

STATISTICAL PHYSICS OF HARD OPTIMIZATION PROBLEMS¹Lenka Zdeborová²*Theoretical Division and Center for Nonlinear Studies,
Los Alamos National Laboratory, NM 87545 USA*

Received 26 May 2009, accepted 3 June 2009

Optimization is fundamental in many areas of science, from computer science and information theory to engineering and statistical physics, as well as to biology or social sciences. It typically involves a large number of variables and a cost function depending on these variables. Optimization problems in the non-deterministic polynomial (NP)-complete class are particularly difficult, it is believed that the number of operations required to minimize the cost function is in the most difficult cases exponential in the system size. However, even in an NP-complete problem the practically arising instances might, in fact, be easy to solve. The principal question we address in this article is: How to recognize if an NP-complete constraint satisfaction problem is typically hard and what are the main reasons for this? We adopt approaches from the statistical physics of disordered systems, in particular the cavity method developed originally to describe glassy systems. We describe new properties of the space of solutions in two of the most studied constraint satisfaction problems - random satisfiability and random graph coloring. We suggest a relation between the existence of the so-called frozen variables and the algorithmic hardness of a problem. Based on these insights, we introduce a new class of problems which we named "locked" constraint satisfaction, where the statistical description is easily solvable, but from the algorithmic point of view they are even more challenging than the canonical satisfiability.

PACS: 89.70.Eg, 75.10.Nr, 64.70.qd, 02.10.Ox, 89.20.Ff

KEYWORDS: Constraint satisfaction problems, Random graphs coloring problem, Average computational complexity, Cavity method, Spin glasses, Replica symmetry breaking, Clustering of solutions, Belief propagation, Satisfiability threshold, Reconstruction on trees

¹Material in this article was presented as a PhD thesis of the author at the University Paris-Sud XI, and at the Charles University in Prague. The work was conducted under supervision of professor Marc Mézard in LPTMS, at University Paris-Sud XI.

²E-mail address: lenka.zdeborova@gmail.com

Contents

1	Hard optimization problems	173
1.1	Importance of optimization problems	173
1.2	Constraint Satisfaction Problems: Setting	174
1.2.1	Definition, factor graph representation	174
1.2.2	List of CSPs discussed in this article	174
1.2.3	Random factor graphs: definition and properties	176
1.3	Computational complexity	178
1.3.1	The worst case complexity	178
1.3.2	The average case hardness	179
1.4	Statistical physics comes to the scene	181
1.4.1	Glance on spin glasses	181
1.4.2	First encounter	181
1.5	The replica symmetric solution	183
1.5.1	Statistical physics description	183
1.5.2	The replica symmetric solution on a single graph	184
1.5.3	Average over the graph ensemble	186
1.5.4	Application for counting matchings	186
1.6	Clustering and Survey propagation	188
1.7	Energetic 1RSB solution	190
1.7.1	Warning Propagation	190
1.7.2	Survey Propagation	192
1.7.3	Application to the exact cover (positive 1-in-3 SAT)	194
1.8	Loose ends	195
1.9	Summary of my contributions to the field	196
2	Clustering	200
2.1	Definition of clustering and the 1RSB approach	200
2.1.1	Properties and equations on trees	202
2.1.2	Back to the sparse random graphs	208
2.1.3	Compendium of the 1RSB cavity equations	210
2.2	Geometrical definitions of clusters	212
2.3	Physical properties of the clustered phase	214
2.4	Is the clustered phase algorithmically hard?	214
3	Condensation	217
3.1	Condensation in a toy model of random subcubes	217
3.2	New in CSPs, well known in spin glasses	218
3.3	Relative sizes of clusters in the condensed phase	220
3.4	Condensed phase in random CSPs	222
3.5	Is the condensed phase algorithmically hard?	223

4	Freezing	225
4.1	Frozen variables	225
4.1.1	Whitening: A way to tell if solutions are frozen	225
4.1.2	Freezing on finite size instances	226
4.1.3	Freezing transition in 3-SAT - exhaustive enumeration	227
4.2	Cavity approach to frozen variables	229
4.2.1	Frozen variables in the entropic 1RSB equations	229
4.2.2	The phase transitions: Rigidity and Freezing	232
4.3	Point like clusters: The locked problems	234
4.3.1	Definition	234
4.3.2	The replica symmetric solution	235
4.3.3	Small noise reconstruction	237
4.3.4	Clustering transition in the locked problems	240
4.4	Freezing - The reason for hardness?	241
4.4.1	Always a trivial whitening core	241
4.4.2	Incremental algorithms	242
4.4.3	Freezing transition and the performance of SP in 3-SAT	243
4.4.4	Locked problems – New extremely challenging CSPs	244
5	Coloring random graphs	246
5.1	Setting	246
5.2	Phase diagram	247
5.3	Large q limit	251
5.3.1	The $2q \log q$ regime: colorability and condensation	252
5.3.2	The $q \log q$ regime: clustering and rigidity	252
5.4	Finite temperature	253
6	Conclusions and perspectives	256
6.1	Key results	256
6.2	Some open problems	256
6.3	Perspectives	257
	Appendices	259
A	1RSB cavity equations at $m = 1$	259
B	Exact entropy for the balanced LOPs	263
B.1	The 1 st moment for occupation models	263
B.2	The 2 nd moment for occupation models	265
B.3	The results	266
C	Stability of the RS solution	268
C.1	Several equivalent methods for RS stability	268
C.2	Stability of the warning propagation	271

D	1RSB stability	272
D.1	Stability of the energetic 1RSB solution	273
D.2	1RSB stability at general m and T	275
E	Populations dynamics	277
E.1	Population dynamics for belief propagation	277
E.2	Population dynamics to solve 1RSB at $m = 1$	278
E.3	Population dynamics with reweighting	279
E.4	Population dynamics with hard and soft fields	281
E.5	The population of populations	282
E.6	How many populations needed?	282
F	Algorithms	284
F.1	Decimation based solvers	284
F.1.1	Unit Clause propagation	284
F.1.2	Belief propagation based decimation	285
F.1.3	Maximal BP decimation on the random coloring	286
F.1.4	Analysis of the uniform exact decimation	286
F.1.5	The Failure of Decimation in the Locked problems	287
F.1.6	Survey propagation based decimation	289
F.2	Search of improvement based solvers	290
F.2.1	Simulated annealing	290
F.2.2	Stochastic local search	290
F.2.3	Belief propagation reinforcement	292
	References	295

1 Hard optimization problems

In this opening chapter we introduce the constraint satisfaction problems and discuss briefly the computer science approach to the computational complexity. We review the studies of the random satisfiability problem in the context of average computational complexity investigations. We describe the connection between spin glasses and random constraint satisfaction problems and highlight the most interesting results coming out from this analogy. We explain the replica symmetric approach to these problems and show its usefulness on the example of counting of matchings [ZM06]. Then we review the survey propagation approach to constraint satisfaction on an example of 1-in-K satisfiability [RSZ07]. Finally we summarize the main contributions of the author to the advances in the statistical physics of hard optimization problems, that are elaborated in the rest of the article.

1.1 Importance of optimization problems

Optimization is a common concept in many areas of human activities. It typically involves a large number of variables, e.g. particles, agents, cells or nodes, and a cost function depending on these variables, such as energy, measure of risk or expenses. The problem consists in finding a state of variables which minimizes the value of the cost function.

In this article we will concentrate on a subset of optimization problems the so-called *constraint satisfaction problems* (CSPs). Constraint satisfaction problems are one of the main building blocks of complex systems studied in computer science, information theory and statistical physics. Their wide range of applicability arises from their very general nature: given a set of N discrete variables subject to M constraints, the CSP consists in deciding whether there exists an assignment of variables which satisfies simultaneously all the constraints. And if such an assignment exists then we aim at finding it.

In computer science, CSPs are at the core of computational complexity studies: the satisfiability of boolean formulas is the canonical example of an intrinsically hard, NP-complete, problem. In information theory, error correcting codes also rely on CSPs. The transmitted information is encoded into a codeword satisfying a set of constraints, so that the information may be retrieved after transmission through a noisy channel, using the knowledge of the constraints satisfied by the codeword. Many other practical problems in scheduling a collection of tasks, in electronic design engineering or artificial intelligence are viewed as CSPs. In statistical physics the interest in CSPs stems from their close relation with the theory of spin glasses. Answering if frustration is avoidable in a system is the first, and sometimes highly nontrivial, step in understanding the low temperature behaviour.

A key point is to understand how difficult it is to solve practical instances of a constraint satisfaction problem. Everyday experience confirms that sometimes it is very hard to find a solution. Many CSPs require a combination of heuristics and combinatorial search methods to be solved in a reasonable time. A key question we address in this article is thus *why* and *when* are some instances of these problems intrinsically hard. Answering this question has, next to its theoretical interest, several practical motivations

- Understanding where the hardness comes from helps to push the performance of CSPs solvers to its limit.

- Understanding which instances are hard helps to avoid them if the nature of the given practical problem permits.
- Finding the very hard problem might be interesting for cryptographic application.

A pivotal step in this direction is the understanding of the onset of hardness in random constraint satisfaction problems. In practice random constraint satisfaction problems are either regarded as extremely hard as there is no obvious structure to be explored or as extremely simple as they permit probabilistic description. Furthermore, random constraint satisfaction models are spin glasses and we shall thus borrow methods from the statistical physics of disordered systems.

1.2 Constraint Satisfaction Problems: Setting

1.2.1 Definition, factor graph representation

Constraint Satisfaction Problem (CSP): Consider N variables $s_1 \dots, s_N$ taking values from the domain $\{0, \dots, q-1\}$, and a set of M constraints. A constraint a concerns a set of k_a different variables which we call ∂a . Constraint a is a function from all possible assignments of the variables ∂a to $\{0, 1\}$. If the constraint evaluates to 1 we say it is satisfied, and if it evaluates to 0 we say it is violated. The constraint satisfaction problem consists in deciding whether there exists an assignment of variables which satisfies simultaneously all the constraints. We call such an assignment a solution of the CSP.

In physics, the variables represent q -state Potts spins (or Ising spins if $q = 2$). The constraints represent very general (non-symmetric) interactions between k_a -tuples of spins. In Boolean constraint satisfaction problems ($q = 2$) a *literal* is a variable or its negation. A *clause* is then a disjunction (logical OR) of literals.

A handy representation for a CSP is the so-called *factor graph*, see [KFL01] for a review. Factor graph is a bipartite graph $G(V, F, E)$ where V is the set of variables (variables nodes, represented by circles) and F is the set of constraints (function nodes, represented by squares). An edge $(ia) \in E$ is present if the constraint $a \in F$ involves the variable $i \in V$. A constraint a is connected to k_a variables, their set is denoted ∂a . A variable i is connected to l_i constraints, their set is denoted ∂i . For clarity we specify the factor graph representation for the graph coloring and exact cover problem in fig. 1.1, both defined in the following section 1.2.2.

1.2.2 List of CSPs discussed in this article

Here we define constraint satisfaction problems which will be discussed in the following. Most of them are discussed in the classical reference book [GJ79]. The most studied constraint satisfaction problems are defined over Boolean variables, $q = 2$, $s_i \in \{0, 1\}$. Sometimes we use equivalently the notation with Ising spins $s_i \in \{-1, +1\}$. CSPs with Boolean variables that we shall discuss in this article are:

- **Satisfiability (SAT) problem:** Constraints are clauses, that is logical disjunctions of literals (i.e., variables or their negations). Example of a satisfiable formula with 3 variables and 4 clauses (constraints) and 10 literals: $(x_1 \vee x_2 \vee \neg x_3) \wedge (x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_3) \wedge (x_1 \vee \neg x_2 \vee x_3)$.




Fig. 1.1. (Color online) Example of a factor graph representation for the coloring (left) and the exact cover (right) problems. The function nodes (squares) in the graph coloring are satisfied if and only if their two neighbours (circles) are in different states (take different colors). The function nodes (squares) in the exact cover problem are satisfied if exactly one variable (circle) around them takes values 1 (full) and the others 0 (empty).

- **K-SAT**: Satisfiability problem where every clause involves K literals, $k_a = K$ for all $a = 1, \dots, M$.
- **Not-All-Equal SAT**: Constraints are satisfied everytime except when all the literals they involve are TRUE or all of them are FALSE.
- **Bicoloring**: Constraints are satisfied except when all variables they involve are equal. Bicoloring is Not-All-Equal SAT without negations.
- **XOR-SAT**: Constraints are logical XORs of literals.
- **Odd (resp. Even) Parity Checks**: A constraint is satisfied if the sum of variables it involves is odd (resp. even). Odd parity checks are XORs without negations.
- **1-in- K SAT**: Constraints are satisfied if exactly one of the K literals they involve is TRUE.
- **Exact Cover, or positive 1-in- K SAT**: Constraints are satisfied if exactly one of the K variables they involve is 1 (occupied). Exact cover, or positive 1-in- K SAT, is 1-in- K SAT without negations.
- **Perfect matching**: Nodes of the original graph become constraints, variables are on edges and determine if the edge is or is not in the matching, see fig. 1.5. Constraints are satisfied if exactly one of the K variables they involve is 1 (belongs to the matching). Note that perfect matching is just a variant of the Exact Cover
- **Occupation problems** are defined by a binary $(K + 1)$ component vector A . All constraints involve K variables, and are satisfied if the sum of variables they involve $r = \sum_{\partial a} s_i$ is such that $A_r = 1$.

- **Locked Occupation Problems (LOPs):** If the vector A is such that $A_i A_{i+1} = 0$ for all $i = 0, \dots, K-1$, and all the variables are present in at least two constraints.

We will also consider in a great detail one CSP with q -ary variables: The **graph coloring** with q colors: Every constraint involves two variables and is satisfied if the two variables are not assigned the same value (color). In physics the q -ary variables are called Potts spins.

1.2.3 Random factor graphs: definition and properties

Given a constraint satisfaction problem with N variables and M constraints, the *constraint density* is defined as $\alpha = M/N$. Denote by $\mathcal{R}(k)$ the probability distribution of the degree of constraints (number of neighbours in the factor graph), and by $\mathcal{Q}(l)$ the probability distribution of the degree of variables. The average connectivity (degree) of constraints is

$$K = \bar{k} = \sum_{k=0}^{\infty} k \mathcal{R}(k). \quad (1.1)$$

The average connectivity of variables is

$$c = \bar{l} = \sum_{l=0}^{\infty} l \mathcal{Q}(l). \quad (1.2)$$

The constraint density is then asymptotically

$$\alpha = \frac{M}{N} = \frac{\bar{l}}{\bar{k}} = \frac{c}{K}. \quad (1.3)$$

A random factor graph with a given N and M is then created as follows: Draw a sequence $\{l_1, \dots, l_N\}$ of N numbers from the distribution $\mathcal{Q}(l)$. Subsequently, draw a sequence $\{k_1, \dots, k_M\}$ of M numbers from the distribution $\mathcal{R}(k)$, such that $\sum_{a=1}^M k_a = \sum_{i=1}^N l_i$. The *random factor graph* is drawn uniformly at random from all the factor graphs with N variables, M constraints and degree sequences $\{l_1, \dots, l_N\}$ and $\{k_1, \dots, k_M\}$.

Another definition leading to a Poissonian degree distribution is used often if the degree of constraints is fixed to K and the number of variables is fixed to N . There are $\binom{N}{K}$ possible positions for a constraint. Each of these positions is taken with probability

$$p = \frac{cN}{K \binom{N}{K}}. \quad (1.4)$$

The number of constraints is then a Poissonian random variable with average $M = cN/K$. The degree of variables is distributed according to a Poissonian law with average c

$$\mathcal{Q}(l) = e^{-c} \frac{c^l}{l!}. \quad (1.5)$$

If $K = 2$ these are the random Erdős-Rényi graphs [ER59]. This definition works also if constraints are changed for variables, that is if the degree of variables and the number of constraints are fixed, as in e.g. the matching problem.

The random factor graphs are called *regular* if both the degrees of constraints and variables are fixed, $\mathcal{R}(k) = \delta(k - K)$ and $\mathcal{Q}(l) = \delta(l - L)$. In section 4.3 we will also use the *truncated Poissonian* degree distribution

$$l \leq 1 : \quad \mathcal{Q}(l) = 0, \quad (1.6a)$$

$$l \geq 2 : \quad \mathcal{Q}(l) = \frac{1}{e^c - (c+1)} \frac{c^l}{l!}. \quad (1.6b)$$

The average connectivity for the truncated Poissonian distribution is then

$$\bar{l} = c \frac{e^c - 1}{e^c - (c+1)}. \quad (1.7)$$

In the cavity approach, the so-called *excess degree distribution* is a crucial quantity. It is defined as follows: Choose an edge (ij) at random and consider the probability distribution of the number of neighbours of i except j . The variables (analogously for constraints) excess degree distribution thus reads

$$q(l) = \frac{(l+1)\mathcal{Q}(l+1)}{\bar{l}}, \quad r(k) = \frac{(k+1)\mathcal{R}(k+1)}{\bar{k}}. \quad (1.8)$$

We will always deal with factor graphs where K and c are of order one, and $N \rightarrow \infty, M \rightarrow \infty$. These are called *sparse random factor graphs*. Concerning the physical properties of sparse random factor graphs the two definitions of a random graph with Poissonian degree distribution are equivalent. Some properties (e.g. the annealed averages) can however depend on the details of the definition.

The tree-like property of sparse random factor graphs — Consider a random variable i in the factor graph. We want to estimate the average length of the shortest cycle going through variable i . Consider a diffusion algorithm spreading into all direction but the one it came from. The probability that this diffusion will arrive back to i in d steps reads

$$1 - \left(1 - \frac{1}{N}\right)^{\sum_{j=1}^d (\gamma_l \gamma_k)^j}, \quad (1.9)$$

where $\gamma_l = \bar{l}^2/\bar{l} - 1$ and $\gamma_k = \bar{k}^2/\bar{k} - 1$ are the mean values of the excess degree distribution (1.8). The probability (1.9) is almost surely zero if

$$d \ll \frac{\log N}{\log \gamma_l \gamma_k}. \quad (1.10)$$

An important property follows: As long as the degree distributions $\mathcal{R}(k)$ and $\mathcal{Q}(l)$ have a finite variance the sparse random factor graphs are locally trees up to a distance scaling as $\log N$ (1.10). We define this as the *tree-like* property.

In this article we consider only degree distributions with a finite variance. A generalization to other cases (e.g. the scale-free networks with long-tail degree distributions) is not straightforward and many of the results which are asymptotically exact on the tree-like structures would be in general only approximative. We observed, see e.g. fig. 2.2, that many of the nontrivial properties predicted asymptotically on the tree-like graphs seems to be reasonably precise even on graphs with about $N = 10^2 - 10^4$ variables. It means that the asymptotic behaviour sets in rather early and does not, in fact, require $\log N \gg 1$.

1.3 Computational complexity

1.3.1 The worst case complexity

Theoretical computer scientists developed the computational complexity theory in order to quantify how hard problems can be in the worst possible case. The most important and discussed complexity classes are the P, NP and NP-complete.

A problem is in the *P (polynomial) class* if there is an algorithm which is able to solve the problem for any input instance of length N in at most cN^k steps, where k and c are constants independent of the input instance. The formal definitions of what is a "problem", its "input instance" and an "algorithm" was formalized in the theory of Turing machines [Pap94], where the definition would be: The complexity class P is the set of decision problems that can be solved by a deterministic Turing machine in polynomial time. A simple example of polynomial problem is sorting a list of N real numbers.

A problem is in the *NP class* if its instance can be stored in memory of polynomial size and if the correctness of a proposed result can be checked in polynomial time. Formally, the complexity class NP is the set of decision problems that can be solved by a non-deterministic Turing machine in polynomial time [Pap94], NP stands for non-deterministic polynomial. Whereas the deterministic Turing machine is basically any of our today computers, the non-deterministic Turing machine can perform unlimited number of parallel computations. Thus, if for finite N there is a finite number of possible solutions all of them can be checked simultaneously. This class contains many problems that we would like to be able to solve efficiently, including the Boolean satisfiability problem, the traveling salesman problem or the graph coloring. Problems which do not belong to the NP class are for example counting the number of solutions in Boolean satisfiability, or the random energy model [Der80, Der81].

All the polynomial problems are in the NP class. It is not known if all the NP problems are polynomial, and it is considered by many to be the most challenging problem in theoretical computer science. It is also one of the seven, and one of the six still open, Millennium Prize Problems that were stated by the Clay Mathematics Institute in 2000 (a correct solution to each of these problems results in a \$1,000,000 prize for the author). A majority of computer scientists, however, believes that the negative answer is the correct one [Gas02].

The concept of NP-complete problems was introduced by Cook in 1971 [Coo71]. All the NP problems can be polynomially reduced to any NP-complete problem, thus if any NP-complete problem would be polynomial then $P=NP$. Cook proved [Coo71] that the Boolean satisfiability problem is NP-complete. Karp soon after added 21 new NP-complete problems to the list [Kar72]. Since then thousands of other problems have been shown to be NP-complete by reductions from other problems previously shown to be NP-complete; many of these are collected in the Garey and Johnson's "Guide to NP-Completeness" [GJ79].

Schaefer in 1978 proved a dichotomy theorem for Boolean ($q = 2$) constraint satisfaction problems. He showed that if the constraint satisfaction problem has one of the following four properties then it is polynomial, otherwise it is NP-complete. (1) All constraints are such that $s_i = 1$ for all i is a solution or $s_i = 0$ for all i is a solution. (2) All constraints concern at most two variables (e.g. in 2-SAT). (3) All constraints are linear equations modulo two (e.g. in XOR-SAT). (4) All constraints are the so-called Horn clauses or all of them are the so-called dual Horn clauses. A Horn clause is a disjunction of variables such that at most one variable is

not negated. A dual Horn clause is when at most one variable is negated. A similar dichotomy theorem exists for 3-state variables, $q = 3$, [Bul02]. Generalization for $q > 3$ is not known.

1.3.2 The average case hardness

Given the present knowledge, it is often said that all the polynomial problems are easy and all the NP-complete problems are very hard. But, independently if $P=NP$ or not, even polynomial problems might be practically very difficult, and some (or even most) instances of the NP-complete problems might be practically very easy.

An example of a still difficult polynomial problem is the primality testing, a first polynomial algorithm was discovered by [AKS04]. But a "proof" of remaining difficulty is the EFF prize [EFF] of \$100,000 to the first individual or group who discovers the first prime number with at least 10,000,000 decimal digits.

And how hard are the NP-complete problems? One way to answer is that under restrictions on the structure an NP-complete problem might become polynomial. Maybe the most famous example is 4-coloring of maps (planar factor graphs) which is polynomial. Moreover, it was a long standing conjecture that every map is colorable with 4 colors, proven by Appel and Haken [AH77b, AH77a]. Interestingly enough 3-coloring of maps is NP-complete [GJ79].

But there are also settings under which the problem stays NP-complete and yet almost every instance can be solved in polynomial time. A historically important example is the Boolean satisfiability where each clause is generated by selecting literals with some fixed probability. Goldberg introduced this random ensemble and showed that the average running time of the Davis-Putnam algorithm [DP60, DLL62] is polynomial for almost all choices of parameter settings [Gol79, GPB82]. Thus in the eighties some computer scientist tended to think that all the NP-complete problems are in fact on average easy and it is hard to find the evil instances which makes them NP-complete.

The breakthrough came at the beginning of the nineties when Cheeseman, Kanefsky and Taylor asked "Where the *really* hard problems are?" in their paper of the same name [CKT91]. Shortly after Mitchell, Selman and Levesque came up with a similar work [MSL92]. Both groups simply took a different random ensemble of the satisfiability (in the second case) and coloring (in the first case) instances: the length of clauses is fixed to be K and they are drawn randomly as described in sec. 1.2.3. They observed that when the density of clauses $\alpha = M/N$ is small the existence of a solution is very likely and if α is large the existence of a solution is very unlikely. And the *really* hard instances were located nearby the critical value originally estimated to be $\alpha_s \approx 4.25$ in the 3-SAT [MSL92]. The hardness was judged from the median running time of the Davis-Putnam-Logemann-Loveland (DPLL) backtracking-based algorithm [DP60, DLL62], see fig. 1.2. This whipped away the thoughts that NP-complete problems might in fact be easy on average. Many other studies and observations followed. The hard instances of random K -satisfiability became very fast important benchmarks for the best algorithms. Moreover, there are some indications that critically constrained instances might appear in real-world applications. One may imagine that in a real world situation the amount of constraints is given by the nature of the problem, and variables usually correspond to something costly, thus the competitive designs contain the smallest possible number of variables.

Given a random K -SAT formula of N variables the probability that it is satisfiable, plotted in fig. 1.3 for 3-SAT, becomes more and more like a step-function as the size N grows. An analogy




Fig. 1.2. (Color online) The easy-hard-easy pattern in the random 3-SAT formulas as the constraint density is changed. Full lines are probabilities that a formula is satisfiable. Dashed lines is the medium running time of the DPLL algorithm. This figure is courtesy of Riccardo Zecchina.

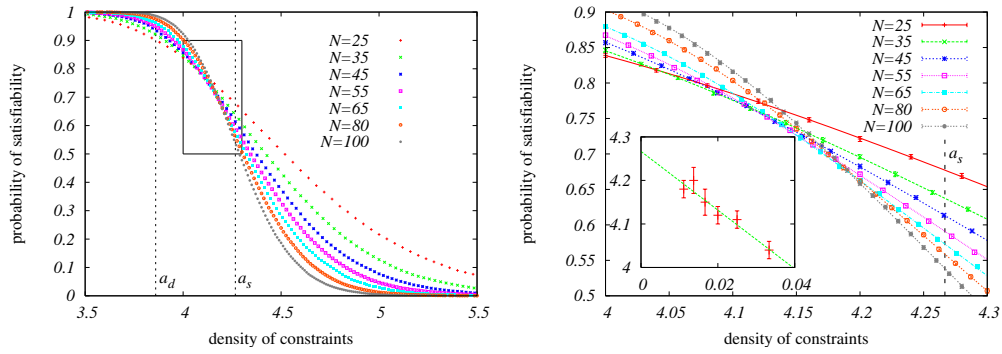


Fig. 1.3. (Color online) Probability that a random 3-SAT formula is satisfiable as a function of the constraint density. In the inset on the left figure is the position of the crossing point between curves corresponding to different sizes as a function of $1/N$. It seems to extrapolate to the analytical value $\alpha_s = 4.267$ [MZ02, MMZ06]. This figure should be put in contrast with fig. 4.1 where the same plot is presented for the freezing transition with a much smaller size of the inset.

with phase transitions in physics cannot be overlooked. The existence and sharpness of the threshold were partially proved [Fri99]. The best known probabilistic bounds of the threshold value in 3-SAT are 3.520 for the lower bound [KKL03, HS03] and 4.506 for the upper bound [DBM00]. Numerical estimates of the asymptotic value of the threshold are $\alpha_s \approx 4.17$ [KS94], $\alpha_s \approx 4.258$ [CA96], $\alpha_s \approx 4.27$ [MZK⁺99b, MZK⁺99a]. The finite size scaling of the curves

in fig. 1.3 is quite involved as the crossing point is moving. That is why the early numerical estimates of the threshold were very inaccurate. The work of Wilson [Wil02], moreover, showed that the experimental sizes are too small and the asymptotic regime for the critical exponent is not reached in any of the current empirical works. The study of XOR-SAT indeed shows a crossover in the critical exponent at sizes which are not accessible for K -SAT [LRTZ01].

The studies of random K -SAT opened up the exciting possibility to connect the hardness with an algorithm-independent property, like the satisfiability phase transition. But what exactly makes the instances near to the threshold hard remained an open question.

1.4 Statistical physics comes to the scene

1.4.1 Glance on spin glasses

Spin glass is one of the most interesting puzzles in statistical physics. An example of a spin glass material is a piece of gold with a small fraction of iron impurities. Physicist, on contrary to the rest of the human population, are interested in the behaviour of these iron impurities and not in the piece of gold itself. A new type of a phase transition was observed from the high temperature paramagnetic phase to the low temperature spin glass phase, where the magnetization of each impurity is *frozen* to a non-zero value, but there is no long range ordering. More than 30 years ago Edwards and Anderson [EA75] introduced a lattice model for such magnetic disordered alloys

$$\mathcal{H} = - \sum_{(ij)} J_{ij} S_i S_j - h \sum_i S_i, \quad (1.11)$$

where $S_i \in \{-1, +1\}$ are Ising spins on a 3-dimensional lattice, the sum runs over all the nearest neighbours, h is the external magnetic field and the interaction J_{ij} is random (usually Gaussian or randomly $\pm J$). The solution of the Edwards-Anderson model stays a largely open problem even today.

The mean field version of the Edwards-Anderson model was introduced by Sherrington and Kirkpatrick [SK75], the sum in the Hamiltonian (1.11) then runs over all pairs (ij) as if the underlying lattice would be fully connected. Sherrington and Kirkpatrick called their paper "Solvable Model of a Spin-Glass". They were indeed right, but the correct solution came only five years later by Parisi [Par80c, Par80b, Par80a]. Parisi's *replica symmetry breaking* (RSB) solution of the Sherrington-Kirkpatrick model gave rise to a whole new theory of the spin glass phase and of the ideal glass transition in structural glasses. The exactness of the Parisi's solution was, however, in doubt till 2000 when Talagrand provided its rigorous proof [Tal06]. The relevance of the RSB picture for the original Edwards-Anderson model is widely discussed but still unknown.

A different mean field version of the Edwards-Anderson model was introduced by Viana and Bray [VB85], the lattice underlying the Hamiltonian (1.11) is then a random graph of fixed average connectivity. The complete solution of the Viana-Bray model is also still an open problem.

1.4.2 First encounter

The Viana-Bray model of spin glasses can also be viewed as random graph bi-partitioning (or bi-coloring at a finite temperature). The peculiarity of the spin glass phase will surely have

some interesting consequences for the optimization problem itself. Indeed, the close connection between optimization problems and spin glass systems brought forward a whole collection of theoretical tools to analyze the structural properties of the optimization problems.

All started in 1985 when Mézard and Parisi realized that the replica theory can be used to solve the bipartite weighted matching problem [MP85]. Let us quote from the introduction of this work: *"This being a kind of pioneering paper, we have decided to present the method {meaning the replica method} on a rather simple problem (a polynomial one) the weighted matching. In this problem one is given $2N$ points $i = 1, \dots, 2N$, with a matrix of distance l_{ij} , and one looks for a matching between the points (a set of N links between two points such that at each point one and only one link arrives) of a minimal length."* Using the replica symmetric (RS) approach they computed the average minimal length, when the elements of the matrix l_{ij} are random identically distributed independent variables.

Shortly after Fu and Anderson [FA86] used the replica method to treat the graph bi-partitioning problem. They were the first to suggest that, possibly, the existence of a phase transition in the average behaviour will affect the actual implementation and performance of local optimization techniques, and that this may also play an important role in the complexity theory. Only later, such a behaviour was indeed discovered empirically by computer scientists [CKT91, MSL92].

The replica method also served to compute the average minimal cost in the random traveling salesman problem [MP86a, MP86b]. Partitioning a dense random graph into more than two groups and the coloring problem of dense random graphs were discussed in [KS87]. Later some of the early results were confirmed rigorously, mainly those concerning the matching problem [Ald01, LW04]. All these early solved models are formulated on dense or even fully connected graph. Thus the replica method and where needed the replica symmetry breaking could be used in its original form. Another example of a "fully connected" optimization problem which was solved with a statistical physics approach is the number partitioning problem [Mer98, Mer00].

And what about our customary random K -satisfiability, which is defined on a sparse graph? Monasson and Zecchina worked out the replica symmetric solution in [MZ96, MZ97]. It was immediately obvious that this solution is not exact as it largely overestimates the satisfiability threshold, the replica symmetry has to be broken in random K -SAT.

An interesting observation was made in [MZK⁺99b]: They defined the backbone of a formula as the set of variables which take the same value in all the ground-state configurations¹. No extensive backbone can exist in the satisfiable phase in the limit of large N . If it would, then adding an infinitesimal fraction of constraints would almost surely cause a contradiction. At the satisfiability threshold an extensive backbone may appear. The authors of [MZK⁺99b] suggested that the problem is computationally hard if the backbone appears discontinuously and easy if it appears continuously. They supported this by replica symmetric solution of the SAT problem with mixed 2-clauses and 3-clauses, the so-called $2+p$ -SAT. Even if the replica symmetric solution is not correct in random K -SAT and even if it overlooks many other important phenomena the concept of backbone is fruitful and we will discuss its generalization in chapter 4.

How to deal with the replica symmetry breaking on a sparse tree-like graph was an open question since 1985, when Viana and Bray [VB85] introduced their model. The solution came only in 2000 when Mézard and Parisi published their paper "Bethe lattice spin glass revisited"

¹In CSPs with a discrete symmetry, e.g. graph coloring, this symmetry has to be taken into account in the definition of the backbone.

[MP01]. They showed how to treat correctly and without approximations the first step of replica symmetry breaking (1RSB) and described how, in the same way, one can in principal deal with more steps of replica symmetry breaking, this extension is however numerically very difficult. But before explaining the 1RSB method we describe the general replica symmetric solutions. And illustrate its usefulness on the problem of counting matchings in graphs [ZM06]. Only then we describe the main results of the 1RSB solution and illustrate the method in the 1-in- K SAT problem [RSZ07]. After we list several "loose ends" which appeared in this approach. Finally we summarize the main contribution of this article. This will be the departure point for the following part of this article which contains most of the original results.

1.5 The replica symmetric solution

The replica symmetric (RS) solution on a locally tree-like graph consists of two steps:

- (1) Compute the partition sum and all the other quantities of interest as if the graph would be a tree.
- (2) The replica symmetric assumption: Assume that the correlations induced by long loops decay fast enough, such that this tree solution is also correct on the only locally tree-like graph.

Equivalent names used in literature for the replica symmetric solution are Bethe-Peierls approximation (in particular in the earlier physics references) or belief propagation (in computer science or when using the iterative equation as an algorithm to estimate the marginal probabilities - magnetizations in physics). Both these conveniently abbreviate to BP.

1.5.1 Statistical physics description

Let $\phi_a(\partial a)$ be the evaluating function for the constraint a depending on the variables neighbouring with a in the factor graph $G(V, F, E)$. A satisfied constraint has $\phi_a(\partial a) = 1$ and violated constraint $\phi_a(\partial a) = 0$. The Hamiltonian can then be written as

$$H_G(\{s\}) = \sum_{a=1}^M [1 - \phi_a(\partial a)] . \quad (1.12)$$

The energy cost is thus one for every violated constraint. The corresponding Boltzmann measure on configurations is:

$$\mu_G(\{s\}, \beta) = \frac{1}{Z_G(\beta)} e^{-\beta H_G(\{s\})} , \quad (1.13)$$

where β is the inverse temperature and $Z_G(\beta)$ is the partition function. The marginals (magnetizations) $\chi_{s_i}^i$ are defined as the probabilities that the variable i takes value s_i

$$\chi_{s_i}^i = \frac{1}{Z_G(\beta)} \sum_{\{s_j\}, j=1, \dots, i-1, i+1, \dots, N} e^{-\beta H_G(\{s_j\}, s_i)} . \quad (1.14)$$




Fig. 1.4. Parts of the factor graph used to compute $\psi_{s_i}^{a \rightarrow i}$ and $\chi_{s_j}^{j \rightarrow a}$.

The goal is to compute the internal energy $E_G(\beta)$ and the entropy $S_G(\beta)$. For $\beta \rightarrow \infty$ (zero temperature limit) these two quantities give the ground state properties. We are interested in the "thermodynamic" limit of large graphs ($N \rightarrow \infty$), and we shall compute expectations over ensembles of graphs of the densities of thermodynamical potentials $\epsilon(\beta) = \mathbb{E}[E_G(\beta)]/N$ and $s(\beta) = \mathbb{E}[S_G(\beta)]/N$, as well as the average free energy density

$$f(\beta) = \frac{-1}{\beta N} \mathbb{E}[\log Z_G(\beta)] = \frac{1}{N} \mathbb{E}[F_G(\beta)] = \epsilon(\beta) - \frac{1}{\beta} s(\beta). \quad (1.15)$$

The reason for this interest is that, for reasonable graph ensembles, $F_G(\beta)$ is self-averaging. This means that the distribution of $F_G(\beta)/N$ becomes more and more sharply peaked around $f(\beta)$ when N increases.

1.5.2 The replica symmetric solution on a single graph

First suppose that the underlying factor graph is a tree, part of this tree is depicted in fig. 1.4. We define messages $\psi_{s_i}^{a \rightarrow i}$ as the probability that node i takes value s_i on a modified graph where all constraints around i apart a were deleted, and $\chi_{s_j}^{j \rightarrow a}$ as the probability that variable j takes value s_j on a modified graph obtained by deleting constraint a . On a tree these messages can be computed recursively as

$$\psi_{s_i}^{a \rightarrow i} = \frac{1}{Z^{a \rightarrow i}} \sum_{\{s_j\}, j \in \partial a - i} \phi_a(\{s\}, s_i, \beta) \prod_{j \in \partial a - i} \chi_{s_j}^{j \rightarrow a} \equiv \mathcal{F}_\psi(\{\chi^{j \rightarrow a}\}), \quad (1.16a)$$

$$\chi_{s_j}^{j \rightarrow a} = \frac{1}{Z^{j \rightarrow a}} \prod_{b \in \partial j - a} \psi_{s_j}^{b \rightarrow j} \equiv \mathcal{F}_\chi(\{\psi^{b \rightarrow j}\}), \quad (1.16b)$$

where $Z^{a \rightarrow i}$ and $Z^{j \rightarrow a}$ are normalization constants, the factor $\phi_a(\{s\}, \beta) = 1$ if the constraint a is satisfied by the configuration $\{s\}$ and $\phi_a(\{s\}, \beta) = e^{-\beta}$ if not. We denote by $\psi^{a \rightarrow i}$ the whole vector $(\psi_0^{a \rightarrow i}, \dots, \psi_{q-1}^{a \rightarrow i})$ and analogically $\chi^{j \rightarrow a} = (\chi_0^{j \rightarrow a}, \dots, \chi_{q-1}^{j \rightarrow a})$. This is one form of the *belief propagation* (BP) equations [KFL01, Pea82], sometimes called sum-product equations. The probabilities ψ, χ are interpreted as messages (beliefs) living on the edges of the factor graph, with the consistency rules (1.16a) and (1.16b) on the function and variable nodes. Equations (1.16) are usually solved by iteration, the name *message passing* is used in this context. In the

following it will be simpler not to consider the "two-levels" equations (1.16) but

$$\psi_{s_i}^{a \rightarrow i} = \frac{1}{Z^{j \rightarrow i}} \sum_{\{s_j\}, j \in \partial a - i} \phi_a(\{s_j\}, s_i, \beta) \prod_{j \in \partial a - i} \prod_{b \in \partial j - a} \psi_{s_j}^{b \rightarrow j} \equiv \mathcal{F}(\{\psi^{b \rightarrow j}\}), \quad (1.17)$$

where $Z^{j \rightarrow i} = Z^{a \rightarrow i} \prod_{j \in \partial a - i} Z^{j \rightarrow a}$. Notice that on simple graphs, i.e., when either $l_i = 2$ for all $i = 1, \dots, N$ or $k_a = 2$ for all $a = 1, \dots, M$, the form (1.17) simplifies further. And on constraint satisfaction problems on simple graphs (e.g. the matching or coloring problems) the "two-levels" equations are almost never used.

Assuming that one has found the fixed point of the belief propagation equations (1.16a-1.16b), one can deduce the various marginal probabilities and the free energy, entropy etc. The marginal probability (1.14) of variable i estimated by the BP equations is

$$\chi_{s_i}^i = \frac{1}{Z^i} \prod_{a \in \partial i} \psi_{s_i}^{a \rightarrow i}. \quad (1.18)$$

To compute the free energy we first define the free energy shift $\Delta F^{a+\partial a}$ after addition of a function node a and all the variables i around it, and the free energy shift ΔF^i after addition of a variable i . These are given in general by:

$$e^{-\beta \Delta F^{a+\partial a}} = Z^{a+\partial a} = \sum_{\{s_i\}, i \in \partial a} \phi_a(\{s_i\}, \beta) \prod_{i \in \partial a} \prod_{b \in \partial i - a} \psi_{s_i}^{b \rightarrow i}, \quad (1.19a)$$

$$e^{-\beta \Delta F^i} = Z^i = \sum_{s_i} \prod_{a \in \partial i} \psi_{s_i}^{a \rightarrow i}. \quad (1.19b)$$

The total free energy is then obtained by summing over all constraints and subtracting the terms counted twice [MP01, YFW03]:

$$F_G(\beta) = \sum_a \Delta F_{a+\partial a} - \sum_i (l_i - 1) \Delta F_i. \quad (1.20)$$

This form of the free energy is variational, i.e., the derivatives $\frac{\partial(\beta F_G(\beta))}{\partial \chi^{i \rightarrow a}}$ and $\frac{\partial(\beta F_G(\beta))}{\partial \psi^{a \rightarrow i}}$ vanish if and only if the probabilities $\chi^{i \rightarrow a}$ and $\psi^{a \rightarrow i}$ satisfy (1.16a-1.16b). This allows to compute easily the internal energy as

$$E_G(\beta) = \frac{\partial \beta F_G(\beta)}{\partial \beta} = - \sum_a \frac{\partial \beta Z^{a+\partial a}}{Z^{a+\partial a}}. \quad (1.21)$$

The entropy is then obtained as

$$S_G(\beta) = \beta [E_G(\beta) - F_G(\beta)]. \quad (1.22)$$

All the equations (1.16)-(1.22) are exact if the graph G is a tree. The replica symmetric approach consists in assuming that all correlations decay fast enough that application of eqs. (1.16)-(1.22) on a large tree-like graph G gives asymptotically exact results. These equations can be used either on a given graph G or to compute the average over the graph (and disorder) ensemble.

1.5.3 Average over the graph ensemble

We now study the typical instances in an ensemble of graphs. We denote the average over the ensemble by $\mathbb{E}(\cdot)$. We assume that the random factor-graph ensemble is given by a prescribed degree distribution $\mathcal{Q}(l)$ for variables and $\mathcal{R}(k)$ for constraints. Let us call $\mathcal{P}(\psi)$ and $\mathcal{O}(\chi)$ the distributions of messages ψ and χ over all the edges of a large typical factor graph from the ensemble. They satisfy the following self-consistent equations

$$\mathcal{P}(\psi) = \sum_{l=1}^{\infty} q(l) \int \prod_{i=1}^l [\mathrm{d}\chi^i \mathcal{O}(\chi^i)] \delta[\psi - \mathcal{F}_{\psi}(\{\chi^i\})], \quad (1.23a)$$

$$\mathcal{O}(\chi) = \sum_{k=1}^{\infty} r(k) \int \prod_{i=1}^k [\mathrm{d}\psi^i \mathcal{P}(\psi^i)] \delta[\chi - \mathcal{F}_{\chi}(\{\psi^i\})], \quad (1.23b)$$

where the functions \mathcal{F}_{ψ} and \mathcal{F}_{χ} represent the BP equations (1.16a-1.16b), $q(l)$ and $r(k)$ are the excess degree distributions defined in (1.8). If there is a disorder in the interaction terms, as e.g. the negations in K -SAT, we average over it at the same place as over the fluctuating degree.

Solving equations (1.23a-1.23b) to obtain the distributions \mathcal{P} and \mathcal{O} is not straightforward. In some cases (on regular factor graphs, at zero temperature, etc.) it can be argued that the distributions \mathcal{P} , \mathcal{O} are sums of Dirac delta functions. Then the solution of eqs. (1.23a-1.23b) can be obtained analytically. But in general distributional equations of this type are not solvable analytically. However, a numerical technique called *population dynamics* [MP01] is very efficient for their resolution. In appendix E we give a pseudo-code describing how the population dynamics technique works.

Once the distributions \mathcal{P} and \mathcal{O} are known the average of the free energy density can be computed by averaging (1.20) over \mathcal{P} . This average expression for the free energy is again in its variational form (see [MP01]), i.e., the functional derivative $\frac{\delta f(\beta)}{\delta \mathcal{P}(h)}$ vanishes if and only if \mathcal{P} satisfies (1.32). The average energy and entropy density are thus expressed again via the partial derivatives.

Factorized solution — As we mentioned, on the ensemble of random regular factor graphs (without disorder in the interactions) the solution of equations (1.23) is very simple: $\mathcal{P}(\psi) = \delta(\psi - \psi^{\text{reg}})$, $\mathcal{O}(\chi) = \delta(\chi - \chi^{\text{reg}})$, where ψ^{reg} and χ^{reg} is a self-consistent solution of (1.16). This is because in the thermodynamical limit an infinite neighbourhood of every variable is exactly identical thus also the marginal probabilities have to be identical in every physical solution.

1.5.4 Application for counting matchings

To demonstrate how the replica symmetric method works to compute the entropy, that is the logarithm of the number of solutions, we review the results for matching on sparse random graphs [ZM06]. The reasoning why the replica symmetric solution is exact for the matching problem is done on the level of self-consistency checks in [ZM06]. And [BN06] have worked out a rigorous proof for graphs with bounded degree and a large girth (length of the smallest loop).

Consider a graph $G(V, E)$ with N vertices ($N = |V|$) and a set of edges E . A *matching* (dimerization) of G is a subset of edges $M \subseteq E$ such that each vertex is incident with at most




Fig. 1.5. On the left, example of a graph with six nodes and six edges. On the right, the corresponding factor graph with six function nodes (squares) and six variable nodes (circles).

one edge in M . In other words the edges in the matching M do not touch each other. The *size of the matching*, $|M|$, is the number of edges in M . Our goal is to compute the entropy of matchings of a given size on a typical large Erdős-Rényi random graph.

We describe a matching by the variables $s_i = s_{(ab)} \in \{0, 1\}$ assigned to each edge $i = (ab)$ of G , with $s_i = 1$ if $i \in M$ and $s_i = 0$ otherwise. The constraints that two edges in a matching cannot touch impose that, on each vertex $a \in V$: $\sum_{b, (ab) \in E} s_{(ab)} \leq 1$. To complete our statistical physics description, we define for each given graph G an energy (or cost) function which gives, for each matching $M = \{s\}$, the number of unmatched vertices:

$$H_G(M = \{s\}) = \sum_a E_a(\{s\}) = N - 2|M|, \quad (1.24)$$

where $E_a = 1 - \sum_{b \in \partial a} s_{(ab)}$.

In the factor graph representation we transform the graph G into a factor graph $F(G)$ as follows (see fig. 1.5): To each edge of G corresponds a variable node (circle) in $F(G)$; to each vertex of G corresponds a function node (square) in $F(G)$. We shall index the variable nodes by indices i, j, k, \dots and function nodes by a, b, c, \dots . The variable i takes value $s_i = 1$ if the corresponding edge is in the matching, and $s_i = 0$ if it is not. The weight of a function node a is

$$\phi_a(\{\partial a\}, \beta) = \mathbb{I} \left(\sum_{i \in \partial a} s_i \leq 1 \right) e^{-\beta(1 - \sum_{i \in \partial a} s_i)}, \quad (1.25)$$

where ∂a is the set of all the variable nodes which are neighbours of the function node a , and the total Boltzmann weight of a configuration is $\frac{1}{Z_G(\beta)} \prod_a \phi_a(\{\partial a\}, \beta)$.

