18

LINEAR EQUATIONS WITH BOOLEAN VARIABLES

Solving a system of linear equations over a finite field $\mathbb{F}$ is arguably one of the most fundamental operations in mathematics. Several algorithms have been devised to accomplish such a task in polynomial time. The best known is Gauss elimination, that has $O(N^3)$ complexity (here N is number of variables in the linear system, and we assume the number of equations to be $M = \Theta(N)$). As a matter of fact, one can improve over Gaussian elimination, and the best existing algorithm for general systems has complexity $O(N^{2.376\cdots})$. Faster methods do also exist for special classes of instances.

The set of solutions of a linear system is an affine subspace of $\mathbb{F}^N$. Despite this apparent simplicity, the geometry of affine or linear subspaces of $\mathbb{F}^N$ can be surprisingly rich. This observation is systematically exploited in coding theory. Linear codes are just linear spaces over finite fields. Nevertheless, they are known to achieve Shannon capacity on memoryless symmetric channels, and their structure is far from trivial, as we already saw in Ch. 11.

From a different point of view, linear systems are a particular example of constraint satisfaction problems. We can associate with a linear system a decision problem (establishing whether it has a solution), a counting problem (counting the number of solutions), an optimization problem (minimize the number of violated equations). While the first two are polynomial, the latter is known to be NP-hard.

In this chapter we consider a specific ensemble of random linear systems over $\mathbb{Z}_2$ (the field of integers modulo 2), and discuss the structure of its set of solutions. The ensemble definition is mainly motivated by its analogy with other random constraint satisfaction problems, which also explains the name XOR-satisfiability (XORSAT).

In the next section we provide the precise definition of the XORSAT ensemble and recall a few elementary properties of linear algebra. We also introduce one of the main objects of study of this chapter: the SAT-UNSAT threshold. Section 18.2 takes a detour into the properties of belief propagation for XORSAT. These are shown to be related to the correlation structure of the uniform measure over solutions and, in Sec. 18.3, to the appearance of a 2-core in the associated factor graph. Sections 18.4 and 18.5 build on these results to compute the SAT-UNSAT threshold and characterize the structure of the solution space. While many results can be derived rigorously, XORSAT offers an ideal playground for understanding the non-rigorous cavity method that will be further developed in the next chapters. This is the object of Sec. 18.6.

18.1 Definitions and general remarks

18.1.1 *Linear systems*

Let $\mathbb{H}$ be a $M \times N$ matrix with entries $H_{ai} \in \{\mathtt{0},\mathtt{1}\}$, $a \in \{1,\dots,M\}$, $i \in \{1,\dots,N\}$, and let $\underline{b}$ be a M-component vector with binary entries $b_a \in \{\mathtt{0},\mathtt{1}\}$. An instance of the **XORSAT** problem is given by a couple $(\mathbb{H},\underline{b})$. The decision problem requires to find a N-component vector $\underline{x}$ with binary entries $x_i \in \{\mathtt{0},\mathtt{1}\}$ which solves the linear system $\mathbb{H}\underline{x} = \underline{b} \mod 2$, or to show that the system has no solution. The name XORSAT comes from the fact that sum modulo 2 is equivalent to the 'exclusive OR' operation: the problem is whether there exists an assignment of the variables $\underline{x}$ which satisfies a set of XOR clauses. We shall thus say that the instance is SAT (resp. UNSAT) whenever the linear system has (resp. doesn't have) a solution.

We shall furthermore be interested in the set of solutions, to be denoted by $\mathcal{S}$, in its size $Z = |\mathcal{S}|$, and in the properties of the uniform measure over $\mathcal{S}$. This is defined by

$$\mu(\underline{x}) = \frac{1}{Z}\,\mathbb{I}(\,\mathbb{H}\underline{x} = \underline{b} \mod 2\,) = \frac{1}{Z}\prod_{a=1}^{M}\psi_a(\underline{x}_{\partial a})\,, \tag{18.1}$$

where $\partial a = (i_a(1),\dots,i_a(K))$ is the set of non-vanishing entries in the a-th row of $\mathbb{H}$, and $\psi_a(\underline{x}_{\partial a})$ is the characteristic function for the a-th equation in the linear system (explicitly $\psi_a(\underline{x}_{\partial a}) = \mathbb{I}(x_{i_1(a)} \oplus \cdots \oplus x_{i_K(a)} = b_a)$, where we denote as usual by $\oplus$ the sum modulo 2). In the following we shall omit to specify that operations are carried $\mod 2$ when clear from the context.

When $\mathbb{H}$ has row weigth p (i.e. each row has p non-vanishing entries), the problem is related to a p-spin glass model. Writing $\sigma_i = 1-2x_i$ and $J_a = 1-2b_a$, we can associate to the XORSAT instance the energy function

$$E(\underline{\sigma}) = \sum_{a=1}^{M}\Big(1 - J_a\prod_{j\in\partial a}\sigma_j\Big)\,, \tag{18.2}$$

which counts (twice) the number of violated equations. This can be regarded as a p-spin glass energy function with binary couplings. The decision XORSAT problem asks whether there exists a spin configuration $\underline{\sigma}$ with zero energy or, in physical jargon, whether the above energy function is 'unfrustrated.' If there exists such a configuration, $\log Z$ is the ground state entropy of the model.

A natural generalization is the MAX-XORSAT problem. This requires to find a configuration which maximizes the number of satisfied equations, i.e. minimizes $E(\underline{\sigma})$. In the following we shall use the language of XORSAT but of course all statements have their direct counterpart in p-spin glasses.

Let us recall a few well known facts of linear algebra that will be useful in the following:

(i) The image of $\mathbb{H}$ is a vector space of dimension rank($\mathbb{H}$) (rank($\mathbb{H}$) is the number of independent lines in $\mathbb{H}$); the kernel of $\mathbb{H}$ (the set $\mathcal{S}_0$ of $\underline{x}$ which solve the homogeneous system $\mathbb{H}\underline{x} = \underline{0}$) is a vector space of dimension $N - \text{rank}(\mathbb{H})$.

(ii) As a consequence, if $M \leq N$ and $\mathbb{H}$ has rank M (all of its lines are independent), then the linear system $\mathbb{H}\underline{x} = \underline{b}$ has a solution for any choice of $\underline{b}$.

(iii) Conversely, if $\text{rank}(\mathbb{H}) < M$, the linear system has a solution if and only if $\underline{b}$ is in the image of $\mathbb{H}$.

If the linear system has at least one solution $\underline{x}_*$, then the set of solutions $\mathcal{S}$ is an affine space of dimension $N - \text{rank}(\mathbb{H})$: one has $\mathcal{S} = \underline{x}_* + \mathcal{S}_0$, and $Z = 2^{N-\text{rank}(\mathbb{H})}$. We shall denote by $\mu_0(\,\cdot\,)$ the uniform measure over the set $\mathcal{S}_0$ of solutions of the homogeneous linear system:

$$\mu_0(\underline{x}) = \frac{1}{Z_0}\,\mathbb{I}(\,\mathbb{H}\underline{x} = \underline{0} \mod 2\,) = \frac{1}{Z_0}\prod_{a=1}^{M}\psi_a^0(\underline{x}_{\partial a}) \tag{18.3}$$

where ψ_a^0 has the same expression as ψ_a but with $b_a = 0$. Notice that μ_0 is always well defined as a probability distribution, because the homogeneous systems has at least the solution $\underline{x} = \underline{0}$, while μ is well defined only for SAT instances. The linear structure has several important consequences.

- If $\underline{y}$ is a solution of the inhomogeneous system, and if $\underline{x}$ is a uniformly random solution of the homogeneous linear system (with distribution μ_0), then $\underline{x}' = \underline{x} \oplus \underline{y}$ is a uniformly random solution of the inhomogeneous system (its probability distribution is μ).
- Under the measure μ_0, there exist only two sorts of variables x_i, those which are 'frozen to 0,' (i.e. take value 0 in all of the solutions) and those which are 'free' (taking value 0 or 1 in one half of the solutions). Under the measure μ (when it exists), a bit can be frozen to 0, frozen to 1, or free. These facts are proved in the next exercise.

Exercise 18.1 Let $f : \{0,1\}^N \to \{0,1\}$ be a linear function (explicitly, $f(\underline{x})$ is the sum of a subset $x_{i(1)}, \dots, x_{i(n)}$ of the bits, $\mod 2$).

(a) If $\underline{x}$ is drawn from the distribution μ_0, $f(\underline{x})$ becomes a random variable taking values in $\{0,1\}$. Show that, if there exists a configuration $\underline{x}$ with $\mu_0(\underline{x}) > 0$ and $f(\underline{x}) = 1$, then $\mathbb{P}\{f(\underline{x}) = 0\} = \mathbb{P}\{f(\underline{x}) = 1\} = 1/2$. In the opposite case, $\mathbb{P}\{f(\underline{x}) = 0\} = 1$.

(b) Suppose that there exists at least one solution to the system $\mathbb{H}\underline{x} = \underline{b}$, so that μ exists. Consider the random variable $f(\underline{x})$ obtained by drawing $\underline{x}$ from the distribution μ. Show that one of the following three cases occurs: $\mathbb{P}\{f(\underline{x}) = 0\} = 1$, $\mathbb{P}\{f(\underline{x}) = 0\} = 1/2$, or $\mathbb{P}\{f(\underline{x}) = 0\} = 0$.

These results apply in particular to the marginal of bit i, using $f(\underline{x}) = x_i$.

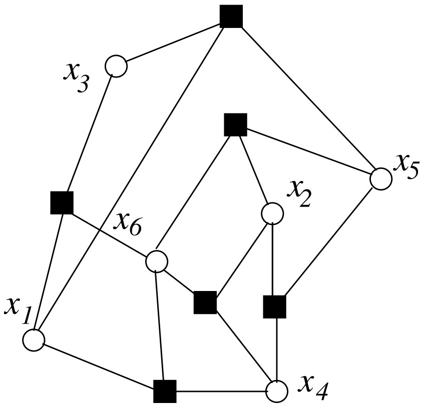

FIG. 18.1. Factor graph for a 3-XORSAT instance with $N = 6$, $M = 6$.

Exercise 18.2 Show that:

(a) If the number of solutions of the homogeneous system is $Z_0 = 2^{N-M}$, then the inhomogeneous system is satisfiable (SAT), and has 2^{N-M} solutions, for any $\underline{b}$.

(b) Conversely, if the number of solutions of the homogeneous system is $Z_0 > 2^{N-M}$, then the inhomogeneous one is SAT only for a fraction $2^{N-M}/Z_0$ of the $\underline{b}$'s.

The distribution μ admits a natural factor graph representation: variable nodes are associated to variables and factor nodes to linear equations, cf. Fig. 18.1. Given a XORSAT formula F (i.e. a pair $\mathbb{H}, \underline{b}$), we denote by $G(F)$ the associated factor graph. It is remarkable that one can identify sub-graphs of $G(F)$ that serve as witnesses of satisfiability or unsatisfiability of F. By this we mean that the existence of such sub-graphs implies satisfiability/unsatisfiability of F. The existence of a simple witness for unsatisfiability is intimately related to the polynomial nature of XORSAT. Such a witness is obtained as follows. Given a subset L of the clauses, draw the factor graph including all the clauses in L, all the adjacent variable nodes, and the edges between them. If this subgraph has *even degree at each of the variable nodes*, and if $\oplus_{a\in L} b_a = 1$, then L is a witness for unsatisfiability. Such a subgraph is sometimes called a frustrated hyper-loop (in analogy with frustrated loops appearing in spin glasses, where function nodes have degree 2).

Exercise 18.3 Consider a 3-XORSAT instance defined through the 6×6 matrix

$$\mathbb{H} = \begin{bmatrix} 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \end{bmatrix} \qquad (18.4)$$

(a) Compute the rank($\mathbb{H}$) and list the solutions of the homogeneous linear system.

(b) Show that the linear system $\mathbb{H}\underline{b} = \underline{0}$ has a solution if and only if $b_1 \oplus b_4 \oplus b_5 \oplus b_6 = 0$. How many solution does it have in this case?

(b) Consider the factor graph associated to this linear system, cf. Fig. 18.1. Show that each solution of the homogeneous system must correspond to a subset U of variable nodes with the following property. The sub-graph induced by U and including all of the adjacent function nodes, has even degree at the function nodes. Find one sub-graph with this property.

18.1.2 *Random XORSAT*

The **random K-XORSAT** ensemble is defined by taking $\underline{b}$ uniformly at random in $\{0, 1\}^M$, and $\mathbb{H}$ uniformly at random among the $N \times M$ matrices with entries in $\{0, 1\}$ which have exactly K non-vanishing elements per row. Each equation thus involves K distinct variables chosen uniformly among the $\binom{N}{K}$ K-uples, and the resulting factor graph is distributed according to the $\mathbb{G}_N(K, M)$ ensemble.

A slightly different ensemble is defined by including each of the $\binom{N}{K}$ possible lines with K non-zero entries independently with probability $p = N\alpha/\binom{N}{K}$. The corresponding factor graph is then distributed according to the $\mathbb{G}_N(K, \alpha)$ ensemble.

Given the relation between homogeneous and inhomogeneous systems described above, it is quite natural to introduce an ensemble of homogeneous linear systems. This is defined by taking $\mathbb{H}$ distributed as above, but with $\underline{b} = \underline{0}$. Since an homogeneous linear system has always at least one solution, this ensemble is sometimes referred to as **SAT K-XORSAT** or, in its spin interpretation, as the **ferromagnetic K-spin model**. Given a K-XORSAT formula F, we shall denote by F_0 the formula corresponding to the homogeneous system.

We are interested in the limit of large systems N, $M \to \infty$ with $\alpha = M/N$ fixed. By applying Friedgut's Theorem, cf. Sec. 10.5, it is possible to show that, **for** $K \geq 3$, the probability for a random formula F to be SAT has a **sharp threshold**. More precisely, there exists $\alpha_{\rm s}^{(N)}(K)$ such that for $\alpha > (1 + \delta)\alpha_{\rm s}^{(N)}(K)$ (respectively $\alpha < (1 - \delta)\alpha_{\rm s}^{(N)}(K)$), $\mathbb{P}\{F \text{ is SAT}\} \to 0$ (respectively $\mathbb{P}\{F \text{ is SAT}\} \to 1$) as $N \to \infty$.

A moment of thought reveals that $\alpha_{\rm s}^{(N)}(K) = \Theta(1)$. Let us give two simple bounds to convince the reader of this statement.

Upper bound: The relation between the homogeneous and the original linear system derived in Exercise 18.2 implies that $\mathbb{P}\{F \text{ is SAT}\} = 2^{N-M}\mathbb{E}\{1/Z_0\}$. As $Z_0 \geq 1$, we get $\mathbb{P}\{F \text{ is SAT}\} \leq 2^{-N(\alpha-1)}$ and therefore $\alpha_{\rm s}^{(N)}(K) \leq 1$.

Lower bound: For $\alpha < 1/K(K-1)$ the factor graph associated with F is formed, with high probability, by finite trees and uni-cyclic components. This corresponds to the matrix $\mathbb{H}$ being decomposable into blocks, each one corresponding to a connected component. The reader can show that, for $K \geq 3$ both a tree formula and a uni-cyclic component correspond to a linear system of full rank. Since each block has full rank, $\mathbb{H}$ has full rank as well. Therefore $\alpha_{\rm s}^{(N)}(K) \geq 1/K(K-1)$.

Exercise 18.4 There is no sharp threshold for $K = 2$.

(a) Let $c(G)$ be the cyclic number of the factor graph G (number of edges minus vertices, plus number of connected components) of a random 2-XORSAT formula. Show that $\mathbb{P}\{F \text{ is SAT}\} = \mathbb{E}\, 2^{-c(G)}$.

(b) Argue that this implies that $\mathbb{P}\{F \text{ is SAT}\}$ is bounded away from 1 for any $\alpha > 0$.

(c) Show that $\mathbb{P}\{F \text{ is SAT}\}$ is bounded away from 0 for any $\alpha < 1/2$.

[Hint: remember the geometrical properties of G discussed in Secs. 9.3.2, 9.4.]

In the next sections we shall show that $\alpha_{\rm s}^{(N)}(K)$ has a limit $\alpha_{\rm c}(K)$ and compute it explicitly. Before dwelling into this, it is instructive to derive two improved bounds.

Exercise 18.5 In order to obtain a better upper bound on $\alpha_{\rm s}^{(N)}(K)$ proceed as follows:

(a) Assume that, for any α, $Z_0 \geq 2^{Nf_K(\alpha)}$ with probability larger than some $\varepsilon > 0$ at large N. Show that $\alpha_{\rm s}^{(N)}(K) \leq \alpha^*(K)$, where $\alpha^*(K)$ is the smallest value of α such that $1 - \alpha - f_K(\alpha) \leq 0$.

(b) Show that the above assumption holds with $f_K(\alpha) = e^{-K\alpha}$, and that this yields $\alpha^*(3) \approx 0.941$. What is the asymptotic behavior of $\alpha^*(K)$ for large K? How can you improve the exponent $f_K(\alpha)$?

Exercise 18.6 A better lower bound on $\alpha_{\rm s}^{(N)}(K)$ can be obtained through a first moment calculation. In order to simplify the calculations we consider here a modified ensemble in which the K variables entering in equation a are chosen independently and uniformly at random (they do not need to be distinct). The scrupulous reader can check at the end that returning to the original ensemble brings only little changes.

(a) Show that for a positive random variable Z, $(\mathbb{E}Z)(\mathbb{E}[1/Z]) \ge 1$. Deduce that $\mathbb{P}\{F \text{ is SAT}\} \ge 2^{N-M}/\mathbb{E}\, Z_{F_0}$.

(b) Prove that

$$\mathbb{E}\, Z_{F_0} = \sum_{w=0}^{N} \binom{N}{w} \left[\frac{1}{2}\left(1 + \left(1 - \frac{2w}{N}\right)^K\right)\right]^M . \tag{18.5}$$

(c) Let $g_K(x) = \mathcal{H}(x) + \alpha \log\left[\frac{1}{2}\left(1 + (1-2x)^K\right)\right]$ and define $\alpha_*(K)$ to be the largest value of α such that the maximum of $g_K(x)$ is achieved at $x = 1/2$. Show that $\alpha_{\rm s}^{(N)}(K) \ge \alpha_*(K)$. One finds $\alpha_*(3) \approx 0.889$.

18.2 Belief propagation

18.2.1 *BP messages and density evolution*

Equation (18.1) provides a representation of the uniform measure over solutions of a XORSAT instance as a graphical model. This suggests to apply message passing techniques. We will describe here belief propagation and analyze its behavior. While this may seem at first sight a detour from the objective of computing $\alpha_{\rm s}^{(N)}(K)$, it will instead provide some important insight.

Let us assume that the linear system $\mathbb{H}\underline{x} = \underline{b}$ admits at least one solution, so that the model (18.1) is well defined. We shall first study the homogeneous version $\mathbb{H}\underline{x} = \mathtt{0}$, i.e. the measure μ_0, and then pass to μ. Applying the general definitions of Ch. 14, the BP update equations (14.14), (14.15) for the homogeneous problem read

$$\nu_{i\to a}^{(t+1)}(x_i) \cong \prod_{b\in\partial i\setminus a} \widehat{\nu}_{b\to i}^{(t)}(x_i)\,, \qquad \widehat{\nu}_{a\to i}^{(t)}(x_i) \cong \sum_{\underline{x}_{\partial a\setminus i}} \psi_a^0(\underline{x}_{\partial a}) \prod_{j\in\partial a\setminus i} \nu_{j\to a}^{(t)}(x_j)\,. \tag{18.6}$$

These equations can be considerably simplified using the linear structure. We have seen that under μ_0,there are two types of variables, those ‘frozen to $\mathtt{0}$’ (i.e. equal to $\mathtt{0}$ in all solutions), and those which are ‘free’ (equally likely to be $\mathtt{0}$ or $\mathtt{1}$). BP aims at determining whether any single bit belongs to one class or the other. Consider now BP messages, which are also distributions over $\{\mathtt{0},\mathtt{1}\}$. Suppose that at time $t = 0$ they also take one of the two possible values that we denote as $*$ (corresponding to the uniform distribution) and $\mathtt{0}$ (distribution

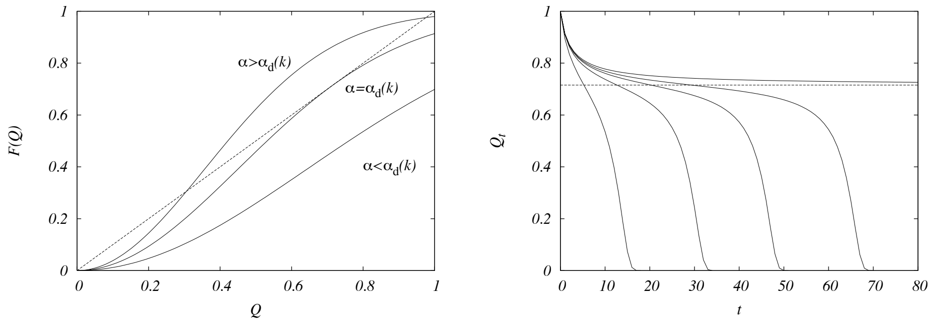

FIG. 18.2. Density evolution for the fraction of 0 messages for 3-XORSAT. On the left: the mapping $F(Q) = 1 - \exp(-K\alpha Q^{K-1})$ below, at and above the critical point $\alpha_{\rm d}(K=3) \approx 0.818468$. On the right: evolution of Q_t for (from bottom to top) $\alpha = 0.75$, 0.8, 0.81, 0.814, 0.818468.

entirely supported on 0). Then, it is not hard to show that this remains true at all subsequent times. The BP update equations (18.6) simplify under this initialization (they reduce to the erasure decoder of Sect. 15.3):

- At a variable node the outgoing message is 0 unless all the incoming are $*$.
- At a function node the outgoing message is $*$ unless all the incoming are 0.

(The message coming out of a degree-1 variable node is always $*$).

These rules preserve a natural partial ordering. Given two sets of messages $\nu = \{\nu_{i\to a}\}$, $\widetilde{\nu} = \{\widetilde{\nu}_{i\to a}\}$, let us say that $\nu^{(t)} \succeq \widetilde{\nu}^{(t)}$ if for each directed edge $i \to a$ where the message $\widetilde{\nu}^{(t)}_{i\to a} = 0$, then $\nu^{(t)}_{i\to a} = 0$ as well. It follows immediately from the update rules that, if for some time t the messages are ordered as $\nu^{(t)} \succeq \widetilde{\nu}^{(t)}$, then this order is preserved at all later times: $\nu^{(s)} \succeq \widetilde{\nu}^{(s)}$ for all $s > t$.

This partial ordering suggests to pay special attention to the two 'extremal' initial conditions, namely $\nu^{(0)}_{i\to a} = *$ for all directed edges $i \to a$, or $\nu^{(0)}_{i\to a} = 0$ for all $i \to a$. The fraction of edges Q_t that carry a message 0 at time t is a deterministic quantity in the $N \to \infty$ limit. It satisfies the recursion:

$$Q_{t+1} = 1 - \exp\{-K\alpha Q_t^{K-1}\}\,, \tag{18.7}$$

with $Q_0 = 1$ (respectively $Q_0 = 0$) for the 0 initial condition (resp. the $*$ initial condition). The density evolution recursion (18.7) is represented pictorially in Fig. 18.2.

Under the $*$ initial condition, we have $Q_t = 0$ at all times t. In fact the all $*$ message configuration is always a fixed point of BP. On the other hand, when $Q_0 = 1$, one finds two possible asymptotic behaviors: $Q_t \to 0$ for $\alpha < \alpha_{\rm d}(K)$, while $Q_t \to Q > 0$ for $\alpha > \alpha_{\rm d}(K)$. Here $Q > 0$ is the largest positive solution of $Q = 1 - \exp\{-K\alpha\, Q^{K-1}\}$. The critical value $\alpha_{\rm d}(K)$ of the density of equations $\alpha = M/N$ separating these two regimes is:

$$\alpha_{\rm d}(K) = \sup \left\{ \alpha \ \text{ such that } \ \forall x \in]0,1] : \ x < 1 - e^{-K\alpha\, x^{K-1}} \right\} . \qquad (18.8)$$

We get for instance $\alpha_{\rm d}(K) \approx 0.818469$, 0.772280, 0.701780 for, respectively, $K = 3$, 4, 5 and $\alpha_{\rm d}(K) = \log K/K[1 + o(1)]$ as $K \to \infty$.

We therefore found two regimes for the homogeneous random XORSAT problem in the large-N limit. For $\alpha < \alpha_{\rm d}(K)$ there is a unique BP fixed point with all messages[25] equal to $*$. The BP prediction for single bit marginals that corresponds to this fixed point is $\nu_i(x_i = \mathtt{0}) = \nu_i(x_i = \mathtt{1}) = 1/2$.