The belief propagation equation (1.16) becomes

$$\chi_{s_i}^{i \rightarrow a} = \frac{1}{Z^{b \rightarrow a}} \sum_{\{s_j\}} \mathbb{I} \left(s_i + \sum_{j \in \partial b - i} s_j \leq 1 \right) e^{-\beta(1 - s_i - \sum_{j \in \partial b - i} s_j)} \prod_{j \in \partial b - i} \chi_{s_j}^{j \rightarrow b}, \quad (1.26)$$

where $Z^{b \rightarrow a}$ is a normalization constant. In statistical physics the more common form of the BP equations uses analog of local magnetic fields instead of probabilities. For every edge between a variable i and a function node a , we define a *cavity field* $h^{i \rightarrow a}$ as

$$e^{-\beta h^{i \rightarrow a}} \equiv \frac{\chi_0^{i \rightarrow a}}{\chi_1^{i \rightarrow a}}. \quad (1.27)$$

The recursion relation between cavity fields is then:

$$h^{i \rightarrow a} = -\frac{1}{\beta} \log \left[e^{-\beta} + \sum_{j \in \partial b - i} e^{\beta h^{j \rightarrow b}} \right]. \quad (1.28)$$

The expectation value (with respect to the Boltzmann distribution) of the occupation number s_i of a given edge $i = (ab)$ is equal to

$$\langle s_i \rangle = \frac{1}{1 + e^{-\beta(h^{i \rightarrow a} + h^{i \rightarrow b})}}. \quad (1.29)$$

The free energy shifts needed to compute the total free energy (1.20) are

$$e^{-\beta \Delta F_{a+i \in \partial a}} = e^{-\beta} + \sum_{i \in a} e^{\beta h^{i \rightarrow a}}, \quad (1.30a)$$

$$e^{-\beta \Delta F_i} = 1 + e^{\beta(h^{i \rightarrow a} + h^{i \rightarrow b})}. \quad (1.30b)$$

The energy, related to the size of the matching via (1.24), is then

$$E_G(\beta) = \sum_a \frac{1}{1 + \sum_{i \in \partial a} e^{\beta(1 + h^{i \rightarrow a})}}. \quad (1.31)$$

This is the sum of the probabilities that node a is not matched.

The distributional equation (1.23) becomes

$$\mathcal{O}(h) = \sum_{k=1}^{\infty} r(k) \int \prod_{i=1}^k [\mathrm{d}h^i \mathcal{O}(h^i)] \delta \left[h + \frac{1}{\beta} \log \left(e^{-\beta} + \sum_i e^{\beta h^i} \right) \right]. \quad (1.32)$$

And the average free energy is explicitly

$$\begin{aligned} f(\beta) &= \frac{\mathbb{E}[F_G(\beta)]}{N} = -\frac{1}{\beta} \sum_{k=0}^{\infty} \mathcal{R}(k) \int \prod_{i=1}^k [\mathrm{d}h^i \mathcal{O}(h^i)] \log \left(e^{-\beta} + \sum_i e^{\beta h^i} \right) \\ &+ \frac{c}{2\beta} \int \mathrm{d}h^1 \mathrm{d}h^2 \mathcal{O}(h^1) \mathcal{O}(h^2) \log \left(1 + e^{\beta(h^1 + h^2)} \right). \end{aligned} \quad (1.33)$$

Where $\mathcal{R}(k)$ is the connectivity distribution of the function nodes, that is the connectivity distribution of the original graph, c is the average connectivity. The distributional equations are solved via the population dynamics method, see appendix E. Figure 1.6 then presents the resulting average entropy as a function of size of the matching.

1.6 Clustering and Survey propagation

As we said previously in the random K -SAT the replica symmetric solution is not generically correct. Mézard and Parisi [MP01] understood how to deal properly and without approximations with the replica symmetry breaking on random sparse graphs, that is how to take into account the correlations induced by long loops. More precisely in their approach only the one-step (at




Fig. 1.6. (Color online) Entropy density $s(m)$ as a function of relative size of the matching $m = |M|/N$ for Erdős-Rényi random graphs with mean degrees $c = 1, 2, 3, 6$. The lower curve is the ground state entropy density for all mean degrees. The curves are obtained by solving eqs. (1.32)-(1.33) with a population dynamics, using a population of sizes $N = 2 \cdot 10^4$ to $2 \cdot 10^5$ and the number of iterations $T = 10000$.

most two-step on the regular graphs) replica symmetry breaking solution is numerically feasible. Anyhow, such a progress opened the door to a better understanding of the optimization problems on sparse graphs. The K -satisfiability played again the prominent role.

To compute the ground state energy within the 1RSB approach we can restrict only to energetic considerations as described in [MP03], we call this approach the *energetic zero temperature limit*. Applying this method to K -satisfiability leads to several outstanding results [MPZ02, MZ02], we describe the three most remarkable ones. Soon after, analog results were obtained for many other optimization problems, for example graph coloring [MPWZ02, BMP⁺03, KPW04], vertex cover [Zho03], bicoloring of hyper-graphs [CNRTZ03], XOR-SAT [FLRTZ01, MRTZ03] or lattice glass models [BM02, RBMM04].

Clustering — It was known already in the “pre-1RSB-cavity era” that replica symmetry broken solution is needed to solve random K -SAT. Such a need is interpreted as the existence of many metastable well-separated states, in the case of highly degenerate ground state this leads to a clustering of solutions in the satisfiable phase [BMW00, MPZ02, MZ02]. The energetic 1RSB cavity method deals with clusters containing frozen variables (clusters with backbones), that is variables which have the same value in all the solutions in the cluster. It predicts how many of such clusters exist at a given energy, the logarithm of this number divided by the system size N defines the complexity function $\Sigma(E)$. According to the energetic cavity method for 3-SAT, clusters exist, $\Sigma(0) \neq 0$, for constraint density $\alpha > \alpha_{SP} = 3.92$ [MPZ02, MZ02].

It was conjectured [MPZ02, MZ02] that there is a link between clustering, ergodicity breaking, existence of many metastable states and the difficulty of finding a ground state via local al-

gorithms. The critical value α_{SP} was called the *dynamical* transition and the region of $\alpha > \alpha_{\text{SP}}$ the *hard-SAT* phase.

Clusters were viewed as a kind of pure states, however, in the view of many a good formal definition was missing. It was also often referred to some sort of geometrical separation between different clusters. A particularly popular one is the following: Clusters are connected components in the graph where solutions are the nodes and two solutions are adjacent if they differ in only d variables. Depending on the model and author the value of d is either one or d is a finite number of d is said to be any sub-extensive number. The notion of x -satisfiability, the existence of pairs of solutions at a distance x , leads to a rigorous proof of existence of exponentially many geometrically separated clusters [MMZ05, DMMZ08, ART06].

The satisfiability threshold computed — The energetic 1RSB cavity method allows to compute the ground state energy and thus also the satisfiability threshold α_s . In 3-SAT its value is $\alpha_s = 4.2667$ [MPZ02, MZ02, MMZ06]. This value is computed as a solution of a closed distributional equation. This time there is an excellent agreement with the empirical estimations. Is the one step of replica symmetry breaking sufficient to locate exactly the satisfiability threshold? The stability of the 1RSB solution was investigated in [MPRT04], the 1RSB energetic cavity was shown to describe correctly the ground state energy for $4.15 < \alpha < 4.39$ in 3-SAT. In particular, it yields the conjecture that the location of the satisfiability threshold is actually exact. From a rigorous point of view it was proven that the 1RSB equations give an upper bound on the satisfiability threshold [FL03, FLT03, PT04].

Survey Propagation: a revolutionary algorithm — The most spectacular result was the development of a new message passing algorithm, the survey propagation [MZ02, BMZ05]. Before the replica and cavity analysis were used to compute the quenched averages of thermodynamical quantities. Using always the self-averaging property that the average of certain (not all) quantities is equal to their value on a large given sample. Mézard and Zecchina applied the energetic 1RSB cavity equations, later called survey propagation, on a single large graph. This resulted in an algorithm which is arguably still the best known for large instances of random 3-SAT near to the satisfiability threshold. And even more interesting than its performance is the conceptual advance this brought into applications of statistical physics to optimization problems.

1.7 Energetic 1RSB solution

In this section we derive the energetic zero-temperature limit of the 1RSB method. When applied to the satisfiability problem this leads, between others, to the calculation of the satisfiability threshold and to the survey propagation equations and algorithm. We illustrate this on the 1-in-3 SAT problem. Before doing so we have to introduce the warning propagation equations, on which the derivation of the survey propagation relies.

1.7.1 Warning Propagation

In general warning propagation (min-sum) is a zero temperature, $\beta \rightarrow \infty$, limit of the belief propagation (sum-product) equations (1.16a-1.16b). It can be used to compute the ground state energy (minimal fraction of violated constraints) at the replica symmetric level. A constraint

satisfaction problem at a finite temperature gives rise to $\phi_a(\{\partial a\}, \beta) = 1$ if the constraint a is satisfied by configuration $\{s_{\partial a}\}$, and $\phi_a(\{\partial a\}, \beta) = e^{-2\beta}$ if a is not satisfied by $\{s_{\partial a}\}$ ². In a general Boolean CSP, with N variables $s_i \in \{-1, 1\}$, the warning propagation can then be obtained from (1.16a-1.16b) by introducing warnings u and h as

$$e^{2\beta h^{i \rightarrow a}} \equiv \frac{\chi_1^{i \rightarrow a}}{\chi_{-1}^{i \rightarrow a}}, \quad e^{2\beta u^{a \rightarrow i}} \equiv \frac{\psi_1^{a \rightarrow i}}{\psi_{-1}^{a \rightarrow i}}. \quad (1.34)$$

This leads in the limit of zero temperature, $\beta \rightarrow \infty$, to

$$h^{i \rightarrow a} = \sum_{b \in \partial i - a} u^{b \rightarrow i}, \quad (1.35a)$$

$$u^{a \rightarrow i} = \frac{1}{2} \left[\max_{\{s_j\}} \left(\sum_{j \in \partial a - i} h^{j \rightarrow a} s_j - 2E_a(\{s_j\}, +1) \right) - \max_{\{s_j\}} \left(\sum_{j \in \partial a - i} h^{j \rightarrow a} s_j - 2E_a(\{s_j\}, -1) \right) \right]. \quad (1.35b)$$

where $E_a(\{s_i\}) = 0$ if the configuration $\{s_i\}$ satisfies the constraint a , and $E_a(\{s_i\}) = 1$ if it does not. The warnings u and h have to be integer numbers, as they can be interpreted as changes in the ground state energy of the cavity subgraphs when the value of variable i is changed from $s_i = 0$ to $s_i = 1$. Given $E_a \in \{0, 1\}$ we have that $h \in \mathbb{Z}$ and $u \in \{-1, 0, +1\}$. The correspondence between values of u and ψ are

$$u = 1 \quad \Leftrightarrow \quad \psi_1 = 1, \quad \psi_{-1} = 0, \quad (1.36a)$$

$$u = -1 \quad \Leftrightarrow \quad \psi_1 = 0, \quad \psi_{-1} = 1, \quad (1.36b)$$

$$u = 0 \quad \Leftrightarrow \quad \psi_1 = \epsilon, \quad \psi_{-1} = 1 - \epsilon, \quad 0 < \epsilon < 1. \quad (1.36c)$$

The warnings u and h can thus be interpreted in the following way

$u^{a \rightarrow i} = -1$	Constraint a tells to variable i : “I think you should be -1 .”
$u^{a \rightarrow i} = 0$	Constraint a tells to variable i : “I can deal with any value you take.”
$u^{a \rightarrow i} = +1$	Constraint a tells to variable i : “I think you should be $+1$.”
$h^{i \rightarrow a} < 0$	Variable i tells to constraint a : “I would prefer to be -1 .”
$h^{i \rightarrow a} = 0$	Variable i tells to constraint a : “I don’t have any strong preferences.”
$h^{i \rightarrow a} > 0$	Variable i tells to constraint a : “I would prefer to be $+1$.”

Given this interpretation the prescriptions (1.35) on how to update the warnings over the graph becomes intuitive, see Tab. 1.1. Variable i collects the preferences from all constraints except a and sends the result to a . Constraint a then decides which value i should take given the preferences of all its other neighbours.

²The factor 2 in the Hamiltonian is introduced for convenience and in agreement with the notation of [RSZ07].

Given the fixed point of the warning propagation (1.35) the total warning of variable i is

$$h^i = \sum_{a \in \partial i} u^{a \rightarrow i}. \quad (1.37)$$

The corresponding energy can be computed as

$$E = \sum_a \Delta E^{a+\partial a} - \sum_i (l_i - 1) \Delta E^i, \quad (1.38)$$

where $\Delta E^{a+\partial a}$ is the number of contradictions created when constraint a and all its neighbours are added to the graph, ΔE^i is the number of contradictions created when variables i is added to the graph. The energy shifts can be computed from (1.19a-1.19b) using (1.34) and taking $\beta \rightarrow \infty$ they read

$$\Delta E^{a+\partial a} = -\max_{\{s_{\partial a}\}} \left[\sum_{i \in \partial a} h^{i \rightarrow a} s_i - E_a(\{s_{\partial a}\}) \right] + \sum_{i \in \partial a} \sum_{b \in \partial i - a} |u^{b \rightarrow i}|; \quad (1.39a)$$

$$\Delta E^i = -\left| \sum_{a \in \partial i} u^{a \rightarrow i} \right| + \sum_{a \in \partial i} |u^{a \rightarrow i}|; \quad (1.39b)$$

To summarize, the warning propagation equations neglect every entropic information in the belief propagation (1.16a-1.16b), thus only the ground state energy can be computed. On the other hand the fact that warnings u and h have a discrete set of possible values simplifies considerably the average over the graph ensemble presented in sec. 1.5.3 as the distribution \mathcal{P} is a sum of three Dirac function, and can be represented by their weights. Deeper interpretations of warning propagation and its fixed points will be given in chapter 4. Note that in the literature the value 0 of warnings is also called * or "joker" [BMWZ03, BZ04].

1.7.2 Survey Propagation

Survey propagation (SP) [MPZ02, MZ02] is a form of belief propagation which aims to count the logarithm of the number of fixed points of warning propagation (1.35) of a given energy

Tab. 1.1. Example of the update (1.35b) in the positive 1-in-3 SAT problem, where exactly one variable in the constraint takes value 1 in order to satisfy the constraint. The first line might seem counter-intuitive, but note that we defined the energy in such a way that configuration (1, 1, 1) is as bad as (1, 1, -1).

$h^{1 \rightarrow a}$	$h^{2 \rightarrow a}$	$u^{a \rightarrow 3}$
+	+	0
+	-	-
+	0	-
0	0	0
-	-	+
-	0	0

(1.38). For the sake of simplicity we present the most basic form of SP which aims to count the logarithm of number of fixed points of the warning propagation with zero energy.

The constraints on values of the warnings assuring that the fixed point of warning propagation corresponds to zero energy are

- For all i and $a \in \partial i$: the warnings $\{u^{b \rightarrow i}\}_{b \in \partial i - a}$ are all non-negative or all non-positive,
- For all a and $i \in \partial a$: the preferred values of all $j \in \partial a - i$ can be realized without violating the constraint a .

We define probabilities that warnings $u^{a \rightarrow i}$ or $h^{i \rightarrow a}$ are positive, negative or null.

$$\mathcal{P}^{a \rightarrow i}(u^{a \rightarrow i}) = q_-^{a \rightarrow i} \delta(u^{a \rightarrow i} + 1) + q_+^{a \rightarrow i} \delta(u^{a \rightarrow i} - 1) + q_0^{a \rightarrow i} \delta(u^{a \rightarrow i}); \quad (1.40a)$$

$$\mathcal{P}^{i \rightarrow a}(h^{i \rightarrow a}) = p_-^{i \rightarrow a} \mu_-(h^{i \rightarrow a}) + p_+^{i \rightarrow a} \mu_+(h^{i \rightarrow a}) + p_0^{i \rightarrow a} \delta(h^{i \rightarrow a}); \quad (1.40b)$$

where $q_-^{a \rightarrow i} + q_+^{a \rightarrow i} + q_0^{a \rightarrow i} = p_-^{i \rightarrow a} + p_+^{i \rightarrow a} + p_0^{i \rightarrow a} = 1$, and $\mu_{\pm}(h)$ are normalized measures with support over \mathbb{Z}^{\pm} . So, to every oriented edge we associate a message $q = (q_-, q_0, q_+)$ or $p = (p_-, p_0, p_+)$ (resp. if oriented towards the variable or the constraint). We call these messages surveys, they are analogous to beliefs $\psi^{a \rightarrow i}$ and $\chi^{i \rightarrow a}$ from (1.16a-1.16b). And thus, if the factor graph is tree, exact iterative equations for q, p can be written. The update of surveys p given incoming q is common for all Boolean CSPs and reads:

$$p_+^{i \rightarrow a} + p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \prod_{b \in \partial i - a} (q_+^{a \rightarrow i} + q_0^{a \rightarrow i}), \quad (1.41a)$$

$$p_-^{i \rightarrow a} + p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \prod_{b \in \partial i - a} (q_-^{b \rightarrow i} + q_0^{b \rightarrow i}), \quad (1.41b)$$

$$p_0^{i \rightarrow a} = \mathcal{N}_{i \rightarrow a}^{-1} \prod_{b \in \partial i - a} q_0^{b \rightarrow i}, \quad (1.41c)$$

where $\mathcal{N}_{i \rightarrow a}$ is the normalization factor. The update of surveys q given the incoming ps depends on the details on the constraint functions. For concreteness we write the equation for the positive 1-in-3 SAT problem. The constraints assuring zero energy then forbids that both the warnings incoming to a constraint a have value $+1$.

$$q_+^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} p_-^{j \rightarrow a} p_-^{k \rightarrow a}, \quad (1.42a)$$

$$q_-^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} [p_+^{j \rightarrow a} (1 - p_+^{k \rightarrow a}) + (1 - p_+^{j \rightarrow a}) p_+^{k \rightarrow a}], \quad (1.42b)$$

$$q_0^{a \rightarrow i} = \mathcal{N}_{a \rightarrow i}^{-1} [p_-^{j \rightarrow a} p_0^{k \rightarrow a} + p_0^{j \rightarrow a} p_-^{k \rightarrow a} + p_0^{j \rightarrow a} p_0^{k \rightarrow a}], \quad (1.42c)$$

where $\mathcal{N}_{a \rightarrow i} = 1 - p_+^{j \rightarrow a} p_+^{k \rightarrow a}$ is the normalization factor, j and k are the other two neighbours of a .

The associated Shannon entropy is called *complexity* [Pal83] (or structural entropy in the context of glasses) and reads [MZ02]

$$\Sigma(E = 0) = \sum_a \log \mathcal{N}^{a + \partial a} - \sum_i (l_i - 1) \log \mathcal{N}^i, \quad (1.43)$$

where $\mathcal{N}^{a+\partial a}$ is the probability that no contradiction is created when the constraint a and all its neighbours are added, \mathcal{N}^i is the probability that no contradiction is created when the variable i is added. Remark the exact analogy with (1.19a-1.19b). We denote $\mathcal{P}_0^i \equiv \prod_{a \in \partial i} q_0^{a \rightarrow i}$ and $\mathcal{P}_\pm^i \equiv \prod_{a \in \partial i} (q_\pm^{a \rightarrow i} + q_0^{a \rightarrow i})$, then

$$\mathcal{N}^i = \mathcal{P}_+^i + \mathcal{P}_-^i - \mathcal{P}_0^i, \quad (1.44a)$$

$$\begin{aligned} \mathcal{N}^{a+\partial a} = & \prod_{i \in \partial a} (\mathcal{P}_+^{i \rightarrow a} + \mathcal{P}_-^{i \rightarrow a} - \mathcal{P}_0^{i \rightarrow a}) - \prod_{i \in \partial a} (\mathcal{P}_-^{i \rightarrow a} - \mathcal{P}_0^{i \rightarrow a}) - \prod_{i \in \partial a} (\mathcal{P}_+^{i \rightarrow a} - \mathcal{P}_0^{i \rightarrow a}) \\ & - \sum_{i \in \partial a} \mathcal{P}_-^{i \rightarrow a} \prod_{j \in \partial a - i} (\mathcal{P}_+^{j \rightarrow a} - \mathcal{P}_0^{j \rightarrow a}). \end{aligned} \quad (1.44b)$$

The second equation collects the contributions from all combinations of arriving surveys except the “contradictory” ones $(+, +, +)$, $(-, -, -)$, $(+, +, 0)$ and $(+, +, -)$ (plus permutations of the latter).

The survey propagation equations (1.41-1.42) and the expression for the complexity function (1.43) are exact on tree graphs. In the spirit of the Bethe approximation, we will assume sufficient decay of correlations and use these equations on a random graph³. To average over the ensemble of random graphs we adopt the same equations as we did for the belief propagation in sec. 1.5.3.

1.7.3 Application to the exact cover (positive 1-in-3 SAT)

The 1-in-3 SAT problem (with probability of negating a variable equal to one-half) is a rare example of an NP-complete problem which is on average algorithmically easy and where the threshold can be computed rigorously [ACIM01]. In particular it was shown that for $\alpha \neq 1$ an instance of the problem can be solved in polynomial time with probability going to one as $N \rightarrow \infty$. This result was generalized into random 1-in-3 SAT where the probability of negating a variable is $p \neq 1/2$ [RSZ07]. In particular we showed that for all $0.273 < p < 0.718$ the RS solution is correct and almost every instance can be solved in polynomial time if the constraint density $\alpha \neq 1/[4p(1-p)]$. When, however, $p < 0.273$ the phase diagram is more complicated, see [RSZ07]. For $p = 0$ the solution of the positive 1-in-3 SAT (exact cover) problem becomes very similar to the one of 3-SAT [MZ02]. The result for the complexity (1.43) in the positive 1-in-3 SAT obtained from the population dynamics method is plotted in fig. 1.7. For more detailed discussion of how the phase diagram changes from the almost-always-easy to the very-hard pattern see [RSZ07].

Up to certain average connectivity of variables $c_{SP} = 1.822$ the only iterative fixed point of the population dynamics gives $q_0^{a \rightarrow i} = p_0^{i \rightarrow a} = 1$ for all (ia) . The associated complexity function is zero. In an interval $(c_{SP}, c_s) = (1.822, 1.879)$ there exist a nontrivial solution giving positive complexity function. There are thus exponentially many different fixed points of the warning propagation. Asymptotically, almost every warning propagation fixed point is associated to a cluster of solutions⁴. Above $c_s = 1.879$ there is a nontrivial solution to the SP equations

³The fact that on a given tree with given boundary conditions the warning propagation has a unique fixed point might seem puzzling at this point. Clarification will be made in the chapter 2.

⁴There might exist fixed points of the warning propagation which are not compatible with any solution, thus do not correspond to a cluster. Such “fake” fixed points are negligible if the 1RSB approach is correct.




Fig. 1.7. (Color online) Average complexity density (logarithm of number of states divided by the number of variables) as a function of the mean degree c for the positive 1-in-3 SAT problem. At $c_{SP} = 1.822$ a nontrivial solution of the survey propagation equations appears, with positive complexity. At $c_s = 1.8789 \pm 0.0002$ the complexity becomes negative: this is the satisfiability transition. At $c_p = 1.992$ the solution at zero energy ceases to exist. The inset magnifies the region where the complexity crosses zero, together with the error bar for the satisfiability transition. Crosses represent results of a population dynamics with $N = 0.5 \cdot 10^5$ elements, squares of $N = 1 \cdot 10^5$, and circles $N = 2 \cdot 10^5$.

giving a negative complexity function. There are thus almost surely no nontrivial fixed points of warning propagation at zero energy.

Before interpreting the survey propagation results, we should check that its application on tree-like random graphs is justified. The method to do this self-consistency check has been developed in [MPRT04] and is discussed in appendix D. For 1-in-3 SAT the result is that SP is stable, thus the results are believed to be correct, for $c \in (1.838, 1.948)$ [RSZ07]. The point c_s belongs to this interval, thus we can interpret it safely as the satisfiability threshold. However, the point c_{SP} has no physical meaning, and some statements that are suggested by its existence are wrong. For example it is not true that there is not exponentially many fixed points of the warning propagation, thus no clustering, for $c < c_{SP}$. This has been remarked in [KMRT⁺07] and a part of chapter 2 will be devoted to understanding this.

1.8 Loose ends

We could summarize the understanding of the subject three years ago in the following way: The 1RSB cavity method was able to compute the satisfiability threshold. The clustered phase was predicted and its existence partially proven. The conjecture that clustering is a key element in understanding of the computational hardness was accepted. The survey propagation inspired decimation algorithm was breath-taking, and the computer science community was getting gradually more and more interested in the concepts which lead to its derivation. It might have seemed

that a real progress can be made only on the mathematical side of the theory, in the analytical analysis of the performance of the message passing algorithms, or in new applications. But several loose ends hanged in the air and the opinions on their resolution were diverse. I will list three of them which I consider to be the most obtruding ones.

(A) The "no man's land", RS unstable but SP trivial — The energetic 1RSB cavity method (survey propagation) predicts the clustering in 3-SAT at $\alpha_{\text{SP}} = 3.92$. But the replica symmetric solution is unstable at already $\alpha_{\text{RS}} = 3.86$, at this point the spin glass susceptibility diverges and equivalently the belief propagation algorithm stops to converge on a single graph, see appendix C. What is the solution in the "no man's land" between α_{RS} and α_{SP} ? The values are even more significant for the 3-coloring or Erdős-Rényi graphs where the corresponding average connectivities are $c_{\text{RS}} = 4$ and $c_{\text{SP}} = 4.42$.

(B) No solutions with nontrivial whitening cores — An iterative procedure called *whitening of a solution* is defined as iteration of the warning propagation equations initialized from a solution. *Whitening core* is the corresponding fixed point. We call white those variables which are assigned the "I do not care" state in the whitening core. A crucial asymptotic property is that if the 1RSB solution is correct then the whitening core of all solutions from one cluster is the same and the non-white variables are the frozen ones in that cluster. Consequently, knowing a solution, the whitening may be used to tell if the solution was or was not in a frozen cluster.

Survey propagation uses information only about frozen cluster. It might seem that every cluster is uniquely described by its whitening core, that is by the set and values of the frozen variables.

Yet, the solutions found by survey propagation have always a trivial, all white, whitening core. This paradox was pointed out in [MMW07] and observed also by the authors of [BZ04]. It was suggested that the concept of whitening might be meaningful only in the thermodynamical limit. But that was not a satisfactory explanation.

(C) Where do the simple local algorithms actually fail — The clustered phase, baptized "Hard" in [MZ02] does not seem to be that hard. There is no local algorithm which would perform well exactly up to $\alpha_{\text{SP}} = 3.92$. For a while it was thought that the 1RSB stability point $\alpha_{\text{II}} = 4.15$, see appendix D, is a better alternative. It was argued that the full-RSB states are more "transparent" for the dynamics than the 1RSB states which should be well defined and separated. Moreover there was at least one empirical result which suggested that the Walk-SAT algorithm stops to work in linear time at that point [AGK04]. But other version of Walk-SAT stopped before or even after, as for example the ASAT which was argued in [AA06] to work in linear time at least up to $\alpha = 4.21$.

1.9 Summary of my contributions to the field

In my first works [ZM06, MMR⁺06, RSZ07] I applied the replica symmetric and the energetic 1RSB method to the matching and the 1-in- K SAT problems. This is why I used these two problems to illustrate the methods in sec. 1.5.4 and 1.7.

The problem of matching on graphs is a common playground for algorithmic and methodological development. I studied the problem of counting maximum matchings in a random graph in [ZM06]. Finding a maximum matching is a well known polynomial problem, while their approximative counting is a much more difficult task. We showed, that the entropy of maximum matchings can be computed using the belief propagation algorithm, a result which was later on partially proved rigorously [BN06].

My interest in the 1-in- K SAT problem stemmed from the work [ACIM01] where the authors computed rigorously the satisfiability threshold and showed that the NP-complete problem is in fact on average algorithmically easy. In [MMR⁺06, RSZ07] we studied the random 1-in-3 SAT in two-parameter space. One parameter is the classical constraint density, the other is the probability p of negating a variable in a constraint ($p = 1/2$ in [ACIM01]). We showed that for $0.2627 < p < 0.7373$ the problem is on average easy and the satisfiability threshold can be computed rigorously. On the other hand for $p < 0.07$ the problem is qualitatively similar to the 3-SAT. We computed the threshold from the energetic 1RSB approach. In the intermediate region the 1RSB approach is not stable, thus it stays an open question how exactly does the problem evolve from an on average easy case to a 3-SAT like case. Qualitatively similar phase diagram was described in the $2 + p$ SAT problem [MZK⁺99a, AKKK01]. We also found an interesting region of the parameter space in the 1-in-3 SAT where the unit clause algorithm provably finds solutions despite the replica symmetric solution being not correct (unstable).

The rest of my works [KMRT⁺07, ZK07, KZ08b, KZ08a, MZ08, AZ08, ZM08] tied up the loose ends from the previous section and mainly addressed the original question of this article: Why are some constraint satisfaction problems intrinsically hard on average and what causes this hardness?

I used the entropic zero temperature 1RSB approach, introduced in [MPR05], to study the structure of solutions in random CSPs. In [KMRT⁺07, ZK07] we discovered that the true clustering (dynamical) transition does not correspond to the onset of a nontrivial solution of the survey propagation equations. We gave a proper definition of the clustering transition and formulated it in terms of extremality of the uniform measure over solutions. The clustering transition happens always before or at the same time as the replica symmetric solution ceases to be stable. This tied up the loose end (A), as in the "no man's land" the energetic 1RSB solution was simply incomplete.

We showed that in general there exist *two* distinct clustered phases below the satisfiable threshold. In the first, *dynamic clustered phase*, an exponentially large number of pure states is needed to cover almost all solutions. However, average properties (such as total entropy) still behave as if the splitting of the measure did not count. In particular, a simple algorithm such as belief propagation gives asymptotically correct estimates of the marginal probabilities. However, the measure over solutions is not extremal and, more importantly, the Monte Carlo equilibration time diverges, thus making the *sampling of solutions* a hard problem. The second kind of clustered phase is the *condensed clustered phase* where a finite number of pure states is sufficient to cover almost all solutions. A number of nontrivial predictions follows: for instance the total entropy has a non-analyticity at the transition to this phase, the marginal probabilities are non-self-averaging and not given anymore by the belief propagation algorithm.

In the context of the coloring problem, i.e. anti-ferromagnetic Potts glass, I also addressed related questions of what does the 1RSB solution predict for the finite temperature phase diagram and when is the 1RSB solutions correct (stable) [ZK07]. We give the full phase diagram for this

model and argue that in the colorable phase for at least 4 colors the 1RSB solutions is stable, and thus believed to be exact.

In order to clarify and substantiate this heuristic picture, we introduced the random subcubes model in [MZ08], a generalization of the random energy model. The random subcubes model is exactly solvable and reproduces the sequence of phase transitions in the real CSPs (clustering, condensation, satisfiability threshold). Its, perhaps, most remarkable property is that it reproduces quantitatively the behaviour of random q -coloring and random K -SAT in the limit of large q and K . We showed that the random subcubes model can also be used as a simple playground for the studies of dynamics in glassy systems.

An important and quite novel phenomena I investigated in [ZK07, KZ08a] is the freezing of variables. A variable is frozen when in all the solutions belonging to one cluster it takes the same value. I discovered that the fraction of such frozen variables undergoes a first order phase transition when the size of states is varied. I introduced the notion of the *rigidity transition* as the point where almost all the dominating clusters become frozen and the *freezing transition* as the point where all the clusters become frozen. The solutions belonging to the frozen clusters can be recognized via the whitening procedure.

We computed the rigidity transition in the random coloring in [ZK07]. And we studied the freezing transition in 3-SAT numerically [AZ08], with the result $\alpha_f = 4.254 \pm 0.009$ (to be compared to the satisfiability threshold $\alpha_s = 4.267$). This study also confirms that the notion of whitening and freezing of variables is meaningful even on relatively small systems.

This allows us to tie up the loose end (B). The survey propagation algorithm describes the most numerous frozen clusters. The range of connectivities where the SP based algorithms are able to find solutions in 3-SAT lies in the phase where most solutions are in fact unfrozen. It is thus much less surprising that the SP based algorithms always find a solution with a trivial whitening.

A very natural question cannot be avoided at this point: What happens in the frozen phase where all the solutions are frozen? We know that such a phase exists, this was shown in [ART06] and numerically in [AZ08]. And we also know from several authors that the known algorithms do not seem to be able to find frozen solutions in polynomial time (that is never for sufficiently large instances). We conjectured in [ZK07] that the freezing is actually a relevant concept for the algorithmical hardness. Thus the answer we suggest to tie up the loose end (C) is that the simple local algorithms stop always before the freezing transition. It is a challenging problem to design an algorithm which would be able to beat this threshold.

In the coloring and satisfiability problems (at reasonably small q and K) the freezing transition is however very near to the satisfiability threshold, see the numbers in [ZK07, AZ08]. It is thus difficult to make strong empirical conclusions about the relation between hardness and freezing. Motivated by the need of problems where the freezing and satisfiability would be well separated I introduced the *locked* constraint satisfaction problems where the freezing transition coincides with the clustering one [ZM08]. The locked CSPs are very interesting from several points of view. The clusters in locked CSPs are point-like, this is why the clustering and freezing coincide. This is also connected with a remarkable technical simplification, as these problems can be fully described on the replica symmetric level.

On the other hand the locked problems are extremely algorithmically challenging. We implemented the best known solvers and showed that they do not find solutions starting very precisely from the clustering (= freezing) transition. At the same time this transition is very well separated

from the satisfiability threshold.

A remarkable point about a subclass of the locked problems which we called *balanced* is that the satisfiability threshold can be obtained exactly from the first and second moment calculation. This adds a huge class of constraint satisfaction problems to a handful of other NP-complete CSPs where the threshold is known rigorously. And it also brings the understanding of which properties of the problem introduce fluctuations which make the second moment method fail.

The numerical work on the 3-SAT problems [AZ08] also addresses another important and almost untouched question: How much are the asymptotic results relevant for systems of practical sizes. We counted the number of clusters in random 3-SAT on instances up to size $N = 150$ and compared to the analytical prediction. We saw that the comparison is strikingly good for already so small systems. This should encourage the application of statistical physics methods to the real world problems.

2 Clustering

In this chapter we introduce the concept of clustering of solutions. First we investigate when does the replica symmetric solution fail. Then we derive the one-step replica symmetry breaking equations on trees and give their interpretation on random graphs. We discuss how several geometrical definitions of clusters might be related to the pure states and review the properties of the clustered phase. Finally, we revise how is the clustering related to the algorithmical hardness and conclude that it is considerably less than previously anticipated. The original contributions to this chapter were published in [KMRT⁺07,ZK07,AZ08].

2.1 Definition of clustering and the 1RSB approach

How to recognize when is the replica symmetric solution correct? First we have to explain what do we precisely mean by "being correct". We obviously require that quantities like the free energy, energy, entropy, marginal probabilities (magnetizations) are asymptotically exact when computed in the replica symmetric approach. But this is not enough, as this is also satisfied in the phase which we will call later the *clustered* (dynamical) 1RSB phase.

A commonly used necessary condition for the validity of the RS solution is referred to as the *local stability towards 1RSB*. It consists in checking that the spin glass susceptibility does not diverge, or equivalently that the belief propagation algorithm converges on a large single graph, or in the probability theory this corresponds to the Kesten-Stigum condition [KS66a,KS66b]. These and other equivalent representations for the replica symmetric stability are discussed in detail in appendix C. If the replica symmetric solution is not stable then it predicts wrong free energy, entropy, correlation functions, etc. But the contrary is far from being true: even if stable, the RS solution might be wrong, and even unphysical (predicting negative entropies in discrete models, negative energies in models with strictly non-negative Hamiltonian function, or discontinuities in functions which physically have to be Lipschitzian).

It is tempting to say: The replica symmetric solution is correct if and only if the assumptions we used when deriving it are correct. In deriving the belief propagation (1.16) and the RS free energy (1.20) we used only one assumption: The neighbours of a variable i are independent random variables, under the Boltzmann measure (1.13), when conditioned on the value of i . As we will see, this assumption is asymptotically correct also in the dynamical 1RSB phase, and thus the RS marginal probabilities, or the free energy function remain asymptotically exact in that phase.

We thus need a different definition for the "RS correctness" which would determine whether the Boltzmann measure (1.13) can be asymptotically described as a single pure state, and whether the equilibration time of a local dynamics is linear in the system size. At the same time we do not want this definition to refer the RSB solution, because obviously we want to justify the need of the RSB solution by the failure of the RS solution.

A definition satisfying the above requirements appeared only recently [MM08,MS05,MS06c], and it can be written in several equivalent ways. From now on we say that the *replica symmetric solution is correct* if and only if one of the following is true.

- (a) The point-to-set correlations decay to zero.
- (b) Reconstruction on the underlying graph is not possible.

- (d) The uniform measure over solutions satisfies the extremality condition.
- (c) The 1RSB equations at $m = 1$, initialized in a completely biased configuration, converge to a trivial fixed point.

In the rest of this section we explain these four statements, and show that they are indeed equivalent, and explain how do they correspond to the existence of a nontrivial 1RSB solution. We should mention that in the so-called locked constraint satisfaction problems this definition have to be slightly changed at zero temperature, we will discuss that in sec. 4.3. The transition from a phase where the RS solution is correct to a phase where it is not is called the *clustering* or the *dynamical transition*.

Gibbs measures and why are the sparse random graphs different — Our goal is to describe the structure of the set of solutions of a constraint satisfaction problem with N variables. Let $\phi_a(\partial a)$ be the constraint function depending on variables $s_i \in \partial a$ involved in the constraint a , $\phi_a(\partial a) = 1$ if the constraint is satisfied, $\phi_a(\partial a) = 0$ if not. The uniform measure over all solutions can be written as

$$\mu(\{s_i\}) = \frac{1}{Z} \prod_{a=1}^M \phi_a(\partial a), \quad (2.1)$$

where Z is the total number of solutions. The uniform measure over solutions is the zero temperature limit, $\beta \rightarrow \infty$, of the Boltzmann measure

$$\mu(\{s_i\}, \beta) = \frac{1}{Z(\beta)} \prod_{a=1}^M e^{-\beta[1-\phi_a(\partial a)]}. \quad (2.2)$$

The above expressions are valid on any given finite factor graph. The theory of Gibbs measures [Geo88] tries to formally define and describe the limiting object to which (2.1-2.2) converge in the thermodynamical limit, $N \rightarrow \infty$. A common way to build this theory is to ask: What is the measure induced in a finite volume Λ when the boundary conditions are fixed? Roughly speaking, the good limiting objects, called the *Gibbs measures* or the *pure states*, are such that boundaries taken from the Gibbs measure induce the same measure inside the finite large volume Λ .

The Ising model on a 2D lattice gives an excellent example of how a phase transition is seen via Gibbs measures. Whereas in the high temperature paramagnetic phase the Gibbs measure is unique, in the ferromagnetic phase there are two extremal measures, one corresponding to the positive average magnetization, the other to the negative average magnetization. Indeed, if a boundary condition is chosen from one of these two then the correct magnetization will be induced in the bulk. In general the bulk in equilibrium can be described by a linear combination of these two *extremal* objects.

In the disordered models the situation might be much more complicated. Indeed the proper definition of the Gibbs measure in the Edwards-Anderson model (1.11) and other glassy models is a widely discussed but still an open problem [Bov06, Tal03, NS92].

The locally tree-like lattices, we are interested in here, are also peculiar from this point of view. The main difference is that in any reasonable definition of the boundary variables, the

boundary has volume comparable to the volume of the interior. Thus again the usual theory of Gibbs measure implies very little. On the other hand the tree structure makes some considerations simpler. We will try to understand what sort of long range correlations might appear on the tree-like graphs by studying the tree graphs with general boundary conditions.

2.1.1 Properties and equations on trees

It is a well known fact that on arbitrary tree, with arbitrary boundary conditions, the belief propagation equations and the Bethe free energy are exact (the thermodynamical limit is not even needed here) [Pea88, KFL01, YFW00].

But what if the boundary conditions are chosen from a complicated measure? Then very little (if anything) is known in general. However, there is a way how to choose the boundary conditions such that the tree is then described by the one-step replica symmetry breaking equations. This is closely linked to the problem of reconstruction on trees, studied in mathematics [EKPS00, Mos01, Mos04]. The link with 1RSB was discovered by Mézard and Montanari [MM06a]. We chose to present the 1RSB equations in this new way, because it opens the door to further mathematical developments. For the original statistical physics derivation we refer to [MP00]. Another recent computer science-like derivation, which is based on the construction of a decorated constraint satisfaction problem and writing belief propagation on such a problem, is presented in [MM08, Mor07].

Reconstruction on trees — We explain the concept of reconstruction on trees [Mos04]. For simplicity we consider q -coloring on a rooted tree with constant branching factor γ (sometimes also called the Cayley tree). A more general situation (with disorder, in the interaction or in the branching factor) is described in appendix A.

Create a rooted tree with branching γ and with L generations. An example of $\gamma = 2$ and $L = 8$ is in fig. 2.1. Assign a color s_0 to the root and broadcast over the edges towards the leaves of the tree in such a way that if a parent node i was assigned color s_i then each of its ancestors is assigned random one of the remaining $q - 1$ colors.

At the end of this broadcasting, every node in the tree is assigned a color, and this assignment corresponds to a proper coloring (neighbours have different colors). Now in an imaginary experiment we forget the colors everywhere but on the leaves. The problem of reconstruction consists in deciding if there is any information left in the values on the leaves (and their correlation) about the original color s_0 of the root in the limit of infinite tree $L \rightarrow \infty$. If the answer is yes then we say that the reconstruction is possible, if the answer is no then the reconstruction is not possible.

Call $\{s\}_l$ the assignment of colors in the l^{th} generation of the tree. Consider formally the probability $\psi_{s_0}(\{s\}_l)$ that a broadcasting process which finished at the configuration $\{s\}_l$ started from the color s_0 at the root. In other words, in what fraction of assignments in the interior of the tree (compatible with the boundary conditions $\{s\}_l$) is the color of the root s_0 ? Reconstruction is possible if and only if

$$\lim_{l \rightarrow \infty} \sum_{r=1}^q \psi_r(\{s\}_l) \log [q \psi_r(\{s\}_l)] > 0. \quad (2.3)$$

Intuitively when the branching γ is small and the number of colors large the information about the root will be lost very fast. If, on the contrary, the branching is large compared to the




Fig. 2.1. (Color online) Illustration of the broadcasting of colors on a binary tree ($\gamma = 2$) for the reconstruction problem.

number of colors some information remains. A simple exercise is to analyze the so-called *naive reconstruction* algorithm [Sem08]. The naive reconstruction is possible if the probability that the leaves determine uniquely the root does not go to zero as the number of generation goes to infinity. We compute the probability η that the far-away boundary is compatible with only one value of the root. Denote η_l the probability that a variable in the l^{th} generation is directly implied conditioned on the value of its parent. The probability η_{l-1} can be computed recursively as

$$\begin{aligned} \eta_{l-1} &= 1 - (q-1) \left(1 - \frac{1}{q-1} \eta_l\right)^\gamma + \frac{(q-1)(q-2)}{2} \left(1 - \frac{2}{q-1} \eta_l\right)^\gamma - \dots \\ &= \sum_{r=0}^{q-1} (-1)^r \binom{q-1}{r} \left(1 - \frac{r}{q-1} \eta_l\right)^\gamma. \end{aligned} \quad (2.4)$$

The terms in this telescopic sum come from probabilities that number r out of the $q-1$ colors are not present in the γ descendants. In the last generation we know the colors by definition of the problem, thus $\eta_\infty = 1$. If the iterative fixed point of (2.4) is positive then the reconstruction is possible.

This simple upper bound on the branching γ for which the reconstruction is possible is actually quite nontrivial and in the limit of large number of colors it coincides with the true threshold at least in the first two orders, see [ZK07] and [Sem08, Sly09]. This upper bound is connected to the presence of frozen variables and will be discussed in a greater detail in chapter 4.

Self-consistent iterative equations for the reconstruction — The iterative equations for the reconstruction problem are equivalent to the one-step replica symmetry breaking equations with

Parisi parameter m , $m = 1$ will apply to the original question of reconstructibility. This was first derived by Mézard and Montanari [MM06a] and it has some deep consequences for the understanding of the RSB solution. We now explain this derivation, still for the coloring problem with a fixed branching γ and q colors. A more general form is presented in appendix A.

For given boundary conditions $\{s\}_l$, constructed as described above, we compute the probability $\psi_{s_i}^{i \rightarrow j}$ (over all broadcasting experiments leading to these boundary conditions) that a variable i had color s_i , where j is the parent of i and the edge (ij) has been cut. Given the probabilities on the descendants of i , which are indexed by $k = 1, \dots, \gamma$, we can write

$$\psi_{s_i}^{i \rightarrow j} = \frac{1}{Z^{i \rightarrow j}} \prod_{k=1}^{\gamma} (1 - \psi_{s_i}^{k \rightarrow i}) \equiv \mathcal{F}_{s_i}(\{\psi^{k \rightarrow i}\}), \quad (2.5)$$

because the descendants can take any other color but s_i . The $Z^{i \rightarrow j}$ is a normalization constant. It should be noticed that this is in fact the belief propagation equation (1.16) for the graph coloring. This equation can also be derived by counting how many assignments are consistent with the boundary conditions $\{s\}_l$. This gives a natural interpretation to $Z^{i \rightarrow j}$

$$Z^{i \rightarrow j} = \frac{Z^{(i)}}{\prod_{k=1}^{\gamma} Z^{(k)}}. \quad (2.6)$$

where $Z^{(i)}$ is the total number of solutions consistent with $\{s\}_l$ if i were the root. Thus $Z^{i \rightarrow j}$ is a change in the number of solutions compatible with the boundary conditions when the γ branches are merged.

Now we consider the distribution over all possible boundary conditions which are achievable by the broadcasting process defined above. We have to specify the probability distribution on the boundary conditions. We consider that the probability of every boundary conditions $\{s\}_l$ is proportional to the power m of the number of ways by which we could create $\{s\}_l$, denote this number $Z(\{s\}_l)$. In other words, the probability of a given boundary condition is proportional to the power m of the number of possible assignments in the bulk of the tree.

$$\mu(\{s\}_l) = \frac{[Z(\{s\}_l)]^m}{\mathcal{Z}(m)}, \quad \text{where} \quad \mathcal{Z}(m) = \sum_{\{s\}_l} [Z(\{s\}_l)]^m. \quad (2.7)$$

The value of $m = 1$ is natural for the original question of reconstruction, because every realization of the broadcasting experiment is then counted in a equiprobable way. We, however, introduced a general power m . The parameter m will play a role of the Legendre parameter, changing its value focuses on boundary conditions compatible with a given number of assignments inside the tree.