For $\alpha > \alpha_{\rm d}(K)$ there exists more than one BP fixed points. We have found two of them: the all-$*$ one, and one with density of $*$'s equal to Q. Other fixed points of the inhomogeneous problem can be constructed as follows for $\alpha \in]\alpha_{\rm d}(K), \alpha_{\rm s}(K)[$. Let $\underline{x}^{(*)}$ be a solution of the inhomogeneous problem, and $\nu, \widehat{\nu}$ be a BP fixed point in the homogeneous case. Then the messages $\nu^{(*)}, \widehat{\nu}^{(*)}$ defined by:

$$\begin{aligned} &\nu^{(*)}_{j\to a}(x_j = \mathtt{0}) = \nu^{(*)}_{j\to a}(x_j = \mathtt{1}) = 1/2 && \text{if } \nu_{j\to a} = *, \\ &\nu^{(*)}_{j\to a}(x_j) = \mathbb{I}(x_j = x^{(*)}_j) && \text{if } \nu_{j\to a} = \mathtt{0}, \end{aligned} \qquad (18.9)$$

(and similarly for $\widehat{\nu}^{(*)}$) are a BP fixed point for the inhomogeneous problem.

For $\alpha < \alpha_{\rm d}(K)$, the inhomogeneous problem admits, with high probability, a unique BP fixed point. This is a consequence of the exercise:

Exercise 18.7 Consider a BP fixed point $\nu^{(*)}, \widehat{\nu}^{(*)}$ for the inhomogeneous problem, and assume all the messages to be of one of three types: $\nu^{(*)}_{j\to a}(x_j = \mathtt{0}) = 1$, $\nu^{(*)}_{j\to a}(x_j = \mathtt{0}) = 1/2$, $\nu^{(*)}_{j\to a}(x_j = \mathtt{0}) = 0$. Assume furthermore that messages are not 'contradictory,' i.e. that there exists no variable node i such that $\widehat{\nu}^{(*)}_{a\to i}(x_i = \mathtt{0}) = 1$ and $\widehat{\nu}^{(*)}_{b\to i}(x_i = \mathtt{0}) = 0$.

Construct a non-trivial BP fixed point for the homogeneous problem.

18.2.2 *Correlation decay*

The BP prediction is that for $\alpha < \alpha_{\rm d}(K)$ the marginal distribution of any bit x_i is uniform under either of the measures μ_0, μ. The fact that the BP estimates do not depend on the initialization is an indication that the prediction is correct. Let us prove that this is indeed the case. To be definite we consider the homogeneous problem (i.e. μ_0). The inhomogeneous case follows, using the general remarks in Sec. 18.1.1.

We start from an alternative interpretation of Q_t. Let $i \in \{1, \dots, N\}$ be a uniformly random variable index and consider the ball of radius t around i in the factor graph G: $\mathsf{B}_{i,t}(G)$. Set to $x_j = \mathtt{0}$ all the variables x_j outside this ball, and let $Q^{(N)}_t$ be the probability that, under this condition, all the solutions of the linear system $\mathbb{H}\underline{x} = \underline{\mathtt{0}}$ have $x_i = \mathtt{0}$. Then the convergence of $\mathsf{B}_{i,t}(G)$ to the tree model

[25] While a vanishing fraction of messages $\nu_{i\to a} = \mathtt{0}$ is not excluded by our argument, it can be ruled out by a slightly lenghtier calculation.

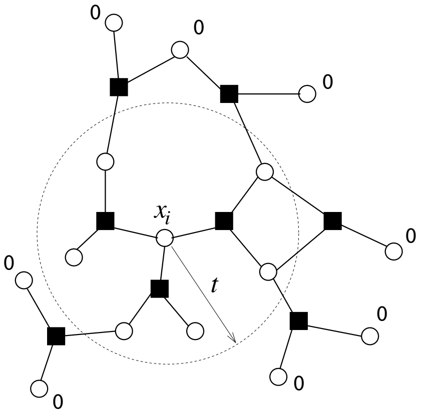

FIG. 18.3. Factor graph for a 3-XORSAT instance with the depth $t = 1$ neighborhood of vertex i, $\mathsf{B}_{i,t}(G)$ indicated. Fixing to $\mathtt{0}$ all the variables outside $\mathsf{B}_{i,t}(G)$ does not imply that x_i must be $\mathtt{0}$ in order to satisfy the homogeneous linear system.

$\mathbb{T}(K, \alpha)$ discussed in Sec. 9.5 implies that, for any given t, $\lim_{N\to\infty} Q_t^{(N)} = Q_t$. It also determines the initial condition to $Q_0 = 1$.

Consider now the marginal distribution $\mu_0(x_i)$. If $x_i = \mathtt{0}$ in all the solutions of $\mathbb{H}\underline{x} = \underline{\mathtt{0}}$, then, a fortiori $x_i = \mathtt{0}$ in all the solutions that fulfill the additional condition $x_j = \mathtt{0}$ for $j \not\in \mathsf{B}_{i,t}(G)$. Therefore we have $\mathbb{P}\{\mu_0(x_i = \mathtt{0}) = 1\} \le Q_t^{(N)}$. By taking the $N \to \infty$ limit we get

$$\lim_{N\to\infty} \mathbb{P}\{\mu_0(x_i = \mathtt{0}) = 1\} \le \lim_{N\to\infty} Q_t^{(N)} = Q_t \,. \tag{18.10}$$

Letting $t \to \infty$ and noticing that the left hand side does not depend on t we get $\mathbb{P}\{\mu_0(x_i = \mathtt{0}) = 1\} \to 0$ as $N \to \infty$. In other words, all but a vanishing fraction of the bits are free for $\alpha < \alpha_{\rm d}(K)$.

The number Q_t also has another interpretation, which generalizes to the inhomogeneous problem. Choose a solution $\underline{x}^{(*)}$ of the homogeneous linear system and, instead of fixing the variables outside the ball of radius t to $\mathtt{0}$, let's fix them to $x_j = x_j^{(*)}$, $j \not\in \mathsf{B}_{i,t}(G)$. Then $Q_t^{(N)}$ is the probability that $x_i = x_i^{(*)}$, under this condition. The same argument holds in the inhomogeneous problem, with the measure μ: if $\underline{x}^{(*)}$ is a solution of $\mathbb{H}\underline{x} = \underline{b}$ and we fix the variables outside $\mathsf{B}_{i,t}(G)$ to $x_j = x_j^{(*)}$, the probability that $x_i = x_i^{(*)}$ under this condition is again $Q_t^{(N)}$. The fact that $\lim_{t\to\infty} Q_t = 0$ when $\alpha < \alpha_{\rm d}(K)$ thus means that a spin decorrelates from the whole set of variables at distance larger than t, when t is large. This formulation of correlation decay is rather specific to XORSAT, because it relies on the dichotomous nature of this problem: Either the 'far away' variables completely determine x_i, or they have no influence on it and it is uniformly random. A more generic formulation of correlation decay, which generalizes to other

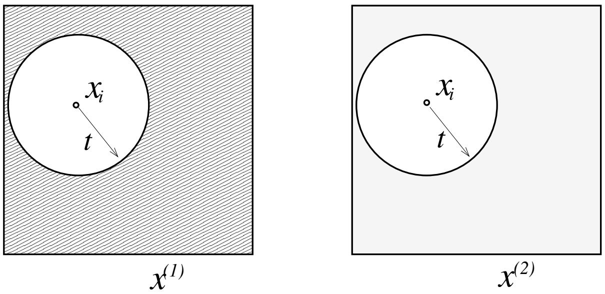

FIG. 18.4. A thought experiment: fix variables 'far' from i to two different assignments and check the influence on x_i. For $\alpha < \alpha_{\rm d}$ there is no influence

problems which don't have this dichotomy property, consists in comparing two different choices $\underline{x}^{(1)}$, $\underline{x}^{(2)}$ of the reference solution (cf. Fig. 18.4). For $\alpha < \alpha_{\rm d}(K)$ the correlations decay even in the worst case:

$$\lim_{N\to\infty} \mathbb{E}\left\{ \sup_{\underline{x}^{(1)},\underline{x}^{(2)}} |\mu(x_i|\underline{x}^{(1)}_{\sim i,t}) - \mu(x_i|\underline{x}^{(2)}_{\sim i,t})| \right\} = Q_t \to 0\,, \tag{18.11}$$

as $t \to \infty$. In Ch. 22 we will discuss weaker (non worst-case) definitions of correlation decay, and their relation to phase transitions.

18.3 Core percolation and BP

18.3.1 2-*core and peeling*

What happens for $\alpha > \alpha_{\rm d}(K)$? A first hint is provided by the instance in Fig. 18.1. In this case, the configuration of messages $\nu^{(t)}_{i\to a} = \mathtt{0}$ on all directed edges $i \to a$ is a fixed point of the BP update for the homogeneous system. A moment of thought shows that this happens because G has the property that each variable node has degree at least 2. We shall now see that, for $\alpha > \alpha_{\rm d}(K)$, G has with high probability a subgraph (called 2-core) with the same property.

We already encountered similar structures in Sec. 15.3, where we identified them as responsible for errors in iterative decoding of LDPC codes over the erasure channel. Let us recall the relevant points[26] from that discussion. Given a factor graph G, a stopping set is a subset of the function nodes such that all the variables have degree larger or equal to 2 in the induced sub-graph. The 2-core is the largest stopping set. It is unique and can be found by the peeling algorithm, which amounts to iterating the following procedure: find a variable node of degree 0 or 1 (a "leaf"), erase it together with the factor node adjacent to it, if there is one. The resulting subgraph, the 2-core, will be denoted as $K_2(G)$.

The peeling algorithm is of direct use for solving the linear system: if a variable has degree 1, the unique equation where it appears allows to express it

[26] Notice that the structure causing decoding errors was the 2-core of the *dual* factor graph that is obtained by exchanging variable and function nodes.

in terms of other variables. It can thus be eliminated from the problem. The 2-core of G is the factor graph associated to the linear system obtained by iterating this procedure, which we shall refer to as the "core system". The original system has a solution if and only if the core does. We shall refer to solutions of the core system as to **core solutions**.

18.3.2 *Clusters*

Core solutions play an important role as the set of solutions can be partitioned according to their core values. Given an assignment $\underline{x}$, denote by $\pi_*(\underline{x})$ its projection onto the core, i.e. the vector of those entries in $\underline{x}$ that corresponds to vertices in the core. Suppose that the factor graph has a non-trivial 2-core, and let $\underline{x}^{(*)}$ be a core solution. We define the **cluster** associated with $\underline{x}^{(*)}$ as the set of solutions to the linear system such that $\pi_*(\underline{x}) =$
$ux^{(*)}$ (the reason for the name cluster will become clear in Sec. 18.5). If the core of G is empty, we shall adopt the convention that the entire set of solutions forms a unique cluster.

Given a solution $\underline{x}^{(*)}$ of the core linear system, we shall denote the corresponding cluster as $\mathcal{S}(\underline{x}^{(*)})$. One can obtain the solutions in $\mathcal{S}(\underline{x}^{(*)})$ by running the peeling algorithm in the reverse direction, starting from $\underline{x}^{(*)}$. In this process one finds variable which are uniquely determined by $\underline{x}^{(*)}$, they form what is called the 'backbone' of the graph. More precisely, we define the **backbone** $B(G)$ as the sub-graph of G that is obtained augmenting $K_2(G)$ as follows. Set $B_0(G) = K_2(G)$. For any $t \geq 0$, pick a function node a which is not in $B_t(G)$ and which has at least $K-1$ of its neighboring variable nodes in $B_t(G)$, and build $B_{t+1}(G)$ by adding a (and its neighborhing variables) to $B_t(G)$. If no such function node exists, set $B(G) = B_t(G)$ and halt the procedure. This definition of $B(G)$ does not depend on the order in which function nodes are added. The backbone contains the 2-core, and is such that any two solutions of the linear system which belong to the same cluster, coincide on the backbone.

We have thus found that the variables in a linear system naturally divide into three possible types: The variables in the 2-core $K_2(G)$, those in $B(G) \setminus K_2(G)$ which are not in the core but are fixed by the core solution, and the variables which are not uniquely determined by $\underline{x}^{(*)}$. This distinction is based on the geometry of the factor graph, i.e. it depends only the matrix $\mathbb{H}$, and not on the value of the right hand side $\underline{b}$ in the linear system. We shall now see how BP finds these structures.

18.3.3 *Core, backbone, and belief propagation*

Consider the homogeneous linear system $\mathbb{H}\underline{x} = \mathtt{0}$, and run BP with initial condition $\nu^{(0)}_{i\to a} = \mathtt{0}$. Denote by $\nu_{i\to a}$, $\widehat{\nu}_{a\to i}$ the fixed point reached by BP (with measure μ_0) under this initialization (the reader is invited to show that such a fixed point is indeed reached after a number of iterations at most equal to the number of messages).

The fixed point messages $\nu_{i\to a}$, $\widehat{\nu}_{a\to i}$ can be exploited to find the 2-core

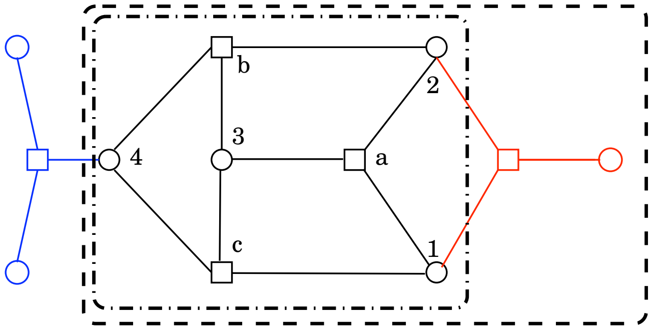

FIG. 18.5. The factor graph of a XORSAT problem, its core (central dash-dotted part) and its backbone (adding one function node and one variable on the right - dashed zone)

$K_2(G)$, using the following properties (which can be proved by induction over t): (i) $\nu_{i\to a} = \widehat{\nu}_{a\to i} = \mathtt{0}$ for each edge (i,a) in $K_2(G)$. (ii) A variable i belongs to the core $K_2(G)$ if and only if it receives messages $\widehat{\nu}_{a\to i} = \mathtt{0}$ from at least two of the neighboring function nodes $a \in \partial i$. (iii) If a function node $a \in \{1,\dots,M\}$ has $\nu_{i\to a} = \mathtt{0}$ for all the neighboring variable nodes $i \in \partial a$, then $a \in K_2(G)$.

The fixed point BP messages also contain information on the backbone: a variable i belongs to the backbone $B(G)$ if and only if it receives at least one message $\widehat{\nu}_{a\to i} = \mathtt{0}$ from its neighboring function nodes $a \in \partial i$.

Exercise 18.8 Consider a XORSAT problem described by the factor graph of Fig. 18.5.

(a) Using the peeling and backbone construction algorithms, check that the core and backbone are those described in the caption.

(b) Compute the BP messages found for the homogeneous problem as a fixed point of BP iteration starting from the all $\mathtt{0}$ configuration. Check the core and backbone that you obtain from these messages.

(c) Consider the general inhomogeneous linear system with the same factor graph. Show that there exist two solutions to the core system: $x_1 = \mathtt{0}, x_2 = b_b \oplus b_c, x_3 = b_a \oplus b_b \oplus b_c, x_4 = b_a \oplus b_b$ and $x_1 = \mathtt{0}, x_2 = b_b \oplus b_c \oplus \mathtt{1}, x_3 = b_a \oplus b_b \oplus b_c, x_4 = b_a \oplus b_b \oplus \mathtt{1}$. Identify the two clusters of solutions.

18.4 The SAT-UNSAT threshold in random XORSAT

We shall now see how a sharp characterization of the core size in random linear systems provides the clue to the determination of the satisfiability threshold. Remarkably, this characterization can again be achieved through an analysis of BP.

18.4.1 *The size of the core*

Consider an homogeneous linear system over N variables drawn from the random K-XORSAT ensemble, and let $\{\nu_{i\to a}^{(t)}\}$ denote the BP messages obtained from the initialization $\nu_{i\to a}^{(0)} = \mathtt{0}$. The density evolution analysis of Sec. 18.2.1 implies that the fraction of edges carrying a message $\mathtt{0}$ at time t, (we called it Q_t) satisfies the recursion equation (18.7). This recursion holds for any given t asymptotically as $N \to \infty$.

It follows from the same analysis that, in the large N limit, the messages $\widehat{\nu}_{a\to i}^{(t)}$ entering a variable node i are i.i.d. with $\mathbb{P}\{\widehat{\nu}_{a\to i}^{(t)} = \mathtt{0}\} = \widehat{Q}_t \equiv Q_t^{K-1}$. Let us for a moment assume that the limits $t \to \infty$ and $N \to \infty$ can be exchanged without much harm. This means that the fixed point messages $\widehat{\nu}_{a\to i}$ entering a variable node i are asymptotically i.i.d. with $\mathbb{P}\{\widehat{\nu}_{a\to i} = \mathtt{0}\} = \widehat{Q} \equiv Q^{K-1}$, where Q is the largest solution of the fixed point equation:

$$Q = 1 - \exp\{-K\alpha\widehat{Q}\}\ , \qquad \widehat{Q} = Q^{K-1}\,. \tag{18.12}$$

The number of incoming messages with $\widehat{\nu}_{a\to i} = \mathtt{0}$ converges therefore to a Poisson random variable with mean $K\alpha\widehat{Q}$. The expected number of variable nodes in the core will be $\mathbb{E}|K_2(G)| = NV(\alpha, K) + o(N)$, where $V(\alpha, K)$ is the probability that such a Poisson random variable is larger or equal to 2, that is

$$V(\alpha, K) = 1 - e^{-K\alpha\widehat{Q}} - K\alpha\widehat{Q}\, e^{-K\alpha\widehat{Q}}\,. \tag{18.13}$$

In Fig. 18.6 we plot $V(\alpha)$ as a function of α. For $\alpha < \alpha_{\rm d}(K)$ the peeling algorithm erases the whole graph, there is no core. The size of the core jumps to some finite value at $\alpha_{\rm d}(K)$ and when $\alpha \to \infty$ the core is the full graph.

Is $K_2(G)$ a random factor graph or does it have any particular structure? By construction it cannot contain variable nodes of degree zero or one. Its expected degree profile (expected fraction of nodes of any given degree) will be asymptotically $\widehat{\Lambda} \equiv \{\widehat{\Lambda}_l\}$, where $\widehat{\Lambda}_l$ is the probability that a Poisson random variable of parameter $K\alpha\widehat{Q}$, conditioned to be at least 2, is equal to l. Explicitly $\widehat{\Lambda}_0 = \widehat{\Lambda}_1 = 0$, and

$$\widehat{\Lambda}_l = \frac{1}{e^{K\alpha\widehat{Q}} - 1 - K\alpha\widehat{Q}}\ \frac{1}{l!}\ (K\alpha\widehat{Q})^l \qquad \text{for } l \geq 2. \tag{18.14}$$

Somewhat surprisingly $K_2(G)$ does not have any more structure than the one determined by its degree profile. This fact is stated more formally in the following theorem.

Theorem 18.1 *Consider a factor graph G from the $\mathbb{G}_N(K, N\alpha)$ ensemble with $K \geq 3$. Then*

(*i*) *$K_2(G) = \emptyset$ with high probability for $\alpha < \alpha_{\rm d}(K)$.*

(*ii*) *For $\alpha > \alpha_{\rm d}(K)$, $|K_2(G)| = NV(\alpha, K) + o(N)$ with high probability.*

(*iii*) *The fraction of vertices of degree l in $K_2(G)$ is between $\widehat{\Lambda}_l - \varepsilon$ and $\widehat{\Lambda}_l + \varepsilon$ with probability greater than $1 - e^{-\Theta(N)}$.*

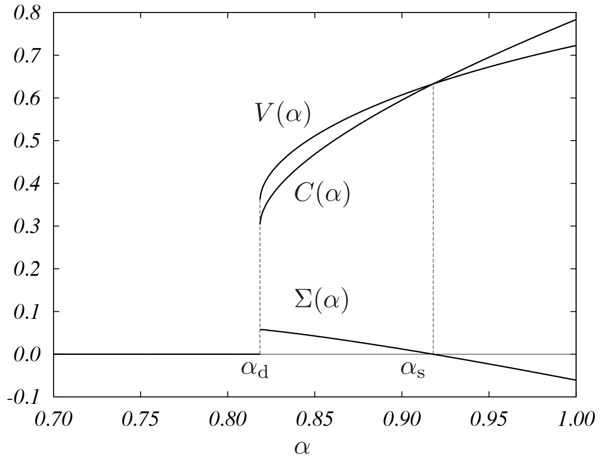

FIG. 18.6. The core of random 3-XORSAT formulae contains $NV(\alpha)$ variables, and $NC(\alpha)$ equations. These numbers are plotted versus the number of equations per variable of the original formula α. The number of solutions to the XORSAT linear system is $\Sigma(\alpha) = V(\alpha) - C(\alpha)$. The core appears for $\alpha \geq \alpha_{\rm d}$, and the system becomes UNSAT for $\alpha > \alpha_{\rm s}$, where $\alpha_{\rm s}$ is determined by $\Sigma(\alpha_{\rm s}) = 0$.

(iv) Conditionally on the number of variable nodes $n = |K_2(G)|$, the degree profile being $\widehat{\Lambda}$, $K_2(G)$ is distributed according to the $\mathbb{D}_n(\widehat{\Lambda}, x^K)$ ensemble.

We will not provide the proof of this theorem. The main ideas have already been presented in the previous pages, except for one important mathematical point: how to exchange the limits $N \to \infty$ and $t \to \infty$. The basic idea is to run BP for a large but fixed number of steps t. At this point the resulting graph is 'almost' a 2-core, and one can show that a sequential peeling procedure stops in less than $N\varepsilon$ steps.

In Fig. 18.7 we compare the statement in this Theorem with numerical simulations. The probability that G contains a 2 core $\mathrm{P}_{\rm core}(\alpha)$ increases from 0 to 1 as α ranges from 0 to ∞, with a threshold becoming sharper and sharper as the size N increases. The threshold behavior can be accurately described using finite size scaling. Setting $\alpha = \alpha_{\rm d}(K) + \beta(K)\, z\, N^{-1/2} + \delta(K)\, N^{-2/3}$ (with properly chosen $\beta(K)$ and $\delta(K)$) one can show that $\mathrm{P}_{\rm core}(\alpha)$ approaches a K-independent non-trivial limit that depends smoothly on z.

18.4.2 *The threshold*

Knowing that the core is a random graph with degree distribution $\widehat{\Lambda}_l$, we can compute the expected number of equations in the core. This is given by the number of vertices times their average degree, divided by K, which yields $NC(\alpha, K) + o(N)$ where

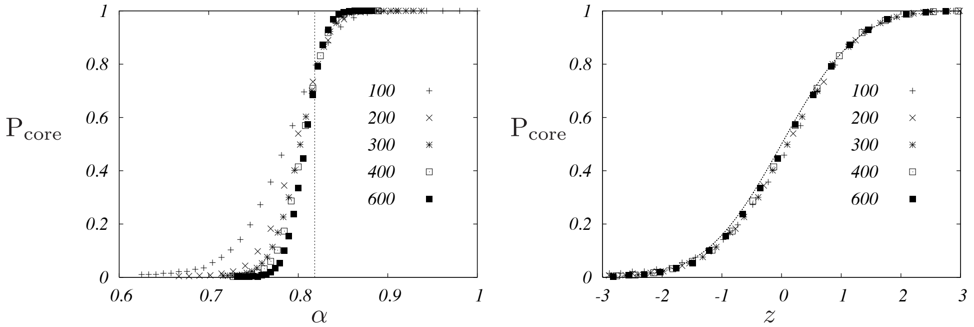

FIG. 18.7. Probability that a random graph from the $\mathbb{G}_N(K,\alpha)$ ensemble with $K = 3$ (equivalently, the factor graph of a random 3-XORSAT formula) contains a 2 core. On the left, the outcome of numerical simulations is compared with the asymptotic threshold $\alpha_{\rm d}(K)$. On the right, scaling plot (see text).

$$C(\alpha, K) = \alpha\widehat{Q}(1 - e^{-K\alpha\widehat{Q}})\,. \tag{18.15}$$

In Fig. 18.6 we plot $C(\alpha, K)$ versus α. If $\alpha < \alpha_{\rm d}(K)$ there is no core. For $\alpha \in]\alpha_{\rm d}, \alpha_{\rm s}[$ the number of equations in the core is smaller than the number of variables $V(\alpha, K)$. Above α_c there are more equations than variables.

A linear system has a solution if and only if the associated core problem has a solution. In a large random XORSAT instance, the core system involves approximately $NC(\alpha, K)$ equations between $NV(\alpha, K)$ variables. We shall show that these equations are, with high probability, linearly independent as long as $C(\alpha, K) < V(\alpha, K)$, which implies the following result

Theorem 18.2. (XORSAT satisfiability threshold.) *For $K \geq 3$, let*

$$\Sigma(K, \alpha) = V(K, \alpha) - C(K, \alpha) = Q - \alpha\widehat{Q}(1 + (K - 1)(1 - Q))\ , \tag{18.16}$$

where $Q, \widehat{Q}$ are the largest solution of Eq. (18.12). Let $\alpha_{\rm s}(K) = \inf\{\alpha\ :\ \Sigma(K, \alpha) < 0\}$. Consider a random K-XORSAT linear system with N variables and $N\alpha$ equations. The following results hold with a probability going to 1 in the large N limit:

(i) The system has a solution when $\alpha < \alpha_{\rm s}(K)$.