Denote $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ the distribution of $\psi^{i \rightarrow j}$, over the measure on the boundary conditions (2.7)

$$P^{i \rightarrow j}(\psi^{i \rightarrow j}) \equiv \sum_{\{s\}_l} \mathbb{I}(\{s\}_l \text{ induce } \psi^{i \rightarrow j}) \frac{[Z^{(i)}(\{s\}_l)]^m}{\mathcal{Z}^{(i)}(m)}. \quad (2.8)$$

Where $Z^{(i)}(\{s\}_l)$ is the number of solutions induced on the subtree rooted in vertex i , $\mathcal{Z}^{(i)}(m)$ is the corresponding normalization. To express the probability distribution $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ as a

function of $P^{k \rightarrow i}(\psi^{k \rightarrow i})$ we need that $\psi^{i \rightarrow j} = \mathcal{F}(\{\psi^{k \rightarrow i}\})$, eq. (2.5). Moreover, $Z^{i \rightarrow j}$ is the increase in the total number of solutions after merging the branches rooted at $k = 1, \dots, \gamma$ into one branch rooted at i . The distributional equation for P is then

$$P^{i \rightarrow j}(\psi^{i \rightarrow j}) = \frac{1}{Z^{i \rightarrow j}} \int \prod_{k=1}^{\gamma} dP^{k \rightarrow i}(\psi^{k \rightarrow i}) (Z^{i \rightarrow j})^m \delta[\psi^{i \rightarrow j} - \mathcal{F}(\{\psi^{k \rightarrow i}\})], \quad (2.9)$$

where \mathcal{F} and $Z^{i \rightarrow j}$ are defined in (2.5), and $Z^{i \rightarrow j}$ is a normalization constant equal to

$$Z^{i \rightarrow j} = \frac{Z^{(i)}(m)}{\prod_{k=1}^{\gamma} Z^{(k)}(m)} = \int \prod_{k=1}^{\gamma} dP^{k \rightarrow i}(\psi^{k \rightarrow i}) (Z^{i \rightarrow j})^m. \quad (2.10)$$

where $Z^{(i)}$ is the normalization from (2.8) if i were the root. Notice that if we start from boundary conditions which are not compatible with any solution then the re-weighting $Z^{i \rightarrow j} = 0$ at the merging where a contradiction is unavoidable. Initially at the leaves the colors of nodes are known. Call δ_r the q -component vector $\psi_{s_i}^{i \rightarrow j} = \delta(s_i, r)$, then the initial distribution is just a sum of singletons

$$P^{\text{init}}(\psi) = \frac{1}{q} \sum_{r=1}^q \delta(\psi - \delta_r). \quad (2.11)$$

Denote $P_0(\psi)$ the distribution created from (2.11) after many iteration of (2.9) with $m = 1$. The reconstruction is possible if and only if $P_0(\psi)$ is nontrivial, that is different from singleton on $\psi_{s_i} = 1/q, \forall s_i$. We define the critical branching factor γ_d in such a way that for $\gamma < \gamma_d$ the reconstruction is not possible, and for $\gamma \geq \gamma_d$ the reconstruction is possible. The critical values $\gamma_d = c_d - 1$ for the coloring problem are reviewed in tab. 5.2.

What are clusters on a tree? If the reconstruction is not possible, then almost all (with respect to (2.7) at $m = 1$) boundary conditions do not contain any information about the original color of the root. However, for rare boundary conditions this might be different. Obviously as long as $\gamma \geq q - 1$ one can always construct boundary conditions which determine uniquely the value of the root (by assigning every of the $q - 1$ colors to the descendants of every node). If $\gamma < q - 1$ then this is no longer possible. And it was proven in [Jon02] that for $\gamma < q - 1$ every boundary conditions lead to an expectation $1/q$ for every color on the root. If the reconstruction is possible, then different boundary conditions may lead to different expectations on the root.

The basic idea of the definition of clusters on a tree is the same as in the classical definition of a Gibbs measure [Geo88]. However, some more work is needed to make the following considerations rigorous. Define a d -neighbourhood of the root as all the nodes up to d^{th} generation, consider $1 \ll d \ll l$. Consider the set \mathcal{S} (resp. \mathcal{S}') of all assignments on the d -neighbourhood compatible with a given boundary condition $\{s\}_l$ (resp. $\{s'\}_l$). Define two boundary conditions $\{s\}_l$ and $\{s'\}_l$ as equivalent if the fraction of elements in which the two sets \mathcal{S} and \mathcal{S}' differ goes to zero as $l, d \rightarrow \infty$. Clusters are then the equivalence classes in the limit $l \rightarrow \infty, d \rightarrow \infty, d \ll l$. The requirement $d \ll l$ comes from the fact that in $l - d$ iterations the equation (2.8) should converge to its iterative fixed point.

As we explained, more than one cluster exists as soon as the branching factor $\gamma \geq q - 1$, but as long as the iterative fixed point of eq. (2.9) at $m = 1$ is trivial all but one clusters are negligible

because they contain an exponentially small fraction of solutions. Indeed, if the reconstruction is not possible it means that the information about the d -neighbourhood is almost surely lost at the l^{th} generation. Thus almost every broadcasting will lead to a boundary condition from the only relevant giant cluster.

Only for $\gamma \geq \gamma_d$, when the reconstruction start to be possible, the total weight of all solutions will be split into many clusters. In every of them the set of expectation values (beliefs) $\psi^{i \rightarrow j}$ will be different. This is related to another derivation of the 1RSB equations where the clusters of solutions on a given graph are identified with fixed points of the belief propagation equations [MM08, Mor07]. There are exponentially many (in the total number of variables N) initial conditions, it is also reasonable to expect that the number of clusters will be exponentially large in N .

The complexity function — The number of solutions compatible with a boundary condition $(\{s\}_l)$ was denoted $Z(\{s\}_l)$ in eq. (2.7). The associated entropy is then, due to interpretation of $Z^{i \rightarrow j}$ (2.6)

$$S(\{s\}_l) \equiv \log [Z(\{s\}_l)] = \sum_i \log Z^{i \rightarrow j}, \quad (2.12)$$

where the sum is over all the vertices i in the tree, if i is a leaf then $Z^{i \rightarrow j} = 1$, if i is the root that j is a imaginary parent of the root. An intuition about this formula is the following: $\log Z^{i \rightarrow j}$ is the change in the entropy when the node i and all edges (ki) , where k are descendants of i , are added. Summing over all i then creates the whole tree.

More commonly, we introduce also messages going from the parents to the descendants and write the expression for the entropy (2.12) in the equivalent Bethe form [YFW03]

$$S(\{s\}_l) = \sum_i \log Z^{i+\partial i} - \sum_{ij} \log Z^{ij}, \quad (2.13)$$

where

$$Z^{i+\partial i} = \sum_{r=1}^q \prod_{k \in \partial i} [1 - \psi_r^{k \rightarrow i}], \quad Z^{ij} = 1 - \sum_{r=1}^q \psi_r^{i \rightarrow j} \psi_r^{j \rightarrow i}, \quad (2.14)$$

where ∂i are all the neighbours (descendants and the parent) of node i . The first sum in (2.13) goes over all the nodes in the tree, the root included, leaves have only one allowed color, thus eq. (2.14) changes correspondingly. Again the meaning of $\log Z^{i+\partial i}$ is the change in the entropy when node i and his neighbourhooding edges are added, each edge is then counted twice, thus the shift in the entropy when an edge (ij) is added, $\log Z^{ij}$, have to be subtracted.

We denote $\Phi(m) \equiv \log \mathcal{Z}(m)$ the thermodynamical potential associated to the measure (2.7). To avoid confusion with the real free energy, associated to the uniform measure over solutions (2.1), we call it the *replicated free entropy*. If a nonzero temperature is involved then $-\Phi(m)/(\beta m)$ is called the *replicated free energy*. The replicated free entropy on a tree can be expressed in totally analogous way as the entropy. From (2.10) we derive

$$\Phi(m) \equiv \log \mathcal{Z}(m) = \sum_i \log \mathcal{Z}^{i \rightarrow j}, \quad (2.15)$$

which is usually written in the equivalent way

$$\Phi(m) = \sum_i \log \mathcal{Z}^{i+\partial i} - \sum_{ij} \log \mathcal{Z}^{ij}, \quad (2.16)$$

where we introduced

$$\mathcal{Z}^{i+\partial i} = \int \prod_{k \in \partial i} dP^{k \rightarrow i}(\psi^{k \rightarrow i}) (Z^{i+\partial i})^m, \quad (2.17a)$$

$$\mathcal{Z}^{ij} = \int dP^{i \rightarrow j}(\psi^{i \rightarrow j}) dP^{j \rightarrow i}(\psi^{j \rightarrow i}) (Z^{ij})^m. \quad (2.17b)$$

We denote $\Sigma(m)$ the Shannon entropy corresponding to measure on the boundary conditions (2.7), and we call it the *complexity* function.

$$\Sigma(m) \equiv - \sum_{\{s\}_l} \mu(\{s\}_l) \log \mu(\{s\}_l) = -mS(m) + \Phi(m), \quad (2.18)$$

where S is the entropy averaged with respect to $\mu(\{s\}_l)$

$$S(m) = \sum_{\{s\}_l} \frac{[Z(\{s\}_l)]^m}{\mathcal{Z}(m)} \log Z(\{s\}_l) = \frac{\partial \Phi(m)}{\partial m}. \quad (2.19)$$

Thus the complexity can also be written as a function of the internal entropy via the Legendre transform of the replicated free entropy $\Phi(m)$

$$\Sigma(S) = -mS + \Phi(m) \quad \text{with} \quad \frac{\partial \Sigma(S)}{\partial S} = -m. \quad (2.20)$$

The reader familiar with the cavity approach surely recognized eqs. (2.9) and (2.15–2.20) as the 1RSB equations.

Interpretation of the complexity function — In the cavity method [MP01] the exponential of the complexity function $\Sigma(m)$ (2.18) counts the number of clusters corresponding to a given value of the parameter m , that is of a given entropy S (2.19). Complexity defined on the full tree is never negative, as it is a Shannon entropy of a discrete random variable. The same is, of course true, about the entropy (2.12).

It is more interesting to consider the complexity (or the entropy) function $\Sigma_d(m)$ on the d -neighbourhood of the root. If the total number of generations of the tree is l we take $1 \ll d \ll l$. And moreover we require $l - d$ to be large enough, such that the distributional iterative equation (2.9) converges to its fixed point in less than $l - d$ iterations. The average complexity function on the d -neighbourhood can then be computed from this fixed point. And it can be both positive or negative. Its negative value then means that the number of clusters is decreasing as we are getting nearer to the root. Two important critical connectivities can be defined

- γ_c : at which the complexity of the "natural" clusters $\Sigma_d(m = 1)$ becomes negative.
- γ_s : at which the maximum of the complexity $\Sigma_d(m = 0)$ becomes negative.

The connectivity γ_s is the tree-analog of the satisfiability threshold. The connectivity γ_c is the tree-analog of the condensation transition on random graphs, see chapter 3.

Strictly speaking, it is not known how to justify the interpretation of the complexity function as the counter of clusters in the derivation we just presented. In the original cavity derivation [MP01] or in the later derivations [MM08, Mor07] this point is well justified. We, however, find the purely tree derivation appealing for further progress on the mathematical side of the theory and that is why we have chosen to present this approach despite this current incompleteness.

2.1.2 Back to the sparse random graphs

We stress that the equations, derived in the previous section, are all exact on a given (even finite) tree and that we have not use any approximation. We were just describing boundary conditions correlated via (2.7). These, in nature recursive, equations are solved via the population dynamics technique, see the appendix E.

To come back to the sparse random graphs, which are only locally tree-like, we can consider equations (2.9-2.26) as an approximation on arbitrary graphs, just as we did with belief propagation. This leads to the one-step replica symmetry breaking (1RSB) approach. Note that on random graphs we will always speak about densities of the entropy, complexity or free-entropy etc. Thus on random graphs: instead of the entropy S defined in (2.12) we consider $s = S/N$. The replicated free entropy Φ (2.15) and complexity Σ (2.18) are also divided by the number of variables. We, however, denote them by the same symbol, as confusion is not possible.

Let us discuss once again, now from the random graph perspective, what are the correlations which make the replica symmetric approach fail. This will finally explain the definition of the *replica symmetric solution being correct* given at the beginning of this section.

Point-to-set correlations — The concept of the point-to-set correlations is common in the theory of glassy systems. Usually it is considered in the phenomenology of the real glassy systems on finite-dimensional lattices, see for example [BB04] and references therein. Here we restrict the discussion to properties relevant for the tree-like lattices.

Call $\overline{B}_d(i)$ all vertices of the graph which are at distance at least d from i , define *point-to-set* correlation function as

$$C_d(i) = \|\mu(i, \overline{B}_d(i)) - \mu(i)\mu(\overline{B}_d(i))\|_{\text{TV}}, \quad (2.21)$$

where $\mu(\cdot)$ is the uniform measure over solutions (2.1), and the total variation distance of two probability distributions is defined as $\|q-p\|_{\text{TV}} = \sum_x |q(x) - p(x)|/2$. The average point-to-set correlation is

$$C_d = \frac{1}{N} \sum_{i=1}^N C_d(i). \quad (2.22)$$

The reconstruction on graphs is then defined via the decay of this correlation function. The reconstruction on tree-like graphs is not in general equivalent to the reconstruction on trees. Roughly said, it is not equivalent in the ferromagnetic models, e.g. the ferromagnetic Ising model, which spontaneously break some of the discrete symmetries. On the other hand on most

of the frustrated models they are equivalent. A general condition, which might be very nontrivial to check, is given in [GM07].

If the point-to-set correlation function decays to zero, $\lim_{d \rightarrow \infty} C_d = 0$, then almost every variable is independent of its far away neighbours. The replica symmetric approach then has to be asymptotically correct on locally tree-like lattices.

On the other hand if the point-to-set correlations do not decay to zero, then the far-away neighbours influence the value of the variable i . And the replica symmetric solution fails to give the correct picture of the properties of the model. The lack of decay of the point-to-set correlations is equivalent to the reconstruction on graphs, and is also equivalent to the existence of a nontrivial solution of the 1RSB equation (2.9) at $m = 1$. This is also equivalent to the extremality condition for the uniform measure (2.1), which was used in definition of [KMRT⁺07] and reads

$$\mathbb{E} \left[\sum_{\bar{B}_d(i)} \mu(\bar{B}_d(i)) \|\mu(i|\bar{B}_d(i)) - \mu(i)\|_{\text{TV}} \right] \xrightarrow{d \rightarrow \infty} 0, \quad (2.23)$$

where the external average is over quenched disorder (in interactions or connectivities).

The point-to-set correlations do not decay to zero for example in the low temperature phase of the ferromagnetic Ising model on a random graph. There it is sufficient to introduce the pure state "up" and the pure state "down" and within these pure states the point-to-set correlations will decay to zero again. On the frustrated models the situation is more complicated but the idea of the resolution is the same: If we manage to split the set of solutions into clusters (pure states) such that within each cluster the point-to-set correlations again decay, the situation is fixed. A statistical description of the properties of clusters can be obtained using the *one-step replica symmetry breaking* (1RSB) equations, derived in the previous section 2.1.1 and summarized in the next section 2.1.3.

However, the correlations might be more complicated and might not be captured fully by the 1RSB approach. In particular the 1RSB approach is correct if and only if the point-to-set correlation decay to zero within clusters and if the replica symmetric statistical description of clusters is correct. In appendix D we will discuss a necessary condition for the 1RSB approach being correct. In case the 1RSB approach does not fully describe the system further steps of replica symmetry breaking might provide a better approximation (that means splitting clusters into sub-clusters or aggregation of clusters) [MP00]. However, on the tree-like lattices, the exact solutions is not known in such cases.

Relation with equilibration time — In glasses, the clustering transition is usually studied at finite temperature and is called the dynamical transition. The clustered phase with $\Sigma(m = 1) > 0$ is called the *dynamical* 1RSB phase. This phase, where most of the static properties do not differ from the replica symmetric (liquid) ones, was first described and discuss in [KT87a, KT87b]. The dynamical transition is associated with a critical slowing down of the dynamical properties, e.g. the equilibration time is expected to diverge at this point. Note that such a purely dynamical phase transition is typical for mean-field models. In the finite dimensional glassy systems the barriers between a metastable and an equilibrium state are finite (independent of the system size). This is because the nucleation length might be large but have to be finite. Thus instead of a sharp dynamical transition in finite dimensional systems we observe only a crossover.

However, even at the mean field level, the exact dynamical description is known only in a few toy models, e.g. the spherical p -spin model [CK93] or the random subcubes model [MZ08]. In general, the dynamical solution is only approximative, still many very interesting results were obtained. For a review see [BCKM98]. In the models on sparse random lattices even the approximation schemes are rather poor, see e.g. [SW04]. Thus the exact general relation between dynamics and the dynamical (clustering) transition is not known.

An important contribution in establishing the link between dynamics and the static solution on random graphs is [MS05, MS06b, MS06c] where the divergence of the point-to-set correlation length is linked with divergence of the equilibration time of the Glauber dynamics. This suggests that beyond the clustering transition the Monte Carlo sampling (or maybe even sampling in general) will be a hard task.

Note also that in the mathematical literature the Glauber dynamics is often studied. Many results exist about the so-called *rapid mixing* of the associated Markov chain [Sin93]. But the rapid mixing questions equilibration in polynomial time, whereas in physics the relevant time scale is linear. Moreover rapid mixing is defined as convergence to the equilibrium measure from any possible initial conditions, whereas in physics of glasses the notion of a typical initial condition should be used instead.

2.1.3 Compendium of the 1RSB cavity equations

We review the 1RSB equations on a general CSP. The order parameter is a probability distribution of the cavity field (BP message) $\psi^{a \rightarrow i} = (\psi_0^{a \rightarrow i}, \dots, \psi_{q-1}^{a \rightarrow i})$. The self-consistent equation for $P^{a \rightarrow i}$ reads

$$P^{a \rightarrow i}(\psi^{a \rightarrow i}) = \frac{1}{Z^{j \rightarrow i}} \int \prod_{j \in \partial a - i} \prod_{b \in \partial j - a} [dP^{b \rightarrow j}(\psi^{b \rightarrow j})] \times (Z^{j \rightarrow i})^m \delta[\psi^{a \rightarrow i} - \mathcal{F}(\{\psi^{b \rightarrow j}\})], \quad (2.24)$$

where the function $\mathcal{F}(\{\psi^{b \rightarrow j}\})$ and the term $Z^{j \rightarrow i}$ are defined by the BP equation (1.17), $Z^{j \rightarrow i}$ is a normalization constant.

The associated thermodynamical potential (2.15) is computed as

$$\Phi(m) = \frac{1}{N} \left[\sum_a \log Z^{a+\partial a} - \sum_i (l_i - 1) \log Z^i \right], \quad (2.25a)$$

$$Z^{a+\partial a} = \int \prod_{i \in \partial a} \prod_{b \in \partial i - a} [dP^{b \rightarrow i}(\psi^{b \rightarrow i})] (Z^{a+\partial a})^m, \quad (2.25b)$$

$$Z^i = \int \prod_{a \in \partial i} [dP^{a \rightarrow i}(\psi^{a \rightarrow i})] (Z^i)^m, \quad (2.25c)$$

where the terms $Z^{a+\partial a}$ and Z^i are the partition sum contributions defined in (1.19).

The logarithm of the number of states divided by the system size defines the complexity function Σ . Inversely the number of states is $e^{N\Sigma}$. At finite temperature the complexity of states with a given internal free energy is a Legendre transformation of the potential $\Phi(m)$

$$\Phi(m) = -\beta m f + \Sigma(f), \quad (2.26)$$

Useful relations between the free energy, complexity and potential Φ are

$$\partial_f \Sigma(f) = \beta m, \quad \partial_m \Phi(m) = -\beta f, \quad m^2 \partial_m \frac{\Phi(m)}{m} = -\Sigma. \quad (2.27)$$

At zero energy, $E = 0$, and zero temperature, $\beta \rightarrow \infty$, the free energy becomes entropy $-\beta f \rightarrow s$. Then the complexity is a function of the internal entropy of states and (2.26) becomes

$$\Phi(m) = ms + \Sigma(s), \quad (2.28)$$

with

$$\partial_s \Sigma(s) = -m, \quad \partial_m \Phi(m) = s, \quad m^2 \partial_m \frac{\Phi(m)}{m} = -\Sigma. \quad (2.29)$$

This is called the *entropic* zero temperature limit. The internal entropy is expressed as

$$s = \frac{1}{N} \left(\sum_a \Delta S^{a+\partial a} - \sum_i (l_i - 1) \Delta S^i \right), \quad (2.30)$$

where $\Delta S^{a+\partial a}$ (ΔS^i resp.) is an internal entropy shift when the constraint a and all its neighbour (the variable i resp.) are added to the graph.

$$\Delta S^{a+\partial a} = \frac{\int \prod_{i \in \partial a} \prod_{b \in \partial i - a} [dP^{b \rightarrow i}(\psi^{b \rightarrow i})] (Z^{a+\partial a})^m \log Z^{a+\partial a}}{\int \prod_{i \in \partial a} \prod_{b \in \partial i - a} [dP^{b \rightarrow i}(\psi^{b \rightarrow i})] (Z^{a+\partial a})^m}, \quad (2.31a)$$

$$\Delta S^i = \frac{\int \prod_{a \in \partial i} [dP^{a \rightarrow i}(\psi^{a \rightarrow i})] (Z^i)^m \log Z^i}{\int \prod_{a \in \partial i} [dP^{a \rightarrow i}(\psi^{a \rightarrow i})] (Z^i)^m}. \quad (2.31b)$$

In the energetic zero temperature limit, described in sec. 1.6 for zero energy, the Parisi parameter $y = \beta m$ is kept constant, thus $m \rightarrow 0$. The free energy then converges to the energy, and (2.26) becomes

$$\Phi(y) = -ye + \Sigma(e), \quad (2.32)$$

where the complexity is this time a function of the energy density e . The survey propagation equations generalized to nonzero y are called the SP- y equations.

Equations (2.24-2.31) are defined on a single instance of the constraint satisfaction problem. Averages \mathcal{P} over the graph ensemble are obtained in a similar manner as in sec. 1.5.3 for the replica symmetric solution.

$$\begin{aligned} \mathcal{P}[P(\psi)] &= \sum_{\{l_i\}} \left[\prod_{l_i} \mathcal{Q}_1(l_i) \right] \int \prod_{i=1}^{K-1} \prod_{j_i=1}^{l_i} \left\{ d\mathcal{P}[P^{j_i}(\psi^{j_i})] \right\} \\ &\times \delta[P(\psi) - \mathcal{F}_2(\{P^{j_i}(\psi^{j_i})\})], \end{aligned} \quad (2.33)$$

where in the sum over $\{l_i\}$, $i \in \{1, \dots, K-1\}$, and the functional \mathcal{F}_2 is defined by (2.24). Analogical expression holds for the average of the complexity or internal entropy. A general method to solve the equation (2.33) is the population of populations described in appendix E.5.

2.2 Geometrical definitions of clusters

Up to now we were describing clusters, i.e., partitions of the space of solutions, in a very abstract way which was defined only in the thermodynamical limit. We showed how to compute the number of clusters of a given size (internal entropy) (2.28), and we argued that the description makes sense if the point-to-set correlation (2.22) decays to zero within almost every cluster of that size. In this last sense clusters are what we would call in statistical physics pure equilibrium states.

On a very intuitive level, clusters are groups of nearby solutions which are in some sense separated from each other. Several geometrical definitions are used in the literature, we want to review the most common ones and state their relation to the definition above. We want to stress that it is not known whether any of the geometric definitions is equivalent to the description given above and used usually in the statistical physics literature.

Strong geometrical separation, x -satisfiability — First rigorous proofs of existence of an exponential number of clusters of solutions in the random K -SAT were based on the concept of x -satisfiability. Two solutions are at distance x if they differ in exactly xN variables. A formula is said x -satisfiable if there is a pair of solutions at distance x , and x -unsatisfiable if there is not.

Mora, Mézard and Zecchina [MMZ05, DMMZ08] managed to prove that for $K \geq 8$ and a constraint density α near enough to the satisfiability threshold the formulas are almost surely x -satisfiable for $x < x_0$, almost surely x -unsatisfiable for $x_1 < x < x_2$, and almost surely x -satisfiable at $x_3 < x < x_4$, where obviously $0 < x_0 < x_1 < x_2 < x_3 < x_4 < 1$. This means that at least two well separated clusters of solutions exist. Proving that there is an exponentially smaller number of pairs of solutions at distances $x < x_1$ than at distances $x > x_2$ leads to the conclusion that an exponential number of well geometrically separated clusters exists [ART06].

However, the x -satisfiability gives too strong conditions of separability. This is illustrated for example in the XOR-SAT problem [MM06b]. It is still an open question if there is or not a gap in the x -satisfiability in the random 3-SAT near to the satisfiability threshold.

Connected-components clusters — Another popular choice of a geometrical definition of clusters is that clusters are connected components in a graph where every solution is a vertex and solutions which differ in d or less variables are connected. The distance d is often said to be any sub-extensive (in the number of variables N) distance, that is $d = o(N)$. However, such a rule is not very practical for numerical investigations.

In K -SAT, in fact, $d = 1$ seems to be a more reasonable choice. There are two reasons: First, clusters defined via $d = 1$ have correct "whitening" properties as we explain in the next paragraph. Second, we numerically investigated the complexity of $d = 1$ connected-components clusters, fig. 2.2 right, and the agreement with the total number of clusters computed from (2.28) at $m = 0$ is strikingly good. In particular, near to the satisfiability threshold $\alpha > 4.15$, where the 1RSB result for the total complexity function is believed to be correct (stable) [MPRT04].

Formally, connected-components clusters have no reason to be equivalent to the notion of pure states. They are not able to reproduce purely entropic separation between clusters, which might exist in models like 3-SAT. However, fig. 2.2 suggests that there is more in this definition than it might seem at a first glance.

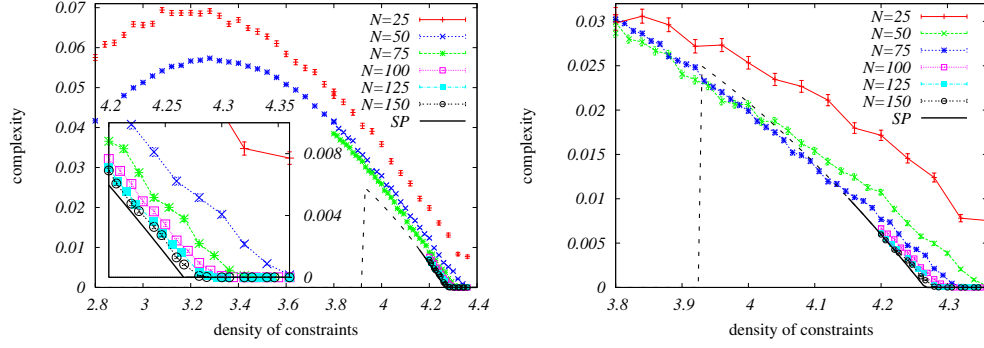


Fig. 2.2. (Color online) Right: Complexity of the connected-components clusters. Left: Complexity of the whitening-core clusters. Both compared to the complexity computed from the survey propagation equations. The data for the SP complexity are courtesy of Stephan Mertens, from [MMZ06].

Whitening-core clusters — We define the *whitening of a solution* as iterations of the warning propagation equations (1.35) initialized in the solution. The fixed point is then called the *whitening core*. Note, that the whitening core is well defined in the sense that the fixed point of the warning propagation initialized in a solution does not depend on the order in which the messages were updated. A whitening core is called trivial if all the warning messages are 0, that is “I do not care”.

The 1RSB equations at $m = 0$, which give the total complexity function, can be derived as belief propagation counting of all possible whitening cores [MZ02, BZ04, MMW07]. Thus another reasonable definition of clusters is that two solutions belong to the same cluster if and only if their whitening core is identical. In fig. 2.2 left we plot numerically computed complexity of the whitening-core clusters compared to the complexity computed from (2.28) at $m = 0$. The agreement is again good, in particular near to the satisfiability threshold, $\alpha > 4.15$, where the SP gives a correct result.

The $d = 1$ connected-components clusters share the property that all the solution from one clusters have the same whitening core. Proof: If this would not be true then there have to exist a pair of solutions which do not have the same whitening core but differ in only one variable, this is not possible because then the whitening could be started in that variable.

Note, however, that the definition of whitening-core clusters put all the solutions with a trivial whitening core into one cluster. This is not correct as, at least near to the clustering threshold, there are many pure states with a trivial whitening core. This is closely connected to the properties of frozen variables which will be discussed in chapter 4.

Enumeration of clusters in 3-SAT: the numerical method — In order to obtain the data in fig. 2.2 we generate instances of the random 3-SAT problem with N variables and M clauses, constraint density is then $\alpha = M/N$. We count number of solutions in $A = 999$ random instances and choose the median one where we count the number of connected-components and whitening-core clusters \mathcal{S} . This is repeated $B = 1000$ times. The average complexity is then computed as $\Sigma = \sum_{i=1}^B \log \mathcal{S}_i / (BN)$, if the median instance was unsatisfiable then we count

zero to the average, that is if all the B instances are unsatisfiable then the complexity is zero. We do such a non-traditional sampling to avoid rare instances with very many solutions, which we would not be able to cluster.

2.3 Physical properties of the clustered phase

Let us give a summary of the properties of the clustered phase, also called the dynamical 1RSB phase. We describe only the situation when $\Sigma(m = 1) > 0$ (2.28), when the opposite is true the properties are completely different as we will discuss in the next chapter 3.

The complexity function computed from (2.28) is the log-number of clusters of a given internal entropy. If a solution is chosen uniformly at random it will almost surely belong to a cluster with entropy s^* such that $\Sigma(s) + s$ is maximized in s^* , $\partial_s \Sigma(s^*) = -1$, that is $m = 1$. At $m = 1$ the total entropy $\Sigma(s^*) + s^* = \Phi(m = 1)$. The replicated free entropy at $\Phi(m = 1)$ is equal to the replica symmetric entropy. Thus the total entropy in the dynamical 1RSB phase is equal to the RS entropy. Also the marginal probabilities at $m = 1$ are equal to the replica symmetric ones

$$\int dP^{i \rightarrow j}(\psi^{i \rightarrow j}) \psi_{s_i}^{i \rightarrow j} = (\psi_{\text{RS}})^{i \rightarrow j}_{s_i} \quad \text{if } m = 1. \quad (2.34)$$

Thus the clustering transition is not a phase transition in the Ehrenfest sense, because the thermodynamical potential, entropy in our case, is analytical at the transition.

The overlap (or here distance) distribution, which is often used to describe the spin glass phase, is also trivial and equal to the replica symmetric one in the dynamical 1RSB phase. Indeed, if exponentially many clusters are needed to cover almost all solutions, then the probability that two solutions happen to belong to the same cluster is zero.

The correlation function between two variables at a distance (shortest path in the graph) d is defined as $\langle s_i s_j \rangle_c = \|\mu(s_i, s_j) - \mu(s_i)\mu(s_j)\|_{\text{TV}}$. The variance of the overlap distribution, which is negligible compared to 1 as we explained, can be expressed as $\sum_{i,j} \langle s_i s_j \rangle_c^2 / N^2$, and thus the two-point correlation have to decay faster with distance than the number on vertices at that distance is growing. This means in particular that two neighbours of a node i are independent if we condition on the value of i , this is again consistent with the fact that the belief propagation equations predict correct total entropy and marginal probabilities.

So far nothing is different from the replica symmetric phase. It is thus not straightforward to recognize the dynamical 1RSB phase based on the original replica computation. Presence of this phase was discovered and discussed in [KT87a, KT87b]. Later purely static methods were developed to identify this phase. The most remarkable is perhaps the ϵ -coupling and the "potential" of [FP95, FP97].

In our setting the main difference between the replica symmetric phase and the dynamical 1RSB phase is that in the later the point-to-set correlations do not decay to zero. Consequently the equilibration time of the local Monte Carlo dynamics diverges and Monte Carlo sampling becomes difficult [MS06b].

2.4 Is the clustered phase algorithmically hard?

Clustering has important implications for the dynamical behaviour. It slows down the equilibration and thus uniform sampling of solutions via local single spin flip Monte Carlo is not possible,




Fig. 2.3. (Color online) The performance of the ASAT algorithm in the 4-coloring of random Erdős-Rényi graphs. Left: The energy density plotted against the number of steps per variable. Right: The average running time (per variable) as a function of the connectivity. The time does not diverge at the clustering transition c_d , but beyond it. The other phase transitions marked are the condensation transition c_c (chap. 3) the rigidity transition c_r (chap. 4) and the colorability threshold c_s

or exponentially slow, beyond the dynamical threshold. But finding one solution is a much simpler problem than sampling.

Analytic arguments — In the 3-coloring of Erdős-Rényi graphs the clustering threshold is $c_d = 4$, as at this point the spin glass susceptibility diverges, see appendix C. In the terms of the reconstruction problem the Kesten-Stigum [KS66a, KS66b] bound is sharp. On the other hand Achlioptas and Moore [AM03] proved that a simple heuristic algorithm is able to find a solution in average polynomial time up to at least $c = 4.03$. This shows that the RSB phase is not necessarily hard.

A similar observation was made in the 1-in-3 SAT problem in [RSZ07]. There is a region in the values of the average density of constraints and the probability of negating a variable in a clause in which the replica symmetric solution is unstable and yet the unit clause propagation algorithm with the short clause heuristics was proven to find a solution in polynomial average time.

We should mention a common contra-argument; which is that in the above mentioned regions the 1RSB approach might not be correct, and the presumably full-RSB phase [Par80c] is more “transparent” for the dynamics of algorithms, see e.g. [MRT04]. However, at least in the 3-coloring, the 1RSB approach seems to be correct in the interval in question, as we argue in appendix D.

Stochastic local search — There is a lot of numerical evidence that relatively simple single spin flip stochastic local search algorithms are able to find solutions in linear time deep in the clustered region. Examples of works where performance of such algorithms was analyzed are [KK07, SAO05, AA06, AAA⁺08]. In fig. 2.3 we give an example of performance of the ASAT algorithm [AA06] in 4-coloring of Erdős-Rényi random graphs [ZK07]. The algorithm

is described in appendix F.2.2. In the 4-coloring ASAT is able to find solutions in linear time beyond the clustering transition $c_d = 8.35$

Simulated annealing — There is no paradox in the observations above. Quantitative statements are, however, difficult to make. Let us describe on an intuitive level the behaviour of an algorithm (dynamics) which satisfies the detailed balance condition and thus in infinite time samples uniformly from the uniform measure (2.1). We think for example about the simulated annealing [KGV83]. Above the dynamical temperature T_d corresponding to an energy E_d the point-to-set correlation function (2.22) decay fast and thus simulated annealing is able to reach the equilibrium. Below temperature T_d this is not the case anymore and the dynamics is stuck for a very long time in one of the clusters, states. But the bottom of this state E_{bottom} lies lower than E_d , thus when lowering the temperature the average energy seen by the simulated annealing also decreases. If $E_{\text{bottom}} = 0$ then the algorithm will find a solution. It is not known how to compute E_{bottom} in general. Sometimes, far from the clustering transition, the *iso-complexity* approach [MRT04] gives a lower bound on E_{bottom} . But in general, as far as we know, there is no argument saying $E_{\text{bottom}} > 0$. This picture can be substantiated for several simple models as the spherical p -spin model [CK93] or the random subcubes model [MZ08]. The connection with the optimization problems was remarked in [KK07].

For the stochastic local search algorithm, which does not satisfy the detailed balanced condition, the situation might be similar. At a point the algorithm is stuck in a cluster, but if this cluster goes down to the zero energy then it might be able to find solutions even in the clustered phase.

However, the current understanding of the dynamics of the mean field glassy systems is far from complete. More studies are needed to understand better the link between the static clustered phase and the dynamical behaviour.

3 Condensation

In this chapter we will describe the so-called condensed clustered phase. Before turning to the models of our interest we present the random subcubes model [MZ08], where the condensation of clusters can be understood on a very elementary probabilistic level. After mentioning that the condensed phase is in fact very well known in spin glasses we describe the Poisson-Dirichlet process which determines the distribution of sizes of clusters in that phase. Further, we discuss general properties of the condensed phase in random CSPs. And finally we address our original question and conclude that the condensation is not much significant for the hardness of finding a solution [ZK07].

3.1 Condensation in a toy model of random subcubes

The random-subcubes model [MZ08] is defined by its solution space $S \subseteq \{0, 1\}^N$; we define S as the union of $\lfloor 2^{(1-\alpha)N} \rfloor$ random clusters (where $\lfloor x \rfloor$ denotes the integer value of x). A random cluster A being defined as:

$$A = \{\sigma \mid \forall i \in \{1, \dots, N\}, \sigma_i \in \pi_i^A\}, \quad (3.1)$$

where π^A is a random mapping:

$$\pi^A : \{1, \dots, N\} \longrightarrow \{\{0\}, \{1\}, \{0, 1\}\}, \quad (3.2)$$

$$i \longmapsto \pi_i^A, \quad (3.3)$$

such that for each variable i , $\pi_i^A = \{0\}$ with probability $p/2$, $\{1\}$ with probability $p/2$, and $\{0, 1\}$ with probability $1 - p$. A cluster is here a random subcube of $\{0, 1\}^N$. If $\pi_i^A = \{0\}$ or $\{1\}$, variable i is said “frozen” in A ; otherwise it is said “free” in A . In this model one given configuration σ might belong to zero, one or several clusters.

We describe the static properties of the set of solutions S in the random-subcubes model in the thermodynamic limit $N \rightarrow \infty$ (the two parameters $0 \leq \alpha \leq 1$ and $0 \leq p \leq 1$ being fixed and independent of N). The internal entropy s of a cluster A is defined as $\frac{1}{N} \log_2 |A|$, i.e., the fraction of free variables in A . The probability $\mathcal{P}(s)$ that a cluster has internal entropy s follows the binomial distribution

$$\mathcal{P}(s) = \binom{N}{sN} (1-p)^{sN} p^{(1-s)N}. \quad (3.4)$$

Then the number of clusters of entropy s , denoted $\mathcal{N}(s)$, is with high probability

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log_2 \mathcal{N}(s) = \begin{cases} \Sigma(s) \equiv 1 - \alpha - D(s \parallel 1 - p) & \text{if } \Sigma(s) \geq 0, \\ -\infty & \text{otherwise,} \end{cases} \quad (3.5)$$

where $D(x \parallel y) \equiv x \log_2 \frac{x}{y} + (1-x) \log_2 \frac{1-x}{1-y}$ is the binary Kullback-Leibler divergence.

We compute the total entropy $s_{\text{tot}} = \frac{1}{N} \log_2 |S|$. First note that a random configuration belongs on average to $2^{N(1-\alpha)} (1 - \frac{p}{2})^N$ clusters. Therefore, if

$$\alpha < \alpha_d \equiv \log_2 (2 - p), \quad (3.6)$$

then with high probability the total entropy is $s_{\text{tot}} = 1$.

Now assume $\alpha > \alpha_d$. The total entropy is given by a saddle-point estimation:

$$\sum_A 2^{s(A)N} = [1 + o(1)]N \int_{\Sigma(s) \geq 0} ds 2^{N[\Sigma(s)+s]}, \quad (3.7)$$

$$\text{whence } s_{\text{tot}} = \max_s [\Sigma(s) + s \mid \Sigma(s) \geq 0]. \quad (3.8)$$

We denote by $s^* = \operatorname{argmax}_s [\Sigma(s) + s \mid \Sigma(s) \geq 0]$ the fraction of free variables in the clusters that dominate the sum. Note that our estimation is valid (there is no double counting) since in every cluster the fraction of solutions belonging to more than one cluster is exponentially small as long as $\alpha > \alpha_d$.

Define $\tilde{s} \equiv 2(1-p)/(2-p)$ such that $\partial_s \Sigma(\tilde{s}) = -1$. The complexity of clusters with entropy \tilde{s} reads:

$$\Sigma(\tilde{s}) = \frac{p}{2-p} + \log_2(2-p) - \alpha. \quad (3.9)$$

\tilde{s} maximizes eq. (3.8) as long as $\Sigma(\tilde{s}) \geq 0$, that is if

$$\alpha \leq \alpha_c \equiv \frac{p}{(2-p)} + \log_2(2-p). \quad (3.10)$$

Then the total entropy reads

$$s_{\text{tot}} = 1 - \alpha + \log_2(2-p) \quad \text{for } \alpha \leq \alpha_c. \quad (3.11)$$

For $\alpha > \alpha_c$, the maximum in (3.8) is realized by the largest possible cluster entropy s_{max} , which is given by the largest root of $\Sigma(s)$. Then $s_{\text{tot}} = s^* = s_{\text{max}}$. We will show in the next section that in such a case almost all solutions belong to only a finite number of largest clusters. This phase is thus called *condensed*, in the sense that almost all solutions are "condensed" in a small number of clusters.

In summary, for a fixed value of the parameter p , and for increasing values of α , four different phases can be distinguished:

- (a) Liquid (replica symmetric) phase, $\alpha < \alpha_d$: almost all configurations are solutions.
- (b) Clustered (dynamical 1RSB) phase with many states, $\alpha_d < \alpha < \alpha_c$: an exponential number of clusters is needed to cover almost all the solutions.
- (c) Condensed clustered phase, $\alpha_c < \alpha < 1$: a finite number of the biggest clusters covers almost all the solutions.
- (d) Unsatisfiable phase, $\alpha > 1$: no cluster, hence no solution, exists.

3.2 New in CSPs, well known in spin glasses

The complexity function $\Sigma(s)$ (2.26) in random CSPs is counting the logarithm of the number of clusters per variable which have internal entropy s per variable. We define *dominating clusters* in the same way as in the random subcubes model, that is clusters of entropy s^* such that

$$s^* = \arg \max_{s, \Sigma(s) > 0} [\Sigma(s) + s]. \quad (3.12)$$




Fig. 3.1. (Color online) The complexity function in the random subcubes model, $\Sigma(s)$ (3.5), for $p = 0.8$ and several values of α . The red dots mark the dominating clusters s^* , $\Sigma(s^*)$. For $p = 0.8$ the dynamical transition $\alpha_d \approx 0.263$ is far away from the plotted values, the condensation transition is $\alpha_c \approx 0.930$, the satisfiability $\alpha_s = 1$.

In chap. 2 we discussed properties of the dynamical 1RSB phase, that is when $\Sigma(s^*) > 0$, in other words when there are exponentially many dominating clusters.

The condensed phase with $\Sigma(s^*) = 0$, described in the random subcubes model, exists also in random CSPs. And in the context of constraint satisfaction problems it was first computed and discussed in [MPR05] and [KMRT⁺07]. However, historically it was the condensed phase where the 1RSB solution was first worked out [Par80c]. A very simple example of condensation can also be found in the random energy model [Der80, Der81]. As we discussed in the previous chapter 2, the dynamical 1RSB phase is well hidden within the replica solution — the total entropy is equal to the replica symmetric entropy, the overlap distribution is trivial and the two-point correlation functions decay to zero etc. All this changes in the condensed phase.

A small digression to the physics of glasses: In structural glasses, the analog of the condensation transition is well known for a long time, its discovery goes back to Kauzmann in 1948 who studied the configurational entropy of glassy materials. Configurational entropy is the difference between the total (experimentally measured) entropy and the entropy of a solid material, this thus corresponds to the complexity function. In the so called fragile structural glasses [Ang95] the extrapolated configurational entropy becomes zero at a positive temperature, nowadays called the Kauzmann temperature. The Kauzmann temperature in the real glasses is, however, only extrapolation. The equilibration time in glasses exceeds the observation time high above the Kauzmann temperature. It is a widely discussed question if there exists a true phase transition at the Kauzmann temperature or not, for a recent discussion see [DS01].

Why does Parisi *maximize* the replicated free energy? As we said, it is the condensed phase which was originally described by Parisi and his one-step replica symmetry breaking solution [Par80c]. Let us now briefly clarify the relation to the replica solution, similar reasoning first appeared in [Mon95]. In sec. 2.1 we called the Legendre transform of the complexity function the replicated free entropy $\Phi(m)$ (2.26). In the replica approach the replicated entropy $\Omega(m) = \Phi(m)/m$ is computed. From (2.28) follows

$$\Omega(m) = s + \frac{\Sigma(s)}{m} \quad \text{where} \quad \frac{\partial \Omega(m)}{\partial m} = -\frac{\Sigma(s)}{m^2}. \quad (3.13)$$

Thus, in the condensed phase, computing the largest root of the function $\Sigma(s)$, in order to maximize the total entropy, is equivalent to extremizing the replicated entropy $\Omega(m)$. Moreover, as the function $\Sigma(s)$ is concave and the parameter m is minus its slope this extrema have to be a *minima*. Thus in the Parisi's replica solution we have to minimize the replicated entropy function with respect to the parameter m . If a temperature is involved then this becomes a maximization of the replicated free energy, this might have seem contra-intuitive in the original solution, but it comes out very naturally in our approach. Other physical interpretation of the maximization was proposed e.g. in [Jan05].

3.3 Relative sizes of clusters in the condensed phase

What is the number of dominating clusters in the condensed phase and what are their relative sizes? So far we know that the entropy per variable of the dominating states is $s^* + o(1)$ and that their number is sub-exponential, $\Sigma(s^*) = 0$. But much more can be said based on purely probabilistic considerations.

Consider that the total number of clusters \mathcal{N} is exponentially large in the system size N , and that $N \rightarrow \infty$. Let the log-number of clusters of a given entropy be distributed according to an analytic function $\Sigma(s)$. Denote $-m^* = \partial_s \Sigma(s^*)$, in the condensed phase $0 < m^* < 1$. Denote the size of the α^{th} largest cluster $e^{Ns^* + \Delta_\alpha}$, $\Delta_\alpha = O(1)$. The probability that there is a cluster of size between $e^{Ns^* + \Delta}$ and $e^{Ns^* + \Delta + d\Delta}$, $\Delta \gg d\Delta$, is $e^{-m^* \Delta} d\Delta$, in other words points Δ_α are constructed from a Poissonian process with rate $e^{-m^* \Delta}$ ¹. Relative size of the α^{th} largest cluster is defined as

$$w_\alpha = \frac{e^{\Delta_\alpha}}{\sum_{\gamma=1}^{\mathcal{N}} e^{\Delta_\gamma}}. \quad (3.14)$$

Point process w_α which is constructed as described above is in mathematics called the Poisson-Dirichlet process [PY97]. The connection between this process and the relative weights of states in the mean field models of spin glasses was (on a non-rigorous level) understood in [MPV85], for more mathematical review see [Tal03]².

Any moment of any w_α can be computed from the generating function [PY97]

$$\mathbb{E}[\exp(-\lambda/w_\alpha)] = e^{-\lambda} \phi_{m^*}(\lambda)^{\alpha-1} \psi_{m^*}(\lambda)^{-\alpha}, \quad (3.15)$$

¹Note that in the random subcubes model the numbers $(Ns^* + \Delta_\alpha) \log(2)$ are integers equal to the number of free variables in the cluster A_α . Then Δ_α are discrete and some of the properties of the resulting process might be different from the Poisson-Dirichlet.

²To avoid confusion, note that the Poisson-Dirichlet process we are interested in is the $\text{PD}(m^*, 0)$ in the notation of [PY97]. In the mathematical literature, it is often referred to the $\text{PD}(0, \theta)$ without indexing by the two parameters.




Fig. 3.2. (Color online) The fractions of solutions covered by the largest clusters as a function of parameter m^* . The lower curve is related to the size of the largest clusters as $1/\mathbb{E}[1/w_1] = 1 - m^*$. The following curves are related to the size of i largest clusters, their distances are $\mathbb{E}[R_\alpha]\mathbb{E}[R_{\alpha-1}] \dots \mathbb{E}[R_1](1 - m^*)$.

where $\lambda \geq 0$ and the functions ϕ_{m^*} and ψ_{m^*} are defined as

$$\phi_{m^*}(\lambda) = m^* \int_1^\infty e^{-\lambda x} x^{-1-m^*} dx, \quad (3.16a)$$

$$\psi_{m^*}(\lambda) = 1 + m^* \int_0^1 (1 - e^{-\lambda x}) x^{-1-m^*} dx. \quad (3.16b)$$

The second moments can be used to express the average probability Y that two random solutions belong to the same cluster

$$Y = \mathbb{E} \left[\sum_{\alpha=1}^{\mathcal{N}} w_\alpha^2 \right] = 1 - m^*. \quad (3.17)$$

This was originally derived in [MPV85]

Another useful relation [PY97] is that the ratio of two consequent points $R_\alpha = w_{\alpha+1}/w_\alpha$, $\alpha = 1, 2, \dots, \mathcal{N}$ is distributed as $\alpha m^* R_\alpha^{\alpha m^* - 1}$. In particular its expectation is $\mathbb{E}[R_\alpha] = \alpha m^* / (1 + \alpha m^*)$ and the random variables R_α are mutually independent. We used these relation to obtain data in figure 3.2.

From the properties of the Poisson-Dirichlet process, it follows that an arbitrary large fraction of the solutions can be covered by a finite number of clusters. When m^* is near to zero, that is near to the satisfiability threshold, the largest cluster covers a large fraction of solutions. On the other side, when m^* is near to one, that is near to the condensation transition, very many (but finite in N) clusters are needed to cover a given fraction of solutions.

3.4 Condensed phase in random CSPs

The total entropy in the condensed phase is strictly smaller than the replica symmetric entropy, $s_{\text{tot}} = s^* < s_{\text{RS}}$. At the condensation transition c_c the total entropy is non-analytic, it has a discontinuity in the second derivative. This can be seen easily for example from the expressions for the random subcubes model. At a finite temperature the discontinuity in the second derivative of the free energy corresponds to a jump in the specific heat. The parameter $m^* = 1$ at the condensation transition and decreases monotonously to $m^* = 0$ at the satisfiability threshold.

Concept of self-averaging — In the physics of disordered systems the self-averaging is a crucial concept. We say that quantity A measured on a system (graph) of N variables is self-averaging if in the limit $N \rightarrow \infty$

$$\frac{\mathbb{E}(A^2) - [\mathbb{E}(A)]^2}{[\mathbb{E}(A^2)]} \rightarrow 0, \quad (3.18)$$

where the average $\mathbb{E}[\cdot]$ is over all the disorder in the system. In other words a quantity is self-averaging if its value on a typical large system is equal to the average value. By computing the average value we thus describe faithfully the typical large system. And also measuring A on a single large system is enough to represent the whole ensemble. On finite-dimensional lattices and off criticality extensive quantities are always self-averaging. This can be shown by building the large lattice from smaller blocks, the additivity of an extensive quantity and the central limit theorem then ensures the self-averaging. At the critical point, on a mean field lattice (fully connected or tree-like) or for non-extensive quantities the answer whether A is self-averaging or not becomes nontrivial.

In the condensed phase quantities which involve the weights of clusters are not self-averaging. This arises from the fact that the dominating clusters are different in every realization of the system. Statistical properties of many quantities of interest can be described from the Poisson-Dirichlet process.

Overlap distribution — The overlap between two solutions is defined as one minus the Hamming distance

$$q(\{s\}, \{s'\}) = \frac{1}{N} \sum_{i=1}^N \delta(s_i, s'_i). \quad (3.19)$$

The overlap between two solutions belonging to two different dominating clusters is q_0 , and between two solutions belonging to the same dominating cluster q_1 . Values q_0 and q_1 are self-averaging. The distribution of overlaps in the limit $N \rightarrow \infty$ can thus be written as

$$P(q) = w \delta(q - q_1) + (1 - w) \delta(q - q_0), \quad (3.20)$$

where the weight w is the probability that two random solutions belong to the same cluster. Thus $w = \sum_{\alpha=1}^{\mathcal{N}} w_{\alpha}^2$, where w_{α} are weights of the clusters (3.14) given by the Poisson-Dirichlet process. The weights change from realization to realization, w is thus not a self-averaging quantity, its typical value fluctuates around the mean $\mathbb{E}(w) = 1 - m^*$ computed in (3.17). The distribution of the random variable w is also known [MPS⁺84].

Two-point correlation functions — The variance of the overlap distribution is

$$\text{var } q = \int q^2 P(q) dq - \left[\int q P(q) dq \right]^2 = w(1-w)(q_1 - q_0)^2. \quad (3.21)$$

At the same time the variance is equal to

$$\begin{aligned} \text{var } q &= \frac{1}{N^2} \sum_{i,j} \sum_{s_i, s_j} |\mu(s_i, s_j) - \mu(s_i)\mu(s_j)| \\ &\approx \frac{1}{N} \sum_i \sum_{s_i, s_0} |\mu(s_i, s_0) - \mu(s_i)\mu(s_0)|, \end{aligned} \quad (3.22)$$

where s_0 is a typical variable in the random graph. If we consider that the two-point correlation function is of order one up to a correlation length ξ and zero after that we get

$$\text{var } q \approx \frac{1}{N} c^\xi, \quad (3.23)$$

where c is approximately the branching factor. In the condensed phase the variance of the overlap is of order one thus the correlation length has to be of order $\log N$. But the shortest path between two random variables is also of order $\log N$ thus the two-point correlations cannot be neglected in the condensed phase.

If two-point correlations cannot be neglected then the derivation of belief propagation equations (1.16a-1.16b) is not valid, because we supposed that the neighbours of a node i are independent when we condition on the value of i . It is thus not surprising that the value to which the BP equations converge (if they do), does not correspond to the true marginal probability. Formally, the BP fixed point corresponds to the 1RSB equations at $m = 1$, but in the condensed phase $m^* < 1$.

In fact, the probability distribution of the true marginal probabilities is another example of a non self-averaging quantity. It again depends on the realization of the Poisson-Dirichlet process.

3.5 Is the condensed phase algorithmically hard?

From the algorithmic point of view the only important difference between the dynamical 1RSB phase and the condensed phase is that in the condensed phase the belief propagation does not estimate correctly the asymptotic marginal probabilities. In the condensed phase, the total entropy cannot be estimated from the BP equations either, thus approximative counting and sampling of solutions will probably be even harder than in the dynamical 1RSB phase.

Concerning the hardness of finding a solution we might expect that the incorrectness of the belief propagation estimates of marginals will play a certain role. However, we used the belief propagation maximal decimation as described in appendix F.1.2 in the 3- and 4-coloring, see fig. 3.3. And this algorithm does not seem to have any problem to pass the condensation transition in both these cases. In particular, in the 3-coloring the gap between the condensation threshold $c_c = 4$ and the limit of performance of the BP decimation $c \approx 4.55$ is huge. The rigidity transition c_r , defined in chapter 4, and the colorability threshold c_s are also marked for comparison in fig. 3.3.

The condensation transition thus does not seem to play any significant role for the computational hardness of finding a solution.




Fig. 3.3. (Color online) The performance of the maximal BP decimation algorithm, described in appendix F.1.2, in the 3-coloring (left) and the 4-coloring (right) of random graphs. This algorithm is able to color random graphs beyond both the clustering c_d and the condensation c_c transitions in 3- and 4-coloring.

4 Freezing

The previous two chapters describe recent contributions to the understanding of the clustering and condensation of solutions in random constraint satisfaction problems. Both these concepts are well known and widely discussed in the mean field theory of glasses and spin glasses for at least a quarter of a century.

The concept of freezing of variables appeared in the studies of optimization problems, that is systems at zero temperature (or infinite pressure). In this chapter we first define the freezing of variables, clusters and solutions, and discuss its properties both in the thermodynamical limit and on finite-size instances. Then we explain how to describe the frozen variables within the one-step replica symmetry breaking approach and we define several possible phase transition associated to the freezing. To simplify the picture we define and solve the "completely frozen" locked constraint satisfaction problem where every cluster contains only one configuration. Finally we give several arguments about connection between the freezing and the average computational hardness. Results of this section are mostly original and were published in [ZK07, KZ08a, AZ08, ZM08].

4.1 Frozen variables

Consider a set of solutions \mathcal{S} of a given instance of a constraint satisfaction problem. Define that a variable i is *frozen* in the set of solutions $A \subset \mathcal{S}$ if it is assigned the same value in all the solutions in the set. If an extensive number of variables is frozen in the set A , then we call A and all the solutions in A *frozen*, otherwise A and all the solutions in A are called *soft* (unfrozen).

A first observation is that the set of all solutions \mathcal{S} is not frozen in the satisfiable phase. If it would be then adding one constraint, i.e., increasing the constraint density by $1/N$, would make the formula unsatisfiable with a finite probability, that would be in a contradiction with the sharpness of the satisfiability threshold. The *backbone* is made of variables frozen in the set of ground states. An extensive backbone can thus exist only in the unsatisfiable phase. Already in [MZK⁺99b] it was argued that there might be a connection between the backbone and the computational hardness of the problem. The suggestion of [MZK⁺99b] was that if the fraction of variables covered by the backbone is discontinuous at the satisfiability transition then it is hard to find satisfying assignments on highly constrained but still satisfiable instances. On the other hand if the backbone appears continuously the problem is easy in the satisfiable phase. This was based on the replica symmetric solution of the random K -SAT which does not describe fully the phase space, in spite of that the relation between the existence of frozen variables inside clusters and the algorithmical hardness seems to be deep and we will develop it in this chapter.

4.1.1 Whitening: A way to tell if solutions are frozen

How to recognize if clusters have frozen variables or not. Or how to recognize if a given solution belongs to a frozen cluster or not. An iterative procedure called *whitening* [Par02a] gives an answer to these questions.

Given a formula of a CSP and one of its solutions $\{s_i\} \in \{-1, 1\}^N$, $i = 1, \dots, N$, the whitening of the solution is defined as iterations of the warning propagation equations (1.35) initialized on the solution. That is, for a binary CSP $h_{\text{init}}^{i \rightarrow a} = s_i$, and $u_{\text{init}}^{a \rightarrow i}$ is computed according

to eq. (1.35b). Note that the fixed point of the whitening does not depend on the order in which the warnings are updated. Indeed, during the iterations the only changes in warnings are from non-zero values to zero values. The fixed point is called the *whitening core* of the solution. The whitening core is called *trivial* if all the warnings are equal to 0, and *nontrivial* otherwise.

In the K -SAT problem whitening can be reformulated in a very natural way: Start with the solution $\{s_i\}$, assign iteratively a "*" (joker) to variables which belong only to clauses which are already satisfied by another variable or already contain a * variable. On a general CSP such procedure is not equivalent to the whitening, and the warning propagation definition has to be used instead in order to obtain all the desired properties and relations to the 1RSB solution.

We now argue that if the 1RSB solution is correct, then frozen variables in the cluster, to which solution $\{s_i\}$ belongs, asymptotically correspond to variables for which in the whitening core the total warning $h^i \neq 0$ (1.37). Thus whitening can be used to decide if the solution $\{s_i\}$ belongs to a frozen cluster without knowing all the solutions in that cluster. The first step to show this property is, as in sec. 2.1.1, to consider the CSP on a tree with given boundary conditions which are compatible with a non-empty set of solutions \mathcal{S} in the interior of the tree. Starting on the leaves we compute iteratively the warnings (1.35) down to the root. Variables which have at least one non-zero incoming warning are frozen in the set \mathcal{S} . The correctness of the 1RSB approach on a tree-like graph means that the picture on a tree captures properly all the asymptotic properties. In particular, the whitening core determines the set of frozen variables on typical large instances of the problem. The correctness of the 1RSB solution is an essential assumption for the above statement. Because all the long-range correlations decay within one cluster the warnings $u^{a \rightarrow i}$ in the whitening core are independent in the absence of i . Thus there truly exist solutions in that cluster in which the variable i takes all the values allowed by the warnings. And on the other hand, if a value is not allowed by the warnings there is no solution where i would be taking this value. For consistency, all solutions in one cluster have to have the same whitening core. However, two different clusters can have the same whitening core. The most important example are all the soft (not frozen) clusters that all have the trivial whitening core.

Whitening, as the iterative fixed point of the warning propagation, may be defined not only for a solution but for any configuration. In this way one may find blocking metastable states. For some preliminary numerical considerations see [SAO05].

4.1.2 Freezing on finite size instances

The definition of whitening is applicable to any (non-random, small, etc.) instance. What does then remain from the asymptotic correspondence between frozen variables and whitening cores?

Consider now clusters as connected components in the graph where all solutions are nodes and where edges are between solutions which differ in only one variable, as in sec. 2.2. Several questions arise about this definition:

- Do all the solutions in the connected-components cluster have the same whitening core? The answer is yes. If there were two solutions with different whitening cores which can be connected by a chain of single-variable flips, then along this chain there would exist a pair of solutions which differ in only one variable i and have different whitening cores. But this is not possible, as the fixed point of the whitening does not depend on the order in which

the warnings were updated, and one could thus start the whitening by setting warnings $h^{i \rightarrow a} = 0$.

- Does the whitening core of a connected-components cluster correspond to the set of frozen variables? The answer is: If in the whitening core $h^i \neq 0$ (1.37) then the variable i is frozen in the connected-components cluster. Proof: If such a variable i is not frozen, then there have to exist a pair of solutions which differ only in the value of this variable. Then all the constraints around i have to be compatible with both these values, this would be in contradiction with $h^i \neq 0$. On the other hand, if in the whitening core $h^i = 0$ then the variable i might still be frozen in the connected-components cluster on a general instance, because correlations which are not considered by the 1RSB solution may play a role.

Consider now clusters as the set of all solutions which share the same whitening core. Whitening-core clusters are aggregations of the connected-components clusters. In particular, all the solutions with a trivial whitening core, which might correspond to exponentially many pure states, are put together.

- What is the set of frozen variables in the whitening-core clusters? The answer is: Again if in the whitening core $h^i \neq 0$ then the variable i is frozen in the whitening-core cluster. In principle, one whitening-core cluster could be an union of several connected-components cluster, but i is frozen to the same value in each of them. The inverse is not correct in general. On finite size instances some variables with a zero warning $h^i = 0$ might be frozen in the whitening-core cluster.
- Can there be a fixed point of the warning propagation (1.35) corresponding to zero energy (1.38) which is not compatible with any solution? The answer is yes. And such fixed points were observed in [BZ04, MMW07, KSS07b]. Again if the 1RSB solution is correct then in the thermodynamical limit these "fake" fixed points are negligible.

4.1.3 Freezing transition in 3-SAT - exhaustive enumeration

Before turning to the cavity description of frozen clusters we investigate the *freezing transition* in the random 3-SAT numerically. We define the freezing transition, α_f , as the smallest density of constraints α such that the whitening core of all solutions is nontrivial, i.e., not made only from zero warnings. We use the whitening core in the definition instead of the real set of frozen variables, because it does not depend on the definition of clusters and it has much smaller finite size effects. The existence of such a frozen phase was proven in the thermodynamical limit for $K \geq 9$ of the K -SAT near to the satisfiability threshold in [ART06].

In order to determine the freezing transition we start with a 3-SAT formula of N variables and all possible clauses, and remove the clauses one by one independently at random¹. We mark the number of clauses M_s where the formula becomes satisfiable as well as the number of clauses $M_f \leq M_s$ where at least one solution starts to have a trivial whitening core. We repeat B -times ($B = 2 \cdot 10^4$ in fig. 4.1) and compute the probabilities that a formula of M clauses is satisfiable $P_s(\alpha, N)$, and unfrozen $P_f(\alpha, N)$ respectively. Due to the memory limitation we could treat

¹In practice we do not start with all the clauses, but as many that in all the repetitions of this procedure the initial instance is unsatisfiable.

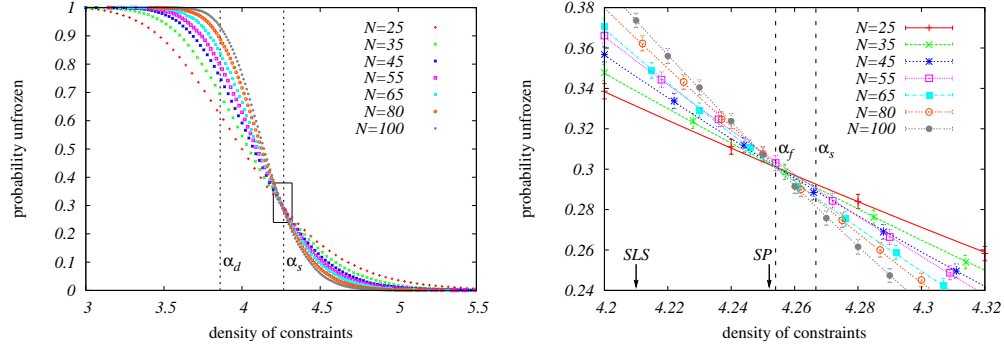


Fig. 4.1. (Color online) Left: Probability that there exists an unfrozen solution as a function of the constraint density α for different system sizes. The clustering [KMRT⁺07] and satisfiability [MPZ02] transitions marked for comparison. Right: A 1:20 zoom on the critical (crossing) point, our estimate for the freezing transition is $\alpha_f = 4.254 \pm 0.009$. The curves are cubic fits in the interval $\alpha \in (4, 4.4)$. The arrows represent estimates of the limits of performance of the best known local search ASAT [AA06] and survey propagation [Par03, CFMZ05] algorithms.

only instances which have less than $5 \cdot 10^7$ solutions which limits us to system sizes $N \leq 100$. The results for the satisfiability threshold are shown in fig. 1.3 and are consistent with previous studies in [KS94, MZK⁺99b, MZK⁺99a]. The probability of being unfrozen, $P_f(\alpha, N)$, is shown in fig. 4.1.

It is tempting to perform a scaling analysis as has been done in [KS94, MZK⁺99b, MZK⁺99a] for the satisfiability threshold. The critical exponent related to the width of the scaling window was defined via rescaling of the constraint density α as $N^{1/\nu_s} [1 - \alpha/\alpha_s(N)]$. Note, however, that the estimate $\nu_s = 1.5 \pm 0.1$ for 3-SAT provided in [MZK⁺99a] is not asymptotically correct. It was proven in [Wil02] that $\nu_s \geq 2$. Indeed, it was shown numerically in [LRTZ01] that a crossover exists at sizes of order $N \approx 10^4$ in the related XOR-SAT problem. A similar situation happens for the scaling of the freezing transition, $P_f(\alpha, N)$, as the proof of [Wil02] applies also here². It would be interesting to investigate the scaling behaviour on an ensemble of instances where the results of [Wil02] do not apply (e.g. graphs without leaves). However, we concentrate instead on the estimation of the critical point, which we do not expect to be influenced by the crossover in the scaling. We are in a much more convenient situation for the freezing transition than for the satisfiability one. The crossing point between functions $P_f(\alpha, N)$ for different system sizes seems to depend very little on N , while for the satisfiability transition it depends very strongly on N , compare the zooms in figs. 1.3 and 4.1.

We determine the value of the freezing transition in random 3-SAT as

$$\alpha_f = 4.254 \pm 0.009, \quad (4.1)$$

which is very near but seems separated from the satisfiability threshold $\alpha_s = 4.267$ [MZ02, MMZ06]. In any case the frozen phase in 3-SAT is very narrow, that is in contrast with the situation in $K \geq 9$ SAT where it covers at least $1/5$ of the large clustered phase [ART06].

²Theorem 1 of [Wil02] applies to the freezing property where the bystander are clauses containing two leaves.

4.2 Cavity approach to frozen variables

In this section we present how to describe the frozen variables within the 1RSB cavity solution. We illustrate the results on an example of the random graph coloring where properties of frozen variables were studied in detail for the first time [ZK07].

The energetic 1RSB (survey propagation), sec. 1.6-1.7, aims to count the total number of frozen clusters. More precisely, it counts the total number of fixed points of the warning propagation (1.35). It can be used to locate the satisfiability threshold or to design survey propagation based solvers [MPZ02, MZ02]. However, as we understood in chapter 2, by neglecting the soft clusters we cannot locate the clustering transition. In chapter 3 we defined the dominant clusters, i.e., those which cover almost all solutions. A natural question arises immediately: Are the dominant clusters frozen or soft? In order to answer the general entropic 1RSB equations (2.24, 2.28) have to be analyzed.

4.2.1 Frozen variables in the entropic 1RSB equations

We remind that in the 1RSB solution of the graph coloring problem the components of the messages (called also the cavity fields) $\psi_{s_i}^{i \rightarrow j}$ are the probabilities that in a given cluster the node i takes the color s_i when the constraint on the edge (ij) is not present. The belief propagation equations, (1.16) in general, (2.5) in coloring, then define the consistency rules between the field $\psi_{s_i}^{i \rightarrow j}$ and fields incoming to i from the other variables than j . In the zero temperature limit we can classify fields $\psi_{s_i}^{i \rightarrow j}$ in the following two categories:

- (i) The *hard (frozen) field* corresponds to the case when all components of $\psi^{i \rightarrow j}$ are strictly zero except the one for color s . This means that in the absence of edge (ij) , variable i takes color s in *all* the solutions from the cluster in question.
- (ii) The *soft field* corresponds to the case when more than one component of $\psi_{s_i}^{i \rightarrow j}$ is nonzero. The variable i is thus not frozen in the absence of edge (ij) , and the colors of all the nonzero components are allowed.

This distinction is also meaningful for the full probabilities $\psi_{s_i}^i$ (1.18). By definition, the variable i is frozen in the cluster if and only if $\psi_{s_i}^i$ is a hard field.

It is important to stress that some of the soft fields on a given instance of the problem might be very small. Some of them might even scale like e^{-N} . We insist on classifying those as the soft fields because they cannot create real contradictions. This subtle distinction becomes important mainly in the implementation of the population dynamics algorithm, see appendix E.

The distribution of fields over clusters $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ (2.24), which is the "order parameter" of the 1RSB equation, can be decomposed into the hard-field part of a weight $\eta_s^{i \rightarrow j}$ and the soft-field part $P_{\text{soft}}^{i \rightarrow j}$ of a weight $\eta_0^{i \rightarrow j} = 1 - \sum_{s=1}^q \eta_s^{i \rightarrow j}$

$$P^{i \rightarrow j}(\psi^{i \rightarrow j}) = \sum_{s=1}^q \eta_s^{i \rightarrow j} \mathbb{I}(\psi^{i \rightarrow j} \text{ frozen into } s) + \eta_0^{i \rightarrow j} P_{\text{soft}}^{i \rightarrow j}(\psi^{i \rightarrow j}). \quad (4.2)$$

Hard fields in the simplest case, $m=0$ — First, we derive equations for the hard fields when the parameter $m=0$ in (2.24). This will, in fact, lead to the survey propagation equations, for

coloring originally derived in [MPWZ02, BMP⁺03] from the energetic 1RSB method (1.6). For simplicity we write the most general form only for the 3-coloring.

We plug (4.2) into eq. (2.24). The reweighting factor $(Z^{i \rightarrow j})^m$ at $m = 0$ is either equal to zero, when the arriving fields are hard and contradictory, or equal to one. This is the origin of a significant simplification. The outgoing field $\psi^{i \rightarrow j}$ might be frozen in direction s if and only if for every other color $r \neq s$ there is at least one incoming field frozen to the color r . The update of probability $\eta_s^{i \rightarrow j}$ that a field is frozen in direction s is for the 3-coloring written as

$$\eta_s^{i \rightarrow j} = \frac{\prod_{k \in i-j} (1 - \eta_s^{k \rightarrow i}) - \sum_{p \neq s} \prod_{k \in i-j} (\eta_0^{k \rightarrow i} + \eta_p^{k \rightarrow i}) + \prod_{k \in i-j} \eta_0^{k \rightarrow i}}{\sum_p \prod_{k \in i-j} (1 - \eta_p^{k \rightarrow i}) - \sum_p \prod_{k \in i-j} (\eta_0^{k \rightarrow i} + \eta_p^{k \rightarrow i}) + \prod_{k \in i-j} \eta_0^{k \rightarrow i}}. \quad (4.3)$$

In the numerator there is a telescopic sum counting the probability that color s and only color s is not forbidden by the incoming fields. In the denominator there is the normalization, i.e., the telescopic sum counting the probability that there is at least one color which is not forbidden. The crucial observation is that at $m = 0$ the self-consistent equations for η do not depend on the soft-fields distribution $P_{\text{soft}}^{i \rightarrow j}(\psi^{i \rightarrow j})$.

If we do not aim at finding of a proper coloring on a single graph but just at computing of the complexity function and similar quantities, we can further simplify eq. (4.3) by imposing the color symmetry. Indeed, the probability that in a given cluster a field is frozen in the direction of a color s has to be independent of s . Then (4.3) becomes, now for general number of colors q :

$$\eta^{i \rightarrow j} = w(\{\eta^{k \rightarrow i}\}) = \frac{\sum_{l=0}^{q-1} (-1)^l \binom{q-1}{l} \prod_{k \in i-j} [1 - (l+1)\eta^{k \rightarrow i}]}{\sum_{l=0}^{q-1} (-1)^l \binom{q}{l+1} \prod_{k \in i-j} [1 - (l+1)\eta^{k \rightarrow i}]}. \quad (4.4)$$

We remind that since $\partial \Sigma(s)/\partial s = -m$ (2.28), the value $m = 0$ corresponds to the point \tilde{s} where the function $\Sigma(s)$ has a zero slope. If a nontrivial solution of (4.3) exists, then $\Sigma(\tilde{s})|_{m=0}$ is the maximum of the curve $\Sigma(s)$. And if the 1RSB solution for clusters at $m = 0$ is correct then it is counting the total log-number of clusters of size \tilde{s} , which is due to the exponential dependence also the total log-number of all clusters, regardless of their size.

Frozen variables at general m , generalized SP — Let us compute how the fraction of hard fields η evolves after one iteration of equation (2.24) at a general value of m . There are two steps in each iteration of (2.24). In the first step, η iterates via eq. (4.4). In the second step we re-weight the fields. Writing $P_m^{\text{hard}}(Z)$ the —unknown— distribution of the reweightings Z^m for the hard fields, one gets

$$\begin{aligned} \eta^{i \rightarrow j} &= \frac{1}{\mathcal{N}^{i \rightarrow j}} \int dZ^{i \rightarrow j} P_m^{\text{hard}}(Z^{i \rightarrow j}) (Z^{i \rightarrow j})^m w(\{\eta^{k \rightarrow i}\}) \\ &= \frac{w(\{\eta^{k \rightarrow i}\})}{\mathcal{N}^{i \rightarrow j}} \int dZ^{i \rightarrow j} P_m^{\text{hard}}(Z^{i \rightarrow j}) (Z^{i \rightarrow j})^m \\ &= \frac{w(\{\eta^{k \rightarrow i}\})}{\mathcal{N}^{i \rightarrow j}} \langle Z_m^{i \rightarrow j} \rangle_{\text{hard}}. \end{aligned} \quad (4.5)$$

A similar equation can formally be written for the soft fields

$$1 - q\eta^{i \rightarrow j} = \frac{1 - qw(\{\eta^{k \rightarrow i}\})}{\mathcal{N}^{i \rightarrow j}} \langle Z_m^{i \rightarrow j} \rangle_{\text{soft}}. \quad (4.6)$$

Writing explicitly the normalization $\mathcal{N}^{i \rightarrow j}$, we finally obtain the generalized survey propagation equations:

$$\eta^{i \rightarrow j} = \frac{w(\{\eta^{k \rightarrow i}\})}{qw(\{\eta^{k \rightarrow i}\}) + [1 - qw(\{\eta^{k \rightarrow i}\})] r(m, \{\eta^{k \rightarrow i}\})}, \quad (4.7)$$

where r is the ratio of average reweighting factors of the soft and hard fields

$$r(m, \{\eta^{k \rightarrow i}\}) = \frac{\langle Z_m^{i \rightarrow j} \rangle_{\text{soft}}}{\langle Z_m^{i \rightarrow j} \rangle_{\text{hard}}}. \quad (4.8)$$

In order to do this recursion, the only nontrivial information needed is the ratio r between soft- and hard-field average reweightings, which depends on the full distribution of soft fields $P_{\text{soft}}^{i \rightarrow j}(\psi^{i \rightarrow j})$. Eq. (4.7) is easy to use in the population dynamics and allows to compute the fraction of frozen variables in typical clusters of a given size (for a given value m).

There are two cases where eq. (4.7) simplifies so that the hard-field recursion becomes *independent* from the soft-field distribution. The first case is, of course, $m = 0$. Then $r = 1$ independently of the edge (ij) , and the equation reduces to the original SP. The second case arises for $m = 1$, where the eq. (4.7) can be written as the equation for the naive reconstruction (2.4). The probability that a variables is frozen at $m = 1$ is the same at the probability that leaves (far away variables) determine uniquely the root in the reconstruction problem, see sec. 2.1.1.

Frozen variables and minimal rearrangements — Montanari and Semerjian [MS05, Sem08] developed a very interesting connection between frozen variables and the so-called *minimal rearrangements*. Given a CSP instance, one of its solutions $\{s_i\}$ and a variable i , find the nearest solution to $\{s_i\}$ where the values of the variable i is changed to $s'_i \neq s_i$. The set of variables on which these two solutions differ is called the *minimal rearrangement*. It was shown in [Sem08] that the size of the average (over variables i , the solution $\{s_i\}$, and the graph ensemble) minimal rearrangement diverges at the rigidity transition (when almost all the dominant clusters become frozen). Indeed, the cavity approach to minimal rearrangements leads to equations analogous to those for frozen variables. Part of the reasoning is the following [SAO05]: Consider a solution of a K -SAT formula and a variable i from its whitening core. By flipping the variable i at least one neighbouring constraint a is made unsatisfied, otherwise the variable would not be in the whitening core. All variables contained in a are also in the whitening core, thus one of them has to be flipped in order to satisfy this constraint. There have to be a chain of flips which can be finished only by closing a loop. The length of the shortest loop going through a typical variable is of order $\log N$. Thus a diverging number of changes is needed to find another solution. Hence the connection between frozen variables and rearrangements is:

- If the variable i is frozen in the cluster to which the solution $\{s_i\}$ belongs, then in order to change the value of i one has to find a solution from a different cluster, thus at an extensive Hamming distance.
- If the variable i is not frozen in the cluster to which the solution $\{s_i\}$ belongs, then the best rearrangement will probably also lie within that cluster and the Hamming distance is finite.

Many more results about rearrangements can be found in [Sem08], they shed light on the onset of frozen variables. An exciting possibility is that the cavity equations for rearrangements might be useful in incremental algorithms for CSPs, like the one of [KK07].

4.2.2 The phase transitions: Rigidity and Freezing

A natural question is: “In which clusters are the hard fields present?” Or more in the terms of the 1RSB solutions: “When does eq. (4.7) have a nontrivial solution $\eta > 0$?” We answer this question in one of the simplest cases, that is for the coloring of random regular graphs of connectivity $c = k + 1$. In tree-like regular graphs the neighbourhood of each node looks identical, thus also the distribution $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ is the same for every edge (ij) . Moreover we search for a color-symmetric solution [ZK07], that is $\eta_s = \eta_r = \eta$ for all $s, r \in \{1, \dots, q\}$. The function $w(\{\eta\})$ in the ensemble of random regular graphs simplifies to

$$w(\eta) = \frac{\sum_{l=0}^{q-1} (-1)^l \binom{q-1}{l} [1 - (l+1)\eta]^k}{\sum_{l=0}^{q-1} (-1)^l \binom{q}{l+1} [1 - (l+1)\eta]^k}. \quad (4.9)$$

First notice that in order to constrain a variable into one color, i.e., create a hard field, one needs at least $q - 1$ incoming fields that forbids all the other colors. It means that the function $w(\{\eta\})$ defined in eq. (4.9) is identically zero for $k < q - 1$ and might be non-zero only for $k \geq q - 1$, where k is the number of incoming fields.

The equation (4.7) also simplifies on a regular graph and η follows a self-consistent relation

$$\eta = w(\eta) \frac{1}{qw(\eta) + [1 - qw(\eta)] r(m)}, \quad (4.10)$$

where $r(m)$ is the average of the reweighting of the soft fields divided by the average of the reweighting of the frozen fields (4.7). The function $r(m)$ is in general not easy to compute, the population dynamics is needed for that. Several properties are, however known:

$$r \rightarrow 0 \quad \text{when} \quad m \rightarrow -\infty, \quad (4.11a)$$

$$r \rightarrow \infty \quad \text{when} \quad m \rightarrow \infty, \quad (4.11b)$$

and $r(m)$ is a monotonous function of m . Moreover, for the internal entropy of clusters $s(m) \rightarrow 0$ when $m \rightarrow -\infty$, and $s(m) \rightarrow \infty$ when $m \rightarrow \infty$, and $s(m)$ is also a monotonous function. We thus solve eq. (4.10) for every possible ratio r . For all $k \geq q - 1$ we compute the solution $\eta(r)$. The result is shown in fig. 4.2 for the 3- and 4-coloring of random regular graphs.

There is a discontinuous phase transition: For $r < r_r$ eq. (4.10) has a solution with a large fraction of frozen fields, $\eta > 0$, whereas for $r > r_r$ the only solution is $\eta = 0$. Note that the index r stands for “rigidity”. In terms of the parameter m , the critical value is $r(m_r) = r_r$. In terms of the internal entropy of clusters $s(m_r) = s_r$. The interpretation is the following:

- Clusters of internal entropy $s < s_r$ are almost all frozen, and the fraction of frozen variables they contain is quite large.
- Clusters of internal entropy $s > s_r$ are almost all soft, meaning the fraction of frozen variables is zero.

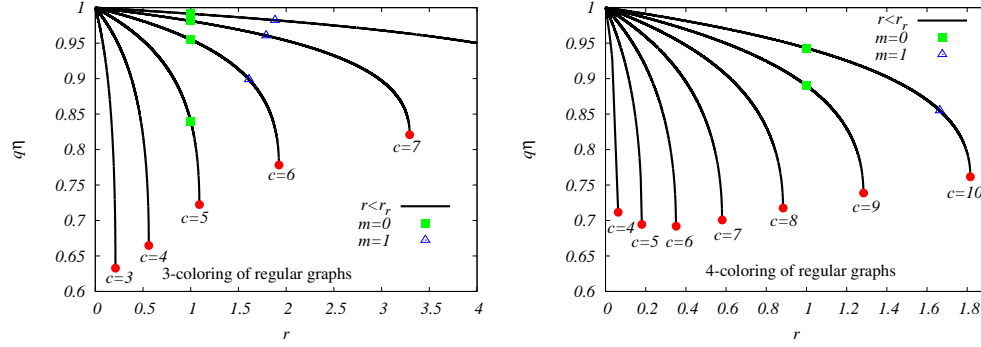


Fig. 4.2. (Color online) The lines are solutions of eq. (4.10) and give the fraction $q\eta$ of hard fields for a given value of ratio $r = \langle Z_m^{i \rightarrow j} \rangle_{\text{soft}} / \langle Z_m^{i \rightarrow j} \rangle_{\text{hard}}$ for $q = 3$ (left) and $q = 4$ (right) in regular random graphs. There is a critical value of the ratio r_r (red point) beyond which only the trivial solution $\eta = 0$ exists. Note that the solutions corresponding to $m = 0$ (green square) and $m = 1$ (blue triangle) only exist for a connectivity large enough.




Fig. 4.3. (Color online) A pictorial sketch of the complexity function of clusters of a given size. Cyan-blue is the complexity of the frozen clusters, magenta of the soft clusters. The total complexity is the envelope, which can be calculated from the entropic 1RSB solution. The black point marks the dominating clusters. Left: In the rigid phase almost all the dominant clusters are frozen, but clusters corresponding to larger entropy might be mostly soft. Middle: In the totally rigid phase almost all clusters of all sizes are frozen, but there still might be exponentially many of soft clusters. Right: The frozen phase where soft clusters almost surely do not exist.

When we change the average constraint density there are at least three interesting phase transitions related to frozen variables. Figure 4.3 sketches the difference between the phases they separate. Recall that s^* is the internal entropy of the dominant clusters, and s_{\max} the internal entropy of the largest clusters $\Sigma(s_{\max}) = 0$.

- The rigidity transition, c_r , at which $s^* = s_r$, separates a phase where a typical dominant cluster is almost surely not frozen from a phase where a typical dominant cluster is almost surely frozen.
- The total rigidity transition, c_{tr} , at which $s_{\max} = s_r$, when almost all clusters of every size become frozen.

- The freezing transition, c_f , separates phase where exponentially many unfrozen cluster exists from a phase where such clusters almost surely do not exist³.

In general it have to be $c_r \leq c_{tr} \leq c_f$. The relation between the rigidity and total rigidity transition is easily obtained from the 1RSB solution. It is thus known that in the q -coloring of Erdős-Rényi graphs $c_r = c_{tr}$ if and only if $q \leq 8$, in K -SAT if and only if $K \leq 5$. For larger q or K the rigidity transition is given by the onset of frozen variables in clusters corresponding to $m = 1$, this is equivalent to the naive reconstruction (2.4).

The relation between the total rigidity transition and the freezing is less known. There are only few studies for the freezing transition in random K -SAT. The first one is the one of [ART06] where they prove that for every $K \geq 9$ the freezing transition is strictly smaller than the satisfiability one $c_f < c_s$. In the large K limit they showed that the frozen phase covers a finite fraction (at least 20%) of the satisfiable region. The second study [MS07] gives a rigorous upper bound on the freezing transition in 3-SAT $\alpha_f < 4.453$, which is slightly better than the best known upper bound on the satisfiability transition in 3-SAT [DBM00]. The third study is numerical [AZ08], presented in fig. 4.1. It shows that in 3-SAT the frozen phase is tiny, about 0.3% of the satisfiable region.

It is not known if the total rigidity transition coincides with the freezing transition. The entropic cavity method describes a typical but not every cluster of a given size. A generalization of the 1RSB equations which would count only the number of soft cluster would answer this question.

To summarize the description of the freezing of variables and clusters in the canonical constraint satisfaction problems, like q -coloring or K -satisfiability, is both numerically and conceptually involved task. Moreover in the experimentally feasible range of q and K the frozen phase is tiny. Thus conclusive statements about the connection between the freezing and the computational hardness are difficult to make. In the next section we introduce the so-called *locked* constraint satisfaction problems where the situation is much more transparent.

4.3 Point like clusters: The locked problems

In order to get a better understanding of the frozen phase we introduce the so-called *locked* constraint satisfaction problems [ZM08]. In these problems the whole clustered phase is at the same time frozen, this is because in the locked problems all the clusters contain only one solution.

4.3.1 Definition

A *locked* constraint satisfaction problem is made of N variables and M *locked* constraints in such a way that every variable is present in at least two constraints. A constraint consisting of $K > 0$ variables is *locked* if and only if for every satisfying assignment of variables changing the value of any (but only one) variable makes the assignment unsatisfying.

A locked constraint of K variables has the property that if $(K - 1)$ variables are assigned then either the constraint cannot be satisfied by any value of the last variable or there is only one value of the last variable which makes the constraint satisfied. All the uniquely extendible

³Note that what is called freezing transition in [Sem08] or in sec. IV.C of [MRTS08] is in fact what we define as the rigidity transition, in agreement with [ZK07].

constraints [Con04, CM04] are locked, XOR-SAT being the most common example. 1-in- K SAT (exact cover) constraint [GJ79] is another common example. On the other hand, the most studied constraint satisfaction problems K -SAT or graph q -coloring ($q > 2$) are not made of locked constraints.

The second important part of the definition of locked constraint satisfaction problems is the requirement that every variable is present in at least two constraints, i.e., leaves are absent. An important property follows: In order to change a satisfying assignment into a different satisfying assignment at least a closed loop of variables have to be changed. If leaves would be allowed changing a path connecting two leaves might be sufficient.

It seems to us that all the random locked constraint satisfaction problems should behave in the way we describe in the following. We, however, investigated in detail only a subclass of the locked problems called *locked occupation problems* (LOP). Occupation constraint satisfaction problem is defined as a problem with binary variables (0-empty, 1-occupied) where each constraint containing K variables is a function of how many of the K variables are occupied. A constraint of the occupation CSP can thus be characterized via a $(K + 1)$ -component vector A , $A_i \in \{0, 1\}$, $i \in 0, \dots, K$. A constraint is satisfied (resp. violated) if it contains r occupied variables where r is such that $A_r = 1$ (resp. $A_r = 0$). For example $A = (0, 1, 0, 0)$ corresponds to the positive 1-in-3 SAT [RSZ07], $A = (0, 1, 1, 0)$ is bicoloring [CNRTZ03], $A = (0, 1, 0, 1, 0)$ is 4-odd parity check (4-XOR-SAT without negations) [MRTZ03].

An occupation problem is locked if all the variables are connected to at least two constraints and the vector A is such that $A_i A_{i+1} = 0$ for all $i = 0, \dots, K - 1$. We study the random ensembles of LOPs where all constraints are identical and the variable degree is either fixed or distributed according to a truncated Poissonian law (1.6).

4.3.2 The replica symmetric solution

The replica symmetric cavity equations, belief propagation (1.16a-1.16b), for the occupation problems read

$$\psi_{s_i}^{a \rightarrow i} = \frac{1}{Z^{a \rightarrow i}} \sum_{\{s_j\}} \delta(A_{s_i + \sum_j s_j} - 1) \prod_{j \in \partial a - i} \chi_{s_j}^{j \rightarrow a}, \quad (4.12a)$$

$$\chi_{s_j}^{j \rightarrow a} = \frac{1}{Z^{j \rightarrow a}} \prod_{b \in \partial j - a} \psi_{s_j}^{b \rightarrow j}, \quad (4.12b)$$

where $\psi_{s_i}^{a \rightarrow i}$ is the probability that the constraint a is satisfied conditioned that the value of the variable i is s_i , and $\chi_{s_j}^{j \rightarrow a}$ is the probability that variable j takes value s_j conditioned that the constraint a was removed from the graph. The normalizations Z have the meaning of the partition function contributions. The replica symmetric entropy s is a zero temperature limit of (1.20)

$$s = \frac{1}{N} \sum_a \log(Z^{a+\partial a}) - \frac{1}{N} \sum_i (l_i - 1) \log(Z^i), \quad (4.13)$$

where the contributions $Z^{a+\partial a}$ (resp. Z^i) are the exponentials of the entropy shifts when the node a and its neighbours (resp. the node i) is added (1.19a-1.19b)

$$Z^{a+\partial a} = \sum_{\{s_i\}} \delta(A_{\sum_i s_i} - 1) \prod_{i \in a} \left(\prod_{b \in i-a} \psi_{s_i}^{b \rightarrow i} \right), \quad (4.14a)$$

$$Z^i = \prod_{a \in i} \psi_0^{a \rightarrow i} + \prod_{a \in i} \psi_1^{a \rightarrow i}. \quad (4.14b)$$

Solving eqs. (4.12a-4.12b) means finding their fixed points. A crucial property of the locked problems is that if $\{s_i\}$ is one of the solutions then

$$\psi_{s_i}^{a \rightarrow i} = 1, \quad \psi_{\neg s_i}^{a \rightarrow i} = 0, \quad (4.15a)$$

$$\chi_{s_i}^{i \rightarrow a} = 1, \quad \chi_{\neg s_i}^{i \rightarrow a} = 0 \quad (4.15b)$$

is a fixed point of eqs. (4.12a-4.12b). The corresponding entropy is then zero, as $Z^i = Z^{a+\partial a} = 1$ for all i, a . In the derivation of [MM08] fixed points of the belief propagation equations correspond to clusters. Thus in the locked problems every solution corresponds to a cluster.

In the satisfiable phase there exist exponentially many solutions (i.e., clusters), thus the iterative fixed point of BP equations (4.12a-4.12b) obtained from a random initialization gives an asymptotically exact value for the total entropy. And the satisfiability threshold coincides with the condensation transition, described in chap. 3. Furthermore, as each cluster contains only one solution the clustered phase is automatically frozen according to the definition in sec. 4.2.2. Interestingly, part of the satisfiable phase is only "fake clustered" meaning that at infinitesimally small temperature there is a single fixed point of the BP equations. This has been discussed e.g. in the context of the perfect matchings in [ZM06]. A general discussion and proper definition of the clustering transition in the locked problems follows in sec. 4.3.3.

Iterative fixed point of eqs. (4.12a-4.14b) averaged over the graph ensemble is in general found via the population dynamics technique, see appendix E. Note that the sum over $\{s_j\}$ in (4.12a) can be computed iteratively in $(K-1)^2$ steps instead of the naive 2^{K-1} steps. Moreover, on the regular graphs ensemble or for some of the symmetric locked problems, such that $A_i = A_{K-i}$ for all $i = 0, \dots, K$, the solutions is *factorized*. In the factorized solution the messages $\chi^{i \rightarrow a}, \psi^{a \rightarrow i}$ are independent of the edge (ia) and the population dynamics is thus not needed.

- For the regular graph ensemble where each variable is present in L constraints the factorized solution is

$$\psi_0 = \frac{1}{Z^{\text{reg}}} \sum_{A_r=1} \binom{K-1}{r} \psi_1^{(L-1)r} \psi_0^{(L-1)(K-1-r)}, \quad (4.16a)$$

$$\psi_1 = \frac{1}{Z^{\text{reg}}} \sum_{Ar+1=1} \binom{K-1}{r} \psi_1^{(L-1)r} \psi_0^{(L-1)(K-1-r)}, \quad (4.16b)$$

and the entropy is

$$s_{\text{reg}} = \frac{L}{K} \log \left[\sum_{A_r=1} \binom{K}{r} \psi_1^{(L-1)r} \psi_0^{(L-1)(K-r)} \right] - (L-1) \log [\psi_0^L + \psi_1^L]. \quad (4.17)$$

- For the symmetric locked problems where the symmetry is not spontaneously broken the solution is also factorized. We call these the *balanced* locked problems. The BP solution is $\psi_1 = \psi_0 = 1/2$ and the corresponding entropy

$$s_{\text{sym}}(\bar{l}) = \log 2 + \frac{\bar{l}}{K} \log \left[2^{-K} \sum_{r=0}^K \delta(A_r - 1) \binom{K}{r} \right], \quad (4.18)$$

where \bar{l} is the average degree of variables. Notably, this result for the entropy can be proven rigorously by computing the first and second moment of the partition sum, i.e., $\langle Z \rangle, \langle Z^2 \rangle$, and using the Chebyshev's inequality. The exact value of the satisfiability threshold is then given by $s_{\text{sym}}(l_s) = 0$. This itself is a remarkable result, because so far the exact threshold was computed in only a handful of the sparse NP-complete CSPs. As far as we know only in the 1-in- K SAT [ACIM01] and [RSZ07], the $2 + p$ -SAT [MZK⁺99a, AKKK01] and the (3, 4)-UE-CSP [CM04]. We dedicate the appendix B to this computation.

The replica symmetric solution might be incorrect if long range correlations are present in the system, as we discussed in detail in chap. 2. A sufficient condition for its correctness is the decay of the point-to-set correlations, which we will discuss in the next section, again in context of the reconstruction problem. A necessary condition for the RS solution to be correct is the non-divergence of the spin glass susceptibility, which can be investigated in several equivalent ways, as described in appendix C. The result for all the locked problems we investigated is that the phase where the entropy (4.13) is positive is always RS stable, whereas part of the phase where the entropy (4.13) is negative might be RS unstable (depending on the parameters and the vector A).

4.3.3 Small noise reconstruction

It is immediate to observe that reconstruction as we defined it in sec. 2.1.1 is always possible for the locked problems. Indeed, if we know $K - 1$ out of K variables around a constraint the last one is given uniquely (no contradiction is possible as we broadcasted a solution). This is related to the fact that at least one closed loop has to be flipped to go from one solution of a given instance of a locked problem to another solution. Typical length of such a minimal loop is of order $\log N$. For very low connectivities, and at infinitesimally low temperature, the BP equations will have a unique fixed point, there the zero temperature $\log N$ clustering is "fake" and will not have a crucial influence on the dynamics and other properties of interest.