(ii) It has no solution when $\alpha > \alpha_{\rm s}(K)$.

(iii) For $\alpha < \alpha_{\rm s}(K)$ the number of solutions is $2^{N(1-\alpha)+o(N)}$, and the number of clusters is $2^{N\Sigma(K,\alpha)+o(N)}$.

Notice that the the last expression in Eq. (18.16) is obtained from Eqs. (18.13) and (18.15) using the fixed point condition (18.12).

The prediction of this theorem is compared with numerical simulations in Fig. 18.8, while Fig. 18.9 summarizes the results on the thresholds for XORSAT.

Proof: We shall convey the basic ideas of the proof and refer to the literature for technical details.

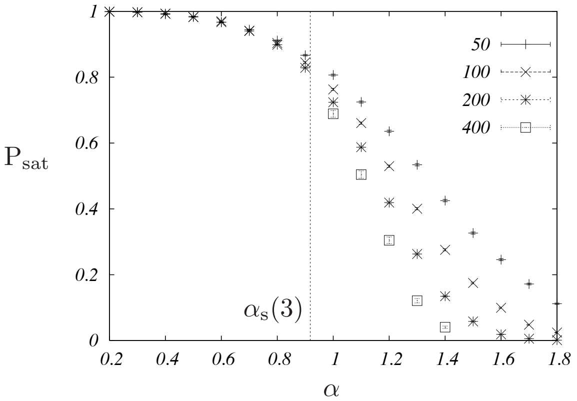

FIG. 18.8. Probability that a random 3-XORSAT formula with N variables and $N\alpha$ equations is SAT, estimated numerically by generating $10^3 \div 10^4$ random instances.

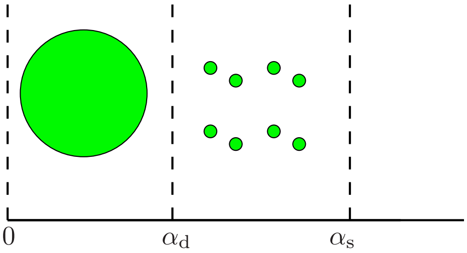

K	3	4	5
$\alpha_{\rm d}$	0.81847	0.77228	0.70178
$\alpha_{\rm s}$	0.91794	0.97677	0.99244

FIG. 18.9. Left: A pictorial view of the phase transitions in random XORSAT systems. The satisfiability threshold is $\alpha_{\rm s}$. In the 'Easy-SAT' phase $\alpha < \alpha_{\rm d}$ there is a single cluster of solutions. In the 'Hard-SAT' phase $\alpha_{\rm d} < \alpha < \alpha_{\rm s}$ the solutions of the linear system are grouped in well separated clusters. Right: The thresholds $\alpha_{\rm d}$, $\alpha_{\rm s}$ for various values of K. At large K one has: $\alpha_{\rm d}(K) \simeq \log K/K$ and $\alpha_{\rm s}(K) = 1 - e^{-K} + O(e^{-2K})$.

Let us start by proving (ii), namely that for $\alpha > \alpha_{\rm s}(K)$ random XORSAT instances are with high probability UNSAT. This follows from a linear algebra argument. Let $\mathbb{H}_*$ denote the $\mathtt{0}-\mathtt{1}$ matrix associated with the core, i.e. the matrix including those rows/columns such that the associated function/variable nodes belong to $K_2(G)$. Notice that if a given row is included in $\mathbb{H}_*$ then all the columns corresponding to non-zero entries of that row are also in $\mathbb{H}_*$. As a consequence, a necessary condition for the rows of $\mathbb{H}$ to be independent is that the rows of $\mathbb{H}_*$ are independent. This is in turn impossible if the number of columns in $\mathbb{H}_*$ is smaller than its number of rows.

Quantitatively, one can show that $M - \text{rank}(\mathbb{H}) \geq \text{rows}(\mathbb{H}_*) - \text{cols}(\mathbb{H}_*)$ (with the obvious meanings of $\text{rows}(\,\cdot\,)$ and $\text{cols}(\,\cdot\,)$). In large random XORSAT systems, Theorem 18.1 implies that $\text{rows}(\mathbb{H}_*) - \text{cols}(\mathbb{H}_*) = -N\Sigma(K,\alpha) + o(N)$ with

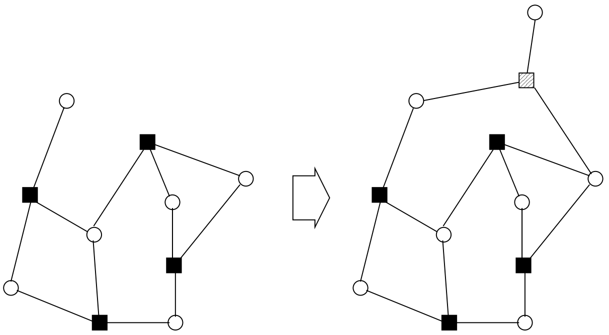

FIG. 18.10. Adding a function nodes involving a variable node of degree one. The corresponding linear equation is independent from the other ones.

high probability. According to our discussion in Sec. 18.1.1, among the 2^M possible choices of the right-hand side vector $\underline{b}$, only $2^{\mathrm{rank}(\mathbb{H})}$ are in the image of $\mathbb{H}$ and thus lead to a solvable system. In other words, conditional on $\mathbb{H}$, the probability that random XORSAT is solvable is $2^{\mathrm{rank}(\mathbb{H})-M}$. By the above argument this is, with high probability, smaller than $2^{N\Sigma(K,\alpha)+o(N)}$. Since $\Sigma(K,\alpha)<0$ for $\alpha>\alpha_{\mathrm{s}}(K)$, it follows that the system is UNSAT with high probability.

In order to show that a random system is satisfiable with high probability when $\alpha<\alpha_{\mathrm{s}}(K)$, one has to prove the following facts: (i) if the core matrix $\mathbb{H}_*$ has maximum rank, then $\mathbb{H}$ has maximum rank as well; (ii) if $\alpha<\alpha_{\mathrm{s}}(K)$, then $\mathbb{H}_*$ has maximum rank with high probability. As a byproduct, the number of solutions is $2^{N-\mathrm{rank}(\mathbb{H})}=2^{N-M}$.

(i) The first step follows from the observation that G can be constructed from $K_2(G)$ through an inverse peeling procedure. At each step one adds a function node which involves at least a degree one variable (see Fig. 18.10). Obviously this newly added equation is linearly independent of the previous ones, and therefore $\mathrm{rank}(\mathbb{H})=\mathrm{rank}(\mathbb{H}_*)+M-\mathrm{rows}(\mathbb{H}_*)$.

(ii) Let $n=\mathrm{cols}(\mathbb{H}_*)$ be the number of variable nodes and $m=\mathrm{rows}(\mathbb{H}_*)$ the number of function nodes in the core $K_2(G)$. Let us consider the homogeneous system on the core, $\mathbb{H}_*\underline{x}=\underline{0}$, and denote by Z_* the number of solutions to this system. We will show that with high probability this number is equal to 2^{n-m}. This means that the dimension of the kernel of $\mathbb{H}_*$ is $n-m$ and therefore $\mathbb{H}_*$ has full rank.

We know from linear algebra that $Z_*\geq 2^{n-m}$. To prove the reverse inequality we use a first moment method. According to Theorem 18.1, the core is a uniformly random factor graph with $n=NV(K,\alpha)+o(N)$ variables and degree profile $\Lambda=\widehat{\Lambda}+o(1)$. Denote by $\mathbb{E}$ the expectation value with respect to this ensemble. We shall use below a first moment analysis to show that, when $\alpha<\alpha_{\mathrm{c}}(K)$:

$$\mathbb{E}\{Z_*\}=2^{n-m}[1+o_N(1)]\,. \tag{18.17}$$

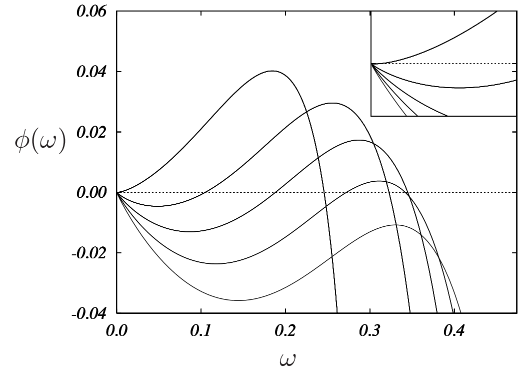

FIG. 18.11. The exponential rate $\phi(\omega)$ of the weight enumerator of the core of a random 3-XORSAT formula. From top to bottom $\alpha = \alpha_{\rm d}(3) \approx 0.818469$, 0.85, 0.88, 0.91, and 0.94 (recall that $\alpha_{\rm s}(3) \approx 0.917935$). Inset: blow up of the small ω region.

Then Markov inequality $\mathbb{P}\{Z_* > 2^{n-m}\} \le 2^{-n+m}\mathbb{E}\{Z_*\}$ implies the bound.

The surprise is that Eq. (18.17) holds, and thus a simple first moment estimate allows to establish that $\mathbb{H}_*$ has full rank. We saw in Exercise 18.6 that the same approach, when applied directly to the original linear system, fails above some $\alpha_*(K)$ which is strictly smaller than $\alpha_{\rm s}(K)$. Reducing the original graph to its two-core has drastically reduced the fluctuations of the number of solutions, thus allowing for a successful application of the first moment method.

We now turn to the proof of Eq. (18.17), and we shall limit ourselves to the computation of $\mathbb{E}\{Z_*\}$ to the leading exponential order, when the core size and degree profiles take their typical values $n = NV(K,\alpha)$, $\Lambda = \widehat{\Lambda}$ and $P(x) = x^K$. This problem is equivalent to computing the expected number of codewords in the LDPC code defined by the core system, which we already did in Sec. 11.2. The result takes the typical form

$$\mathbb{E}\{Z_*\} \doteq \exp\left\{N \sup_{\omega\in[0,V(K,\alpha)]} \phi(\omega)\right\}. \tag{18.18}$$

Here $\phi(\omega)$ is the exponential rate for the number of solutions with weight $N\omega$. Adapting Eq. (11.18) to the present case, we obtain the parametric expression:

$$\begin{aligned}\phi(\omega) = &-\omega\log x - \eta(1-e^{-\eta})\log(1+yz) + \qquad (18.19)\\ &+\sum_{l\ge 2} e^{-\eta}\frac{\eta^l}{l!}\log(1+xy^l) + \frac{\eta}{K}(1-e^{-\eta})\log q_K(z)\,,\\ \omega = &\sum_{l\ge 2} e^{-\eta}\,\frac{\eta^l}{l!}\,\frac{xy^l}{1+xy^l}\,. \qquad (18.20)\end{aligned}$$

where $\eta = K\alpha\widehat{Q}_*$, $q_K(z) = [(1+z)^K + (1-z)^K]/2$ and $y = y(x)$, $z = z(x)$ are the solution of

$$z = \frac{\sum_{l\geq 1}[\eta^l/l!]\;[xy^{l-1}/(1+xy^l)]}{\sum_{l\geq 1}[\eta^l/l!]\;[1/(1+xy^l)]}\,, \qquad y = \frac{(1+z)^{K-1}-(1-z)^{K-1}}{(1+z)^{K-1}+(1-z)^{K-1}}\,. \quad (18.21)$$

With a little work one sees that $\omega_* = V(K,\alpha)/2$ is a local maximum of $\phi(\omega)$, with $\phi(\omega_*) = \Sigma(K,\alpha)\log 2$. As long as ω_* is a global maximum, $\mathbb{E}\{Z_*|n,\Lambda\} \doteq \exp\{N\phi(\omega_*)\} \doteq 2^{n-m}$. It turns out, cf. Fig. 18.11, that the only other local maximum is at $\omega = 0$ corresponding to $\phi(0) = 0$. Therefore $\mathbb{E}\{Z_*|n,\Lambda\} \doteq 2^{n-m}$ as long as $\phi(\omega_*) = \Sigma(K,\alpha) > 0$, i.e. for any $\alpha < \alpha_{\rm s}(K)$

Notice that the actual proof of Eq. (18.17) is more complicate because it requires estimating the sub-exponential factors. Nevertheless it can be carried out successfully. □

18.5 The Hard-SAT phase: clusters of solutions

In random XORSAT, the whole regime $\alpha < \alpha_{\rm s}(K)$ is SAT. This means that, with high probability there exist solutions to the random linear system, and the number of solutions is in fact $Z \doteq e^{N(1-\alpha)}$. Notice that the number of solutions does not present any precursor of the SAT-UNSAT transition at $\alpha_{\rm s}(K)$ (recall that $\alpha_{\rm s}(K) < 1$), nor does it carry any trace of the sudden appearence of a non-empty two core at $\alpha_{\rm d}(K)$.