Thus for the locked problem it is useful to modify the definition of the clustering transition presented in chap. 2. In order to do that we need to introduce the *small noise (SN) reconstruction*. Construct an infinite tree hyper-graph, assign a value 1 or 0 to its root and iteratively assign its offsprings uniformly at random but in such a way that the constraints are satisfied (constraints play the role of noiseless channels). At the end of the procedure forget the values of all variables in the bulk but also of an infinitesimal fraction ϵ of leaves. If the remaining $1 - \epsilon$ leaves contain some information about the original value on the root then we say that the small noise reconstruction is possible, if they do not the small noise reconstruction is not possible. The phase where the SN reconstruction is not possible is then only "fake clustered" and is more similar to the liquid

phase. Whereas the phase where the SN reconstruction is possible has all the properties of the clustered phase, except that each of the clusters contains only one configuration⁴.

All the equations we derived in sec. 2.1.1 for the reconstruction apply also for the SN reconstruction. Except the specification of the initial conditions (2.11) which for the SN reconstruction is instead

$$P^{\text{init}}(\vec{\psi}) = \frac{1-\epsilon}{2} [\delta(\vec{\psi} - \delta_0) + \delta(\vec{\psi} - \delta_1)] + \epsilon \delta(\psi_0 - \frac{1}{2}) \delta(\psi_1 - \frac{1}{2}), \quad (4.19)$$

where $\epsilon \ll 1$. The second term accounts for the fraction of leaves on which the value of the variable has been forgotten. The fixed point of the 1RSB equation (2.24) is then either trivial (corresponding to the replica symmetric solution) or nontrivial describing solutions as an ensemble of totally frozen clusters. This has several interesting consequences: The threshold for the naive SN reconstruction (i.e., the one taking into account only the frozen variables) coincide with the true threshold for SN reconstruction. The solution of the 1RSB equation (2.24) in the locked problem does not depend on the value of the parameter m .

A general form of the 1RSB equations at $m = 1$ for occupation problems is derived in appendix A. First we consider only problems where the replica symmetric solution is factorized. We define μ_1 (resp. μ_0) as the probability that a variable which in the broadcasting had value 1 (resp. 0) is uniquely determined by the boundary conditions. Based on the general eq. (A.10), we derive self-consistent equations for μ_1, μ_0 on regular graphs ensemble of connectivity of variables L :

$$\begin{aligned} \mu_1 &= \frac{1}{\psi_1 Z^{\text{reg}}} \sum_{A_{r+1}=1, A_r=0} \binom{k}{r} (\psi_1)^{lr} (\psi_0)^{l(k-r)} \sum_{s=0}^{s_1} \binom{r}{s} \\ &\times [1 - (1 - \mu_0)^l]^{k-r} [1 - (1 - \mu_1)^l]^{r-s} (1 - \mu_1)^{ls}, \end{aligned} \quad (4.20a)$$

$$\begin{aligned} \mu_0 &= \frac{1}{\psi_0 Z^{\text{reg}}} \sum_{A_r=1, A_{r+1}=0} \binom{k}{r} (\psi_1)^{lr} (\psi_0)^{l(k-r)} \sum_{s=0}^{s_0} \binom{k-r}{s} \\ &\times [1 - (1 - \mu_1)^l]^r [1 - (1 - \mu_0)^l]^{k-r-s} (1 - \mu_0)^{ls}, \end{aligned} \quad (4.20b)$$

where $l = L - 1, k = K - 1$. The indices s_1, s_0 in the second sum of both equations are the largest possible but such that $s_1 \leq r, s_0 \leq K - 1 - r$, and $\sum_{s=0}^{s_1} A_{r-s} = 0, \sum_{s=0}^{s_0} A_{r+1+s} = 0$. The values ψ_0, ψ_1 are the fixed point of eqs. (4.16a-4.16b), and Z^{reg} is the corresponding normalization. These lengthy equations have in fact a simple meaning. The first sum is over the possible numbers of occupied variables on the descendants in the broadcasting. The sums over s is over the number of variables which were not implied by at least one constraint but still such that the set of incoming implied variables implies the outgoing value. The term $1 - (1 - \mu)^l$ is the probability that at least one constraint implies the variable, $(1 - \mu)^l$ is the probability that none of the constraints implies the variable.

The second case where the BP equations are factorized are the *balanced* locked problems. That is LOPs with symmetric vector A where the symmetry is not spontaneously broken. Then

⁴Note that a rigorous study of a related robust reconstruction exists [JM04]. In robust reconstruction, however, one allows ϵ to be arbitrarily near to one.

$\psi_0 = \psi_1 = 1/2$ and thus also $\mu_0 = \mu_1 = \mu$. For the ensemble of graphs with truncated Poissonian degree distribution of coefficient c we derive from (A.10)

$$\mu = \frac{2}{g_A} \sum_{A_{r+1}=1} \binom{k}{r} \sum_{s=0}^{s_1} \binom{r}{s} \left(\frac{1 - e^{-c\mu}}{1 - e^{-c}} \right)^{k-s} \left(\frac{e^{-c\mu} - e^{-c}}{1 - e^{-c}} \right)^s, \quad (4.21)$$

where $k = K - 1$, and $g_A = \sum_{r, A_{r+1}=1} \binom{k}{r} + \sum_{r, A_r=1} \binom{k}{r}$ and the value s is, as before, the number of descendants which were not directly implied.

In both these cases, there are two solutions to eqs. (4.20a-4.20b) and (4.21). One is $\mu = 0$ and the other $\mu = 1$. The small noise reconstruction is investigated by the iterative stability of the solution $\mu = 1$. If it is stable then the SN reconstruction is possible, all variables are almost surely directly implied. If it is not stable then the only other solution is $\mu = 0$. Few observations are immediate, for example if $L \geq 3$ then the solution $\mu_1 = \mu_0 = 1$ of (4.20a-4.20b) is always iteratively stable. Iterative stability of (4.21) gives for the balanced locked problems, marked by * in tab. 4.1:

$$\frac{e^{c_d} - 1}{c_d} = K - 1 - \frac{\sum_{r=0}^{K-2} r \delta(A_{r+1} - 1) \delta(A_{r-1}) \delta(A_r) \binom{K-1}{r}}{\sum_{r=0}^{K-2} \delta(A_{r+1} - 1) \binom{K-1}{r}}. \quad (4.22)$$

Tab. 4.1. The locked cases of the occupation CSPs for $K \leq 6$ (cases with a trivial ferromagnetic solution are omitted). In the regular graphs ensemble the phase is clustered for $L \geq L_d = 3$, and unsatisfiable for $L \geq L_s$. Values c are the critical parameters of the truncated Poissonian ensemble (1.6), the corresponding average connectivities \bar{l} are given via eq. (1.7). All these problems are RS stable at least up to the satisfiability threshold. For the balanced cases, marked as *, the dynamical threshold follows from (4.21), and the satisfiability threshold, which can be computed rigorously, app. B, from (4.18).

A	name	L_s	c_d	c_s	l_d	l_s
0100	1-in-3 SAT	3	0.685(3)	0.946(4)	2.256(3)	2.368(4)
01000	1-in-4 SAT	3	1.108(3)	1.541(4)	2.442(3)	2.657(4)
00100*	2-in-4 SAT	3	1.256	1.853	2.513	2.827
01010*	4-odd-PC	5	1.904	3.594	2.856	4
010000	1-in-5 SAT	3	1.419(3)	1.982(6)	2.594(3)	2.901(6)
001000	2-in-5 SAT	4	1.604(3)	2.439(6)	2.690(3)	3.180(6)
010100	1-or-3-in-5 SAT	5	2.261(3)	4.482(6)	3.068(3)	4.724(6)
010010	1-or-4-in-5 SAT	4	1.035(3)	2.399(6)	2.408(3)	3.155(6)
0100000	1-in-6 SAT	3	1.666(3)	2.332(4)	2.723(3)	3.113(4)
0101000	1-or-3-in-6 SAT	6	2.519(3)	5.123(6)	3.232(3)	5.285(6)
0100100	1-or-4-in-6 SAT	4	1.646(3)	3.366(6)	2.712(3)	3.827(6)
0100010	1-or-5-in-6 SAT	4	1.594(3)	2.404(6)	2.685(3)	3.158(6)
0010000	2-in-6 SAT	4	1.868(3)	2.885(4)	2.835(3)	3.479(4)
0010100*	2-or-4-in-6 SAT	6	2.561	5.349	3.260	5.489
0001000*	3-in-6 SAT	4	1.904	3.023	2.856	3.576
0101010*	6-odd-PC	7	2.660	5.903	3.325	6

4.3.4 Clustering transition in the locked problems

In the locked problem where the replica symmetric solution is not factorized there is another equivalent way to locate the clustering transition, which is simpler than solving eq. (A.10). It is the investigation of the iterative stability of the nontrivial fixed point of the survey propagation. In LOPs the survey propagation equations consist of eqs. (1.41) and

$$q_1^{a \rightarrow i} = \frac{1}{\mathcal{N}^{a \rightarrow i}} \left[\sum_{\{r_j\}} C_1(\{r_j\}) \prod_{j \in a-i} p_{r_j}^{j \rightarrow a} \right], \quad (4.23a)$$

$$q_{-1}^{a \rightarrow i} = \frac{1}{\mathcal{N}^{a \rightarrow i}} \left[\sum_{\{r_j\}} C_{-1}(\{r_j\}) \prod_{j \in a-i} p_{r_j}^{j \rightarrow a} \right], \quad (4.23b)$$

$$q_0^{a \rightarrow i} = \frac{1}{\mathcal{N}^{a \rightarrow i}} \left[\sum_{\{r_j\}} C_0(\{r_j\}) \prod_{j \in a-i} p_{r_j}^{j \rightarrow a} \right], \quad (4.23c)$$

where the indexes $r_j \in \{1, -1, 0\}$, $\mathcal{N}^{a \rightarrow i}$ is the normalization constant. The C_1/C_{-1} (resp. C_0) takes values 1 if and only if the incoming set of $\{r_j\}$ forces the variable i to be occupied/empty (resp. let the variable i free), in all other cases the C 's are zero. Let us call s_1, s_{-1}, s_0 the number of indexes 1, $-1, 0$ in the set $\{r_j\}$ then

- $C_1 = 1$ if and only if $A_{s_1+s_0+1} = 1$ and $A_{s_1+n} = 0$ for all $n = 0 \dots s_0$;
- $C_{-1} = 1$ if and only if $A_{s_1} = 1$ and $A_{s_1+1+n} = 0$ for all $n = 0 \dots s_0$;
- $C_0 = 1$ if and only if there exists $m, n = 0 \dots s_0$ such that $A_{s_1+n} = A_{s_1+m+1} = 1$.

The SP equations in LOPs have two different fixed points:

- The trivial one: $q_0^{a \rightarrow i} = p_0^{i \rightarrow a} = 1, q_1^{a \rightarrow i} = p_1^{i \rightarrow a} = q_{-1}^{a \rightarrow i} = p_{-1}^{i \rightarrow a} = 0$ for all edges (ai) .
- The BP-like one: $q_0^{a \rightarrow i} = p_0^{i \rightarrow a} = 0, q^{a \rightarrow i} = \psi^{a \rightarrow i}, p^{i \rightarrow a} = \chi^{i \rightarrow a}$ for all edges (ai) , where ψ and χ is the solution of the BP equations (4.12a-4.12b).

The small noise reconstruction is then investigated, using the population dynamics, from the iterative stability of the BP-like fixed point. If it is stable then the SN reconstruction is possible and the phase is clustered. If it is not stable then we are in the liquid phase. Of course, this approach gives the same critical connectivity l_d as the previous one, because for the locked problems the solutions of the 1RSB equation (2.24) is independent of the parameter m .

We remind at this point that in a general CSP, where the sizes of clusters fluctuate, the SP equations are not related to the reconstruction problem, more technically said the 1RSB solutions at $m = 0$ and at $m = 1$ are different. The solution of the locked problems is sometimes called frozen 1RSB [MMR04, MMR05].

4.4 Freezing - The reason for hardness?

We describe several strong evidences that it is hard to find frozen solutions. We also give several arguments for why it is so. However, the precise mechanism stays an open question and strictly speaking the freezing of variables might just be going along with a true yet unknown reason. Or even there might be an algorithm which is able to find the frozen solutions efficiently waiting for a discovery. But in any case, we show that freezing of variables is an important new aspect in the search of the origin of the average computational hardness.

4.4.1 Always a trivial whitening core

Several studies of the random 3-SAT problem [MMW07, BZ04, SAO05] showed that all known algorithms on large instances systematically find only solutions with a trivial whitening core (defined in sec. 4.1.1). On small instances of the problem solutions with a nontrivial whitening core can be found as observed by several authors, and studied systematically in sec. 4.1.3.

For solutions found by the stochastic local search algorithms, see appendix F, this observation is reasonable, as argued already in [SAO05]. Consider that a stochastic local search finds a configuration which is not a solution, but its whitening core is nontrivial. Then a diverging number of variables have to be rearranged in order to satisfy one of the unsatisfied constraints [Sem08]. In the clusters with a trivial whitening core the rearrangements are finite [Sem08] and thus stochastic local dynamics might be able to find them more easily.

The fact of finding only the "white" solutions is, however, quite surprising for the survey propagation algorithm. The SP equations compute probabilities (over clusters) that a variables is frozen in a certain value. This information is then used in a decimation, reinforcement, etc. algorithms, see appendix F. Thus SP is explicitly exploring the information about nontrivial whitening cores and in spite of that it finishes finding solutions with trivial whitening cores.

A related, and rather surprising, result was shown in [DRZ08]. The authors considered the random bi-coloring problem in the rigid, but not frozen, phase. That is a phase where most solutions are frozen, but rare unfrozen ones still exist. They showed that belief propagation reinforcement solver, see appendix F, is in some cases able to find these exponentially rare, but unfrozen, solutions.

We have observed the same phenomena in one of the non-locked occupation problem $A = (0110100)$, that is 1-or-2-or-4-in-6 SAT. On regular factor graphs this problem is in the liquid phase for $L \leq 6$, in the rigid phase for $7 \leq L \leq 9$, where almost all the solutions are frozen, and it is unsatisfiable for $L \geq 10$. In fig. 4.4 we show that belief propagation reinforcement finds almost always solutions for $L = 8$, but as the size of instances is growing the fraction of cases in which the solution is frozen goes to zero.

We listed this paradox, that only the all-white solutions can be found, as one of the loose ends in sec. 1.8. The resolution we suggest here, and substantiate in the following, is that every known algorithm is able to find efficiently (in polynomial - but more often in experiments we mean linear or quadratic - time) only the unfrozen solutions. The frozen solutions are intrinsically hard to find and all the known algorithms have to run for an exponential time to find them.




Fig. 4.4. (Color online) Algorithmic performance in the rigid phase of the 1-or-2-or-4-in-6 SAT at $L = 8$. In red is the rate of success of the belief propagation reinforcement algorithms as a function of system size (out of 100 trials). The algorithm basically always succeeds to find a solution. In blue is the fraction of solutions which were frozen (had a nontrivial whitening core). Almost all solutions are frozen in this problem, yet it is algorithmically easier to find the rare unfrozen solutions, in particular in instances of larger size.

4.4.2 Incremental algorithms

Adopted from [KK07]: Consider an instance of a constraint satisfaction problem of N variables and M constraints. Order randomly the set of constraints and remove all of them. Without constraints any configuration is a solution. In each step: First, add back one of the constraints. Second, if needed rearrange the configuration in such a way that it satisfies the new and all the previous constraints. Repeat until there are some constraints left. We call such strategy the *incremental algorithm* for CSPs. And one can ask about its computational complexity. The way by which the rearrangement is found in the second step needs to be specified. But independently of this specification we know that if the new constraint connects frozen and contradictory variables then the size of the minimal rearrangement diverges [Sem08], thus in the frozen phase the incremental algorithm has to be at best super-linear.

Another understanding of the situation is gained by imagining the space of solutions at a given constraint density. As we are adding the constraints some solutions are disappearing and none are appearing. At the clustering transition the space of solutions splits into exponentially many clusters. As more constraints are added the clusters are becoming smaller, they may split into several smaller ones and some may completely disappear. However, only the frozen clusters can disappear, if a constraint is added between two frozen and contradictory variables. Note also that each frozen cluster will almost surely disappear before an infinitesimally small fraction of constraints is added. An unfrozen cluster, on the other hand, may only become smaller or split. Indeed, if a constraint is added any solution belonging to an unfrozen cluster may be rearranged

in a finite number of steps [Sem08]. The incremental algorithm in this setting works as a non-intelligent animal would be escaping from a rising ocean on a Pacific hilly island [KK07]. As the water starts to rise the animal would step away from it. As the water keeps rising at a point the animal would be blocked in one of the many smaller islands. This island will be getting smaller and smaller and it will disappear at a point and the animal will have to learn how to swim. But at this point there might still be many small higher island. All of them will disappear eventually. For sure the animal will be in trouble before all the clusters (island) start to contain frozen variables.

Moreover, if the sequence of constraints to be added is not known in advance there is no way to choose the best cluster, because which cluster is the best depends completely on the constraints to be added. This proves that no incremental algorithm is able to work in linear time in the frozen phase. On the other hand it was shown experimentally in [KK07] for the coloring problem that such algorithms work in linear time in part of the clustered (or even the condensed) phase.

4.4.3 Freezing transition and the performance of SP in 3-SAT

How does the freezing transition in 3-SAT, $\alpha_f = 4.254 \pm 0.009$ fig. 4.1, compare to the performance of the best known random 3-SAT solver — the survey propagation? We are aware of two studies where the performance of SP is investigated systematically and with a reasonable precision, [Par03] and [CFMZ05].

In [Par03] the survey propagation decimation is studied. The SP fixed point is found on the decimated graph and the variable having the largest bias is fixed as long as the SP fixed point is nontrivial. When the SP fixed point becomes trivial the Walk-SAT algorithm finishes the search for a solution. In [Par03] the residual complexity is measured on the partially decimated graph. It is observed that if the residual complexity becomes negative then solutions are never found, if on the other hand the residual complexity is positive just before the survey propagation fixed point becomes trivial then solutions are found. The value of complexity in the last step before the fixed point becomes trivial is extrapolated, fig. 2 of [Par03] for system size $N = 3 \cdot 10^5$, to zero at a constraint density $\alpha = 4.252 \pm 0.003$ (we estimated the error bar based on data from [Par03]).

In [CFMZ05] the survey propagation reinforcement is studied. The rate of success is plotted as a function of the complexity function. From fig. 8 of [CFMZ05] it is estimated that SP reinforcement (more precisely its implementation presented in [CFMZ05]) finds solution in more than 50% of trials if $\Sigma > 0.0013$. The data do not really concentrate on this point, thus it is difficult to obtain a reliable error bar of this value, our educated guess is 0.0013 ± 0.0003 this would correspond to a constraint density $\alpha = 4.252 \pm 0.004$.

The striking agreement between our value for the freezing transition and the performance limit of the survey propagation supports the suggestion that the frozen phase is hard for any known algorithm. The trouble for a better study of the frozen phase in 3-SAT is its size, it covers only 0.3% of the satisfiable phase. In K -SAT with large K the frozen phase becomes wider, but as K grows the constraint density of the satisfiability threshold grows like $2^K \log K$, empirical study thus becomes infeasible very fast. It is also not very easy to compute the freezing transition or to check if the 1RSB solution is correct in the frozen phase. Thus K -SAT (and q -coloring) are not very suitable problems for understanding better how exactly the freezing influences the search for a solution.

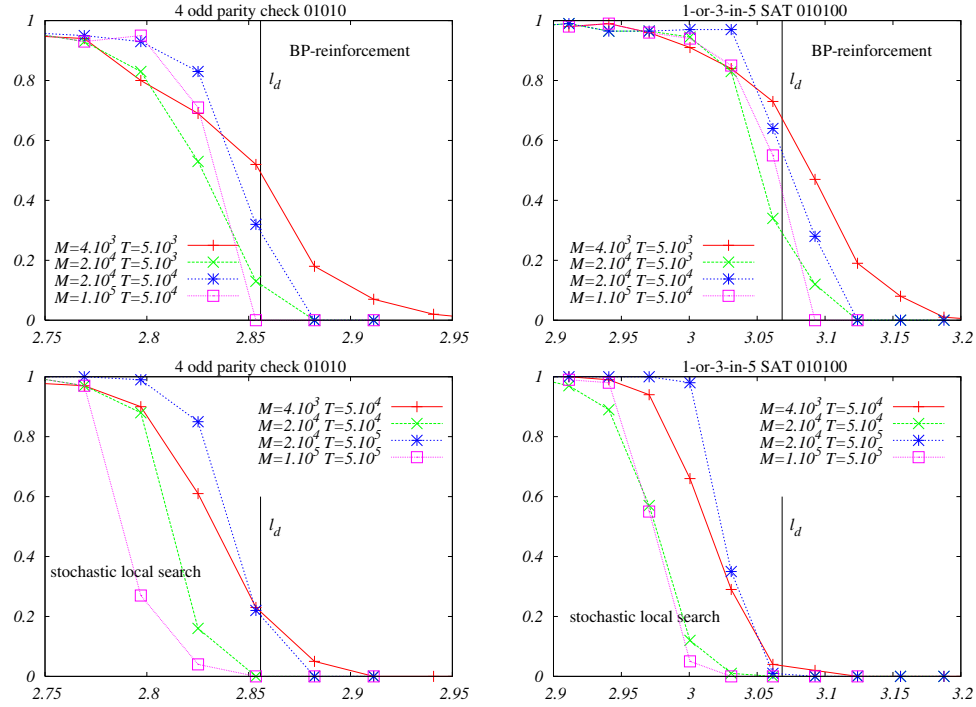


Fig. 4.5. (Color online) The probability of success of the BP-REINFORCEMENT (top) and the stochastic local search ASAT (bottom) plotted against the average connectivity for two of the locked occupation problems. The clustering transition is marked by a vertical line, the satisfiability threshold is $l_s = 4$ for the 4-odd parity checks, and $l_s = 4.72$ for the 1-or-3-in-5 SAT. The challenging task is to design an algorithm which would work also in the clustered phase of the NP-complete locked problems.

4.4.4 Locked problems – New extremely challenging CSPs

We introduced the locked problems to challenge the suggestion about hardness of the frozen phase [ZM08]. It is rather easy to compute the freezing transition here, it coincides with the clustering transition l_d . Moreover, the frozen phase is wide, taking more than 50% of the satisfiable phase for some of the locked problems, see table 4.1. As in the locked problems every cluster consists of one solution, all the variables are frozen. Consequently the replica symmetric approach describes correctly the phase diagram. From this point of view the locked problems seems extremely easy compared to K -SAT.

On the other hand, experiments with the best known solvers of random CSPs show that the frozen phase of locked problems is very hard. And some of the very good solvers, e.g. the belief propagation based decimation, do not work at all even at the lowest connectivities (for an explanation see appendix F).

In fig. 4.5 we show the performance of the BP-REINFORCEMENT and the stochastic local search ASAT algorithms. Both the algorithms are described in appendix F, they are the best

we were able to find for the locked CSPs. The greediness parameter in the stochastic local search ASAT we evaluated as the most optimal is $p = 5 \cdot 10^{-5}$ for the 4-odd parity check, and $p = 3 \cdot 10^{-5}$ for the 1-or-3-in-5 SAT. In the BP-REINFORCEMENT the optimal forcing parameter π changes slightly with the connectivity. For the 1-or-3-in-5 SAT we used $\pi = 0.42$ for $2.9 \leq \bar{l} < 3.0$ and $\pi = 0.43$ for $3.0 \leq \bar{l} \leq 3.2$. For the 4-odd parity checks we used $\pi = 0.44$ for $2.75 \leq \bar{l} \leq 2.95$.

Of course, the parity check problem is an exceptional locked problem, as it is not NP-complete and can be solve via Gaussian elimination. However, our study shows that algorithms which do not use directly the linearity of the problem fail in the same way as they do in the NP-complete cases. Instances of the regular XOR-SAT indeed belong between the hardest benchmarks for all the best known satisfiability solvers which do not explore linearity of the problem, see e.g. [HJKN06].

Figure 4.5 puts in the evidence that in all the random locked problems the best known algorithms stop to be able to find solutions (in linear time) at the clustering transition. This supports the conjecture about freezing being relevant for algorithmical hardness. The locked problems are thus (at least until they are "unlocked") the new benchmarks of hard constraint satisfaction problems.

5 Coloring random graphs

In the previous three chapters we developed tools for describing the structure of solution and the phase diagram of random constraint satisfaction problems. These tools were applied to the problem of coloring random graphs in a series of works [KMRT⁺07, ZK07, KZ08b, KZ08a]. In this section we summarize the results.

5.1 Setting

Coloring of a graph is an assignment of colors to the vertices of the graph such that two adjacent vertices do not have the same color. The question is if on a given graph a coloring with q colors exists. Figure 5.1 gives an example of 3-coloring of a graphs with $N = 22$ vertices and $M = 27$ edges, the average connectivity is $c = 2M/N \approx 2.45$.

It is immediate to realize that the q -coloring problem is equivalent to the question of determining if the ground-state energy of a Potts anti-ferromagnet on a random graph is zero or not [KS87]. Consider indeed a graph $G = (V, E)$ defined by its vertices $V = \{1, \dots, N\}$ and edges $(i, j) \in E$ which connect pairs of vertices $i, j \in V$; and the Hamiltonian

$$\mathcal{H}(\{s\}) = \sum_{(i,j)} \delta(s_i, s_j). \quad (5.1)$$

With this choice there is no energy contribution for neighbours with different colors, but a positive contribution otherwise. The ground state energy is thus zero *if and only if* the graph is q -colorable. This transforms the coloring problem into a well-defined statistical physics model.

Studies of coloring of sparse random graphs have a long history in mathematics and computer science, see [ZK07] for some references. From the statistical physics perspective it was first studied in [vMS02], where the replica symmetric solution was worked out, and the replica symmetric stability was investigated numerically. Results were compared to Monte Carlo simulations and simulated annealing was used as a solver for coloring. The energetic 1RSB solution and the survey propagation algorithm for graph coloring were developed in [MPWZ02, BMP⁺03]. Subsequently [KPW04] studied the stability of the 1RSB solution and its large q limit. The entropic 1RSB solution was studies in [MPR05] for 3-coloring of Erdős-Rényi graphs. The entropic




Fig. 5.1. (Color online) Example of a proper 3-coloring of a small graph.

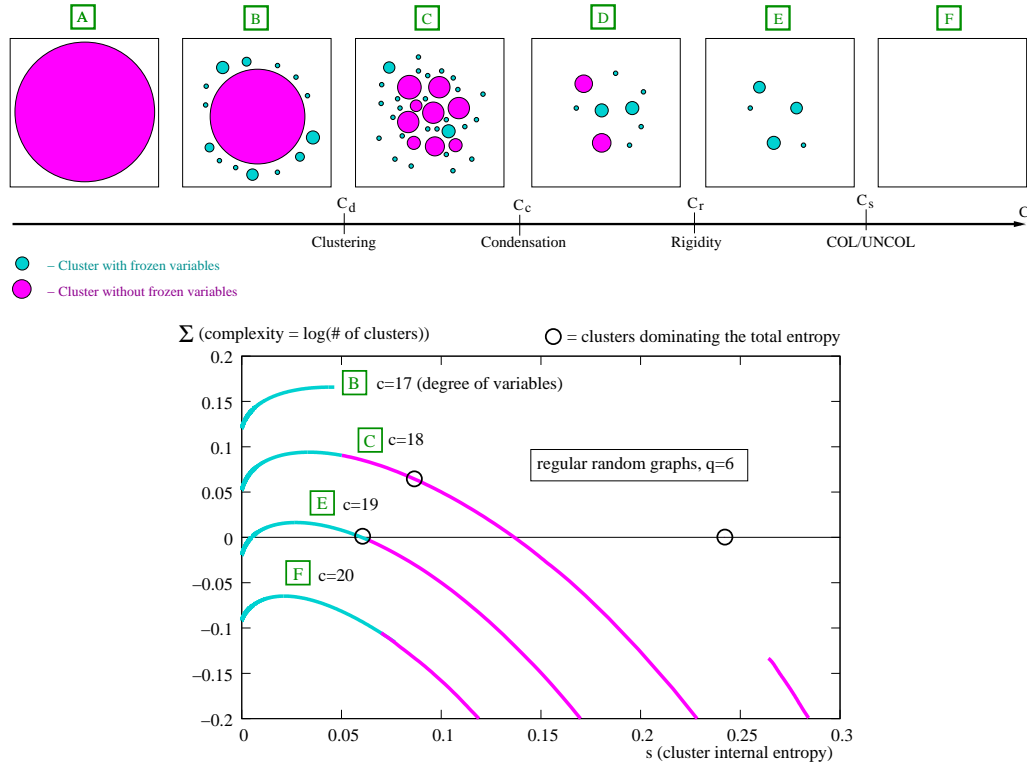


Fig. 5.2. (Color online) Up: Sketch of the structure of solutions in the random coloring problem. The depicted phase transitions arrive in the above order on Erdős-Rényi graphs for number of colors $4 \leq q \leq 8$. Down: Complexity (log-number) of clusters of a given entropy, $\Sigma(s)$, for 6-coloring or random regular graphs. The circles mark the dominating clusters, i.e., those which cover almost all solutions.

1RSB solution was, however, fully exploited only in [KMRT⁺07, ZK07, KZ08b, KZ08a] and the resulting phase diagram is discussed here.

5.2 Phase diagram

Figure 5.2 summarizes how does the structure of solutions of the coloring problem change when the average connectivity is increased, (A)→(F). In fig. 5.2 up, each colored "pixel" corresponds to one solution, and each circle to one cluster. As the average connectivity is increased, some solutions disappear and the overall structure of clusters changes. This is depicted in the six snapshots (A)→(F). The magenta clusters are the unfrozen ones, the cyan-blue clusters are the frozen ones. Figure 5.2 down, the corresponding complexity (log-number) of clusters of a given entropy, $\Sigma(s)$, computed from the 1RSB approach (2.28) for the 6-coloring of random regular graphs. More detailed description of the different phases for q -coloring follows.

- (A) **A unique cluster exists:** For connectivities low enough, all the proper colorings are found in a single cluster, where it is easy to “move” from one solution to another. Only one possible—and trivial—fixed point of the BP equations exists at this stage (as can be proved rigorously in some cases [BG06]). The entropy can be computed and reads in the large graph size limit

$$s = \frac{\log \mathcal{N}_{\text{sol}}}{N} = \log q + \frac{c}{2} \log \left(1 - \frac{1}{q} \right). \quad (5.2)$$

- (B) **Some (irrelevant) clusters appear:** As the connectivity is slightly increased, the phase space of solutions decomposes into a large (exponential) number of different clusters. It is tempting to identify that as the clustering transition. But in this phase all but one of these clusters contain relatively very few solutions, as compare to the whole set. Thus almost all proper colorings still belong to one single giant cluster, and the replica symmetric solution is correct, eq. (5.2) gives the correct entropy.
- (C) **The clustered phase:** For larger connectivities, the large single cluster decomposes into an exponential number of smaller ones: this now defines the genuine clustering threshold c_d . Beyond this threshold, a local algorithm that tries to move in the space of solutions will remain prisoner of a cluster of solutions for a diverging time [MS06c]. Interestingly, it can be shown that the total number of solutions is still given by eq. (5.2). Thus the free energy (entropy) has no singularity at the clustering transition, which is therefore not a phase transition in the sense of Ehrenfest. Only a diverging length scale (point-to-set correlation length) and time scale (the equilibration time) when c_d is approached justify the name “phase transition”.
- (D) **The condensed phase:** As the connectivity is increased further, another phase transition arises at the condensation threshold, c_c , where most of the solutions are found in a finite number of the largest clusters. Total entropy in the condensed phase is strictly smaller than (5.2). It has a non-analyticity at c_c therefore this is a genuine static phase transition. The condensation transition can be observed from the two-point correlation functions or from the overlap distribution.
- (E) **The rigid phase:** As explained in chapter 4, two different types of clusters exist. In the first type, the *unfrozen* ones, magenta in fig. 5.2, all variables can take at least two different colors. In the second type, *frozen* clusters, cyan in fig. 5.2, a finite fraction of variables is allowed only one color within the cluster and is thus “frozen” into this color. In the rigid phase, a random proper coloring belongs almost surely to a frozen cluster. Depending on the value of q , this transition may arise before or after the condensation transition (see tab. 5.1).
- (F) **The uncolorable phase:** Eventually, the connectivity c_s is reached beyond which no more solutions exist. The ground state energy is zero for $c < c_s$ and then grows continuously for $c > c_s$.

In table 5.1 we present all the critical values for coloring of Erdős-Rényi graphs, in table 5.2 for random regular graphs. Notice the special role of 3-coloring where the clustering and

condensation transitions coincide and are given by the local stability of the replica symmetric solution, see app. C. Notice also that for $q \geq 9$ in Erdős-Rényi graphs and $q \geq 8$ in regular graph the rigidity transition arrives before the condensation transition.

Few more words about the rigidity transition and the rigid phase in coloring. In sec. 4.2.2, next to the rigid phase, we also defined the *totally rigid* phase where almost all the clusters of every size become frozen. And the frozen phase where strictly all clusters become frozen. Note

Tab. 5.1. Critical connectivities c_d (dynamical, clustering), c_r (rigidity), c_c (condensation, Kauzmann) and c_s (colorability) for the phase transitions in the coloring problem on Erdős-Rényi graphs. The connectivities c_{SP} (where the first non trivial solution of SP appears) and $c_{r(m=1)}$ (where hard fields appear at $m = 1$) are also given. The error bars consist of the numerical precision on evaluation of the critical connectivities by the population dynamics technique, see appendix E.

q	c_d	c_r	c_c	c_s	c_{SP}	$c_{r(m=1)}$
3	4	4.66(1)	4	4.687(2)	4.42(1)	4.911
4	8.353(3)	8.83(2)	8.46(1)	8.901(2)	8.09(1)	9.267
5	12.837(3)	13.55(2)	13.23(1)	13.669(2)	12.11(2)	14.036
6	17.645(5)	18.68(2)	18.44(1)	18.880(2)	16.42(2)	19.112
7	22.705(5)	24.16(2)	24.01(1)	24.455(5)	20.97(2)	24.435
8	27.95(5)	29.93(3)	29.90(1)	30.335(5)	25.71(2)	29.960
9	33.45(5)	35.658	36.08(5)	36.490(5)	30.62(2)	35.658
10	39.0(1)	41.508	42.50(5)	42.93(1)	35.69(3)	41.508

Tab. 5.2. The transition thresholds for regular random graphs: c_{SP} is the smallest connectivity with a nontrivial solution at $m = 0$; the clustering threshold c_d is the smallest connectivity with a nontrivial solution at $m = 1$; the rigidity threshold c_r is the smallest connectivity at which hard fields are present in the dominant states, the condensation c_c is the smallest connectivity for which the complexity at $m = 1$ is negative and c_s the smallest uncolorable connectivity. Note that 3-coloring of 5-regular graphs is exactly critical for that $c_d = 5^+$. The rigidity transition may not exist due to the discreteness of the connectivities.

q	c_{SP}	c_d	c_r	c_c	c_s
3	5	5^+	-	6	6
4	9	9	-	10	10
5	13	14	14	14	15
6	17	18	19	19	20
7	21	23	-	25	25
8	26	29	30	31	31
9	31	34	36	37	37
10	36	39	42	43	44
20	91	101	105	116	117




Fig. 5.3. (Color online) Entropies and complexities as a function of the average connectivity for the 5-coloring of Erdős-Rényi graphs. The replica symmetric entropy is in dashed black, the total entropy in red. The complexity of dominant clusters in red. The total complexity, computed from the survey propagation, is in dashed blue.

that in the random graph coloring the rigidity transition coincides with the total rigidity transition for $q \leq 8$ for Erdős-Rényi graphs and for $q \leq 7$ for regular graphs. For larger values of q the rigidity transition is given by the $m = 1$ computation. We have not computed the total rigidity transition for larger q , but it is accessible from the present method. The freezing transition is, however, not accessible for the entropic 1RSB cavity approach. We cannot exclude that in the totally rigid phase there might still be some rare unfrozen clusters.

Note also an interesting feature about the 1RSB entropic solution; in fig. 5.2 down, for the connectivity $c = 17$ the function $\Sigma(s)$ consists of two branches. The low-entropy branch with frozen clusters, and the high-entropy branch with soft clusters. Note that the soft branch may also exist for positive values of complexity, e.g. in 4-coloring of Erdős-Rényi graphs. We interpreted the gap as the nonexistence of clusters of the corresponding size. The gap might, however, be an artifact of the 1RSB approximation which most likely does not describe correctly clusters of the corresponding size. For the discussion of correctness of the 1RSB solutions see appendix D.

To make the picture complete we plot the important complexities and entropies as a function of the average connectivity, for 5-coloring of Erdős-Rényi graphs see fig. 5.3. We plotted in dashed black the replica symmetric entropy (5.2), which in coloring is equal to the annealed one s_{ann} . The correct total entropy s_{tot} (in red) differs from the replica symmetric one in the condensed and uncolorable phase. The complexity of the dominating clusters (those covering almost all solutions) Σ_{dom} (in red, computed at $m = 1$) is non-zero between the clustering and the condensation transition. The total complexity Σ_{max} (in blue), maximum of the curves $\Sigma(s)$, can be computed in the region where survey propagation gives a nontrivial result. The colorability threshold corresponds to $\Sigma_{\text{max}} = 0$. We call c_{SP} the smallest connectivity at which survey propagation gives a nontrivial result, i.e., the part of the curve $\Sigma(s)$ with a zero slope




Fig. 5.4. (Color online) The fraction of solutions covered by the largest cluster as a function of the average connectivity for 4-coloring of Erdős-Rényi graphs. In the condensed phase the fraction covered by the largest cluster is not self-averaging and is determined by the Poisson-Dirichlet process with parameter m^* .

exists. Clusters exist also for $c < c_{SP}$, but computing their total complexity is more involved and we have not done it. The rigidity transition c_r cannot be determined from these quantities.

In fig. 5.4 we sketch what fraction of solutions is covered by the largest cluster as the average connectivity increases for 4-coloring of Erdős-Rényi graphs. In the replica symmetric phase $c < c_d$ the largest cluster covers almost all solutions. In the dynamical 1RSB phase the largest cluster covers an exponentially small fraction of solutions. In the condensed phase the largest state covers fraction of about $1 - m^*$ of solutions¹, but this part of the curve is not self-averaging. In the uncolorable phase there are no clusters of solutions, the ground state is made from one cluster.

5.3 Large q limit

The coloring of random graphs in the limit of large number of colors might seem a very unpractical and artificial problem. However, it allows many simplifications in the statistical description (rigorous or not) and a lot of insight can be obtained from this limit.

It is known from the cavity method, but also from a rigorous lower [ANP05] and upper [Luc91] bound that the colorability threshold for large number of colors scales like $2q \log q$. At the same time a very naive algorithm: Pick at random an uncolored vertex and assign it at random a color which is not assigned to any of its neighbours, was shown to work in polynomial (linear) time up to a connectivity scaling as $q \log q$. In other words this algorithm uses about twice as many colors than needed. Such a performance is not very surprising, a very naive algorithm

¹More precisely from the properties of the Poisson-Dirichlet process, described in sec. 3.3, if the fraction of solutions covered by the largest state is w then $1 - m^* = 1/\mathbb{E}(1/w)$.

performs half as good as possible. The surprise comes with the fact that it is an open problem if there is a polynomial algorithm which would work at connectivity $(1 + \epsilon)q \log q$ for an arbitrarily small positive ϵ .

5.3.1 The $2q \log q$ regime: colorability and condensation

The complexity function $\Sigma(s)$ at connectivity

$$c = 2q \log q - \log q + \gamma \quad (5.3)$$

where $\gamma = \Theta(1)$ was computed in [ZK07] and reads

$$\Sigma(s) = \frac{s}{\log 2} \left[1 - \log \frac{s}{\varepsilon \log 2} \right] - \varepsilon(2 + \gamma) + o(\varepsilon). \quad (5.4)$$

where $\varepsilon = 1/2q$. From this expression it is easy to see that the coloring threshold corresponds to

$$\gamma_s = -1. \quad (5.5)$$

and the condensation transition

$$\gamma_c = -2 \log 2. \quad (5.6)$$

Notice, as in [MZ08], that the complexity of the random subcubes model (3.5), sec. 3.1, gives exactly the expression (5.4) if we take the parameters of the random subcubes model as ²

$$p = 1 - \varepsilon, \quad \alpha = 1 + \varepsilon \frac{1 + \gamma}{\log 2}. \quad (5.7)$$

This is a striking property of the coloring problem in the limit of large number of colors near to the colorability threshold. The $1 - \varepsilon$ is a fraction of frozen variables in each cluster. Almost all the soft variables can take only one of two colors. The expression (5.4) means that the soft variables are mutually almost independent and the clusters have shape of small hypercubes. And the other way around, this property makes the random subcubes model more than just a pedagogical example of the condensation transition.

5.3.2 The $q \log q$ regime: clustering and rigidity

Another interesting scaling regime is defined as

$$c = q(\log q + \log \log q + \alpha), \quad (5.8)$$

where $\alpha = \Theta(1)$ is of order one. The large q scaling of the rigidity transition ($m = 1$) is easily expressed from (2.4):

$$\alpha_r = 1. \quad (5.9)$$

²We remind that in the section 3.1 entropies were logarithms of base 2 whereas everywhere else they are natural logarithms.




Fig. 5.5. (Color online) We plotted the difference $(c_r - c_d)/q = \alpha_r - \alpha_d$ and $(c_r - c_{SP})/q = \alpha_r - \alpha_{SP}$. The data are taken from table 5.1. The difference $\alpha_r - \alpha_{SP}$ indeed seems to converge to the theoretical $\log 2$, the difference $\alpha_r - \alpha_d$ seems to converge to around $1/4$.

This was originally computed in [KMRT⁺07, ZK07] and [Sem08]. The onset of a nontrivial solution for the survey propagation corresponds to the rigidity transition at $m = 0$ and reads [KPW04]

$$\alpha_{SP} = 1 - \log 2. \quad (5.10)$$

An empirical observation is that for $q = 3$ the threshold for survey propagation is smaller than the rigidity at $m = 1$, but for $q \geq 4$ the order changes and the distances between the two threshold grows with q . Based on this observation we conjectured that the clustering transition is

$$1 - \log 2 \geq \alpha_d \geq 1. \quad (5.11)$$

Note that recently the dynamical transition was proved to be $1 - \log 2 \geq \alpha_d$ [Sly09]. Figure 5.5 actually suggest that $\alpha_d \approx 1/4$. Its precise location is actually an interesting problem because it could shed light on the way soft fields converge to hard fields in the cavity approach.

Concerning the total rigidity transition, where almost all the clusters of all sizes become frozen, we have not manage to compute it in the large q limit. It is not even clear if the relevant scaling is as (5.8). The same is true for the even more interesting freezing transition, where all the clusters become frozen.

5.4 Finite temperature

It is interesting to study how does the antiferromagnetic Potts model, coloring at zero temperature, behave at finite temperature. In particular which of the zero temperature phase transitions survive to positive temperatures and what do they correspond to in the phenomenology of glasses. This has been done in [KZ08b] and we summarize the main results here.

The belief propagation equation for coloring (2.5) generalizes at finite temperature to

$$\psi_{s_i}^{i \rightarrow j} = \frac{1}{Z^{i \rightarrow j}} \prod_{k \in \partial i - j} [1 - (1 - e^{-\beta}) \psi_{s_i}^{k \rightarrow i}] \equiv \mathcal{F}_{s_i}(\{\psi^{k \rightarrow i}\}). \quad (5.12)$$

The distributional 1RSB equation (2.24) is the same.

- **The clustering transition** — becomes the dynamical phase transition T_d at positive temperature. The notion of reconstruction on trees, introduced in sec. 2.1.1, generalizes to positive temperatures. Constraints then play the role of noisy channels in the broadcasting. The dynamical temperature T_d is then defined via divergence of the point-to-set correlations (2.22). Or equivalently via the onset of a nontrivial solution of the 1RSB equations at $m = 1$. At the dynamical transition the point-to-set correlation length and the equilibration time diverge. There is however no non-analyticity in the free energy, Ehrenfest might thus not call it a phase transition.
- **The condensation transition** — becomes the Kauzmann phase transition T_K at positive temperature. The point at which the complexity function at $m = 1$ (structural entropy) becomes negative defines the Kauzmann temperature [Kau48]. At the Kauzmann temperature the free energy has a discontinuity in the second derivative. This corresponds to the discontinuity in the specific heat. Kauzmann transition is thus genuine even in the sense of Ehrenfest.
- **The rigidity transition** — is a purely zero temperature phase transition. At positive temperature the fields $\psi_{s_i}^{i \rightarrow j}$ (5.12) cannot be hard.
- **The colorability transition** — is a purely zero temperature phase transition. At the colorability threshold the ground state energy becomes positive (it has discontinuity in the first derivative). At a finite temperature, however, there is no corresponding non-analyticity.

Figure 5.6 shows the temperature phase diagram of 3- (left) and 4-coloring (right) on both Erdős-Rényi (up) and regular (down) random graphs. The dynamical temperature is in blue, the Kauzmann temperature in black.

The temperature at which the replica symmetric solution becomes locally unstable, see appendix C, is called T_{local} . In the terms of reconstruction on trees this is the Kesten-Stigum bound [KS66a, KS66b]. This temperature is a lower bound on the dynamical temperature T_d , but also on the Kauzmann temperature T_K . This is because below T_{local} the two-point correlations do not decay, which is possible only below T_K . Note that in the 3-coloring $T_d = T_K = T_{\text{local}}$ and this phase transition is continuous in the order parameter $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ (2.24). For $q \geq 4$ colors we find instead $T_d > T_K > T_{\text{local}}$ and the dynamical transition is discontinuous. At large connectivity, however, the three temperatures are very close, see fig. 5.6 where the T_{local} is in pink.

Correctness of the 1RSB solution — The last question concerns correctness of the 1RSB solutions itself. The local stabilities of the 1RSB solution are discussed in appendix D. The temperature at which the 1RSB solutions becomes type II locally unstable, see appendix D, is

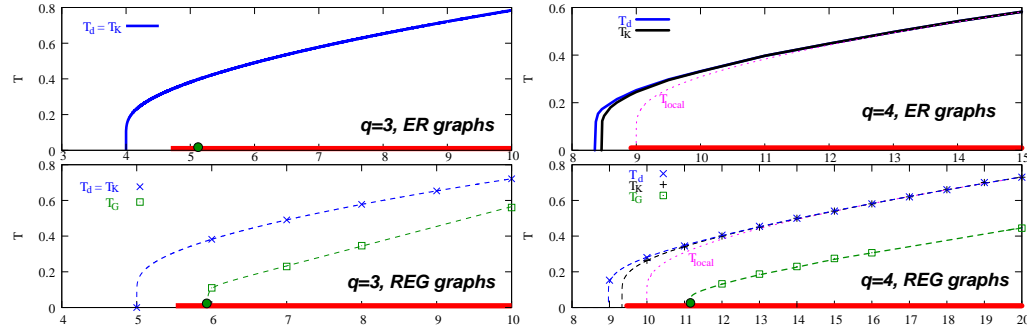


Fig. 5.6. (Color online) Phase diagrams for the 3-state (left) and 4-state (right) anti-ferromagnetic Potts glass on Erdős-Rényi graphs of average degree c (top) and regular graphs of degree c (bottom). For $q = 3$ the transition is continuous $T_d = T_K = T_{\text{local}}$. For $q = 4$, we find that $T_d > T_K > T_{\text{local}}$, while for larger connectivities these three critical temperatures become almost equal. The Gardner temperature T_G for regular graphs is also shown (green), below T_G the 1RSB solution is not correct anymore (for Erdős-Rényi graph we expect this curve to look similar). The bold (red) lines at zero temperature represent the uncolorable connectivities $c > c_s$.

called the Gardner temperature T_G [Gar85]. We computed it only on the ensemble of random regular graphs, see fig. 5.6, the T_G is in green. We do not know how to compute the stability of the type I, but we argued that the corresponding critical temperature should be smaller than the T_{local} . An important consequence is that in the colorable region the 1RSB solution is stable for $q \geq 4$ coloring.

Coloring with three colors is a bit special, as $T_{\text{local}} = T_d = T_K$. However, at small temperatures the stability of type I can be investigated from the energetic approach, again discussed in app. D. It follows that at least in interval $c \in (c_s, c_G) = (4.69, 5.08)$ the 1RSB solution is stable at low temperature. For $c > c_G$ on contrary the Gardner temperature is strictly positive. We cannot exclude that part of the colorable phase is unstable, but in such a case the unstable region would have a sort of re-entrant behaviour. Moreover the ferromagnetic fully connected 3-state Potts model has also a continuous dynamical transition $T_d = T_{\text{local}}$ yet it is 1RSB stable near to T_d [GKS85]. We thus find more likely that also the colorable phase of 3-coloring is 1RSB stable.

Finally, the local stability is only a necessary condition. The full correctness of the 1RSB approach have to be investigated from the 2RSB approach. We implemented the 2RSB on the regular coloring, the results are not conclusive, as the numerics is involved. but we have not found any sign for a nontrivial 2RSB solution in the colorable region.

6 Conclusions and perspectives

In this final section we highlight the, in our view, most interesting results of this thesis. More complete overview of the original contributions is presented in sec. 1.9. Scientific research is such that every answered question raises a number of new questions to be answered. We thus bring up a list of open problems which we find particularly intrinsic. Finally we give a brief personal view on the perspective applications of the results obtained in this work.

6.1 Key results

The main question underlying this study is: How to recognize if an NP-complete problem is typically hard and what are the main reasons for this?

In order to approach the answer we studied the structure of solutions in random constraint satisfaction problem - mainly in the graph coloring. We did not neglect the entropic contributions, as was common in previous studies, and this led to much more complete description of the phase diagram and associated phase transitions, see summarizing fig. 5.2.

The most interesting concept in these new findings was the freezing of variables. We pursued its study and investigated its relation to the average computational hardness. We introduced the *locked* constraint satisfaction, where the statistical description is easily solvable and the clustered phase is automatically frozen. We indeed observed empirically that these problems are much harder than the canonical K-satisfiability. They should thus become a new challenge for algorithmical development. As we mention in the perspectives, we also anticipate that the locked constraint satisfaction problems are of a more general interest.

6.2 Some open problems

(A) Clusters and their counting on trees — In sec. 2 we derived the 1RSB equations on purely tree graphs. Our derivation was, however, not complete as it is not straightforward why the complexity function should be counting the clusters as we defined them on trees. More physically founded derivations are for example the original one [MP00]. And also the one presented in [MM08] where the complexity is shown to count the fixed points of the belief propagation. We are, however, persuaded that the purely tree approach is more appealing from the probabilistic point of view, as treating correlations in the boundary conditions on trees is easier than treating the random graphs directly, for a recent progress see e.g. [Sly09, GM07, DM08]. This is why we chose to present this derivation despite its incompleteness.

In general we should say that creating better mathematical grounds for the replica symmetry breaking approach is a very important and challenging task.

(B) What is the meaning of the gap in the $\Sigma(s)$ function — We computed the number of clusters of a given entropy via the 1RSB method. For some intervals of parameters there is no solution corresponding to certain intermediate sizes. In other words there is a gap in the 1RSB function $\Sigma(s)$. See e.g. fig. 5.2, we observed such a gap in many other cases. Does this gap mean that there are truly no clusters of corresponding sizes or does it mean that the 1RSB method is wrong in that region or is there another explanation?

(C) Analysis of dynamical processes — In this thesis we described in quite a detail the static (equilibrium) properties of the constraint satisfaction problems. Very little is known about the dynamical properties – here we mean both the physical dynamics (with detailed balance) and the dynamics of algorithms. Focusing on results described here: the dynamics of the random subcubes model can be solved [MZ08], and the uniform belief propagation decimation can be analyzed [MRTS07], see also appendix F.1.4. However in general even the performance of simulated annealing as a solver is not known. And the understanding of why the survey propagation decimation works so well in 3-SAT and not that well in other problems, e.g. the locked problems or for larger K , is also very pure.

The most exciting conjecture of this work is the connection between the algorithmical hardness and freezing of variables. Several indirect arguments and empirical results were explained in sec. 4.4 to support this conjecture. It is, however, not very clear what is the detailed origin of the connection between presence of frozen variables in solutions and the fact that dynamics (of a solver) does not seem to be able to find them.

(D) Beyond random graphs and the thermodynamical limit — For practical application the perhaps most important point is to understand what is the relevance of our results for instances which are not random or not infinite. For example fig. 2.2 suggests that even on small random instances the clustering can be observed and is thus probably relevant. We also observed that the solutions-related quantities seems to have stronger finite size effects than the clusters-related properties, compare e.g. fig. 1.3 with fig. 4.1. This is an interesting point and it should be pursued.

6.3 Perspectives

This work should have a practical impact on the design of new solvers of constraint satisfaction problems. Instances with only frozen solutions should be used as new benchmarks for SAT solvers. At the same time where the design allows such instances should be avoided.

More concretely, the belief propagation algorithm is used as a standard approximative inference technique in artificial intelligence and information theory. One of the important problems with applications of the belief propagation is the fact that in many cases it does not converge. Many converging modifications were introduced. It might be interesting to investigate in this context the reinforced belief propagation, see appendix F.2.3, which sometimes converges towards a fixed point when the standard belief propagation does not. As the reinforcement algorithm seems to be very efficient, robust and is not theoretically well understood different variants of the implementation should be studied empirically. It would be interesting to see if this algorithm performs well on non-random graphs, or if it can provide information useful for the practical solvers. Several other concepts enhanced in this thesis might show up useful in algorithmic applications. We feel that the whitening of solutions might be one of them.

We introduced the locked constraint satisfaction problems as a new algorithmical challenge. Moreover the simplicity of their statistical description makes accessible several quantities which are difficult to compute in the K -SAT problem. For example the weight enumerator function or the x -satisfiability threshold. But these new models are exciting from many other points of view. Their hardness might be appealing for noise tolerant cryptographic applications. Planted

ensemble of the locked problems might be a very good one-way functions. The fact that the solutions of the locked problems are well separated makes them excellent candidates for nonlinear error correcting codes. It will be interesting to investigate if they can be advantageous over the standard linear low-density-parity-check codes [Gal62, Gal68, MN95, Mon01].

Clusters of solutions come up naturally in the pattern recognition and machine learning problems. There each cluster corresponds to a pattern which should be learned or recognized. Similarly the different phenotypes of a cell might be viewed as clusters of fixed points of the corresponding gene regulation network. The methods developed in this thesis might thus have impact also in these exciting fields.

Appendices

A 1RSB cavity equations at $m = 1$

Here we derive how the 1RSB equation (2.24) simplifies at $m = 1$ for the problems where the replica symmetric solution is not factorized. We restrict to the occupation models, but a generalization to other models is straightforward. Advantage of these equations is that the unknown object is not a functional of functionals but only a single functional. Moreover, the final self-consistent equation does not contain the reweighting term. This simplification makes implementation of the population dynamics at $m = 1$ much simpler, and thus the computation of the clustering and condensation transitions easier. Derivation of the corresponding equations for the K -SAT problem can be found in [MRTS08].

Write the RS equation (1.17) for the occupation problems in the form

$$\psi_{s_i}^{a \rightarrow i} = \frac{1}{Z^{j \rightarrow i}} \sum_{\{s_j\}} C_a(\{s_j\}, s_i) \prod_{j \in \partial a - i} \left(\prod_{b \in \partial j - a} \psi_{s_j}^{b \rightarrow j} \right) \equiv \mathcal{F}_{s_i}(\{\psi^{b \rightarrow j}\}), \quad (\text{A.1})$$

where the constraints $C_a(\{s_j\}, s_i) = 1$ if $\sum_j s_j + s_i \in A$, and 0 otherwise. Let $\mathcal{P}_{\text{RS}}(\psi)$ be the distribution of RS fields over the graph.

The 1RSB equations (2.24) at $m = 1$ are

$$P^{a \rightarrow i}(\psi^{a \rightarrow i}) = \frac{1}{Z^{j \rightarrow i}} \int \prod_{j \in \partial a - i} \prod_{b \in \partial j - a} [\mathrm{d}\psi^{b \rightarrow j} P^{b \rightarrow j}(\psi^{b \rightarrow j})] Z^{j \rightarrow i}(\{\psi^{b \rightarrow j}\}) \delta[\psi^{a \rightarrow i} - \mathcal{F}(\{\psi^{b \rightarrow j}\})] \equiv \mathcal{F}_2(\{P^{b \rightarrow j}\}). \quad (\text{A.2})$$

The averages over states

$$\overline{\psi}_{s_i}^{a \rightarrow i} = \int \mathrm{d}\psi_{s_i}^{a \rightarrow i} P^{a \rightarrow i}(\psi_{s_i}^{a \rightarrow i}) \psi_{s_i}^{a \rightarrow i} \quad (\text{A.3})$$

satisfy the RS equation (A.1). And consequently the RS and 1RSB normalizations are equal $Z^{j \rightarrow i} = \overline{Z}^{j \rightarrow i}$. The full order parameter is the probability distribution of P 's over the graph, it follow the self-consistent equation

$$\begin{aligned} \mathcal{P}_{\text{1RSB}}[P(\psi)] &= \sum_{l_1, \dots, l_{K-1}} q(l_1, \dots, l_{K-1}) \\ &\int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \left\{ \mathrm{d}P^j(\psi^j) \mathcal{P}_{\text{1RSB}}^j[P^j(\psi^j)] \right\} \delta[P(\psi) - \mathcal{F}_2(\{P^j\})]. \end{aligned} \quad (\text{A.4})$$

We define the average distribution $\overline{P}(\psi|\overline{\psi})$ on those edges where the RS field is equal to a given value $\overline{\psi}$

$$\overline{P}(\psi|\overline{\psi}) \mathcal{P}_{\text{RS}}(\overline{\psi}) \equiv \int \mathrm{d}P(\psi) \mathcal{P}_{\text{1RSB}}[P(\psi)] P(\psi) \delta\left[\overline{\psi} - \int \mathrm{d}\psi P(\psi) \psi\right]. \quad (\text{A.5})$$

Now we rewrite all the terms on the right hand side using the incoming fields and distributions, i.e., using first eq. (A.4) and then (A.2).

$$\begin{aligned}
\bar{P}(\psi|\bar{\psi})\mathcal{P}_{\text{RS}}(\bar{\psi}) &= \sum_{\{l\}} q(\{l\}) \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \left\{ dP^j(\psi^j) \mathcal{P}_{\text{1RSB}}^j[P^j(\psi^j)] \right\} \\
&\quad \mathcal{F}_2(\{P^j\}) \delta \left[\bar{\psi} - \int d\psi \mathcal{F}_2(\{P^j\}) \psi \right] \\
&= \sum_{\{l\}} q(\{l\}) \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \left\{ dP^j(\psi^j) \mathcal{P}_{\text{1RSB}}^j[P^j(\psi^j)] \right\} \\
&\quad \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} [d\psi^j P^j(\psi^j)] \frac{Z(\{\psi^j\})}{\mathcal{Z}} \delta [\psi - \mathcal{F}(\{\psi^j\})] \delta [\bar{\psi} - \mathcal{F}(\{\bar{\psi}^j\})] \\
&= \sum_{\{l\}} q(\{l\}) \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} [d\bar{\psi}^j \mathcal{P}_{\text{RS}}(\bar{\psi}^j)] \delta [\bar{\psi} - \mathcal{F}(\{\bar{\psi}^j\})] \\
&\quad \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} [d\psi^j \bar{P}^j(\psi^j|\bar{\psi}^j)] \frac{Z(\{\psi^j\})}{Z(\{\bar{\psi}^j\})} \delta [\psi - \mathcal{F}(\{\psi^j\})] , \tag{A.6}
\end{aligned}$$

where the original Dirac function was rewritten using

$$\begin{aligned}
\int d\psi \mathcal{F}_2(\{P^j\}) \psi &= \frac{1}{\mathcal{Z}} \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} [d\psi^j P^j(\psi^j)] Z(\{\psi^j\}) \int d\psi \psi \delta [\psi - \mathcal{F}(\{\psi^j\})] \\
&= \frac{1}{\mathcal{Z}} \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} [d\psi^j P^j(\psi^j)] Z(\{\psi^j\}) \mathcal{F}(\{\psi^j\}) \\
&= \mathcal{F}(\{\bar{\psi}^j\}) , \tag{A.7}
\end{aligned}$$

and in last equality was obtained using the integral of eq. (A.5)

$$\int d\bar{\psi} \bar{P}(\psi|\bar{\psi}) \mathcal{P}_{\text{RS}}(\bar{\psi}) = \int dP(\psi) \mathcal{P}_{\text{1RSB}}[P(\psi)] P(\psi) . \tag{A.8}$$

To simplify the equations further, in particular to get rid of the reweighting term $Z(\{\psi^j\})$, we define a distribution \bar{P}_s

$$\bar{\psi}_s \bar{P}_s(\psi|\bar{\psi}) \equiv \psi_s \bar{P}(\psi|\bar{\psi}) \quad \Rightarrow \quad \bar{P}(\psi|\bar{\psi}) = \sum_s \bar{\psi}_s \bar{P}_s(\psi|\bar{\psi}) , \tag{A.9}$$

then by factorizing the sum over components s we get

$$\begin{aligned} \bar{\psi}_s \bar{P}_s(\psi|\bar{\psi}) \mathcal{P}_{\text{RS}}(\bar{\psi}) &= \sum_{\{l\}} q(\{l\}) \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \left[d\bar{\psi}^j \mathcal{P}_{\text{RS}}(\bar{\psi}^j) \right] \delta \left[\bar{\psi} - \mathcal{F}(\{\bar{\psi}^j\}) \right] \\ &\quad \sum_{\{s_i\}} C(\{s_i\}, s) \frac{\prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \bar{\psi}_{s_i}^j}{Z(\{\bar{\psi}^j\})} \\ &\quad \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \left[d\psi^j \bar{P}_{s_i}^j(\psi^j|\bar{\psi}^j) \right] \delta \left[\psi - \mathcal{F}(\{\psi^j\}) \right]. \end{aligned} \quad (\text{A.10})$$

This final equation might look more complicated than the original one, but, in fact, it is much easier to solve. It could seem that we need a population of populations to represent the distribution $\bar{P}_s(\psi|\bar{\psi}) \mathcal{P}_{\text{RS}}(\bar{\psi})$. But keeping in mind that the proper initial conditions are

$$\bar{P}_1(\psi_1 = 1|\bar{\psi}) = 1, \quad \bar{P}_0(\psi_0 = 1|\bar{\psi}) = 1, \quad (\text{A.11})$$

independently of the RS field $\bar{\psi}$ we see that the probability distribution $\bar{P}_s(\psi|\bar{\psi}) \mathcal{P}_{\text{RS}}(\bar{\psi})$ may be represented by a population of triplets of fields - the first one corresponding to the RS field $\bar{\psi}$ and the other two corresponding to the two components (A.11).

In the population dynamics we first equilibrate the RS distribution $\mathcal{P}_{\text{RS}}(\bar{\psi})$ and then initialize the other two components according to (A.11). In every step of the update we first fix randomly the set of indexes $\{j\}$ and compute the new $\bar{\psi}$, then given the value s we choose the set of indexes $\{s_i\}$ according to a probability law given by the first line of eq. (A.10), then we compute the new ψ for $s = 0$ and $s = 1$ and change a random triplet in the population for the new values. In summary, eq. (A.10) allows to reduce the double-functional equations at $m = 1$ into a simple-functional form, which is much easier to solve.

The internal entropy $s = s_{\text{RS}} - \Sigma$, and thus also the complexity function, may be computed by making very similar manipulations as

$$\begin{aligned} s &= \alpha \sum_{\{l\}} q(\{l\}) \int \prod_{i=1}^K \prod_{j=1}^{l_i} \left[d\bar{\psi}^j \mathcal{P}_{\text{RS}}(\bar{\psi}^j) \right] \frac{\sum_{\{s_i\}} C(\{s_i\}) \prod_{i=1}^K \prod_{j=1}^{l_i} \bar{\psi}_{s_i}^j}{Z^{a+\partial a}(\{\bar{\psi}^j\})} \\ &\quad \int \prod_{i=1}^K \prod_{j=1}^{l_i} \left[d\psi^j \bar{P}_{s_i}^j(\psi^j|\bar{\psi}^j) \right] \log Z^{a+\partial a}(\{\psi^j\}) \\ &- \sum_l \mathcal{Q}(l)(l-1) \int \prod_{i=1}^l \left[d\bar{\psi}^i \mathcal{P}_{\text{RS}}(\bar{\psi}^i) \right] \frac{\sum_{s_i} \prod_{i=1}^l \bar{\psi}_{s_i}^i}{Z^i(\{\bar{\psi}^i\})} \\ &\quad \int \prod_{i=1}^l \left[d\psi^i \bar{P}_{s_i}^i(\psi^i|\bar{\psi}^i) \right] \log Z^i(\{\psi^i\}). \end{aligned} \quad (\text{A.12})$$

We can also express other quantities, e.g. the inter $q_0 = q_{RS}$ and intra q_1 state overlaps.

$$\begin{aligned} q_1 &= \int dP(\psi) \mathcal{P}_{\text{IRSB}} \int d\psi P(\psi) \sum_{\sigma} \psi_{\sigma} \\ &= \sum_{\sigma, s} \int d\bar{\psi} \mathcal{P}_{\text{RS}}(\bar{\psi}) \bar{\psi}_s \int d\psi \bar{P}_s(\psi | \bar{\psi}) \psi_{\sigma}^2. \end{aligned} \quad (\text{A.13})$$

Factorized RS solution — Several times, see e.g. sec. 4.3.3, we used the equations at $m = 1$ for problems with factorized RS solution, $\mathcal{P}_{\text{RS}}(\psi) = \delta(\psi - \bar{\psi})$. The derivation is straightforward from (A.10)

$$\begin{aligned} \bar{P}_s(\psi) &= \sum_{\{l\}} q(\{l\}) \frac{1}{\bar{\psi}_s Z} \sum_{\{s_i\}} C(\{s_i\}, s) \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} \bar{\psi}_{s_i}^j \\ &\quad \int \prod_{i=1}^{K-1} \prod_{j=1}^{l_i} d\bar{P}_{s_i}(\psi^j) \delta(\psi - \mathcal{F}(\psi^j)). \end{aligned} \quad (\text{A.14})$$

Proper initial conditions for the population dynamics resolution of (A.14) is $\bar{P}_s(\psi_s = 1) = 1$.

At zero temperature the distributions can be written as the sum of the frozen and soft part

$$\bar{P}_1(\psi) = \mu_1 \delta\left(\psi - \begin{pmatrix} 1 \\ 0 \end{pmatrix}\right) + (1 - \mu_1) \tilde{P}_1(\psi), \quad (\text{A.15a})$$

$$\bar{P}_0(\psi) = \mu_0 \delta\left(\psi - \begin{pmatrix} 0 \\ 1 \end{pmatrix}\right) + (1 - \mu_0) \tilde{P}_0(\psi). \quad (\text{A.15b})$$

Self-consistent equations for the fractions of hard fields μ_1, μ_0 (4.20a-4.20b) follow from (A.14).

B Exact entropy for the balanced LOPs

Rigorous results about the entropy and the satisfiability threshold can be obtain comparing the first and second moment of the number of solutions, that is: If a number of solution on a graph G is \mathcal{N}_G then the first moment is average over the graph ensemble:

$$\langle \mathcal{N}_G \rangle = \sum_{\{\sigma\}} \text{Prob}(\{\sigma\} \text{ is SAT}) . \quad (\text{B.1})$$

The second moment is

$$\langle \mathcal{N}_G^2 \rangle = \sum_{\{\sigma_1\}, \{\sigma_2\}} \text{Prob}(\{\sigma_1\} \text{ and } \{\sigma_2\} \text{ are both SAT}) . \quad (\text{B.2})$$

The Markov inequality then gives an upper bound on the entropy and the satisfiability threshold

$$\text{Prob}(\mathcal{N}_G > 0) \leq \langle \mathcal{N}_G \rangle . \quad (\text{B.3})$$

The Chebyshev's inequality gives a lower bound via

$$\text{Prob}(\mathcal{N}_G > 0) \geq \frac{\langle \mathcal{N}_G \rangle^2}{\langle \mathcal{N}_G^2 \rangle} . \quad (\text{B.4})$$

B.1 The 1st moment for occupation models

Let us remind that the occupation models are defined via a $(K + 1)$ -component vector A , such that $A_i = 1$ if and only if there can be i occupied particles around a constraint of K variables. We consider by default $A_0 = A_K = 0$, i.e., that everybody full of empty is not a solution. We also consider all the M constraints are the same. We have $Q(l)N$ variables of connectivity l , where $\sum_{l=0}^{\infty} Q(l) = 1$ and $\bar{l} = \sum_{l=0}^{\infty} lQ(l) = KM/N$.

In order to compute the first moment we divide variables into groups according to their connectivity and in each groups we choose fraction t_l of occupied variables. Number of ways in which this is possible is then multiplied by a probability that such a configuration satisfies simultaneously all the constraints.

$$\begin{aligned} \langle \mathcal{N}_G \rangle &= \int_0^1 dt \sum_{\{t_l\}} \prod_l \binom{Q(l)N}{t_l Q(l)N} \sum_{r_1, \dots, r_M=1}^K \prod_{a=1}^M \delta(A_{r_a} - 1) \\ &\times \binom{N \sum_l l(1-t_l)Q(l)}{(K-r_1) \dots (K-r_M)} \binom{N \sum_l l t_l Q(l)}{r_1 \dots r_M} \left[\binom{\bar{l}N}{K \dots K} \right]^{-1} \\ &\times \delta \left(\sum_{a=1}^M r_a - \bar{l}tN \right) \delta \left(t\bar{l}N - \sum_l l t_l Q(l)N \right), \end{aligned} \quad (\text{B.5})$$

where t is the total fraction of occupied variables, this variable might seem ambiguous, as it can be integrated out, but we will appreciate its usefulness later, r_a is a number of occupied variables in a constraint a .

We develop expression (B.5) in the exponential order. In order to do so we exchange the last two delta functions by their Fourier transforms, introducing two complex Lagrange parameters $\log x$ and $\log u$.

$$\begin{aligned} \langle \mathcal{N}_G \rangle \approx & \int dt \int \prod_l dt_l \int dx \int du \exp N \left\{ - \sum_l Q(l) [t_l \log t_l + (1 - t_l) \log (1 - t_l)] \right. \\ & + \bar{l} [t \log t + (1 - t) \log (1 - t)] + \log u \left[\sum_l l t_l Q(l) - t \bar{l} \right] \\ & \left. + \frac{\bar{l}}{K} \log \left[\sum_{r=1}^K \delta(A_r - 1) \binom{K}{r} x^r \right] - t \bar{l} \log x \right\}. \end{aligned} \quad (\text{B.6})$$

Saddle point with respect to parameters t_l gives us

$$t_l = \frac{u^l}{1 + u^l}, \quad (\text{B.7})$$

and we call $p_A(x) = \sum_{r=1}^K \delta(A_r - 1) \binom{K}{r} x^r$. Using this we have

$$\begin{aligned} \langle \mathcal{N}_G \rangle \approx & \int dt dx du \exp N \left\{ \frac{\bar{l}}{K} \log p_A(x) - t \bar{l} \log x \right. \\ & \left. + \sum_l Q(l) \log (1 + u^l) - t \bar{l} \log u + \bar{l} [t \log t + (1 - t) \log (1 - t)] \right\}. \end{aligned} \quad (\text{B.8})$$

The saddle point equations read

$$\partial_u : \quad t = \frac{1}{\bar{l}} \sum_l l Q(l) \frac{u^l}{1 + u^l}, \quad (\text{B.9a})$$

$$\partial_x : \quad t = \frac{x \partial_x p_A(x)}{K p_A(x)}, \quad (\text{B.9b})$$

$$\partial_t : \quad t = \frac{xu}{1 + xu}, \quad (\text{B.9c})$$

As the parameter t is the only physically meaningful from the three, the goal is to express the annealed entropy as a function of t and find its maxima. We do that by inverting numerically (B.9a) and plugging (B.9c) in (B.8). Eq. (B.9c) then express the saddle point with respect to the parameter t . We can write

$$s_{\text{ann}}(t) = \sum_l Q(l) \log [1 + u(t)^l] + \frac{\bar{l}}{K} \log p_A(t), \quad (\text{B.10})$$

where

$$p_A(t) = \sum_{r=1}^K \delta(A_r - 1) \binom{K}{r} \left(\frac{t}{u(t)} \right)^r (1 - t)^{K-r}, \quad (\text{B.11})$$

where $u(t)$ is an inverse of (B.9a).

For the regular graphs $Q(l) = \delta(l - L)$ the inverse of (B.9a) is explicit $u = [t/(1 - t)]^{1/L}$ and thus

$$s_{\text{ann reg}}(t) = \frac{L}{K} \log \left\{ \sum_{r=1}^K \delta(A_r - 1) \binom{K}{r} [t^r (1 - t)^{K-r}]^{\frac{L-1}{L}} \right\}. \quad (\text{B.12})$$

B.2 The 2nd moment for occupation models

The second moment is computed in a similar manner. First we fix that in a fraction $t_{x,l}$ of nodes of connectivity l the variable is occupied in both the solutions σ_1, σ_2 in (B.2). In a fraction $t_{y,l}$ the variable is occupied in σ_1 and empty in σ_2 and the other way round for $t_{z,l}$. We sum over all possible combinations of $0 \leq t_{x,l}, t_{y,l}, t_{z,l}$ such that $\sum_{w=x,y,z} t_{w,l} \leq 1$. All this is multiplied by the probability that such two configurations σ_1, σ_2 both satisfy all the constraints.

$$\begin{aligned} \langle \mathcal{N}_G^2 \rangle &= \int dt_x dt_y dt_z \sum_{\{t_{x,l}\}, \{t_{y,l}\}, \{t_{z,l}\}} \prod_l \binom{Q(l)N}{(t_{x,l}Q(l)N) (t_{y,l}Q(l)N) (t_{z,l}Q(l)N)} \\ &\quad \sum_{r_{x,1}, \dots, r_{x,M}} \sum_{r_{y,1}, \dots, r_{y,M}} \sum_{r_{z,1}, \dots, r_{z,M}} \prod_{a=1}^M \delta(A_{r_{x,a}+r_{y,a}} - 1) \delta(A_{r_{x,a}+r_{z,a}} - 1) \\ &\quad \left(\binom{N \sum_l l (1 - \sum_{w=x,y,z} t_{w,l}) Q(l)}{(K - \sum_{w=x,y,z} r_{w,1}) \dots (K - \sum_{w=x,y,z} r_{w,M})} \right) \\ &\quad \prod_{w=x,y,z} \binom{N \sum_l l t_{w,l} Q(l)}{r_{w,1} \dots r_{w,M}} \left[\binom{\bar{l}N}{K \dots K} \right]^{-1} \\ &\quad \prod_{w=x,y,z} \delta \left(\sum_{a=1}^M r_{w,a} - \bar{l} t_{w,l} N \right) \delta \left(t_w \bar{l} N - \sum_l l t_{w,l} Q(l) N \right). \quad (\text{B.13}) \end{aligned}$$

We introduce Fourier transforms at a place of both the Dirac functions, the conjugated parameters are $\log x, \log y, \log z$ for the first Dirac function, and $\log u_x, \log u_y, \log u_z$ for the second one. After that we suppress the parameters $t_{w,l}$ in the same manner as we did for the first moment. We obtain for the second moment entropy

$$\begin{aligned} s_{2\text{nd}} &= \bar{l} [t_x \log t_x + t_y \log t_y + t_z \log t_z + (1 - t_x - t_y - t_z) \log (1 - t_x - t_y - t_z)] \\ &\quad - \bar{l} (t_x \log x + t_y \log y + t_z \log z) + \frac{\bar{l}}{K} \log p_A(x, y, z) \\ &\quad + \sum_l Q(l) \log (1 + u_x^l + u_y^l + u_z^l) - \bar{l} (t_x \log u_x + t_y \log u_y + t_z \log u_z), \quad (\text{B.14}) \end{aligned}$$

where

$$\begin{aligned} p_A(x, y, z) &= \sum_{r_1, r_2=0}^K \delta(A_{r_1} A_{r_2} - 1) \sum_{s=\max(0, r_1+r_2-K)}^{\min(r_1, r_2)} \binom{K}{(r_1-s)(r_2-s)s} \\ &\quad \times x^s y^{(r_1-s)} z^{(r_2-s)}, \quad (\text{B.15}) \end{aligned}$$

and the saddle point with respect to t_w , w and u_w ($w = x, y, z$) is

$$\partial_{t_w} : \quad t_w = \frac{1}{\bar{l}} \sum_l l Q(l) \frac{u_w^l}{1 + u_x^l + u_y^l + u_z^l}, \quad w = x, y, z, \quad (\text{B.16a})$$

$$\partial_w : \quad t_w = \frac{w \partial_w p_A(x, y, z)}{K p_A(x, y, z)}, \quad w = x, y, z, \quad (\text{B.16b})$$

$$\partial_{u_w} : \quad w u_w = \frac{t_w}{1 - t_x - t_y - t_z}, \quad w = x, y, z. \quad (\text{B.16c})$$

Once again the parameters t_w are physically meaningful, so we want to express $s_{2\text{nd}}$ as a function of these. We thus need to inverse (B.16a), note that such an inverse is well defined, and using (B.16c) we obtain

$$\begin{aligned} s_{2\text{nd}}(t_x, t_y, t_z) &= \frac{\bar{l}}{K} \log p_A(t_x, t_y, t_z) \\ &+ \sum_l Q(l) \log \left\{ 1 + \sum_{w \in \{x, y, z\}} [u_w(t_x, t_y, t_z)]^l \right\}, \end{aligned} \quad (\text{B.17})$$

where

$$\begin{aligned} p_A(t_x, t_y, t_z) &= \sum_{r_1, r_2=0}^K \delta(A_{r_1} A_{r_2} - 1) \sum_{s=\max(0, r_1+r_2-K)}^{\min(r_1, r_2)} \binom{K}{(r_1-s)(r_2-s)s} \\ &\left(\frac{t_x}{u_x(t_x, t_y, t_z)} \right)^s \left(\frac{t_y}{u_y(t_x, t_y, t_z)} \right)^{(r_1-s)} \\ &\left(\frac{t_z}{u_z(t_x, t_y, t_z)} \right)^{(r_2-s)} (1 - t_x - t_y - t_z)^{(K-r_1-r_2+s)}. \end{aligned} \quad (\text{B.18})$$

The global maximum with respect to t_x, t_y, t_z needs to be found.

For the regular ensemble $Q(l) = \delta(l - L)$ the function (B.16a) is explicitly reversible and the final expression for the second moment entropy simplifies significantly

$$\begin{aligned} s_{2\text{nd}, \text{reg}}(t_x, t_y, t_z) &= \frac{L}{K} \log \left\{ \sum_{r_1, r_2, s} \frac{K! \delta(A_{r_1} - 1) \delta(A_{r_2} - 1)}{(r_1 - s)! (r_2 - s)! s! (K - r_1 - r_2 + s)!} \right. \\ &\left. \left[t_x^s t_y^{(r_1-s)} t_z^{(r_2-s)} (1 - \sum_w t_w)^{(K-r_1-r_2+s)} \right]^{\frac{L-1}{L}} \right\}, \end{aligned} \quad (\text{B.19})$$

where the range of summations is the same as in (B.18).

B.3 The results

The main result is that for some of the symmetric ($A_{K-r} = A_r$ for all $r = 0, \dots, K$) and locked occupation problems ($Q(0) = Q(1) = 0$) the first and second moments computation leads the exact entropy of solutions (4.18). And thus also the exact satisfiability threshold. The cases

where this statement holds are marked by a * in tab. 4.1, and we call them *balanced* LOPs. We observed that some of the balanced problems A are created iteratively starting from 010 or 01010 and adding

$$A^{K+2} = 0A^K0, \quad A^{K+4} = 01A^K10. \quad (\text{B.20})$$

We, however, found also other balanced cases than (B.20). The simplest example of symmetric locked problem which is not balanced is $A = 010010$, and many others of higher K .

Let us now show this result. For all the symmetric occupation problems:

- The annealed entropy (B.10) has a stationary point at $t = 1/2$ ($u = 1, x = 1$). At this stationary the entropy evaluates to (4.18).
- The second moments entropy (B.17) has a stationary point at $t_x = t_y = t_z = 1/4$ ($u_x = u_y = u_z = 1, x = y = z = 1$). At this stationary point the second moment entropy evaluates to twice the (4.18). To prove this statement observe that for the symmetric problems $p_A(1/4, 1/4, 1/4) = [p_A(1/2)]^2$. This last identity can be derived from the Vandermonde's combinatorial identity

$$\binom{K}{r_2} = \sum_{s=0}^{r_1} \binom{r_1}{s} \binom{K-r_1}{r_2-s}. \quad (\text{B.21})$$

- The second moment entropy has another stationary point at $t_x = 1/2, t_y = t_z = 0$ or $t_x = 0, t_y = t_z = 1/2$. This stationary point is equal to the first moment entropy at $t = 1/2$.

In the problems where one of the above stationary points is the global maximum the annealed entropy is exact and the satisfiability threshold easily calculable from (4.18).

In the symmetric problems with leaves ($Q(1) > 0$), or those which are not locked (e.g. 0110) or not balanced (e.g. 010010) another competing maximum of the second moment entropy appears before the annealed entropy goes to zero.

We investigated numerically that this does not happen for the balanced problems described by the recursion (B.20). So far we were not able to prove this last point analytically. This is, however, a technical problem, much simpler than the original one.

The main message of this analysis is what are the ingredients of the model which make the satisfiability threshold accessible to the second moment computations. Here we showed that it is on one hand the (unbroken) symmetry of the problem and on the other hand the point-like clusters. Such a general result might be surprising because otherwise the satisfiability threshold is known exactly in only a handful of the NP-complete problems [ACIM01, MZK⁺99a, AKKK01, CM04].

C Stability of the RS solution

In chapter 2 we argued in detail that the replica symmetric solution is correct if and only if the point-to-set correlations decay to zero, or equivalently if the reconstruction is not possible. Failure of the RS solution may (but does not have to) manifest itself via the divergence of the spin glass susceptibility. In a system with Ising variables $s_i \in \{-1, +1\}$ this is defined as

$$\chi_{\text{SG}} = \frac{1}{N} \sum_{i,j} \langle s_i s_j \rangle_c^2, \quad (\text{C.1})$$

where $\langle \cdot \rangle_c$ is the connected expectation with respect to the Boltzmann measure.

Originally the replica symmetric instability was investigated from the spectrum of the Hessian matrix in a celebrated paper by de Almeida and Thouless [dAT78]. Equivalence between the RS stability and the convergence of the belief propagation equations on a single large graph is also often stated. In the reconstruction on trees this corresponds to the Kesten-Stigum condition [KS66a, KS66b]. It is not straightforward to see that all these statements are equivalent. We thus try to put a bit of order to the different ways of expressing the stability of the RS solution¹.

C.1 Several equivalent methods for RS stability

Susceptibility chains — Perhaps the most direct way how to investigate the divergence of the spin glass susceptibility (C.1) is to write

$$\chi_{\text{SG}} \approx \sum_i \mathbb{E}(\langle s_i s_0 \rangle_c^2) \approx \sum_d \gamma^d \mathbb{E}(\langle s_d s_0 \rangle_c^2), \quad (\text{C.2})$$

where s_0 is a typical variable (the origin), s_d is a variable at distance d from s_0 , and γ^d is the typical number of variables at distance d from s_0 ($\gamma = \bar{l}^2 / \bar{l} - 1$). The average $\mathbb{E}(\cdot)$ is over the randomness of the graph. The spin glass susceptibility diverges if and only if $\lambda > 1$ where

$$\lambda = \gamma \lim_{d \rightarrow \infty} \left[\mathbb{E}(\langle s_d s_0 \rangle_c^2) \right]^{\frac{1}{d}} \quad (\text{C.3})$$

Using the fluctuation dissipation theorem we can rewrite

$$\mathbb{E}(\langle s_0 s_d \rangle_c^2) \approx \mathbb{E} \left[\left(\frac{\partial h_0}{\partial h_d} \right)^2 \right] = \mathbb{E} \left[\prod_{i=1}^d \left(\frac{\partial h_{i-1}}{\partial h_i} \right)^2 \right], \quad (\text{C.4})$$

where h_0, \dots, h_d is a sequence of cavity fields (1.34) on the shortest path from s_0 to s_d . The dependence of the cavity field h_i on h_{i-1} is given by the belief propagation equations. This method to investigate the RS stability was used e.g. in [MMR05] or [ZM06]. It is numerically involved and not very precise as in practice d can be taken only at maximum 10 – 20.

¹This overview has been worked out in collaboration with F. Krzakala and F. Ricci-Tersenghi.

Noise propagation — Call v_d^0 the contribution to the spin glass susceptibility from the layer of variables at a distance d from 0

$$v_d^0 = \sum_{k, |k,0|=d} \left(\frac{\partial h_0}{\partial h_k} \right)^2 = \sum_{i \in \partial 0} \left(\frac{\partial h_0}{\partial h_i} \right)^2 \sum_{k, |k,i|=d-1} \left(\frac{\partial h_i}{\partial h_k} \right)^2 = \sum_{i \in \partial 0} \left(\frac{\partial h_0}{\partial h_i} \right)^2 v_{d-1}^i, \quad (\text{C.5})$$

where h_k are cavity fields at distance d from h_0 , and the sum is over all the cavity fields needed to compute h_0 . The spin glass susceptibility diverges if and only if the numbers v_d are on average growing with the distance d .

The evolution of numbers v can be followed via the population dynamics method. Next to the population of fields h we keep also a population of positive numbers v . When a field h_0 is updated according to the belief propagation equations, we update also the number v^0 according to (C.5). The RS solution is stable if and only if the overall sum $\sum_i v^i$ is decreasing during the population dynamics updates. This method was implemented e.g. in [MS06a] or [RSZ07]. It is simple and numerically very precise.

Deviation of two replicas — Consider a general form of the belief propagation equations $h = f(\{h_i\})$. After averaging over the graph ensemble we obtain distributional equations (1.23a-1.23b) which are solved via the population dynamics technique. Consider now two replicas of the resulting population, each element i differs by δh_i . Keep running the population dynamics on both these replicas and record how the differences δh_i are changing

$$\delta h_0 = \sum_{i \in \partial 0} \frac{\partial h_0}{\partial h_i} \delta h_i. \quad (\text{C.6})$$

The differences δh can be negative and positive. Take $v = (\delta h)^2$ then

$$v_0 = \left(\sum_{i \in \partial 0} \frac{\partial h_0}{\partial h_i} \delta h_i \right)^2 = \sum_{i \in \partial 0} \left(\frac{\partial h_0}{\partial h_i} \right)^2 v_i + \sum_{i \neq j} \frac{\partial h_0}{\partial h_i} \frac{\partial h_0}{\partial h_j} \delta h_i \delta h_j. \quad (\text{C.7})$$

The second term can be neglected because the terms δh_i and δh_j are independent. This brings us back to the equation (C.5).

Thus the replica symmetric solutions is stable if and only if the two infinitesimally different replicas do not deviate one from another. This method is very fast to implement and is thus useful for preliminary checks of the RS stability.

Convergence of the belief propagation — The stability of replica symmetric solutions is equivalent to the convergence of the belief propagation equations on a large random graph. This fact follows directly from the previous paragraph. Eq. (C.6) gives the rate of convergence (divergence) of two nearby trajectories of the dynamical map defined by the BP iterative equations.

Variance propagation — Often a "variance" formulation of the stability is described. Assume that instead of a value h_i on every link, there is a narrow distribution of values $g(h_i)$ parameter-

ized by a mean \bar{h}_i and a small variance v_i . How does \bar{h} and v evolve? We have now

$$\bar{h} = \int dh g(h) h = \int \prod_i [dh_i g_i(h_i)] f(\{h_i\}), \quad (\text{C.8})$$

$$v = \int dh g(h) (h - \bar{h})^2 = \int \prod_i [dh_i g_i(h_i)] f^2(\{h_i\}) - (\bar{h})^2, \quad (\text{C.9})$$

where $h = f(\{h_i\})$ is the belief propagation equation. However, since the variance is infinitesimal, the variation of h_i around \bar{h}_i is very small, so that

$$f(\{h_i\}) = f(\{\bar{h}_i\}) + \sum_i (h_i - \bar{h}_i) \left. \frac{\partial f(\{h_i\})}{\partial h_i} \right|_{\bar{h}_i}, \quad (\text{C.10})$$

and therefore one obtains $\bar{h} = f(\{\bar{h}_i\})$ and

$$v = \sum_i v_i \left(\left. \frac{\partial f(\{h_i\})}{\partial h_i} \right|_{\bar{h}_i} \right)^2, \quad (\text{C.11})$$

which is nothing else then equations (C.5).

Numerical instability towards the 1RSB solution — The RS stability can also be investigated from the numerical stability of the trivial solution of the 1RSB equations. Indeed if the distribution of fields over states is regarded the probability distribution of a small variance $g(h)$ then the 1RSB equation (2.24) gives for a p^{th} moment of $g(h)$

$$\bar{h}^p = \frac{1}{Z} \int \prod_i [dh_i g_i(h_i)] Z^m(\{h_i\}) f^p(\{h_i\}), \quad (\text{C.12})$$

where Z is the normalization of the BP equations and its m^{th} power is the reweighting factor. Expansion gives

$$Z^m(\{h_i\}) = Z^m(\{\bar{h}_i\}) + m Z^{m-1}(\{h_i\}) \sum_i (h_i - \bar{h}_i) \left. \frac{\partial Z(\{h_i\})}{\partial h_i} \right|_{\bar{h}_i}. \quad (\text{C.13})$$

The equations for the variances (C.11) does not depend on the second term from (C.13), as this is of a smaller order. As a consequence the condition for stability is independent of the parameter m .

It is quite remarkable fact that the divergence of the spin glass susceptibility corresponds to the appearance of a nontrivial solution of the 1RSB equation at *all* the values of m . In particular because we observed that when the instability is not present the onset of a nontrivial 1RSB solution is m dependent, see e.g. fig. D.2.

The eigenvalues of the Hessian — The replica symmetric solution is a minimum of the Gibbs free energy. This is often investigated from the spectra of the matrix of second derivatives called Hessian. The equivalence between this approach and the divergence of the spin glass susceptibility is a classical result, see e.g. the book of Fischer and Hertz [FH91], page 98-100.

C.2 Stability of the warning propagation

At zero temperature the necessary (but not sufficient) condition for the replica symmetric solution to be stable is the convergence of the warning propagation on a single graph. Obviously if the warning propagation does not converge then BP does not either, and convergence of the BP is equivalent to the replica symmetric stability. Advantage of the investigation of the warning propagation convergence is that it can be treated analytically, without using the population dynamics method.

Consider a model with Ising spins where the warnings u (1.35b) can take only three possible values $u \in \{-1, 0, 1\}$. Consider warning u and one of the warnings on which u depends, say u_0 . Except u_0 the warning u depends also on u_1, \dots, u_k , where k is distributed according to $\mathcal{Q}(k)$. The degree distribution conditioned on the presence of two edges is

$$\tilde{\mathcal{Q}}(k-2) = \frac{k(k-1)}{k^2 - \bar{k}} \mathcal{Q}(k). \quad (\text{C.14})$$

Call $P(a \rightarrow b | c \rightarrow d)$ the probability that the warning u changes from value a to value b provided that the warning u_0 was changed from value c to value d . This probability can be always computed from the probabilities p_-, p_0, p_+ that a warning $u = -1, 0, +1$

$$P(a \rightarrow b | c \rightarrow d) = \sum_k \tilde{\mathcal{Q}}(k) P_k(p_-, p_0, p_+; a \rightarrow b | c \rightarrow d), \quad (\text{C.15})$$

where the function P_k depends on the model in consideration. This probability describes a proliferation of a "bug" in the warning propagation. We define a *bug proliferation* matrix P_{ij} of dimension 6, $i \equiv a \rightarrow b, j \equiv c \rightarrow d$. The stability of the warning propagation is then governed by the largest (in absolute value) eigenvalue of this matrix λ_{\max} . The warning propagation is stable if and only if

$$\gamma \lambda_{\max} < 1, \quad (\text{C.16})$$

where $\gamma = \bar{k}^2 / \bar{k} - 1$ is the growth rate of the tree (γ^d is the typical number of vertices at distance d from the root). This analysis is often called bug proliferation [KPW04, MMZ06] (mostly in the context of the 1RSB stability). This investigation of the warning propagation stability was used e.g. in [ZM06] or [CKRT05].

An example where the warning propagation is stable, however, the belief propagation is not, can be found in [RSZ07] for the 1-in-K SAT problem. In 1-in-K SAT the warning propagation stability threshold corresponds to the unit clause propagation upper bound [RSZ07].

D 1RSB stability

Concerning the correctness of the 1RSB solution: the Boltzmann measure is split into clusters. This leads to an exact description of the system if and only if both the following conditions are satisfied.

- Condition of type I — the point-to-set correlation with respect to the measure over clusters decay to zero. The statistics over clusters may be described on the replica symmetric (tree) level. Clusters do not tend to aggregate.
- Condition of type II — the point-to-set correlations within the dominating clusters decay to zero. The interior of these clusters may be described on the replica symmetric (tree) level. Clusters do not tend to fragment into smaller ones.

Within the cavity approach these conditions can be checked from the 2RSB equation

$$P_2^{i \rightarrow j} [P^{i \rightarrow j}] = \frac{1}{\mathcal{Z}_2^{i \rightarrow j}} \int \prod_{k \in \partial i - j} dP_2^{i \rightarrow j} [P^{k \rightarrow i}] (\mathcal{Z}^{i \rightarrow j})^{m_2} \delta [P^{i \rightarrow j} - \mathcal{F}_2(\{P^{k \rightarrow i}\})] \quad (\text{D.1})$$

where the functional \mathcal{F}_2 is given by the 1RSB equation (2.24). We call the solution of (D.1) trivial if either $P_2^{i \rightarrow j} [P^{i \rightarrow j}] = \delta [P^{i \rightarrow j}]$ or each $P^{i \rightarrow j}(\psi^{i \rightarrow j}) = \delta(\psi^{i \rightarrow j} - \bar{\psi}^{i \rightarrow j})$, where the $P^{i \rightarrow j}$ is the solution of (2.24). If and only if the (population dynamics) solution of the 2RSB equation at $m = m^*$, $m_2 = 1$ and at $m = 1$, $m_2 = m^*$ is trivial then the two conditions are satisfied, and the 1RSB solution at m^* is correct.

Solving the 2RSB equation is, however, numerically involved. Even on random regular graphs the population dynamics of populations is needed, see app. E.5. Moreover the reweighting taking in account the term $(\mathcal{Z}^{i \rightarrow j})^{m_2}$ is costly. It is thus extremely useful to check the local stability of the 1RSB solution in the lines of the appendix C. The two types of local stability follow.