On the other hand the threshold $\alpha_{\rm d}(K)$ separates two phases, that we will call **'Easy-SAT'** (for $\alpha < \alpha_{\rm d}(K)$) and **'Hard-SAT' phase** (for $\alpha \in]\alpha_{\rm d}(K), \alpha_{\rm s}(K)[$). These two phases differ in the structure of the solution space, as well as in the behavior of some simple algorithms.

In the Easy-SAT phase there is no core, solutions can be found in (expected) linear time using the peeling algorithm and they form a unique cluster. In the Hard-SAT the factor graph has a large 2-core, and no algorithm is known that finds a solution in linear time. Solutions are partitioned in $2^{N\Sigma(K,\alpha)+o(N)}$ clusters. Until now the name 'cluster' has been pretty arbitrary, and only denoted a subset of solutions that coincide in the core. The next result shows that distinct clusters are 'far apart' in Hamming space.

Proposition 18.3 *In the Hard-SAT phase there exists $\delta(K,\alpha) > 0$ such that, with high probability, any two solutions in distinct clusters have Hamming distance larger than $N\delta(K,\alpha)$.*

Proof: The proof follows from the computation of the weight enumerator exponent $\phi(\omega)$, cf. Eq. (18.20) and Fig. 18.11. One can see that for any $\alpha > \alpha_{\rm d}(K)$, $\phi'(0) < 0$, and, as a consequence there exists $\delta(K,\alpha) > 0$ such that $\phi(\omega) < 0$ for $0 < \omega < \delta(K,\alpha)$. This implies that if $\underline{x}_*$, $\underline{x}'_*$ are two distinct solution of the core linear system, then either $d(\underline{x}_*, \underline{x}'_*) = o(N)$ or $d(\underline{x}, \underline{x}') > N\delta(K,\alpha)$. It turns out that the first case can be excluded along the lines of the minimal distance calculation of Sec. 11.2. Therefore, if $\underline{x}$, $\underline{x}'$ are two solutions belonging to distinct clusters $d(\underline{x}, \underline{x}') \geq d(\pi_*(\underline{x}), \pi_*(\underline{x}')) \geq N\delta(K,\alpha)$. □

This result suggests to regard clusters as 'lumps' of solutions well separated from each other. One aspect which is conjectured, but not proved, concerns the fact that clusters form 'well connected components.' By this we mean that any two solutions in the a cluster can be joined by a sequence of other solutions, whereby two successive solutions in the sequence differ in at most s_N variables, with $s_N = o(N)$ (a reasonable expectation is $s_N = \Theta(\log N)$).

18.6 An alternative approach: the cavity method

The analysis of random XORSAT in the previous sections relied heavily on the linear structure of the problem, as well as on the very simple instance distribution. This section describes an alternative approach that is potentially generalizable to more complex situations. The price to pay is that this second derivation relies on some assumptions on the structure of the solution space. The observation that our final results coincide with the ones obtained in the previous section gives some credibility to these assumptions.

The starting point is the remark that BP correctly computes the marginals of $\mu(\,\cdot\,)$ (the uniform measure over the solution space) for $\alpha < \alpha_{\rm d}(K)$, i.e. as long as the set of solutions forms a single cluster. We want to extend its domain of validity to $\alpha > \alpha_{\rm d}(K)$. If we index by $n \in \{1, \dots, \mathcal{N}\}$ the clusters, the uniform measure $\mu(\,\cdot\,)$ can be decomposed into the convex combination of uniform measures over each single cluster:

$$\mu(\,\cdot\,) = \sum_{n=1}^{\mathcal{N}} w_n\, \mu^n(\,\cdot\,)\,. \tag{18.22}$$

Notice that in the present case $w_n = 1/\mathcal{N}$ is independent of n and the measures $\mu^n(\,\cdot\,)$ are obtained from each other via a translation, but this will not be true in more general situations.

Consider an inhomogeneous XORSAT linear system and denote by $\underline{x}^{(*)}$ one of its solutions in cluster n. The distribution μ^n has single variable marginals $\mu^n(x_i) = \mathbb{I}(x_i = x_i^{(*)})$ if node i belongs to the backbone, and $\mu^n(x_i = \mathtt{0}) = \mu^n(x_i = \mathtt{1}) = 1/2$ on the other nodes.

In fact we can associate to each solution $\underline{x}^{(*)}$ a fixed point of the BP equation. We already described this in Section 18.2.1, cf. Eq. (18.9). On this fixed point messages take one of the following three values: $\nu^{(*)}_{i\to a}(x_i) = \mathbb{I}(x_i = \mathtt{0})$ (that we will denote as $\nu^{(*)}_{i\to a} = \mathtt{0}$), $\nu^{(*)}_{i\to a}(x_i) = \mathbb{I}(x_i = \mathtt{1})$ (denoted $\nu^{(*)}_{i\to a} = \mathtt{1}$), $\nu^{(*)}_{i\to a}(x_i = \mathtt{0}) = \nu^{(*)}_{i\to a}(x_i = \mathtt{1}) = 1/2$ (denoted $\nu^{(*)}_{i\to a} = *$). Analogous notations hold for function-to-variable node messages. The solution can be written most easily in terms of the latter

$$\widehat{\nu}^{(*)}_{a\to i} = \begin{cases} \mathtt{1} \text{ if } x_i^{(*)} = \mathtt{1} \text{ and } i, a \in B(G), \\ \mathtt{0} \text{ if } x_i^{(*)} = \mathtt{0} \text{ and } i, a \in B(G), \\ * \text{ otherwise.} \end{cases} \tag{18.23}$$

Notice that these messages only depend on the value of $x_i^{(*)}$ on the backbone of G, hence they depend on $\underline{x}^{(*)}$ only through the cluster it belongs to. Reciprocally, for any two distinct clusters, the above definition gives two distinct fixed points. Because of this remark we shall denote these fixed points as $\{\nu_{i\to a}^{(n)}, \widehat{\nu}_{a\to i}^{(n)}\}$, where n is a cluster index.

Let us recall the BP fixed point condition:

$$\nu_{i\to a} = \begin{cases} * & \text{if } \widehat{\nu}_{b\to i} = * \text{ for all } b \in \partial i\backslash a, \\ \text{any `non $*$' } \widehat{\nu}_{b\to i} & \text{otherwise.} \end{cases} \tag{18.24}$$

$$\widehat{\nu}_{a\to i} = \begin{cases} * & \text{if } \exists j \in \partial a\backslash i \text{ s.t. } \widehat{\nu}_{j\to a} = *, \\ b_a \oplus \nu_{j_1\to a} \oplus \cdots \oplus \nu_{j_l\to a} & \text{otherwise.} \end{cases} \tag{18.25}$$

Below we shall denote symbolically these equations as

$$\nu_{i\to a} = \mathsf{f}\{\widehat{\nu}_{b\to i}\}\,, \qquad \widehat{\nu}_{a\to i} = \hat{\mathsf{f}}\{\nu_{j\to a}\}\,. \tag{18.26}$$

Let us summarize our findings.

Proposition 18.4 *To each cluster n we can associate a distinct fixed point of the BP equations (18.25) $\{\nu_{i\to a}^{(n)}, \widehat{\nu}_{a\to i}^{(n)}\}$, such that $\widehat{\nu}_{a\to i}^{(n)} \in \{\mathtt{0}, \mathtt{1}\}$ if i, a are in the backbone and $\widehat{\nu}_{a\to i}^{(n)} = *$ otherwise.*

Note that the converse of this proposition is false: there may exist solutions to the BP equations which are not of the previous type. One of them is the all $*$ solution. Nontrivial solutions exist as well as shown in Fig. 18.12.

An introduction to the 1RSB cavity method in the general case will be presented in Ch. 19. Here we give a short informal preview in the special case of the XORSAT: the reader will find a more formal presentation in the next chapter. The first two assumptions of the 1RSB cavity method can be summarized as follows (all statements are understood to hold with high probability).

Assumption 1 *In a large random XORSAT instance, for each cluster 'n' of solutions, the BP solution $\nu^{(n)}, \widehat{\nu}^{(n)}$ provides an accurate 'local' description of the measure $\mu^n(\,\cdot\,)$.*

This means that for instance the one point marginals are given by $\mu^n(x_j) \cong \prod_{a\in\partial j} \widehat{\nu}_{a\to j}^{(n)}(x_j) + o(1)$, but also that local marginals inside any finite cavity are well approximated by formula (14.18).

Assumption 2 *For a large random XORSAT instance in the Hard-SAT phase, the number of clusters $e^{N\Sigma}$ is exponential in the number of variables. Further, the number of solutions of the BP equations (18.25) is, to the leading exponential order, the same as the number of clusters. In particular it is the same as the number of solutions constructed in Proposition 18.4.*

A priori one might have hoped to identify the set of messages $\{\nu_{i\to a}^{(n)}\}$ for each cluster. The cavity method gives up this ambitious objective and aims to

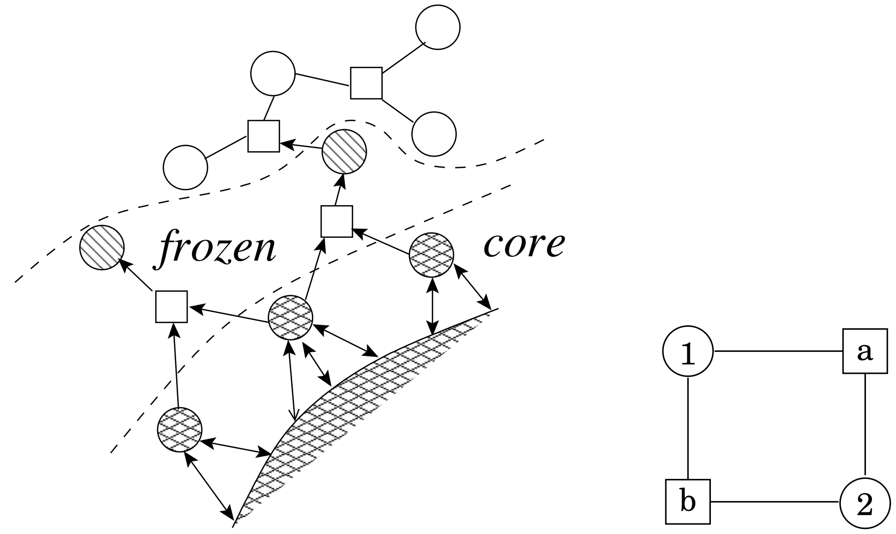

FIG. 18.12. Left: A set of BP messages associated with one cluster (cluster number n) of solutions. An arrow along an edge means that the corresponding message (either $\nu_{i\to a}^{(n)}$ or $\widehat{\nu}_{a\to i}^{(n)}$) takes value in $\{\mathtt{0},\mathtt{1}\}$. The other messages are equal to $*$. Right: A small XORSAT instance. The core is the whole graph. In the homogeneous problem there are two solutions, which form two clusters: $x_1 = x_2 = 0$ and $x_1 = x_2 = 1$. Beside the two corresponding BP fixed points described in Proposition 18.4, and the all-$*$ fixed point, there exist other fixed points such as $\widehat{\nu}_{a\to 1} = \nu_{1\to b} = \widehat{\nu}_{b\to 2} = \nu_{2\to a} = \mathtt{0}$, $\widehat{\nu}_{a\to 2} = \nu_{2\to b} = \widehat{\nu}_{b\to 1} = \nu_{1\to a} = *$.

compute the distribution of $\nu_{i\to a}^{(n)}$ for any fixed edge $i \to a$, when n is a cluster index drawn with distribution $\{w_n\}$. We thus want to compute the quantities:

$$Q_{i\to a}(\nu) = \mathbb{P}\left\{\nu_{i\to a}^{(n)} = \nu\right\}\,, \qquad \widehat{Q}_{a\to i}(\widehat{\nu}) = \mathbb{P}\left\{\widehat{\nu}_{a\to i}^{(n)} = \widehat{\nu}\right\}\,. \tag{18.27}$$

for $\nu, \widehat{\nu} \in \{\mathtt{0},\mathtt{1},*\}$. Computing these probabilities rigorously is still a challenging task. In order to proceed, we make some assumption on the joint distribution of the messages $\nu_{i\to a}^{(n)}$ when n is a random cluster index (chosen from the probability w_n).