- Stability of type I — the inter-cluster spin glass susceptibility does not diverge.

$$\chi_{\text{SG}}^{\text{inter}} = \frac{1}{N} \sum_{i,j} (\overline{\langle s_i \rangle \langle s_j \rangle} - \overline{\langle s_i \rangle} \overline{\langle s_j \rangle})^2, \quad (\text{D.2})$$

where the overline denotes an average over clusters

$$\overline{x(\psi^{i \rightarrow j})} = \int x(\psi^{i \rightarrow j}) dP^{i \rightarrow j}(\psi^{i \rightarrow j}). \quad (\text{D.3})$$

- Stability of type II — the intra-cluster spin glass susceptibility does not diverge.

$$\chi_{\text{SG}}^{\text{intra}} = \frac{1}{N} \sum_{i,j} \langle s_i s_j \rangle_c^2. \quad (\text{D.4})$$

The instability of second type is sometimes called the Gardner instability due to [Gar85].

Again, there are several equivalent ways how to investigate the 1RSB stability. This time we first describe the zero temperature - frozen fields - version before turning to the general formalism.

D.1 Stability of the energetic 1RSB solution

In the energetic zero temperature limit the 1RSB distribution $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ can be split into the frozen and soft part as in (4.2). Moreover the self-consistency equations on the weights of the frozen fields, called the SP- y equations, do not depend on the details of the soft part. The methods for stability investigation of the SP- y equations were developed in [Par02b, MRT03, MPRT04, RBMM04].

Type I — SP- y convergence — The divergence of the inter-cluster spin glass susceptibility is in general equivalent to the non-convergence of the 1RSB equations (2.24) on a single graph. The reason is exactly the same as for the equivalence of the non-divergence of the spin glass susceptibility and the convergence of the belief propagation equations, which we explained in app. C.1. In the energetic zero temperature limit the convergence of the general 1RSB equations becomes convergence of the SP- y equations on a single graph. All the methods described in app. C.1 for the stability of the belief propagation equations can be used directly.

Remark in particular that the chain method (C.3), used e.g. in [RBMM04, KPW04], is not the simplest choice. The chains of length $d \rightarrow \infty$ have to be considered numerically, and the treatable values are only $d \approx 10-20$. This leads to an imprecision for a relatively large numerical effort. It is much more precise to use for example the noise propagation (C.5) as e.g. in [RSZ07].

Type II — Bug proliferation — The intra-state susceptibility is investigated in exactly the same manner as the replica symmetric stability. The only difference is that the average over clusters have to be taken properly. The energetic 1RSB solution is based on the warning propagation equations averaged properly over the clusters. Thus the 1RSB stability of the type II leads to the bug proliferation, as in app. C.2, averaged over the clusters.

Roughly explained, if we consider a model with Ising spins, we have the three components surveys $p = (p_-, p_0, p_+)$ on each edge. Where p_s is the probability over clusters that the warning on this edge takes the value s . Consider, as in app. C.2, a warning u and one of the incoming warnings u_0 , the remaining incoming warnings are indexed by $i = 1, \dots, k$ where k is distributed according to $\tilde{Q}(k)$ (C.14). Define $P_k(a \rightarrow b | c \rightarrow d)$ as the probability, over clusters, that the warning u changes from a value a to a value b provided that the warning u_0 was changed from a value c to a value d . Consider $P_k(a \rightarrow b | c \rightarrow d)$ as a matrix of dimension 6. And consider a chain of edges of length d . The proliferation of an instability "bug" is given by the product of matrices P_k along this chain. The product is averaged over the realizations of disorder (in degrees, etc.). We define the stability parameter as

$$\lambda_{II}(d) = \gamma \left(\text{Tr} \langle P_{k_1}^1 \dots P_{k_d}^d \rangle \right)^{\frac{1}{d}}. \quad (\text{D.5})$$

The SP- y is 1RSB stable if and only if $\lim_{d \rightarrow \infty} \lambda_{II}(d) < 1$. For more detailed presentation of the 1RSB bug proliferation method or concrete examples see e.g. [RBMM04, MMZ06, KPW04] and [RSZ07]. In all the implementations of this method the chain of $d \rightarrow \infty$ edges was used. Unlike in the type I stability, it is not know if this can be avoided in general.

Some results — The investigation of the 1RSB stability as we just described can be very simply incorporated to the population dynamics method used to solve the survey propagation

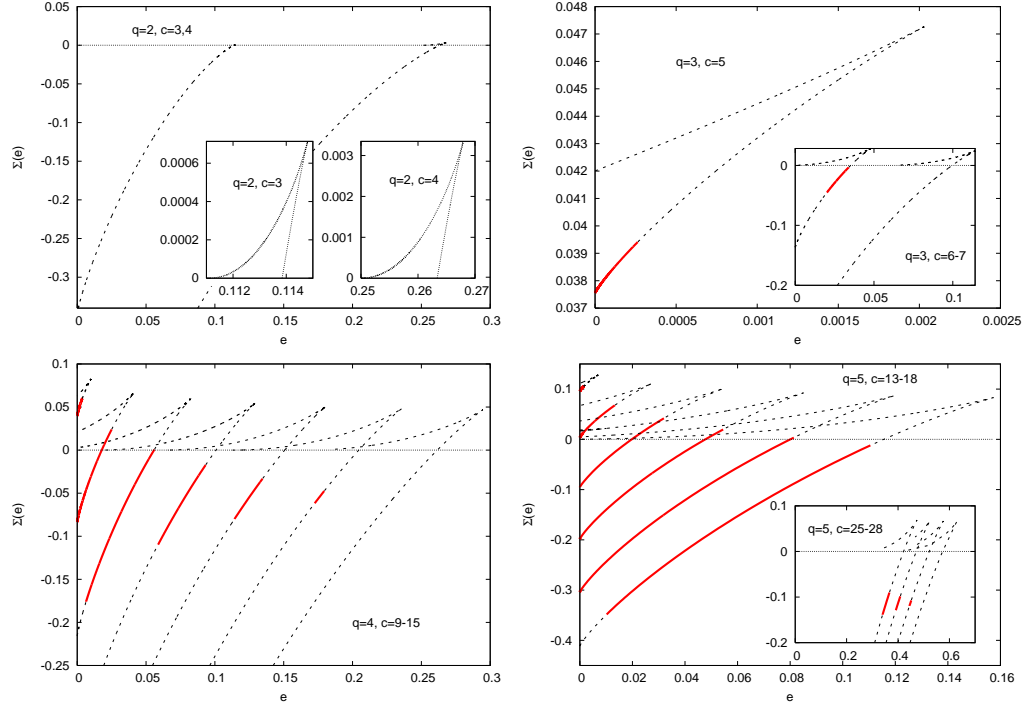


Fig. D.1. (Color online) The complexity as a function of energy for the coloring of random regular graphs. The 1RSB stable parts of the curves are in bold red.

equations. This means that on random regular graphs the stability equations become algebraic, as the values of surveys do not depend on the index of the edge. In fig. D.1 we present the result for coloring of random regular graphs.

On all the parts of fig. D.1 the complexity function is plotted against energy, $\Sigma(e)$ (2.32). This function is the main output of the 1RSB energetic method, the SP- y equations. The parameter y corresponds to the slope of the complexity function $y = \partial \Sigma(e) / \partial e$. Note that only the concave parts of the curves are physical.

The red parts of the $\Sigma(e)$ curves are the 1RSB stable parts. It seems to be a general fact that the instability of type I happens first for large values of y , and the instability of the type II for small values of y . The unphysical (convex) branch is always type II unstable. The instability of type I is sometimes completely absent.

An important observation is that the stability of the 1RSB energetic solution does not guarantee the stability of the full 1RSB solution. Differently said, the soft fields can destabilize the full solution. On the other hand also the opposite is true — the instability of the clusters corresponding to $m = 0$ does not imply the instability of the dominating clusters at m^* . We thus want to stress that the results of [MPRT04, RBMM04, MMZ06, KPW04] and others have to be taken with these two facts in mind.

D.2 1RSB stability at general m and T

The stability of the full 1RSB equations at a general value of the parameter m and of the temperature T is a more difficult task. We are not aware of any study where this would be practically considered for models on random graphs, apart from [KZ08b]. We review shortly the main findings and difficulties.

Type I — Divergence of the inter-cluster spin glass susceptibility (D.2) is equivalent to the non-convergence of the probability distributions $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ (2.24). But here arrives the biggest problem, how to judge if a probability distribution converges? The probability distribution $P^{i \rightarrow j}(\psi^{i \rightarrow j})$ is represented by a population of random elements picked from this distribution. How to decouple the randomness coming from this sampling and the one coming from the eventual non-convergence? Of course, provided that the numerical difficulty does not rise to the level of directly solving the 2RSB equations. This is not known in general and it is a technical but important open problem in the subject.

One interesting observation can be made, however: If the RS solution is instable then the 1RSB solution at $m = 1$ is type I instable. Indeed, if the mean value of the probability distribution does not converge then the 1RSB solution is type I instable. At the value $m = 1$ the mean (A.3) satisfies the simple belief propagation equations, as explained in app. A.

Type II — Divergence of the intra-cluster spin glass susceptibility (D.4) is much easier to investigate on a general level. It is equivalent to checking if the 1RSB iteration are stable against small changes in the probabilities ψ . Arguably the simplest way to do so is the *deviation of two replicas* method, described for the RS stability in app. C.1. We first find a fixed point of the 1RSB equations (2.24) using the population dynamics method. Then we create a second copy of the populations representing the distributions $P^{i \rightarrow j}(\psi^{i \rightarrow j})$. We perturb infinitesimally every of its elements $\psi^{i \rightarrow j}$. The 1RSB is type II stable if and only if the two copies converge to the same point. The noise propagation and other methods from C.1 can be used equivalently.

Some results and connection to the SP-y stability — Figure D.2 depicts the results for the stability of type II in the space of the parameters m and temperature T . The 1RSB solution is type II stable above the red curve m_{II} .

It is interesting to state the connection between the general m, T stability and the energetic zero temperature limit. The parameter $m = yT$ when $T \rightarrow 0$, thus when the stability of the frozen fields is relevant for the full stability the parameter $y_{\text{II}}T$ gives the slope of $m_{\text{II}}(T)$ near to zero T . This indeed seems to be the case, as shown in fig. D.2.

Based on the arguments above, it seems reasonable that the following assumptions are correct:

- (i) The stability of the energetic method gives the full stability for small m and T .
- (ii) If the RS solutions is stable then the 1RSB is stable type I at $m = 1$.
- (iii) If the 1RSB at a given temperature is type I (II resp.) stable at a given m , then it is type I (II resp.) stable for all smaller (larger resp.) m .

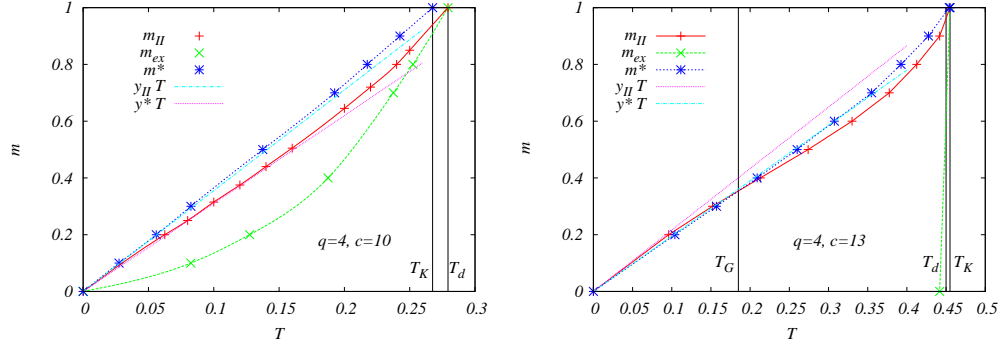


Fig. D.2. (Color online) Example of m - T diagrams for the 4-state anti-ferromagnetic Potts model on c -regular random graphs (left $c = 10$, right $c = 13$). A nontrivial solution of the 1RSB eq. (2.24) exists above the curve m_{ex} (green). The curves in blue m^* represent the thermodynamic value of the parameter m . The red curve m_{II} is the lower border of the type II stable region. The straight lines ($y_{\text{II}}T$ and y^*T) represent the slopes corresponding to the energetic 1RSB solution $yT = m$ in $T \rightarrow 0$. The energetic 1RSB solution is type II stable for $y > y_{\text{II}}$. The line $y_{\text{II}}T$ seems to give correctly the slope of m_{II} . This suggests that the stability of frozen variables is equivalent to the full stability for small m and T . Other examples of diagrams of this type are presented in [KZ08b].

Assuming as above, the stability of the 1RSB solution in the region where the RS solution is stable is given by the type II (Gardner) stability, which we know how to investigate. The result is depicted e.g. in fig. 5.6. This would mean that the stability of type II is always more important for the thermodynamical solution. And in particular that in the random coloring problem for $q \geq 4$ the 1RSB solution is stable in all the colorable phase.

The situation for 3-coloring is more subtle as 3-coloring is not RS stable for $c \geq c_d$. However, from assumption (i) follows that the interval of connectivities $(c_s, c_G) = (4.69, 5.08)$ is 1RSB stable at small temperatures. Thus we expect also all the colorable phase to be 1RSB stable (otherwise the phase diagram at fig. 5.6 would have to present a sort of re-entrant behaviour). This would also be in agreement with the situation in the fully connected ferromagnetic 3-state Potts model [GKS85]¹.

¹This is a contra-example to the common claim that in the systems with continuous dynamical transition ($T_d = T_{\text{local}}$) the 1RSB solution is not stable.

E Populations dynamics

Population dynamics is a numerical method to solve efficiently distributional equations of type (1.32) or (2.24) and compute observables of type (1.33) or (2.25a). In this context it was developed in [MP01]. As the form of the 1RSB equations was more or less known before, and they were solved approximatively using various forms of the variational ansatz, see e.g. [BMW00], it may be argued that the population dynamics technique was the crucial ingredient which made the spin glass models on random graphs solvable. Recently rigorous versions of this method were developed to analyze the performance of decoding algorithms [RU01], the name *density evolution* is often used in this context.

The main idea is to represent the probability distribution by a *population* (sample) of N elements drawn independently at random from this distribution. The algorithm starts from a random list and it mimics T iterations of the distributional equations and (hopefully) converges to a good representation of the desired fixed point. Several generalizations or subtleties are encountered and we describe some of them in the following. Consider the a random constraint satisfaction model specified by degree distribution $\mathcal{R}(k)$ of constraints, and $\mathcal{Q}(l)$ of variables, the excess degree distributions $r(k)$ and $q(l)$ are given by (1.8).

E.1 Population dynamics for belief propagation

The simplest version of the population dynamics is used to solve

- Belief propagation distributional equations (1.23a-1.23b) and compute the corresponding average free energy (1.20), entropy, etc. The complete replica symmetric solution is obtained this way.
- Survey propagation distributional equations, obtained from (1.41-1.42), and compute the average complexity function (1.43). The satisfiability transition is obtained this way.

The pseudocode for the procedures POPULATION-DYNAMICS and ONE-MEASUREMENT follows. To compute the observable Φ (free energy, entropy, complexity, etc.) we first call procedure POPULATION-DYNAMICS with $T = T_{\text{equil}}$ (equilibration time) and sufficiently large N . After we repeat ONE-MEASUREMENT plus POPULATION-DYNAMICS with $T = T_{\text{rand}}$ (randomization time) and M sufficiently large, but smaller than N . And finally we compute averages and error bars of these measurements.

In some problems the constraints are themselves random (negations in K -SAT, interactions in a spin glass etc.). The choice of this quenched randomness is then done at line 9 of POPULATION-DYNAMICS, and at line 8 of ONE-MEASUREMENT.

The population $\{\psi\}$ is randomly initialized to a random assignment at line 1 of POPULATION-DYNAMICS. That is all the zero components of the surveys (1.41-1.42) are zero, and the beliefs are completely biased, i.e., either $(1, 0)$ or $(0, 1)$. Such a choice is justified from the analogy with the reconstruction on trees where the proper initial condition is given by (2.11).

Satisfactory results are usually obtained with the population sizes and times of order $N \approx 10^4 - 10^5$, $T_{\text{equil}} \approx 10^3 - 10^4$, $T_{\text{rand}} \approx 10$, $M \approx N$. But these values may change problem from problem and a special care have to be taken about the numerics every time as basically no convergence theorems are known for a general case.

POPULATION-DYNAMICS($r(k), q(l), N, T$)

```

1  Initialize randomly  $N$ -component array  $\{\psi\}$ ;
2  for  $t = 1, \dots, T$ :
3      do for  $i = 1, \dots, N$ :
4          do Draw an integer  $k$  from the distribution  $r(k)$ ;
5              for  $d = 1, \dots, k$ :
6                  do Draw an integer  $l$  from the distribution  $q(l)$ ;
7                      Draw indexes  $j_1, \dots, j_l$  uniformly in  $\{1, \dots, N\}$ ;
8                      Compute  $\chi_d$  from  $\{\psi_{j_1}, \dots, \psi_{j_l}\}$  according to eq. (1.16b);
9                      Compute  $\psi_{\text{new}}$  from  $\{\chi_1, \dots, \chi_k\}$  according to eq. (1.16a);
10                      $\psi_i \leftarrow \psi_{\text{new}}$ ;
11  return array  $\{\psi\}$ ;

```

ONE-MEASUREMENT($\mathcal{R}(k), \mathcal{Q}(l), q(l), N, M$)

```

1  Initialize  $\Phi_{\text{constraint}} = 0; \Phi_{\text{variable}} = 0$ ;
2  for  $i = 1, \dots, M$ :  $\triangleright$  Compute the constraint part.
3      do Draw an integer  $k$  from the distribution  $\mathcal{R}(k)$ ;
4          for  $d = 1, \dots, k$ :
5              do Draw an integer  $l$  from the distribution  $q(l)$ ;
6                  Draw indexes  $j_1, \dots, j_l$  uniformly in  $\{1, \dots, N\}$ ;
7                  Compute  $\chi_d = \prod_{n=1}^l \psi_{j_n}$ ;
8                  Compute  $Z_{\text{new}}$  from  $\{\chi_1, \dots, \chi_k\}$  according to eq. (1.19a);
9                   $\Phi_{\text{constraint}} \leftarrow \Phi_{\text{constraint}} + \log Z_{\text{new}}$ ;
10 for  $i = 1, \dots, M$ :  $\triangleright$  Compute the variable part.
11     do Draw an integer  $l$  from the distribution  $\mathcal{Q}(l)$ ;
12         Draw indexes  $j_1, \dots, j_l$  uniformly in  $\{1, \dots, N\}$ ;
13         Compute  $Z_{\text{new}}$  from  $\{\psi_{j_1}, \dots, \psi_{j_l}\}$  according to eq. (1.19b);
14          $\Phi_{\text{variable}} \leftarrow \Phi_{\text{variable}} + (l - 1) \log Z_{\text{new}}$ ;
15 return  $(\alpha \Phi_{\text{constraint}} - \Phi_{\text{variable}})/M$ ;

```

E.2 Population dynamics to solve 1RSB at $m = 1$

The general 1RSB equations for general random graph ensemble require a population dynamics with population of populations. We will explain this in sec. E.5. Treating the population of populations requires a lot of CPU time and it is not very precise, thus anytime we have the opportunity to avoid this we have to take it. One such opportunity is the simplification of the 1RSB equations at $m = 1$ explained in appendix A. Conveniently, both the clustering and the condensation transitions are obtained this way.

The population dynamics method have to be adapted to solve eq. (A.10) and to measure the entropy of states (A.11). We give the $m = 1$ generalization of the procedure POPULATION-DYNAMICS, the changes in ONE-MEASUREMENT are then straightforward. Note that lines 11 and 13 take in general 2^k steps as we need to compute probability of every combination of the set $\{s_1, \dots, s_k\}$.

```

PD-( $m = 1$ )-GENERALIZATION( $r(k), q(l), N, T$ )
1   $\{\psi^{\text{RS}}\} \leftarrow \text{POPULATION-DYNAMICS}(r(k), q(l), N, T)$ ;
2  Initialize  $N$ -component arrays  $\{\psi^1 \leftarrow 1\}$  and  $\{\psi^0 \leftarrow 0\}$ ;
3  for  $t = 1, \dots, T$ :
4  for  $i = 1, \dots, N$ :
5      do Draw an integer  $k$  from the distribution  $r(k)$ ;
6      for  $d = 1, \dots, k$ :
7          do Draw an integer  $l_d$  from the distribution  $q(l)$ ;
8              Draw indexes  $j(d, 1), \dots, j(d, l_d)$  uniformly in  $\{1, \dots, N\}$ ;
9              Compute  $\chi_d^{\text{RS}}$  from  $\{\psi_{j(d,1)}^{\text{RS}}, \dots, \psi_{j(d,l_d)}^{\text{RS}}\}$  according to eq. (1.16b);
10          $s \leftarrow 1$ ;
11         Choose  $\{s_1, \dots, s_k\}$  with prob. given by the 2nd line of eq. (A.10);
12          $s \leftarrow 0$ ;
13         Choose  $\{r_1, \dots, r_k\}$  with prob. given by the 2nd line of eq. (A.10);
14         for  $d = 1, \dots, k$ :
15             do Compute  $\chi_d^1$  from  $\{\psi_{j(d,1)}^{s_d}, \dots, \psi_{j(d,l_d)}^{s_d}\}$  according to eq. (1.16b);
16             Compute  $\chi_d^0$  from  $\{\psi_{j(d,1)}^{r_d}, \dots, \psi_{j(d,l_d)}^{r_d}\}$  according to eq. (1.16b);
17             Compute  $\psi_{\text{new}}^{\text{RS}}$  from  $\{\chi_1^{\text{RS}}, \dots, \chi_k^{\text{RS}}\}$  according to eq. (1.16a);
18             Compute  $\psi_{\text{new}}^1$  from  $\{\chi_1^1, \dots, \chi_k^1\}$  according to eq. (1.16a);
19             Compute  $\psi_{\text{new}}^0$  from  $\{\chi_1^0, \dots, \chi_k^0\}$  according to eq. (1.16a);
20              $\psi_i^{\text{RS}} \leftarrow \psi_{\text{new}}^{\text{RS}}$ ;
21              $\psi_i^1 \leftarrow \psi_{\text{new}}^1$ ;
22              $\psi_i^0 \leftarrow \psi_{\text{new}}^0$ ;
23 return arrays  $\{\psi^{\text{RS}}\}, \{\psi^1\}, \{\psi^0\}$ ;

```

E.3 Population dynamics with reweighting

A simplification of the 1RSB equations (2.24) arises for the ensemble of random regular graphs, there the distribution $\mathcal{P}^{i \rightarrow j}(\psi^{i \rightarrow j})$ over clusters is the same for every edge (ij) . In the corresponding population dynamics a special care have to be taken about the reweighting term $(Z^{i \rightarrow j})^m$.

We describe two different strategies to deal with the reweighting. In the first one REWEIGHTING-FASTER the elements of the population have all the same weight and thus in each sweep the population needs to be re-sampled and some less probable elements might be lost. In the second strategy REGULAR-REWEIGHTING-PRECISE each element has its own weight, no re-sampling is needed, but the search of a random element, at the line 10, takes $\log N$ steps. Thus the first strategy is faster, the second one is slightly more precise. Which one is eventually better seems to be problem specific.

Consider a population $\{\psi\}$ where each element ψ_i has weight w_i . The weights are computed from the BP update (1.16a-1.16a) as $w_i = \left(Z^{a \rightarrow i} \prod_{j \in \partial a - i} Z^{j \rightarrow a} \right)^m$.

REWEIGHTING-FASTER($N, \{\psi\}, \{w\}$)

```

1   $w_{\text{tot}} \leftarrow 0$ ;
2  for  $i = 1, \dots, N$ :
3      do  $w_{\text{tot}} \leftarrow w_{\text{tot}} + w_i$ ;
4   $\triangleright z_i$  is the cumulative distribution of indexes  $i$ ;
5   $z_0 = 0$ ;
6  for  $i = 1, \dots, N$ :
7      do  $z_i \leftarrow z_{i-1} + w_i/w_{\text{tot}}$ 
8   $\triangleright$  Trick to make a list of ordered random numbers  $n_i$  in  $O(N)$  steps.
9   $G \leftarrow 0$ ;
10 for  $i = 1, \dots, N$ :
11     do  $n_i \leftarrow -\log \text{RAND}$ ;
12      $\triangleright$  RAND outputs a random number in the interval  $(0, 1)$ .
13      $G \leftarrow G + n_i$ ;
14  $G \leftarrow G - \log \text{RAND}$ ;
15  $n_1 \leftarrow n_1/G$ ;
16 for  $i = 2, \dots, N$ :
17     do  $n_i \leftarrow n_i/G$ ;
18      $n_i \leftarrow n_i + n_{i-1}$ ;
19  $\triangleright$  Finally making the new population.
20  $p \leftarrow 0$ ;
21 for  $i = 1, \dots, N$ 
22     do while ( $n_i > z_p$ )  $p \leftarrow p + 1$ ;
23      $\psi_i^{\text{new}} \leftarrow \psi_p$ ;
24 return array  $\{\psi^{\text{new}}\}$ ;
```

REGULAR-REWEIGHTING-PRECISE($r(k), q(l), N, T, m$)

```

1  Initialize randomly  $N$ -component arrays  $\{\psi\}$  and  $\{w\}$ ;
2  for  $t = 1, \dots, T$ :
3      for  $i = 1, \dots, N$ :
4          do Draw an integer  $k$  from the distribution  $r(k)$ ;
5               $Z_{\text{new}} \leftarrow 1$ ;
6              for  $d = 1, \dots, k$ :
7                  do Draw an integer  $l$  from the distribution  $q(l)$ ;
8                      for  $n = 1, \dots, l$ :
9                          do Create cumulative probability distribution from weights  $\{w\}$ ;
10                             Draw index  $j_n$  from this cumulative distribution;
11                             Compute  $\chi_d$  from  $\{\psi_{j_1}, \dots, \psi_{j_l}\}$  according to eq. (1.16b);
12                              $Z_{\text{new}} \leftarrow Z_{\text{new}} \cdot Z_d$ , where  $Z_d$  is the norm. from eq. (1.16b);
13                             Compute  $\psi_{\text{new}}$  from  $\{\chi_1, \dots, \chi_k\}$  according to eq. (1.16a);
14                              $Z_{\text{new}} \leftarrow Z_{\text{new}} \cdot Z_d$ , where  $Z_d$  is the norm. from eq. (1.16a);
15                              $\psi_i \leftarrow \psi_{\text{new}}$ ;
16                              $w_i \leftarrow (Z_{\text{new}})^m$ ;
17 return array  $\{\psi\}$ , weights  $\{w\}$ ;
```


E.4 Population dynamics with hard and soft fields

Fraction of frozen variables (again on random regular graphs for simplicity) can be obtained by solving equation (4.10). To compute the value $r(m)$ a population needs to be kept for the soft part of the distribution P_{soft} , eq. (4.2). It is important to stress that when evaluating the **if** conditions on lines 15,19 and 23 we consider as frozen only the incoming fields created at line 13.

PD-HARD-SOFT($r(k), q(l), N, T, m$)

```

1  Initialize randomly  $N$ -component array  $\{\psi \leftarrow \text{RAND}\}$ ;
2   $\eta \leftarrow 1$ ;
3  for  $t = 1, \dots, T$ :
4      do  $i \leftarrow 1$ ;
5           $h \leftarrow 0$ ;  $Z_{\text{hard}} \leftarrow 0$ ;  $Z_{\text{soft}} \leftarrow 0$ ;
6          while  $i \leq N$ :
7              do Draw an integer  $k$  from the distribution  $r(k)$ ;
8                   $Z_{\text{new}} \leftarrow 1$ ;
9                  for  $d = 1, \dots, k$ :
10                     do Draw an integer  $l$  from the distribution  $q(l)$ ;
11                         for  $r = 1, \dots, l$ :
12                             do if  $\text{RAND} < \eta$ 
13                                 then Set  $\psi_r$  to be a frozen field;
14                                 else Draw  $\psi_r$  uniformly from  $\{\psi\}$ ;
15                             if No contradiction between the frozen fields in  $\{\psi_1, \dots, \psi_l\}$ 
16                                 then Compute  $\chi_d$  from  $\{\psi_1, \dots, \psi_l\}$  using eq. (1.16b);
17                                      $Z_{\text{new}} \leftarrow Z_{\text{new}} \cdot Z_d$ ,  $Z_d$  is the norm. from (1.16b);
18                                 else goto line 7;
19                             if No contradiction between the frozen fields in  $\{\chi_1, \dots, \chi_k\}$ 
20                                 then Compute  $\psi_{\text{new}}$  from  $\{\chi_1, \dots, \chi_k\}$  according to eq. (1.16a);
21                                      $Z_{\text{new}} \leftarrow Z_{\text{new}} \cdot Z_d$ , where  $Z_d$  is the norm. from eq. (1.16a);
22                                 else goto line 7;
23                             if  $\psi_{\text{new}}$  is a frozen field
24                                 then  $Z_{\text{hard}} \leftarrow Z_{\text{hard}} + (Z_{\text{new}})^m$ ;
25                                      $h \leftarrow h + 1$ ;
26                                 else  $Z_{\text{soft}} \leftarrow Z_{\text{soft}} + (Z_{\text{new}})^m$ ;
27                                      $\psi_i \leftarrow \psi_{\text{new}}$ ;
28                                      $w_i \leftarrow (Z_{\text{new}})^m$ ;
29                                      $i \leftarrow i + 1$ ;
30           $r \leftarrow (Z_{\text{soft}} h) / (Z_{\text{hard}} N)$ ;
31          Update  $\eta$  according to eq. (4.10);
32           $\{\psi\} \leftarrow \text{REWEIGHTING-FASTER}(N, \{\psi\}, \{w\})$ ;
33  return array  $\{\psi\}, \eta$ ;
```

E.5 The population of populations

The general 1RSB equations take form (2.33), the order parameter $\mathcal{P}[P(\psi)]$ is a distribution (over the graph ensemble) of distributions (over the clusters). It can be represented by a population $\{\{\psi\}\}$ of N -component populations $\{\psi\}_i$, where $i = 1, \dots, M$. We sketch here the corresponding population dynamics of populations. Again this has been first described in [MP01].

POPULATION-OF-POPULATIONS($r(k), q(l), N, M, T, m$)

```

1  Initialize randomly  $M \times N$ -component array  $\{\{\psi\}\}$ ;
2  for  $t = 1, \dots, T$ :
3      do for  $i = 1, \dots, M$ :
4          do Draw an integer  $k$  from the distribution  $r(k)$ ;
5              for  $d = 1, \dots, k$ :
6                  do Draw an integer  $l_d$  from the distribution  $q(l)$ ;
7                      Draw indexes  $i(d, 1), \dots, i(d, l_d)$  uniformly in  $\{1, \dots, M\}$ ;
8                       $\{\psi\}_{\text{new}} \leftarrow \text{ONE-STEP}(\{\{\psi\}\}, \{i(1, 1), \dots, i(k, l_k)\}, \{l\}, k, N, m)$ ;
9                       $\{\psi\}_i \leftarrow \{\psi\}_{\text{new}}$ ;
10 return array  $\{\{\psi\}\}$ ;
```

ONE-STEP($\{\{\psi\}\}, \{i(1, 1), \dots, i(k, l_k)\}, \{l\}, k, N, m$)

```

1  for  $j = 1, \dots, N$ :
2      do  $Z_{\text{new}} \leftarrow 1$ ;
3          for  $d = 1, \dots, k$ :
4              do Draw indexes  $j(d, 1), \dots, j(d, l_d)$  uniformly in  $\{1, \dots, N\}$ ;
5                  Compute  $\chi_d$  from  $\{\psi_{i(d, 1), j(d, 1)}, \dots, \psi_{i(d, l_d), j(d, l_d)}\}$  using (1.16b);
6                   $Z_{\text{new}} \leftarrow Z_{\text{new}} \cdot Z_d$ ,  $Z_d$  is the norm. from (1.16b);
7                  Compute  $\psi_{\text{new}}$  from  $\{\chi_1, \dots, \chi_k\}$  according to eq. (1.16a);
8                   $Z_{\text{new}} \leftarrow Z_{\text{new}} \cdot Z_d$ ,  $Z_d$  is the norm from (1.16a);
9                   $w_j \leftarrow (Z_{\text{new}})^m$ ;
10                  $\psi_j \leftarrow \psi_{\text{new}}$ ;
11  $\{\psi\} \leftarrow \text{REWEIGHTING-FASTER}(N, \{\psi\}, \{w\})$ ;
12 return array  $\{\psi\}$ ;
```

Depending on the problem we are about to solve the population of populations might also be combined with the reweighting of populations or the separation of the frozen and soft fields, see e.g. appendix D of [ZK07].

E.6 How many populations needed?

We make a summary of which level of the population dynamics technique is needed depending on the problem. References are just examples and are biased towards works presented in this thesis.

- Analytical solution
 - Belief propagation on regular graphs [ZM06, ZK07, ZM08].

- General warning propagation with integer warnings [ZM06,RSZ07].
- Frozen variables at $m = 1$ [ZK07,ZM08].
- Survey propagation on regular graphs (frozen variables at $m = 0$, energetic cavity) [KPW04] or [ZK07,ZM08].
- Single population
 - General belief propagation in models with discrete variables [ZM06,ZK07,ZM08].
 - General survey propagation (1RSB at $m = 0$, energetic cavity) on model with integer warnings [RSZ07,ZM08], or very precise numerics in [MMZ06].
 - 1RSB at $m = 1$ [MM06a,MRTS08] or [KMRT⁺07,ZK07].
 - 1RSB on random regular graphs [KMRT⁺07,ZK07].
 - 2RSB at $m = 0$ (energetic cavity) on regular graphs [Riv05].
- Population of populations
 - General 1RSB (also finite temperature) [MP01,MPR05,MRTS08] or [KMRT⁺07,ZK07,KZ08b].
 - 2RSB of random regular graphs [KZ08b].
 - 2RSB at $m = 0$ (energetic cavity).
 - 3RSB at $m = 0$ (energetic cavity) on regular graphs.

We are not aware on any work where the last two points would be implemented. More levels of replica symmetry breaking would require more levels of populations. We are not aware of any work where more than population of populations would be treated. Rather than pushing the numerics in this direction new theoretical works are needed for models where the 1RSB solution is not correct.

F Algorithms

Here we do not aim to provide a complete summary of algorithms used to solve the random constraint satisfaction problems. We just define and briefly discuss algorithms which were used, generalized or tested in the context of this thesis. Strictly speaking we are almost always dealing with incomplete solvers, that is algorithms which might find a solution but never provide a certificate of unsatisfiability. It is an open and interesting questions if the methods presented in this thesis can imply something for certification of unsatisfiability.

F.1 Decimation based solvers

A large class of algorithms for CSPs is based on the following iterative scheme:

DECIMATION

- 1 **repeat** Choose a variable i ;
- 2 Choose a value s_i ;
- 3 Assign i the value s_i and simplify the formula;
- 4 **until** Solution or contradiction is found;

The nontrivial part is how to choose a variable in step 1 and how to choose its value in step 2. In the following we describe several more or less sophisticated or efficient strategies.

Note that all these strategies can be improved by *backtracking*, that is if a contradiction was found we return to the last variable where another value than the one we chose was possible and make this choice instead.

F.1.1 Unit Clause propagation

One of the simplest (and obvious) strategies is to choose and assign a variable which is present in a constraint which is compatible with only one value of that variable. In K-SAT this is equivalent to assigning variables belonging to clauses which contain only this variable, hence the name *unit clause*. If no such variable exists one possibility (the random heuristics) is to choose an arbitrary variable and assign it a random value from the available ones. The unit clause propagation combined with the random heuristics (without backtracking) is not very efficient solver of K-SAT. But the situation is more fortunate for some other constraint satisfaction problems. The most interesting example being perhaps the 1-in-K SAT [ACIM01] and [RSZ07]. The random 1-in-K SAT exhibits a sharp satisfiability phase transition for $K \geq 3$. Moreover, if the probability of negation of variables lies in the interval $(0.2726, 0.7274)$ (for $K = 3$) then:

- In the satisfiable phase the unit clause propagation combined with the random heuristics finds a solution with finite probability in every run.
- In the unsatisfiable phase every run of the unit clause propagation leads to a contradiction with finite probability after the assignment of the very first variable.

Hence, with random restarts the random 1-in-3 SAT is almost surely solvable in polynomial time in the whole phase space (given the probability of a negation is as specified above). At the same

time the 1-in-3 SAT is an NP-complete problem, it thus provides a rare example of an on average easy NP-complete problem with a satisfiability phase transition.

Unit clause propagation is the main element of all the exact solvers of constraint satisfaction problems. The most studied example being the Davis-Putnam-Logemann-Loveland (DPLL) algorithm [DP60, DLL62] for K-SAT which combines the unit clause propagation with the pure literal elimination (pure literal appears either only negated or non-negated) with backtracking. It was mostly this algorithm which was used when the connection between the algorithmical hardness and phase transitions was being discovered [MSL92, CKT91]. Moreover, all the modern complete solvers of the satisfiability problem follow a similar, more elaborated, path.

F.1.2 Belief propagation based decimation

Belief propagation [Pea82] computes, or on general graphs approximates, marginal probabilities. These can then be used to find an actual solution. In some problems the marginals give the solution directly, e.g. in the error correcting codes [Gal68], in the matching [BSS05, BSS06], or the random field Ising model at zero temperature [KW05, Che08] etc. In constraint satisfaction problems, typically, marginals do not give a direct information about a solution. For example in coloring of random graphs, the BP equations always converge to all marginals being equal to $1/q$. Belief propagation based decimation strategies have been studied recently.

In every cycle of the algorithm DECIMATION, the belief propagation equations are updated until they converge or a maximal number of updates per variable T_{\max} is reached. At least two strategies how to choose the decimated variable and its value were tested and studied, see e.g. [KMRT⁺07] and [MRTS07]:

- Uniform BP decimation – Choose a variable at random and assign its value according to the marginal probability estimated by BP.
- Maximal BP decimation – Find the variable with the most biased BP marginal and assign it the most probable value.

The other two combinations where a random variables is assigned its most probable value or when the most biased variable is assigned random value according to its marginal probability can be think of. The BP decimation, as described above, runs in quadratic time. In eventual practical implementations a small fraction of variables should be decimated at each step, thus reducing the computational complexity to linear (or log-linear if the maximum convergence time increases as $\log N$).

The empirically best strategy is the maximal BP decimation. This can be understood from the fact that this strategy aims to destroy the smallest possible number of solutions in every step, as argued on a more quantitative level in [Par03]. We gave as an example the performance of the maximal BP decimation in the 3- and 4-coloring of random Erdős-Rényi graphs [ZK07] in fig. 3.3.

The uniform BP decimation is less successful, because it aims not only to find a solution but also to sample solutions uniformly at random. Indeed, if an exact calculation of marginal probabilities would be used instead of the BP estimates the uniform *exact* decimation would lead to a perfect sampling. The uniform exact decimation is a process which can be analyzed using the cavity method. The result then sheds light on the limitations of the BP decimation. This analysis

was developed in [MRTS07], and we give an example for the factorized occupation problems in the following.

F.1.3 Maximal BP decimation on the random coloring

We implemented the maximal BP decimation algorithm on the random graph coloring. We chose $T_{\max} = 10$, if a solution is not found we restart with $T_{\max} = 20$ and eventually once again with $T_{\max} = 40$. The fraction of successful runs is plotted in fig. 3.3 and we see that this algorithm works even in condensed phase where the BP marginals are not asymptotically correct, or in a phase where the equations do not even converge. The non-convergence of the belief propagation equations is ignored (in 3-coloring from the beginning, in 4-coloring after a small fraction, typically around 10%, of variables was fixed). It thus seems that in coloring the BP decimation is a very robust algorithm.

What is the reason for the failure of the maximal BP decimation at higher connectivities? A straightforward suggestion would be that it should not work in the condensed phase where the BP marginals are not asymptotically correct. But we do not observe anything particular in the performance curves at the condensation transition. A second natural suggestion would be that BP should converge in order that the algorithm works, this also does not seem to be the case, as BP does not converge in the 3-coloring for connectivity $c > 4$ and yet the algorithm is perfectly able to find solutions. Moreover, even in 4-coloring where the BP equations converge on large formulas in all the satisfiable phase, after a certain (rather small) fraction of variables is decimated the convergence is lost. As we argued in appendix C the non-convergence of BP is equivalent to the local instability of the replica symmetric solution. It thus seems that the reduced problem, after a certain fraction of variable was fixed, is even harder from the statistical physics perspective than the original problem. Yet, this does not seem to be fatal for the finding of solutions. Finally, in the region where the BP decimation algorithm really does not succeed we observed that a precursor of the failure exists. The normalizations in the BP equations (1.16a-1.16b) gradually decrease to zero, meaning that the incoming beliefs become almost contradictory.

F.1.4 Analysis of the uniform exact decimation

The uniform exact decimation after θN steps is equivalent to taking a solution uniformly at random and fixing its first θN variables. Such a procedure can be analyzed [MRTS07] and conclusions made about the influence of small errors in the BP estimates of marginals.

Given an instance of the CSP, consider a solution $\{s\}$ taken uniformly at random and reveal the value of each variable with probability θ . Denote Φ the fraction of variables which were either revealed or are directly implied by the revealed ones. To compute $\Phi(\theta)$ we derive the cavity equations on a tree. Denote $\Phi_s^{i \rightarrow b}$ the probability that a variable i is fixed conditioned on the value s of the variable i and on the absence of the edge (ib) :

$$\Phi_s^{i \rightarrow b} = \theta + (1 - \theta) \left[1 - \prod_{a \in \partial i - b} (1 - q_s^{a \rightarrow i}) \right]. \quad (\text{F.1})$$

Meaning that the variable i was either revealed or not, and if not it is implied if at least one of the incoming constraints implies it. The $q_s^{a \rightarrow i}$ is a probability that constraint a implies variable i to

be s conditioned on: 1) variable i taking the value $s \in \{s\}$ in the solution we chose, 2) variable i was not revealed directly and 3) the edge (ai) is absent.

We write the expression for $q_s^{a \rightarrow i}$ only for random occupation CSPs on random regular graphs where the replica symmetric equation is factorized. Then also $q_s^{a \rightarrow i}$ and $\Phi_s^{i \rightarrow b}$ are factorized, that is independent of a, b, i . The conditioned probability q_s is the ratio of the probability that variable i takes the value s and is implied by the constraint a and probability that variable i takes the value s :

$$q_1 = \frac{1}{\psi_1 Z^{\text{reg}}} \sum_{\substack{A_r=0 \\ A_{r+1}=1}} \binom{k}{r} (\psi_1)^{lr} (\psi_0)^{l(k-r)} \sum_{s=0}^{s_1} \binom{r}{s} \Phi_0^{k-r} \Phi_1^{r-s} (1 - \Phi_1)^s, \quad (\text{F.2a})$$

$$q_0 = \frac{1}{\psi_0 Z^{\text{reg}}} \sum_{\substack{A_r=1 \\ A_{r+1}=0}} \binom{k}{r} (\psi_1)^{lr} (\psi_0)^{l(k-r)} \sum_{s=0}^{s_0} \binom{k-r}{s} \Phi_1^r \Phi_0^{k-r-s} (1 - \Phi_0)^s, \quad (\text{F.2b})$$

where $l = L - 1$, $k = K - 1$. The indexes s_1, s_0 in the second sum of both equations are the largest possible but such that $s_1 \leq r$, $s_0 \leq K - 1 - r$, and $\sum_{s=0}^{s_1} A_{r-s} = 0$, $\sum_{s=0}^{s_0} A_{r+1+s} = 0$. The terms $\Phi_1^r \Phi_0^{k-r-s} (1 - \Phi_0)^s$ and $\Phi_1^{r-s} \Phi_0^{K-r-1} (1 - \Phi_1)^s$ are the probabilities that a sufficient number of incoming variables was revealed such that the out-coming variable is implied (not conditioned on its value). The first sum goes over all the possible numbers of 1's being assigned on the incoming variables, r . The term $\psi_1^{lr} \psi_0^{l(k-r)}$ is then the probability that such a configuration took place. The cavity probabilities that the corresponding variable takes value 0/1, ψ_0, ψ_1 are taken from the BP equations (4.16a-4.16b), Z^{reg} is the normalization in (4.16a-4.16b). The first condition on r takes care about the values of the incoming neighbours being compatible with the value of the variable i on which is conditioned, the second condition on r is satisfied if and only if the value of the variable i is implied by the incoming configuration.

Once a solution for q_s is found (from initial conditions $\Phi = \theta$) the total probability that a variable is fixed is computed as

$$\Phi(\theta) = \theta + (1 - \theta) \{ \mu_1 [1 - (1 - q_1)^L] + \mu_0 [1 - (1 - q_0)^L] \}, \quad (\text{F.3})$$

where μ_0, μ_1 are the total BP marginals, $\mu_s = \psi_s^L / (\psi_0^L + \psi_1^L)$.

Notice the complete analogy between eqs. (F.2b-F.2a) and the equations for hard fields at $m = 1$ (4.20b-4.20a). To compute the function $\Phi(\theta)$ for a general CSP on a general graph ensemble a derivation in the lines of app. A have to be adapted, see also [MRTS07]. Finally note that as the probabilities ψ_1, ψ_0 are taken from the belief propagation equations the form (F.2b-F.2a) is not correct in the condensed phase (but in the locked problems the satisfiable phase is never condensed).

F.1.5 The Failure of Decimation in the Locked problems

In the locked problems, see sec. 4.3, the BP decimation algorithm does not succeed to find a satisfying assignment even at the lowest possible connectivity. To give an example in the 1-or-3-in-5 SAT on truncated Poissonian graphs the maximal BP decimation succeeds to find a solution in only about 25% at the lowest average connectivity $\bar{l} = 2$, and this fraction drops down to less than 5% at already $\bar{l} = 2.3$ (to be compared with the clustering threshold $l_d = 3.07$, or the satisfiability threshold $l_s = 4.72$).

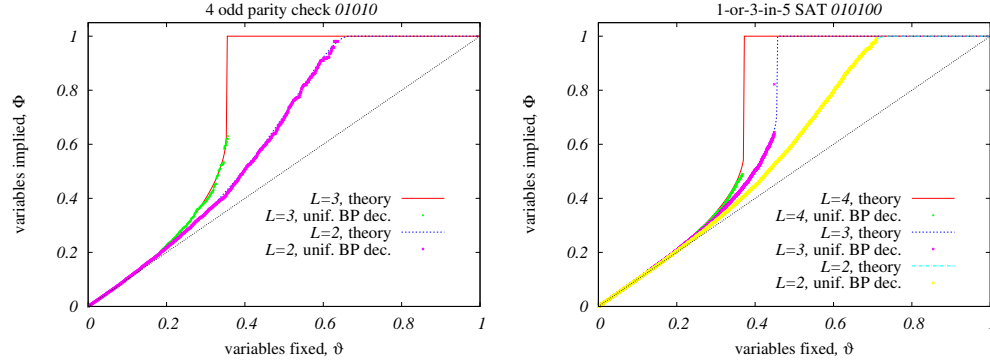


Fig. F.1. (Color online) Analytical analysis of the BP inspired uniform decimation. Number of variables directly implied $\Phi(\theta)$ plotted against number of variables fixed θ .

Interestingly, the precursors of the failure of the BP decimation algorithm observed in the graph coloring are not present in the locked problems. In particular the BP equations converge during all the process and the normalizations in the BP equations (1.16a-1.16b) stays finite. However, the above analysis of the function $\Phi(\theta)$ sheds light on the origin of the failure.

In fig. F.1 we compare the function $\Phi(\theta)$ (F.3) with the experimental performance of the uniform BP decimation. Before the failure of the algorithm (when a contradiction is encountered) the two curves collapse perfectly. The reason why the algorithm fails to find solutions is now transparent.

- **Avalanche of direct implications** – In some cases the function $\Phi(\theta)$ has a discontinuity at a certain spinodal point θ_s ($\theta_s \approx 0.46$ at $L = 3$ of the 1-or-3-in-5 SAT). Before θ_s after fixing one variable there is a finite number of direct implications. As the loops are of order $\log N$ these implications never lead to a contradiction. At the spinodal point θ_s after fixing one more variable and extensive avalanche of direct implications follows. Small (order $1/N$) errors in the previously used BP marginals may thus lead to a contradiction. This indeed happens in almost all the runs we have done. For more detailed discussion see [MRTS07].
- **No more free variables** – The second reason for the failure is specific to the locked problems, more precisely to the problems where $\Phi = 1$ is a solutions of (F.2a-F.2b). In these cases function $\Phi(\theta) \rightarrow 1$ at some $\theta_1 < 1$ ($\theta_1 \approx 0.73$ at $L = 4$ of 1-or-3-in-5 SAT). In other words if we reveal a fraction $\theta > \theta_1$ of variables from a random solution, the reduced problem will be compatible with only that given solution. Again a little error in the previously fixed variables and the BP uniform decimation ends up in a contradiction. If on the contrary the function $\Phi(\theta)$ reaches value 1 only for $\theta = 1$ then the residual entropy is positive and there should everytime be some space to correct previous small errors, demonstrated on a non-locked problem in fig. F.2.

These two reasons of failure of the BP uniform decimation seems quite different. But they have one property in common. As the point of failure is approached we observe a divergence

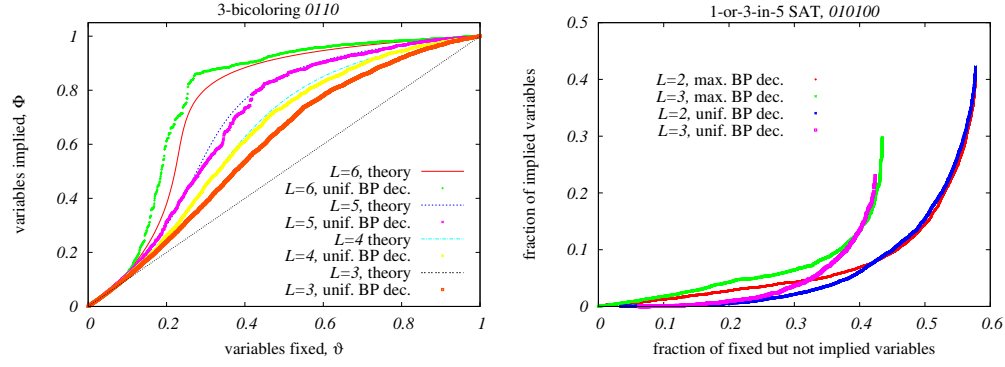


Fig. F.2. (Color online) Left: For comparison, the BP uniform decimation works well on the non-locked problems, the example is for bicoloring. Right: Comparison of the maximal and uniform decimation. Number of directly implied variables is plotted against number of variables which were free just before being fixed. Behaviour of the two decimation strategies is similar.

of the ratio between the number of variables which were not directly implied before being fixed and the number of directly implied variables, see fig. F.2. This ratio can also be computed for the maximal BP decimation and no quantitative difference is observed for the locked problems, thus the two reasons above explain also the failure of the, otherwise more efficient, maximal BP decimation.

F.1.6 Survey propagation based decimation

The seminal works [MPZ02, MZ02] not only derived the survey propagation equations, but also suggested it as a base for a decimation algorithm for random 3-SAT. The performance is spectacular, near to the satisfiability threshold on large random 3-SAT formulas it works faster than any other known algorithm. SP based decimation seem to be able to find solutions in $O(N \log N)$ time up to the connectivity $\alpha = 4.252$ in 3-SAT [Par03] (to be compared with the satisfiability threshold $\alpha_s = 4.267$).

Survey propagation equations (1.41-1.42) aim to compute the probability (over clusters) that a certain variables is frozen to take a certain value. This information can then be used to design a strategy for the DECIMATION algorithm. In particular, as long as the result of survey propagation is nontrivial (not all $p_0^{i \rightarrow a} = 1$) the variable with the largest bias $|p_+^i - p_-^i|$ is chosen and is assigned the more probable value. After a certain fraction of variables is decimated the fixed point of the survey propagation on the reduced formula is trivial. The suggestion of [MPZ02, MZ02] is that such a reduced formula is easily satisfiable and some of the well known heuristic algorithms may be used to solve it (Walk-SAT, see the next section F.2.2, was used in the original implementation). Note also that the original implementation of [MPZ02, MZ02] decimated a fraction of variables at each DECIMATION step, thus reducing significantly the computational time.

Originally, the success of the survey propagation based algorithm was contributed to the fact that survey propagation equations take into account the clustering of solutions. This was,

however, put in doubt since. To give an example, in the locked problems, see sec. 4.3, the survey propagation equations give an identical fixed point as the belief propagation and as we argued in the previous section F.1.2 the maximal BP decimation fails to find solutions in the locked problem in the whole range of connectivities.

The true reason for the high performance of survey propagation in 3-SAT thus stays an open problem. For example, and unlike with BP, there are usually no problems with SP convergence during the decimation. Two very interesting observations were made in [KSS07a] for SP the decimation algorithm on K -SAT. First, the SP decimation indeed makes the formula gradually simpler for local search algorithms, see sec. F.2.2, again in contrast with BP decimation. Second, the SP decimation on K -SAT does not create any (or a very small number) of direct implications (unit clauses) during the process. Given that creation of direct implication makes the decimation fail in the locked problems, as we just showed, this might be a promising direction for a new understanding.

F.2 Search of improvement based solvers

Here we describe another large class of CSPs solvers, the *search of improvement algorithms*. All these algorithms start with a random assignment of variables. Then different rules are adopted to gradually improve this assignment and eventually to find a solution. The most typical example of that strategy is the simulated annealing [KGV83] or stochastic local search algorithms like Walk-SAT [SLM92, SKC94].

F.2.1 Simulated annealing

In physics simulated annealing is a popular and very universal solver of optimization problems. It is based on running the Metropolis [MRR⁺53] (or other Monte Carlo) algorithm and gradually decreasing the temperature-like parameter. Simulated annealing algorithm respects the detailed balance condition, after large time it thus converges to the equilibrium state, and it is thus guaranteed to find the optimal state in a finite time for a finite system size. In general, the time can of course depend exponentially on the system size, and in such a case it is not really of practical interest.

We argued in chap. 2 that at the clustering (dynamical) transition the equilibration time of a detailed balance local dynamics diverges. However, the clusters which appear at the dynamical energy $E_d > 0$ have bottom at an energy $E_{\text{bottom}} \leq E_d$ and numerical performance of the simulated annealing in the 3-coloring of random graphs [vMS02] suggests that E_{bottom} might be zero even if E_d is positive. More precise numerical investigation of this point is, however, needed.

F.2.2 Stochastic local search

Solving K -SAT by a pure random walk was suggested in [Pap91]:

PURE-RANDOM-WALK-SAT(T_{\max})

```

1 Draw a random assignment of variables;
2  $T \leftarrow 0$ ;
3 repeat Draw a random unsatisfied constraint  $a$ ;
4     Flip a random variable  $i$  belonging to  $a$ ;
5      $T \leftarrow T + 1$ ;
6 until Solution is found or  $T > NT_{\max}$ ;

```

In random 3-SAT this simple strategy seems to work in linear time up to $\alpha_{\text{RW}} \approx 2.7$ [SM03]. Improvements of the PURE-RANDOM-WALK-SAT have led to a large class of so-called stochastic local search algorithms. All are based on a random walk in the configurational space with more complicated rules about which variables would be flipped. The version called WALKSAT introduced in [SKC94, SKC96] became, next to the DPLL-based exact solvers, the most widely used solver of practical SAT instances. In random 3-SAT the Walk-SAT with $p = 0.5$ was shown to work in linear time up to about $\alpha_{\text{WS}} = 4.15$ [AGK04].

WALKSAT(T_{\max}, p)

```

1 Draw a random assignment of variables;
2  $T \leftarrow 0$ ;
3 repeat Pick a random unsatisfied constraint  $a$ ;
4     if Exists a variable  $i$  in  $a$  that is not necessary in any other constraint;
5         then Flip this variable  $i$ ;
6     else if RAND <  $p$ ;
7         then Flip a random variable  $i$  belonging to  $a$ ;
8     else Flip  $i$  (from  $a$ ) that minimizes the # of unsat. constraints;
9      $T \leftarrow T + 1$ ;
10 until Solution is found or  $T > NT_{\max}$ ;

```

Several other variants of stochastic local search on random 3-SAT were studied in [SAO05] showing that with a proper tuning of parameters like p the linear performance can be extended up to at least $\alpha \approx 4.20$. Finally a version of the stochastic local search called ASAT was introduced in [AA06]. In random 3-SAT ASAT works in a linear time at least up to $\alpha = 4.21$ [AA06]. We adapted the implementation of ASAT and studied its performance in coloring and on the occupation CSPs.

ASAT(T_{\max}, p)

```

1 Draw a random assignment of variables;
2  $T \leftarrow 0$ ;
3 Create the list  $\{v\}$  of variables which are present in unsatisfied constraints.
4 repeat Pick a random variable  $i$  from the list  $\{v\}$ ;
5     Compute the change of energy  $\Delta E$  if the value of  $i$  is flipped.
6     if  $\Delta E \leq 0$ ;
7         then Flip  $i$ ;
8     else if RAND <  $p$ ;
9         then Flip  $i$ ;
10    else Do nothing;

```

- 11 Update list $\{v\}$ of variables which are present in unsatisfied constraints.
 12 $T \leftarrow T + 1$;
 13 **until** Solution is found or $T > NT_{\max}$;

In the coloring problem where variables take one from more than two possible values, the only modification of ASAT is that we choose a random value into which the variable is flipped on line 5. The performance for the 4-coloring of Erdős-Rényi graphs was sketched in fig. 2.3.

There are two free parameters in the ASAT algorithm, the maximal number of steps per variable T_{\max} and, more importantly, the greediness (temperature-like) parameter p , which need to be optimized. In [AA06] and [ZK07] it was observed that in the random K-SAT and random coloring problems the optimal value of p does not depend on the system size N , neither very strongly on the constraint density α . But these observation might be model dependent, as it indeed seems to be the case for the locked problems.

F.2.3 Belief propagation reinforcement

A "search of improvement" solver can also be based on the belief propagation equations. The idea of the *belief propagation reinforcement*, introduced in [CFMZ05]¹, is to write belief propagation equations with an external "magnetic" field (site potential) $\mu_{s_i}^i$

$$\psi_{s_i}^{a \rightarrow i} = \frac{1}{Z^{a \rightarrow i}} \sum_{A_{s_i} + \sum_{s_j} s_j = 1} \prod_{j \in \partial a - i} \chi_{s_j}^{j \rightarrow a}, \quad (\text{F.4a})$$

$$\chi_{s_i}^{i \rightarrow a} = \frac{1}{Z^{i \rightarrow a}} \mu_{s_i}^i \prod_{b \in \partial i - a} \psi_{s_i}^{b \rightarrow i}, \quad (\text{F.4b})$$

and then iteratively update this field in order to make the procedure converge to a solution given by the direction of the external field $r_i = \text{argmax}_{s_i} \mu_{s_i}^i$. At every step the configuration given by the direction of the external field is regarded as the current configuration which is being improved.

The question is how to update the external field. The basic idea is to choose the local potential $\mu_{s_i}^i$ in some way proportional to the current value of the total marginal probability $\chi_{s_i}^i$, which is computed without the external fields as

$$\chi_{s_i}^i = \frac{1}{Z^i} \prod_{b \in \partial i} \psi_{s_i}^{b \rightarrow i}. \quad (\text{F.5})$$

How exactly, and how often should the value of local potential be updated is open to many different implementations, some of them can be found in [BZ06, DRZ08]. The same as in the local search algorithm it is not well understood, beyond a purely experimental level, how the details of the implementation influence the final performance. We tried several ways and the best performing seemed to be the following

$$\mu_1^i = (\pi)^{l_i - 1}, \quad \mu_0^i = (1 - \pi)^{l_i - 1}, \quad \text{if } \xi_0^i > \xi_1^i, \quad (\text{F.6a})$$

$$\mu_1^i = (1 - \pi)^{l_i - 1}, \quad \mu_0^i = (\pi)^{l_i - 1}, \quad \text{if } \xi_0^i \leq \xi_1^i, \quad (\text{F.6b})$$

¹Strictly speaking the reinforcement strategy was first introduced for the survey propagation equations, but the concept is the same for belief propagation.

where $0 \leq \pi \leq 1/2$, l_i is the degree of variable i and the auxiliary variable $\xi_{s_i}^i$ is computed before updating the field μ^i

$$\xi_{s_i}^i = (\mu_{s_i}^i)^{\frac{1}{l_i-1}} \chi_{s_i}^i. \quad (\text{F.7})$$

BP-REINFORCEMENT(T_{\max}, n, π)

- 1 Initialize $\mu_{s_i}^i$ and $\psi_{s_i}^{a \rightarrow i}$ randomly;
- 2 $T \leftarrow 0$;
- 3 Compute the current configuration $r_i = \operatorname{argmax}_{s_i} \mu_{s_i}^i$;
- 4 **repeat** Make n sweeps of the BP iterations (F.4a-F.4b);
- 5 Update all the local fields $\mu_{s_i}^i$ according to (F.6a-F.6b);
- 6 Update $r_i = \operatorname{argmax}_{s_i} \mu_{s_i}^i$;
- 7 $T \leftarrow T + 1$;
- 8 **until** $\{r\}$ is a solution or $T > T_{\max}$;

How should the strength of the forcing π be chosen? Empirically we observed three different regimes:

- a) $\pi_{\text{BP-like}} < \pi < 0.5$: When the forcing is weak the BP-REINFORCEMENT converges very fast to a BP-like fixed point, the values of the local fields do not point towards any solution. On contrary many constraints are violated by the final configuration $\{r_i\}$.
- b) $\pi_{\text{conv}} < \pi < \pi_{\text{BP-like}}$: The BP-REINFORCEMENT converges to a solution $\{r_i\}$.
- c) $0 < \pi < \pi_{\text{conv}}$: When the forcing is too strong the BP-REINFORCEMENT does not converge. And many constraints are violated by the configuration $\{r_i\}$ which is reached after T_{\max} steps.

When the constraint density in the CSP is large the regime b) disappears and $\pi_{\text{conv}} = \pi_{\text{BP-like}}$. For an obvious reason our goal is to find $\pi_{\text{conv}} < \pi < \pi_{\text{BP-like}}$. The point $\pi_{\text{BP-like}}$ is very easy to find, because for larger π the convergence of BP-REINFORCEMENT to a BP-like fixed point happens in just several sweeps. Thus in all the runs we chose π to be just bellow $\pi_{\text{BP-like}}$, that is to hit the possible gap between $\pi_{\text{BP-like}}$ and π_{conv} . The value of π chosen in this way does not seem to depend on the size of the system, it, however, depends slightly on the constraint density.

Experimentally it seems that the optimal number of BP sweeps on line 4 of BP-REINFORCEMENT is very small, typically $n = 2$, in agreement with [CFMZ05]. We observed with a surprise that when n is much larger not only the total running time is larger but the overall performance of the algorithm is worse.

In the regime where the BP-REINFORCEMENT algorithm performs well the median running time T seems to be independent of the size, leading to an overall linear time complexity. The total CPU time is comparable to the time achieved by the stochastic local search ASAT.

There is an imperfection of our implementation of the BP-REINFORCEMENT, because in small fraction of cases, for all connectivities, the algorithm is blocked in a configuration with only 1-3 violated constraints. If this happens we reinforce stronger the problematic variables which sometimes shifts the problem to a different part of the graph, where it might be resolved. Also a restart leads to a solution.

We tested the BP-REINFORCEMENT algorithm mainly in the occupation CSPs, the results are shown in sec. 4.3. Survey propagation reinforcement can be implemented in a similar way, as was done originally in [CFMZ05].

References

- [AA06] John Ardelius and Erik Aurell. Behavior of heuristics on large and hard satisfiability problems. *Phys. Rev. E*, 74:037702, 2006.
- [AAA⁺08] Mikko Alava, John Ardelius, Erik Aurell, Petteri Kaski, Supriya Krishnamurthy, Pekka Orponen, and Sakari Seitz. Circumspect descent prevails in solving random constraint satisfaction problems. *Proc. Nat. Acad. Sc. U.S.A.*, 105:15253–15257, 2008.
- [ACIM01] Dimitris Achlioptas, Arthur Chtcherba, Gabriel Istrate, and Cristopher Moore. The phase transition in 1-in-k sat and nae 3-sat. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 721–722, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics.
- [AGK04] Erik Aurell, Uri Gordon, and Scott Kirkpatrick. Comparing beliefs, surveys and random walks. In *Proc. of 17th NIPS*, page 804, 2004.
- [AH77a] K. Appel and W. Haken. Every planar map is four colorable. ii. reducibility. *Illinois J. Math.*, 21, 1977.
- [AH77b] K. Appel and W. Haken. Every planar map is four colorable. part i. discharging. *Illinois J. Math.*, 21, 1977.
- [AKKK01] Dimitris Achlioptas, Lefteris M. Kirousis, Evangelos Kranakis, and Danny Krizanc. Rigorous results for random (2+p)-sat. *Theoretical Computer Science*, 256(1-2):109–129, 2001.
- [AKS04] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. Primes is in p. *Annals of Mathematics*, 160(2):781–793, 2004.
- [Ald01] D. J. Aldous. The $\zeta(2)$ limit in the random assignment problem. *Rand. Struct. Algo.*, 18:381–418, 2001.
- [AM03] D. Achlioptas and C. Moore. Almost all graphs with average degree 4 are 3-colorable. *J. Comput. Syst. Sci.*, 67:441, 2003.
- [Ang95] C. A. Angell. Formation of glasses from liquids and biopolymers. *Science*, 267(5206):1924–1935, 1995.
- [ANP05] D. Achlioptas, A. Naor, and Y. Peres. Rigorous location of phase transitions in hard optimization problems. *Nature*, 435:759–764, 2005.
- [ART06] Dimitris Achlioptas and Federico Ricci-Tersenghi. On the solution-space geometry of random constraint satisfaction problems. In *Proc. of 38th STOC*, pages 130–139, New York, NY, USA, 2006. ACM.
- [AZ08] John Ardelius and Lenka Zdeborová. Exhaustive enumeration unveils clustering and freezing in random 3-sat. *Phys. Rev. E*, 78:040101(R), 2008.
- [BB04] J. P. Bouchaud and G. Biroli. On the Adam-Gibbs-Kirkpatrick-Thirumalai-Wolynes scenario for the viscosity increase of classes. *J. Chem. Phys.*, 121:7347–7354, 2004.
- [BCKM98] J.-P. Bouchaud, L. Cugliandolo, J. Kurchan, and M. Mézard. Out of equilibrium dynamics in spin glasses and other glassy systems. In A. P. Young, editor, *Spin Glasses and Random Fields*. World Scientific, Singapore, 1998.
- [BG06] Antar Bandyopadhyay and David Gamarnik. Counting without sampling: new algorithms for enumeration problems using statistical physics. In *Proc. of the 17th ACM-SIAM Symposium on Discrete Algorithms*, pages 890 – 899, New York, USA, 2006. ACM Press.
- [BM02] G. Biroli and M. Mézard. Lattice glass models. *Phys. Rev. Lett.*, 88:025501, 2002.
- [BMP⁺03] A. Braunstein, R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina. Polynomial iterative algorithms for coloring and analyzing random graphs. *Phys. Rev. E*, 68:036702, 2003.

- [BMW00] G. Biroli, R. Monasson, and M. Weigt. A variational description of the ground state structure in random satisfiability problems. *Eur. Phys. J. B*, 14:551, 2000.
- [BMWZ03] A. Braunstein, M. Mézard, M. Weigt, and R. Zecchina. Constraint satisfaction by survey propagation. In Allon Percus, Gabriel Istrate, and Cristopher Moore, editors, *Computational Complexity and Statistical Physics*, page 107. Oxford University Press, 2003.
- [BMZ05] A. Braunstein, M. Mézard, and R. Zecchina. Survey propagation: An algorithm for satisfiability. *Random Struct. Algorithms*, 27(2):201–226, 2005.
- [BN06] Mohsen Bayati and Chandra Nair. A rigorous proof of the cavity method for counting matchings. arXiv:cond-mat/0607290v2 [cond-mat.dis-nn], 2006.
- [Bov06] Anton Bovier. *Statistical Mechanics of Disordered Systems: A Mathematical Perspective*. Cambridge University Press, 2006.
- [BSS05] M. Bayati, D. Shah, and M. Sharma. Maximum weight matching via max-product belief propagation. In *Proc. IEEE Int. Symp. Information Theory*, 2005.
- [BSS06] M. Bayati, D. Shah, and M. Sharma. A simpler max-product maximum weight matching algorithm and the auction algorithm. In *Proc. IEEE Int. Symp. Information Theory*, 2006.
- [Bul02] Andrei A. Bulatov. A dichotomy theorem for constraints on a three-element set. *Proc. of FOCS 2002*, page 649, 2002.
- [BZ04] A. Braunstein and R. Zecchina. Survey propagation as local equilibrium equations. *Journal of Statistical Mechanics: Theory and Experiment*, page P06007, 2004.
- [BZ06] A. Braunstein and R. Zecchina. Learning by Message Passing in Networks of Discrete Synapses. *Physical Review Letters*, 96(3):030201, 2006.
- [CA96] James M. Crawford and Larry D. Auton. Experimental results on the crossover point in random 3-sat. *Artif. Intell.*, 81(1-2):31–57, 1996.
- [CFMZ05] Joel Chavas, Cyril Furtlehner, Marc Mézard, and Riccardo Zecchina. Survey-propagation decimation through distributed local computations. *J. Stat. Mech.*, page P11016, 2005.
- [Che08] M. Chertkov. Exactness of belief propagation for some graphical models with loops. *J. Stat. Mech.*, page P10016, 2008.
- [CK93] L. F. Cugliandolo and J. Kurchan. Analytical solution of the off-equilibrium dynamics of a long-range spin glass model. *Phys. Rev. Lett.*, 71:173, 1993.
- [CKRT05] Tommaso Castellani, Florent Krzakala, and Federico Ricci-Tersenghi. Spin glass models with ferromagnetically biased couplings on the bethe lattice: analytic solution and numerical simulations. *Eur. Phys. J. B*, 47:99, 2005.
- [CKT91] Peter Cheeseman, Bob Kanefsky, and William M. Taylor. Where the Really Hard Problems Are. In *Proc. 12th IJCAI*, pages 331–337, San Mateo, CA, USA, 1991. Morgan Kaufmann.
- [CM04] H. Connamacher and M. Molloy. The exact satisfiability threshold for a potentially intractable random constraint satisfaction problem. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*. IEEE Computer Society, 2004.
- [CNRTZ03] Tommaso Castellani, Vincenzo Napolano, Federico Ricci-Tersenghi, and Riccardo Zecchina. Bicoloring random hypergraphs. *J. Phys. A*, 36:11037, 2003.
- [Con04] H. Connamacher. A random constraint satisfaction problem that seems hard for dpll. In *SAT 2004 - The Seventh International Conference on Theory and Applications of Satisfiability Testing, 10-13 May 2004, Vancouver, BC, Canada, Online Proceedings*, 2004.
- [Coo71] Stephen A. Cook. The complexity of theorem-proving procedures. In *Proc. 3rd STOC*, pages 151–158, New York, NY, USA, 1971. ACM.

- [dAT78] J. R. L. de Almeida and D. J. Thouless. Stability of the Sherrington-Kirkpatrick solution of a spin-glass model. *J. Phys. A*, 11:983–990, 1978.
- [DBM00] Olivier Dubois, Yacine Boufkhad, and Jacques Mandler. Typical random 3-sat formulae and the satisfiability threshold. In *SODA '00: Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 126–127, Philadelphia, PA, USA, 2000. Society for Industrial and Applied Mathematics.
- [Der80] B. Derrida. Random-energy model: Limit of a family of disordered models. *Phys. Rev. Lett.*, 45:79–82, 1980.
- [Der81] B. Derrida. Random-energy model: An exactly solvable model of disordered systems. *Phys. Rev. B*, 24:2613–2626, 1981.
- [DLL62] Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem-proving. *Commun. ACM*, 5(7):394–397, 1962.
- [DM08] A. Dembo and A. Montanari. Ising models on locally tree-like graphs. arXiv:0804.4726v2 [math.PR], 2008.
- [DMMZ08] H. Daudé, T. Mora, M. Mézard, and R. Zecchina. Pairs of sat assignments and clustering in random boolean formulae. *Theoretical Computer Science*, 393:260–279, 2008.
- [DP60] Martin Davis and Hillary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.
- [DRZ08] L. Dall'Asta, A. Ramezanpour, and R. Zecchina. Entropy landscape and non-gibbs solutions in constraint satisfaction problems. *Phys. Rev. E*, 77:031118, 2008.
- [DS01] Debenedetti and Stillinger. Supercooled liquids and the glass transition. *Nature*, 410(6825):259–267, 2001.
- [EA75] S. F. Edwards and P. W. Anderson. Theory of spin-glasses. *J. Phys. F*, 5:965–974, 1975.
- [EFF] <http://w2.eff.org/awards/coop-prime-rules.php>.
- [EKPS00] William Evans, Claire Kenyon, Yuval Peres, and Leonard J. Schulman. Broadcasting on trees and the Ising model. *Ann. Appl. Probab.*, 10:410–433, 2000.
- [ER59] P. Erdős and A. Rényi. On random graphs. *Publ. Math. Debrecen*, 6:290–297, 1959.
- [FA86] Y. Fu and P. W. Anderson. Application of statistical mechanics to NP-complete problems in combinatorial optimization. *J. Phys. A*, 19:1605–1620, 1986.
- [FH91] K. H. Fischer and J. A. Hertz. *Spin-Glasses*, volume 1 of *Cambridge Studies in Magnetism*. Cambridge University Press, Cambridge, 1991.
- [FL03] S. Franz and M. Leone. Replica bounds for optimization problems and diluted spin systems. *J. Stat. Phys.*, 3-4:535–564, 2003.
- [FLRTZ01] Silvio Franz, Michele Leone, Federico Ricci-Tersenghi, and Riccardo Zecchina. Exact solutions for diluted spin glasses and optimization problems. *Phys. Rev. Lett.*, 87(12):127209, Aug 2001.
- [FLT03] Silvio Franz, Michele Leone, and Fabio Lucio Toninelli. Replica bounds for diluted non-poissonian spin systems. *J. Phys. A: Math. Gen.*, 36:10967–10985, 2003.
- [FP95] S. Franz and G. Parisi. Recipes for Metastable States in Spin Glasses. *Journal de Physique I*, 5:1401–1415, November 1995.
- [FP97] Silvio Franz and Giorgio Parisi. Phase diagram of coupled glassy systems: A mean-field study. *Phys. Rev. Lett.*, 79(13):2486–2489, Sep 1997.
- [Fri99] E. Friedgut. Sharp thresholds of graph properties, and the k -sat problem. *J. Amer. Math. Soc.*, 12, 1999.
- [Gal62] Robert G. Gallager. Low-density parity check codes. *IEEE Trans. Inform. Theory*, 8:21–28, 1962.

- [Gal68] R. G. Gallager. *Information theory and reliable communication*. John Wiley and Sons, New York, 1968.
- [Gar85] E. Gardner. Spin glasses with p-spin interactions. *Nuclear Physics B*, 257:747–765, 1985.
- [Gas02] William I. Gasarch. The $p = np$ poll. *SIGACT News*, 33(2):34–47, 2002.
- [Geo88] H.-O. Georgii. *Gibbs Measures and Phase Transitions*. De Gruyter, Berlin, 1988.
- [GJ79] M.R. Garey and D.S. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. Freeman, San Francisco, 1979.
- [GKS85] D. J. Gross, I. Kanter, and H. Sompolinsky. Mean-field theory of the potts glass. *Phys. Rev. Lett.*, 55(3):304–307, Jul 1985.
- [GM07] A. Gerschenfeld and A. Montanari. Reconstruction for models on random graphs. In *Proc. of 48th FOCS*, pages 194–204. IEEE Computer Society, 2007.
- [Gol79] A. Goldberg. On the complexity of the satisfiability problem. In *Courant Computer Science Report*, volume 16, New York, NY, USA, 1979.
- [GPB82] A. Goldberg, Jr. P.W. Purdom, and C.A. Brown. Average time analysis of simplified davis-putnam procedure. *Information Process. Lett.*, 15(2):72–75, 1982. see also Errata, vol. 16, 1983, p. 213.
- [HJKN06] Harri Haanpää, Matti Järvisalo, Petteri Kaski, and Ilkka Niemelä. Hard satisfiable clause sets for benchmarking equivalence reasoning techniques. *Journal on Satisfiability, Boolean Modeling and Computations*, 2:27–46, 2006.
- [HS03] M. Hajiaghayi and G. B. Sorkin. The Satisfiability Threshold of Random 3-SAT Is at Least 3.52. arXiv: math/0310193, 2003.
- [Jan05] V. Janiš. Stability of solutions of the sherrington-kirkpatrick model with respect to replications of the phase space. *Phys. Rev. B*, 71:214403, 2005.
- [JM04] Svante Janson and Elchanan Mossel. Robust reconstruction on trees is determined by the second eigenvalue. *Ann. Probab.*, 32:2630–2649, 2004.
- [Jon02] J. Jonasson. Uniqueness of uniform random colorings of regular trees. *Statistics and Probability Letters*, 57:243–248, 2002.
- [Kar72] R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New-York, 1972.
- [Kau48] W. Kauzmann. The nature of the glassy state and the behavior of liquids at low temperatures. *Chem. Rev.*, 43:219, 1948.
- [KFL01] F. R. Kschischang, B. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Inform. Theory*, 47(2):498–519, 2001.
- [KGV83] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
- [KK07] F. Krzakala and J. Kurchan. A landscape analysis of constraint satisfaction problems. *Phys. Rev. E*, 76:021122, 2007.
- [KKL03] A. Kaporis, L. Kirousis, and E. Lalas. Selecting complementary pairs of literals. In *Proc. LICS'03 Workshop on Typical Case Complexity and Phase Transitions*, 2003.
- [KMRT⁺07] Florent Krzakala, Andrea Montanari, Federico Ricci-Tersenghi, Guilhem Semerjian, and Lenka Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *Proc. Natl. Acad. Sci. U.S.A.*, 104:10318, 2007.
- [KPW04] F. Krzakala, A. Pagnani, and M. Weigt. Threshold values, stability analysis and high- q asymptotics for the coloring problem on random graphs. *Phys. Rev. E*, 70:046705, 2004.
- [KS66a] H. Kesten and B. P. Stigum. Additional limit theorems for indecomposable multidimensional galton-watson processes. *The Annals of Mathematical Statistics*, 37:1463, 1966.

- [KS66b] H. Kesten and B. P. Stigum. Limit theorems for decomposable multi-dimensional galton-watson processes. *J. Math. Anal. Appl.*, 17:309, 1966.
- [KS87] I. Kanter and H. Sompolinsky. Graph optimisation problems nad the potts glass. *J. Phys. A: Math. Gen.*, 20:L673–L679, 1987.
- [KS94] S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random boolean expression. *Science*, 264:1297–1301, 1994.
- [KSS07a] Lukas Kroc, Ashish Sabharwal, and Bart Selman. Decimation strategies: Surveys, beliefs, and local information. in preparation, 2007.
- [KSS07b] Lukas Kroc, Ashish Sabharwal, and Bart Selman. Survey propagation revisited. In *Proc. of 23rd AUA*, pages 217–226, Arlington, Virginia, USA, 2007. AUA Press.
- [KT87a] T.R. Kirkpatrick and D. Thirumalai. Dynamics of the structural glass transition and the p -spin-interaction spin-glass model. *Phys. Rev. Lett.*, 58:2091, 1987.
- [KT87b] T.R. Kirkpatrick and D. Thirumalai. p -spin-interaction spin-glass models: Connections with the structural glass problem. *Phys. Rev. B*, 36:5388, 1987.
- [KW05] Vladimir Kolmogorov and Martin Wainwright. On the optimality of tree-reweighted max-product message passing. In *In 21st Conference on Uncertainty in Artificial Intelligence (UAI)*, 2005.
- [KZ08a] Florent Krzakala and Lenka Zdeborová. Phase transitions and computational difficulty in random constraint satisfaction problems. *J. Phys.: Conf. Ser.*, 95:012012, 2008.
- [KZ08b] Florent Krzakala and Lenka Zdeborová. Potts glass on random graphs. *Eur. Phys. Lett.*, 81:57005, 2008.
- [LRTZ01] M. Leone, F. Ricci-Tersenghi, and R. Zecchina. Phase coexistence and finite-size scaling in random combinatorial problems. *J. Phys. A*, 34:4615, 2001.
- [Luc91] T. Luczak. The chromatic number of random graphs. *Combinatorica*, 11:45, 1991.
- [LW04] S. Linusson and J. Wästlund. A proof of Parisi’s conjecture on the random assignment problem. *Probability Theory and Related Fields*, 128:419–440, 2004.
- [Mer98] Stephan Mertens. Phase transition in the number partitioning problem. *Phys. Rev. Lett.*, 81(20):4281–4284, Nov 1998.
- [Mer00] Stephan Mertens. Random costs in combinatorial optimization. *Phys. Rev. Lett.*, 84(6):1347–1350, Feb 2000.
- [MM06a] Marc Mézard and Andrea Montanari. Reconstruction on trees and spin glass transition. *J. Stat. Phys.*, 124:1317–1350, september 2006.
- [MM06b] T. Mora and M. Mézard. Geometrical organization of solutions to random linear Boolean equations. *Journal of Statistical Mechanics: Theory and Experiment*, 10:P10007, October 2006.
- [MM08] M. Mézard and A. Montanari. *Information, Physics, Computation: Probabilistic approaches*. Cambridge University Press, Cambridge, 2008. In preparation: www.lptms.u-psud.fr/membres/mezard/.
- [MMR04] O. C. Martin, M. Mézard, and O. Rivoire. Frozen glass phase in the multi-index matching problem. *Phys. Rev. Lett.*, 93:217205, 2004.
- [MMR05] O. C. Martin, M. Mézard, and O. Rivoire. Random multi-index matching problems. *J. Stat. Mech.*, 2005.
- [MMR⁺06] E. Maneva, T. Meltzer, J. Raymond, A. Sportiello, and L. Zdeborová. A hike in the phases of the 1-in-3 satisfiability. In *Les Houches Summer School, Session LXXXV 2006 on Complex Systems*. 2006.

- [MMW07] Elitza N. Maneva, Elchanan Mossel, and Martin J. Wainwright. A new look at survey propagation and its generalizations. *J. ACM*, 54(4), 2007.
- [MMZ05] M. Mézard, T. Mora, and R. Zecchina. Clustering of solutions in the random satisfiability problem. *Physical Review Letters*, 94:197205, 2005.
- [MMZ06] Stephan Mertens, Marc Mézard, and Riccardo Zecchina. Threshold values of random k-sat from the cavity method. *Random Struct. Algorithms*, 28(3):340–373, 2006.
- [MN95] David J. C. MacKay and R. M. Neal. Good codes based on very sparse matrices. In *Proceedings of the 5th IMA Conference on Cryptography and Coding*, pages 100–111, London, UK, 1995. Springer-Verlag.
- [Mon95] R. Monasson. Structural glass transition and the entropy of the metastable states. *Phys. Rev. Lett.*, 75:2847, 1995.
- [Mon01] A. Montanari. The glassy phase of Gallager codes. *Eur. Phys. J. B.*, 23:121–136, 2001.
- [Mor07] T. Mora. *Géométrie et inférence dans l’optimisation et en théorie de l’information*. PhD thesis, Université Paris-Sud, 2007. <http://tel.archives-ouvertes.fr/tel-00175221/en/>.
- [Mos01] Elchanan Mossel. Reconstruction on trees: Beating the second eigenvalue. *Ann. Appl. Probab.*, 11(1):285–300, 2001.
- [Mos04] E. Mossel. Survey: Information flow on trees. In J. Neštril and P. Winkler, editors, *Graphs, Morphisms and Statistical Physics*, DIMACS series in discrete mathematics and theoretical computer science, pages 155–170, 2004.
- [MP85] M. Mézard and G. Parisi. Replicas and optimization. *J. Physique*, 46:L771–L778, 1985.
- [MP86a] M. Mézard and G. Parisi. Mean-field equations for the matching and the travelling salesman problem. *Europhys. Lett.*, 2:913–918, 1986.
- [MP86b] M. Mézard and G. Parisi. A replica analysis of the travelling salesman problem. *J. Physique*, 47:1285–1296, 1986.
- [MP00] M. Mézard and G. Parisi. Statistical physics of structural glasses. *J. Phys.: Condens. Matter*, 12:6655–6673, 2000.
- [MP01] M. Mézard and G. Parisi. The bethe lattice spin glass revisited. *Eur. Phys. J. B*, 20:217, 2001.
- [MP03] M. Mézard and G. Parisi. The cavity method at zero temperature. *J. Stat. Phys.*, 111:1–34, 2003.
- [MPR05] M. Mézard, M. Palassini, and O. Rivoire. Landscape of solutions in constraint satisfaction problems. *Phys. Rev. Lett.*, 95:200202, 2005.
- [MPRT04] A. Montanari, G. Parisi, and F. Ricci-Tersenghi. Instability of one-step replica-symmetry-broken phase in satisfiability problems. *J. Phys. A*, 37:2073, 2004.
- [MPS⁺84] M. Mézard, G. Parisi, N. Sourlas, G. Toulouse, and M. A. Virasoro. Replica symmetry breaking and the nature of the spin-glass phase. *J. Physique*, 45:843–854, 1984.
- [MPV85] M. Mézard, G. Parisi, and M. A. Virasoro. Random free energies in spin-glasses. *J. Physique Lett.*, 46:L217–L222, 1985.
- [MPWZ02] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina. Coloring random graphs. *Phys. Rev. Lett.*, 89:268701, 2002.
- [MPZ02] M. Mézard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297:812–815, 2002.
- [MRR⁺53] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of State Calculations by Fast Computing Machines. *The Journal of Chemical Physics*, 21:1087–1092, June 1953.
- [MRT03] A. Montanari and F. Ricci-Tersenghi. On the nature of the low-temperature phase in discontinuous mean-field spin glasses. *Eur. Phys. J. B*, 33:339, 2003.

- [MRT04] Andrea Montanari and Federico Ricci-Tersenghi. Cooling-schedule dependence of the dynamics of mean-field glasses. *Phys. Rev. B*, 70(13):134406, 2004.
- [MRTS07] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Solving constraint satisfaction problems through belief propagation-guided decimation. arXiv:0709.1667v1 [cs.AI], 2007.
- [MRTS08] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian. Clusters of solutions and replica symmetry breaking in random k -satisfiability. *J. Stat. Mech.*, page P04004, 2008.
- [MRTZ03] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina. Alternative solutions to diluted p -spin models and XORSAT problems. *J. Stat. Phys.*, 111:505, 2003.
- [MS05] A. Montanari and G. Semerjian. From large scale rearrangements to mode coupling phenomenology. *Phys. Rev. Lett.*, 94:247201, 2005.
- [MS06a] E. Marinari and G. Semerjian. On the number of circuits in random graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 6:P06019, June 2006.
- [MS06b] A. Montanari and G. Semerjian. On the dynamics of the glass transition on bethe lattices. *J. Stat. Phys.*, 124:103–189, 2006.
- [MS06c] A. Montanari and G. Semerjian. Rigorous inequalities between length and time scales in glassy systems. *J. Stat. Phys.*, 125:23, 2006.
- [MS07] Elitza Maneva and Alistair Sinclair. On the satisfiability threshold and clustering of solutions of random 3-sat formulas. arXiv:0710.0805v1 [cs.CC], 2007.
- [MSL92] David G. Mitchell, Bart Selman, and Hector J. Levesque. Hard and easy distributions for SAT problems. In *Proc. 10th AAAI*, pages 459–465, Menlo Park, California, 1992. AAAI Press.
- [MZ96] R. Monasson and R. Zecchina. Entropy of the K -satisfiability problem. *Phys. Rev. Lett.*, 76:3881–3885, 1996.
- [MZ97] Rémi Monasson and Riccardo Zecchina. Statistical mechanics of the random k -satisfiability model. *Phys. Rev. E*, 56(2):1357–1370, Aug 1997.
- [MZ02] M. Mézard and R. Zecchina. Random k -satisfiability problem: From an analytic solution to an efficient algorithm. *Phys. Rev. E*, 66:056126, 2002.
- [MZ08] T. Mora and L. Zdeborová. Random subcubes as a toy model for constraint satisfaction problems. *J. Stat. Phys.*, 131:1121–1138, 2008.
- [MZK⁺99a] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. 2+p-sat: Relation of typical-case complexity to the nature of the phase transition. *Random Structures and Algorithms*, 15:414, 1999.
- [MZK⁺99b] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic phase transitions. *Nature*, 400:133–137, 1999.
- [NS92] C. M. Newman and D. L. Stein. Multiple states and thermodynamic limits in short-ranged ising spin-glass models. *Phys. Rev. B*, 46(2):973–982, Jul 1992.
- [Pal83] R.G. Palmer. In *Proc. of the Heidelberg Colloquium on spin glasses, Lecture Notes in Physics 192*, Berlin, 1983. Springer.
- [Pap91] Christos H. Papadimitriou. On selecting a satisfying truth assignment (extended abstract). In *Proceedings of the 32nd annual symposium on Foundations of computer science*, pages 163–169, Los Alamitos, CA, USA, 1991. IEEE Computer Society Press.
- [Pap94] C. H. Papadimitriou. *Computational complexity*. Addison-Wesley, 1994.
- [Par80a] G. Parisi. Magnetic properties of spin-glasses in a new mean-field theory. *J. Phys. A*, 13:1887–1895, 1980.
- [Par80b] G. Parisi. The order parameter for spin-glasses: A function on the interval 0–1. *J. Phys. A*, 13:1101–1112, 1980.

- [Par80c] G. Parisi. A sequence of approximated solutions to the SK model for spin-glasses. *J. Phys. A Lett.*, 13:L115–L121, 1980.
- [Par02a] G. Parisi. On local equilibrium equations for clustering states. arXiv:cs.CC/0212047, 2002.
- [Par02b] G. Parisi. On the survey-propagation equations for the random k-satisfiability problem. arXiv:cs.CC/0212009, 2002.
- [Par03] G. Parisi. Some remarks on the survey decimation algorithm for k-satisfiability. arXiv:cs/0301015, 2003.
- [Pea82] J. Pearl. Reverend bayes on inference engines: A distributed hierarchical approach. In *Proceedings American Association of Artificial Intelligence National Conference on AI*, pages 133–136, Pittsburgh, PA, USA, 1982.
- [Pea88] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.
- [PT04] Dmitry Panchenko and Michel Talagrand. Bounds for diluted mean-fields spin glass models. *Probability Theory and Related Fields*, 130(3):319–336, 2004.
- [PY97] J. Pitman and M. Yor. The two-parameter poisson-dirichlet distribution derived from a stable subordinator. *Ann. Probab.*, 25:855–900, 1997.
- [RBMM04] O. Rivoire, G. Biroli, O. C. Martin, and M. Mézard. Glass models on bethe lattices. *Eur. Phys. J. B*, 37:55–78, 2004.
- [Riv05] Olivier Rivoire. *Phases vitreuses, optimisation et grandes déviations*. PhD thesis, Université Paris-Sud, 2005. <http://tel.archives-ouvertes.fr/tel-00009956/en/>.
- [RSZ07] J. Raymond, A. Sportiello, and L. Zdeborová. The Phase Diagram of 1-in-3 Satisfiability Problem. *Phys. Rev. E*, 76:011101, February 2007.
- [RU01] T. Richardson and R. Urbanke. The capacity of low-density parity-check codes under message-passing decoding. *IEEE Trans. Inform. Theory*, 47, 2001.
- [SAO05] Sakari Seitz, Mikko Alava, and Pekka Orponen. Focused local search for random 3-satisfiability. *J. Stat. Mech.*, page P06006, 2005.
- [Sem08] Guilhem Semerjian. On the freezing of variables in random constraint satisfaction problems. *J. Stat. Phys.*, 130:251, 2008.
- [Sin93] A. Sinclair. *Algorithms for Random Generation and Counting: A Markov Chain Approach*. Birkhäuser, Boston-Basel-Berlin, 1993.
- [SK75] D. Sherrington and S. Kirkpatrick. Solvable model of a spin-glass. *Phys. Rev. Lett.*, 35:1792–1796, 1975.
- [SKC94] Bart Selman, Henry A. Kautz, and Bram Cohen. Noise strategies for improving local search. In *Proc. 12th AAAI*, pages 337–343, Menlo Park, CA, USA, 1994. AAAI Press.
- [SKC96] Bart Selman, Henry A. Kautz, and Bram Cohen. Local search strategies for satisfiability testing. In Michael Trick and David Stifler Johnson, editors, *Proceedings of the Second DIMACS Challenge on Cliques, Coloring, and Satisfiability*, Providence RI, 1996.
- [SLM92] Bart Selman, Hector J. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In Paul Rosenbloom and Peter Szolovits, editors, *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 440–446, Menlo Park, California, 1992. AAAI Press.
- [Sly09] Allan Sly. Reconstruction of random colourings. *Communications in Mathematical Physics*, 288:943–961, 2009.
- [SM03] Guilhem Semerjian and Rémi Monasson. Relaxation and metastability in a local search procedure for the random satisfiability problem. *Phys. Rev. E*, 67(6):066103, Jun 2003.

- [SW04] G. Semerjian and M. Weigt. Approximation schemes for the dynamics of diluted spin models: the Ising ferromagnet on a Bethe lattice. *Journal of Physics A Mathematical General*, 37:5525–5546, May 2004.
- [Tal03] M. Talagrand. *Spin glasses : a challenge for mathematicians. Cavity and mean field models*. Springer-Verlag, New-York, 2003.
- [Tal06] M. Talagrand. The parisi formula. *Ann. Math.*, 163:221–263, 2006.
- [VB85] L. Viana and A. J. Bray. Phase diagrams for dilute spin-glasses. *J. Phys. C*, 18:3037–3051, 1985.
- [vMS02] J. van Mourik and D. Saad. Random graph coloring: Statistical physics approach. *Phys. Rev. E*, 66:056120, 2002.
- [Wil02] David B. Wilson. On the critical exponents of random k-sat. *Random Structures and Algorithms*, 21:182–195, 2002.
- [YFW00] J.S. Yedidia, W.T. Freeman, and Y. Weiss. Generalized belief propagation. In *Advances in Neural Information Processing Systems (NIPS)*, volume 13, pages 689–695, 2000.
- [YFW03] J.S. Yedidia, W.T. Freeman, and Y. Weiss. Understanding belief propagation and its generalizations. In *Exploring Artificial Intelligence in the New Millennium*, pages 239–236. Science & Technology Books, 2003.
- [Zho03] H. Zhou. Vertex cover problem studied by cavity method: Analytics and population dynamics. *Eur. Phys. J. B*, 32:265–20, 2003.
- [ZK07] L. Zdeborová and F. Krzakala. Phase transitions in the coloring of random graphs. *Phys. Rev. E*, 76:031131, 2007.
- [ZM06] L. Zdeborová and M. Mézard. The number of matchings in random graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 5:P05003, May 2006.
- [ZM08] L. Zdeborová and M. Mézard. Hard constraint satisfaction problems. *Phys. Rev. Lett.*, 101:078702, 2008.