The simplest idea would be to assume that messages on 'distant' edges are independent. For instance let us consider the set of messages entering a given variable node i. Their only correlations are induced through BP equations along the loops to which i belongs. Since in random K-XORSAT formulae such loops have, with high probability, length of order $\log N$, one might think that messages incoming a given node are asymptotically independent. Unfortunately this assumption is false. The reason is easily understood if we assume that $\widehat{Q}_{a\to i}(\mathtt{0}), \widehat{Q}_{a\to i}(\mathtt{1}) > 0$ for at least two of the function nodes a adjacent to a given variable node i. This would imply that, with positive probability a randomy sampled cluster has $\nu_{a\to i}^{(n)} = \mathtt{0}$, and $\nu_{b\to i}^{(n)} = \mathtt{1}$. But there does not exist any such cluster, because in such a situation there is no consistent prescription for the marginal distribution of x_i under $\mu^n(\,\cdot\,)$.

Our assumption will be that the next simplest thing happens: messages are independent conditional to the fact that they do not contradict each other.

Assumption 3 *Consider the Hard-SAT phase of a random XORSAT problem. Denote by $i \in G$ a uniformly random node, by n a random cluster index with distribution $\{w_n\}$, and let ℓ be an integer ≥ 1. Then the messages $\{\nu_{j\to b}^{(n)}\}$, where (j, b) are all the edges at distance ℓ from i and directed towards i, are asymptotically independent under the condition of being* ***compatible****.*

Here 'compatible' means the following. Consider the linear system $\mathbb{H}_{i,\ell}\underline{x}_{i,\ell} = \underline{\mathtt{0}}$ for the neighborhood of radius ℓ around node i. If this admits a solution under the boundary condition $x_j = \nu_{j\to b}$ for all the boundary edges (j, b) on which $\{\nu_{j\to b}\} \in \{\mathtt{0}, \mathtt{1}\}$, then the messages $\{\nu_{j\to b}\}$ are said to be compatible.

Given the messages $\nu_{j\to b}$ at the boundary of a radius-ℓ neighborhood, the BP equations (18.24) and (18.25) allow to determine the messages inside this neighborhood. Consider in particular two nested neighborhoods at distance ℓ and $\ell + 1$ from i. The inwards messages on the boundary of the largest neighborhood completely determines the ones on the boundary of the smallest one. A little thought shows that, if the messages on the outer boundary are distributed according to Assumption 3, then the distribution of the resulting messages on the inner boundary also satisfies the same assumption. Further, the distributions are consistent if and only if the following 'survey propagation' equations are satisfied by the one-message marginals:

$$Q_{i\to a}(\nu) \cong \sum_{\{\widehat{\nu}_b\}} \prod_{b\in\partial i\setminus a} \widehat{Q}_{b\to i}(\widehat{\nu}_b)\;\; \mathbb{I}(\nu = \mathsf{f}\{\widehat{\nu}_b\})\; \mathbb{I}(\{\widehat{\nu}_b\}_{b\in\partial i\setminus a} \in \mathsf{COMP})\;, \tag{18.28}$$

$$\widehat{Q}_{a\to i}(\widehat{\nu}) = \sum_{\{\nu_j\}} \prod_{j\in\partial a\setminus i} Q_{j\to a}(\nu_j)\;\; \mathbb{I}(\widehat{\nu} = \hat{\mathsf{f}}\{\nu_j\})\,. \tag{18.29}$$

Here and $\{\widehat{\nu}_b\} \in \mathsf{COMP}$ only if the messages are compatible (i.e. they do not contain both a $\mathtt{0}$ and a $\mathtt{1}$). Since Assumptions 1, 2, 3 above hold only with high probability and asymptotically in the system size, the equalities in (18.28), (18.29) must also be interpreted as approximate. The equations should be satisfied within any given accuracy ε, with high probability as $N \to \infty$.

Exercise 18.9 Show that Eqs. (18.28), (18.29) can be written explicitly as

$$Q_{i\to a}(\mathtt{0}) \cong \prod_{b\in\partial i\setminus a} (\widehat{Q}_{b\to i}(\mathtt{0}) + \widehat{Q}_{b\to i}(*)) - \prod_{b\in\partial i\setminus a} \widehat{Q}_{b\to i}(*)\,, \tag{18.30}$$

$$Q_{i\to a}(\mathtt{1}) \cong \prod_{b\in\partial i\setminus a} (\widehat{Q}_{b\to i}(\mathtt{1}) + \widehat{Q}_{b\to i}(*)) - \prod_{b\in\partial i\setminus a} \widehat{Q}_{b\to i}(*) \tag{18.31}$$

$$Q_{i\to a}(*) \cong \prod_{b\in\partial i\setminus a} \widehat{Q}_{b\to i}(*)\,, \tag{18.32}$$

where the $\cong$ symbol hides a global normalization constant, and

$$\widehat{Q}_{a\to i}(\mathtt{0}) = \frac{1}{2}\left\{\prod_{j\in\partial a\setminus i}(Q_{j\to a}(\mathtt{0}) + Q_{j\to a}(\mathtt{1})) + \prod_{j\in\partial a\setminus i}(Q_{j\to a}(\mathtt{0}) - Q_{j\to a}(\mathtt{1}))\right\}\,, \tag{18.33}$$

$$\widehat{Q}_{a\to i}(\mathtt{1}) = \frac{1}{2}\left\{\prod_{j\in\partial a\setminus i}(Q_{j\to a}(\mathtt{0}) + Q_{j\to a}(\mathtt{1})) - \prod_{j\in\partial a\setminus i}(Q_{j\to a}(\mathtt{0}) - Q_{j\to a}(\mathtt{1}))\right\}\,, \tag{18.34}$$

$$\widehat{Q}_{a\to i}(*) = 1 - \prod_{j\in\partial a\setminus i}(Q_{j\to a}(\mathtt{0}) + Q_{j\to a}(\mathtt{1}))\,. \tag{18.35}$$

The final step of the 1RSB cavity method consists in looking for a solution of Eqs. (18.28), (18.29). There are no rigorous results on the existence or number of such solutions. Further, since these equations are only approximate, approximate solutions should be considered as well. In the present case a very simple (and somewhat degenerate) solution can be found that yields the correct predictions for all the quantities of interest. In this solution, the message distributions take one of two possible forms: on some edges one has $Q_{i\to a}(\mathtt{0}) = Q_{i\to a}(\mathtt{1}) = 1/2$ (with an abuse of notation we shall write $Q_{i\to a} = \mathtt{0}$ in this case), on some other edges $Q_{i\to a}(*) = 1$ (we will then write $Q_{i\to a} = *$). Analogous forms hold for $\widehat{Q}_{a\to i}$. A little algebra shows that this is a solution if and only if the η's satisfy

$$Q_{i\to a} = \begin{cases} * & \text{if } \widehat{Q}_{b\to i} = * \text{ for all } b\in\partial i\setminus a, \\ \mathtt{0} & \text{otherwise.}\end{cases} \tag{18.36}$$

$$\widehat{Q}_{a\to i} = \begin{cases} * & \text{if } \exists j\in\partial a\setminus i \text{ s.t. } \widehat{Q}_{j\to a} = *, \\ \mathtt{0} & \text{otherwise.}\end{cases} \tag{18.37}$$

These equations are identical to the original BP equations for the homogeneous problem (this feature is very specific to XORSAT and will not generalize to more advanced applications of the method). However the interpretation is now completely different. On the edges where $Q_{i\to a} = \mathtt{0}$ the corresponding message $\nu^{(n)}_{i\to a}$ depend on the cluster n and $\nu^{(n)}_{i\to a} = \mathtt{0}$ (respectively $= \mathtt{1}$) in half of the clusters. These edges are those inside the core, or in the backbone but directed 'outward' with respect to the core, as shown in Fig.18.12. On the other edges, the message does not depend upon the cluster and $\nu^{(n)}_{i\to a} = *$ for all n's.

A concrete interpretation of these results is obtained if we consider the one bit marginals $\mu^n(x_i)$ under the single cluster measure. According to Assumption 1 above, we have $\mu^n(x_i = \mathtt{0}) = \mu^n(x_i = \mathtt{1}) = 1/2$ if $\widehat{\nu}^{(n)}_{a\to i} = *$ for all $a \in \partial i$. If on the other hand $\widehat{\nu}^{(n)}_{a\to i} = \mathtt{0}$ (respectively $= \mathtt{1}$) for at least one $a \in \partial i$, then $\mu^n(x_i = \mathtt{0}) = 1$ (respectively $\mu^n(x_i = \mathtt{0}) = 0$). We thus recover the full solution discussed in the previous sections: inside a given cluster n, the variables in the backbone are completely frozen, either to $\mathtt{0}$ or to $\mathtt{1}$. The other variables have equal probability to be $\mathtt{0}$ or $\mathtt{1}$ under the measure μ^n.

The cavity approach allows to compute the complexity $\Sigma(K, \alpha)$ as well as many other properties of the measure $\mu(\,\cdot\,)$. We will see this in the next chapter.

Notes

Random XORSAT formulae were first studied as a simple example of random satisfiability in (Creignou and Daudé, 1999). This work considered the case of 'dense formulae' where each clause includes $O(N)$ variables. In this case the SAT-UNSAT threshold is at $\alpha = 1$. In coding theory this model had been characterized since the work of Elias in the fifties (Elias, 1955), cf. Ch. 6.

The case of sparse formulae was addressed using moment bounds in (Creignou, Daudé and Dubois, 2003). The replica method was used in (Ricci-Tersenghi, Weigt and Zecchina, 2001; Franz, Leone, Ricci-Tersenghi and Zecchina, 2001*a*; Franz, Mézard, Ricci-Tersenghi, Weigt and Zecchina, 2001*b*) to derive the clustering picture, determine the SAT-UNSAT threshold, and study the glassy properties of the clustered phase.

The fact that, after reducing the linear system to its core, the first moment method provides a sharp characterization of the SAT-UNSAT threshold was discovered independently by two groups: (Cocco, Dubois, Mandler and Monasson, 2003) and (Mézard, Ricci-Tersenghi and Zecchina, 2003). The latter also discusses the application of the cavity method to the problem. The full second moment calculation that completes the proof can be found for the case $K = 3$ in (Dubois and Mandler, 2002).

The papers (Montanari and Semerjian, 2005; Montanari and Semerjian, 2006*a*; Mora and Mézard, 2006) were devoted to finer geometrical properties of the set of solutions of random K-XORSAT formulae. Despite these efforts, it remains to be proved that clusters of solutions are indeed 'well connected.'

Since the locations of various transitions are known rigorously, a natural question is to study the critical window. Finite size scaling of the SAT-UNSAT transition was investigated numerically in (Leone, Ricci-Tersenghi and Zecchina, 2001). A sharp characterization of finite-size scaling for the appearence of a 2-

core, corresponding to the clustering transition, was achieved in (Dembo and Montanari, 2008*a*).

19

THE 1RSB CAVITY METHOD

The effectiveness of belief propagation depends on one basic assumption: when a function node is pruned from the factor graph, the adjacent variables become weakly correlated with respect to the resulting distribution. This hypothesis may break down either because of the existence of small loops in the factor graph, or because variables are correlated on large distances. In factor graphs with a locally tree-like structure, the second scenario is responsible for the failure of BP. The emergence of such long range correlations is a signature of a phase transition separating a 'weakly correlated' and a 'highly correlated' phase. The latter is often characterized by the decomposition of the (Boltzmann) probability distribution into well separated 'lumps' (pure Gibbs states).

We considered a simple example of this phenomenon in our study of random XORSAT. A similar scenario holds in a variety of problems from random graph coloring to random satisfiability and spin glasses. The reader should be warned that the structure and organization of pure states in such systems is far from being fully understood. Furthermore, the connection between long range correlations and pure states decomposition is more subtle than suggested by the above remarks.

Despite these complications, physicists have developed a non-rigorous approach to deal with this phenomenon: the "one step replica symmetry breaking" (1RSB) cavity method. The method postulates a few properties of the pure state decomposition, and, on this basis, allows to derive a number of quantitative predictions ('conjectures' from a mathematics point of view). Examples include the satisfiability threshold for random K-SAT and other random constraint satisfaction problems.

The method is rich enough to allow for some self-consistency checks of such assumptions. In several cases in which the 1RSB cavity method passed this test, its predictions have been confirmed by rigorous arguments (and there is no case in which they have been falsified so far). These successes encourage the quest for a mathematical theory of Gibbs states on sparse random graphs.

This chapter explains the 1RSB cavity method. It alternates between a general presentation and a concrete illustration on the XORSAT problem. We strongly encourage the reader to read the previous chapter on XORSAT before the present one. This should help her to gain some intuition of the whole scenario.

We start with a general description of the 1RSB glass phase, and the decomposition in pure states, in Sec. 19.1. Section 19.2 introduces an auxiliary constraint satisfaction problem to count the number of solutions of BP equations. The 1RSB analysis amounts to applying belief propagation to this auxiliary problem. One can then apply the methods of Ch. 14 (for instance, density evolution) to the auxiliary problem. Section 19.3 illustrates the approach on the XORSAT problem and shows how the 1RSB cavity method recovers the rigorous results of the previous chapter.

In Sec. 19.4 we show how the 1RSB formalism, which in general is rather complicated, simplifies considerably when the temperature of the auxiliary constraint satisfaction problem takes the value $\mathtt{x} = 1$. Section 19.5 explains how to apply it to optimization problems (leveraging on the min-sum algorithm) leading to the Survey Propagation algorithm. The concluding section 19.6 describes the physical intuition which underlies the whole method. The appendix 19.6.3 contains some technical aspects of the survey propagation equations applied to XORSAT, and their statistical analysis.

19.1 Beyond BP: many states

19.1.1 *Bethe measures*

The main lesson of the previous chapters is that in many cases, the probability distribution specified by graphical models with a locally tree-like structure takes a relatively simple form, that we shall call a Bethe measure (or Bethe state). Let us first define precisely what we mean by this, before we proceed to discuss what kinds of other scenarios can be encountered.

As in Ch. 14, we consider a factor graph $G = (V, F, E)$, with variable nodes $V = \{1, \cdots, N\}$, factor nodes $F = \{1, \cdots, M\}$ and edges E. The joint probability distribution over the variables $\underline{x} = (x_1, \dots, x_N) \in \mathcal{X}^N$ takes the form

$$\mu(\underline{x}) = \frac{1}{Z} \prod_{a=1}^{M} \psi_a(\underline{x}_{\partial a}) \, . \tag{19.1}$$

Given a subset of variable nodes $U \subseteq V$ (which we shall call a ‘cavity’), the **induced subgraph** $G_U = (U, F_U, E_U)$ is defined as the factor graph that includes all the factor nodes a such that $\partial a \subseteq U$, and the adjacent edges. We also write $(i, a) \in \partial U$ if $i \in U$ and $a \in F \setminus F_U$. Finally, a **set of messages** $\{\widehat{\nu}_{a\to i}\}$ is a set of probability distributions over $\mathcal{X}$, indexed by directed edges $a \to i$ in E with $a \in F$, $i \in V$.

Definition 19.1. (Informal) *The probability distribution μ is a* ***Bethe measure*** *(or* ***Bethe state****) if there exists a set of messages $\{\widehat{\nu}_{a\to i}\}$, such that, for ‘almost all’ the ‘finite size’ cavities U, the distribution $\mu_U(\,\cdot\,)$ of the variables in U is approximated as*

$$\mu_U(\underline{x}_U) \cong \prod_{a\in F_U} \psi_a(\underline{x}_{\partial a}) \prod_{(ia)\in\partial U} \widehat{\nu}_{a\to i}(x_i) \;+\; \mathsf{err}(\underline{x}_U) \, , \tag{19.2}$$

where $\mathsf{err}(\underline{x}_U)$ is a ‘small’ error term, and $\cong$ denotes as usual equality up to a normalization.

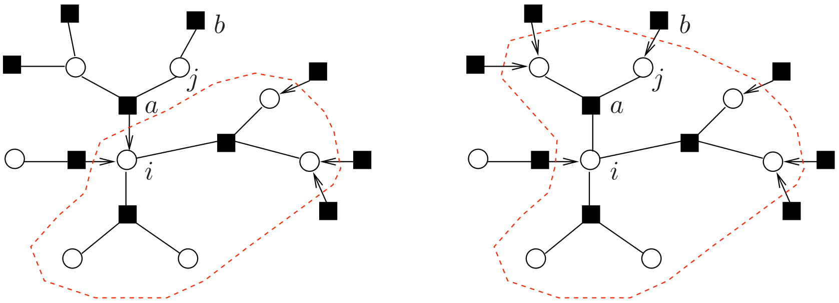

FIG. 19.1. Two examples of cavities. The right hand one is obtained by adding the extra function node a. The consistency of the Bethe measure in these two cavities implies the BP equation for $\widehat{\nu}_{a\to i}$, see Exercise 19.1.

A formal definition should specify what is meant by 'almost all', 'finite size' and 'small.' This can be done by introducing a tolerance ϵ_N (with $\epsilon_N \downarrow 0$ as $N \to \infty$) and a size L_N (where L_N is bounded as $N \to \infty$). One then requires that some norm of $\mathsf{err}(\,\cdot\,)$ (e.g. an L_p norm) is smaller than ϵ_N for a fraction larger than $1 - \epsilon_N$ of all possible cavities U of size $|U| < L_N$. The underlying intuition is that the measure $\mu(\,\cdot\,)$ is well approximated locally by the given set of messages. In the following we shall follow physicists' habit of leaving implicit the various approximation errors.

Notice that the above definition does not make use of the fact that μ factorizes as in Eq. (19.1). It thus apply to any distribution over $\underline{x} = \{x_i : i \in V\}$.

If $\mu(\,\cdot\,)$ is a Bethe measure with respect to the message set $\{\widehat{\nu}_{a\to i}\}$, then the consistency of Eq. (19.2) for different choices of U implies some non-trivial constraints on the messages. In particular if the loops in the factor graph G are not too small (and under some technical condition on the functions $\psi_a(\,\cdot\,)$) then the messages must be close to satisfying BP equations. More precisely, we define a **quasi-solution** of BP equations as a set of messages which satisfy almost all the equations within some accuracy. The reader is invited to prove this statement in the exercise below.

Exercise 19.1 Assume that $G = (V, F, E)$ has girth larger than 2, and that $\mu(\,\cdot\,)$ is a Bethe measure with respect to the message set $\{\widehat{\nu}_{a\to i}\}$ where $\widehat{\nu}_{a\to i}(x_i) > 0$ for any $(i,a) \in E$, and $\psi_a(\underline{x}_{\partial a}) > 0$ for any $a \in F$. For $U \subseteq V$, and $(i,a) \in \partial U$, define a new subset of variable nodes as $W = U \cup \partial a$ (see Fig. 19.1).

Applying Eq. (19.2) to the subsets of variables U and W, show that the message must satisfy (up to an error term of the same order as $\mathsf{err}(\,\cdot\,)$):

$$\widehat{\nu}_{a\to i}(x_i) \cong \sum_{\underline{x}_{\partial a\setminus i}} \psi_a(\underline{x}_{\partial a}) \prod_{j\in\partial a\setminus i} \Big\{ \prod_{b\in\partial j\setminus a} \widehat{\nu}_{b\to j}(x_j) \Big\}\,. \qquad (19.3)$$

Show that these are equivalent to the BP equations (14.14), (14.15).
[Hint: Define, for $k \in V$, $c \in F$, $(k,c) \in E$, $\nu_{k\to c}(x_k) \cong \prod_{d\in\partial k\setminus c} \widehat{\nu}_{d\to k}(x_k)$].

It would be pleasant if the converse was true, i.e. if each quasi-solution of BP equations corresponded to a distinct Bethe measure. In fact such a relation will be at the heart of the assumptions of the 1RSB method. However one should keep in mind that this is not always true, as the following example shows:

Example 19.2 Let G be a factor graph with the same degree $K \geq 3$ both at factor and variable nodes. Consider binary variables, $\mathcal{X} = \{\mathtt{0}, \mathtt{1}\}$, and, for each $a \in F$, let

$$\psi_a(x_{i_1(a)}, \dots, x_{i_K(a)}) = \mathbb{I}(x_{i_1(a)} \oplus \cdots \oplus x_{i_K(a)} = \mathtt{0})\,. \qquad (19.4)$$

Given a perfect matching $\mathsf{M} \subseteq E$, a solution of BP equations can be constructed as follows. If $(i,a) \in \mathsf{M}$, then let $\widehat{\nu}_{a\to i}(x_i) = \mathbb{I}(x_i = \mathtt{0})$ and $\nu_{i\to a}(\mathtt{0}) = \nu_{i\to a}(\mathtt{1}) = 1/2$. If on the other hand $(i,a) \notin \mathsf{M}$, then let $\widehat{\nu}_{a\to i}(\mathtt{0}) = \widehat{\nu}_{a\to i}(\mathtt{1}) = 1/2$ and $\nu_{i\to a}(\mathtt{0}) = \mathbb{I}(x_i = \mathtt{0})$ (variable to factor node).

Check that this is a solution of BP equations and that all the resulting local marginals coincide with the ones of the measure $\mu(\underline{x}) \cong \mathbb{I}(\underline{x} = \underline{\mathtt{0}})$, independently of M. If one takes for instance G to be a random regular graph with degree $K \geq 3$, both at factor nodes and variable nodes, then the number of perfect matchings of G is, with high probability, exponential in the number of nodes. Therefore we have constructed an exponential number of solutions of BP equations that describe the same Bethe measure.

19.1.2 *A few generic scenarios*

Bethe measures are a conceptual tool for describing distributions of the form (19.1). Inspired by the study of glassy phases (see Sec. 12.3), statistical mechanics studies have singled out a few generic scenarios in this respect, that we informally describe below.

- RS (replica symmetric). This is the simplest possible scenario: the distribution $\mu(\,\cdot\,)$ is a Bethe measure.

A slightly more complicated situation (that we still ascribe to the 'replica symmetric' family) arises when $\mu(\,\cdot\,)$ decomposes into a finite set of Bethe measures related by 'global symmetries', as in the Ising ferromagnet discussed in Sec. 17.3.

d1RSB (dynamic one-step replica symmetry breaking). There exists an exponentially large (in the system size N) number of Bethe measures. The measure μ decomposes into a convex combination of these Bethe measures:

$$\mu(\underline{x}) = \sum_n w_n\, \mu^n(\underline{x})\,, \tag{19.5}$$

with weights w_n exponentially small in N. Furthermore $\mu(\,\cdot\,)$ is itself a Bethe measure.

s1RSB (static one-step replica symmetry breaking). As in the d1RSB case, there exists an exponential number of Bethe measures, and μ decomposes into a convex combination of such states. However, a finite number of the weights w_n is of order 1 as $N \to \infty$, and (unlike in the previous case) μ is not itself a Bethe measure.

In the following we shall focus on the d1RSB and s1RSB scenarios, that are particularly interesting, and can be treated in a unified framework (we shall sometimes refer to both of them as 1RSB). More complicate scenarios, such as 'full RSB', are also possible. We do not discuss such scenarios here because, so far, one has a relatively poor control of them in sparse graphical models.

In order to proceed further, we shall make a series of assumptions on the structure of Bethe states in the 1RSB case. While further research work is required to formalize completely these assumptions, they are precise enough for deriving several interesting quantitative predictions.

To avoid technical complications, we assume that the compatibility functions $\psi_a(\,\cdot\,)$ are strictly positive. (The cases with $\psi_a(\,\cdot\,) = 0$ should be treated as limit cases of such models). Let us index by n the various quasi-solutions $\{\nu^n_{i\to a}, \widehat{\nu}^n_{a\to i}\}$ of the BP equations. To each of them we can associate a Bethe measure, and we can compute the corresponding Bethe free-entropy $\mathbb{F}_n = \mathbb{F}(\underline{\nu}^n)$. The three postulates of the 1RSB scenario are listed below.

Assumption 1 *There exist exponentially many quasi-solutions of BP equations. The number of such solutions with free-entropy* $\mathbb{F}(\underline{\nu}^n) \approx N\phi$ *is (to leading exponential order)* $\exp\{N\Sigma(\phi)\}$*, where* $\Sigma(\,\cdot\,)$ *is the* ***complexity*** *function*[27] .

This can be expressed more formally as follows. There exists a function $\Sigma : \mathbb{R} \to \mathbb{R}_+$ (the complexity) such that, for any interval $[\phi_1, \phi_2]$, the number of quasi-solutions of BP equations with $\mathbb{F}(\underline{\nu}^n) \in [N\phi_1, N\phi_2]$ is $\exp\{N\Sigma_* + o(N)\}$ where $\Sigma_* = \sup\{\Sigma(\phi) : \quad \phi_1 \le \phi \le \phi_2\}$. We shall also assume in the following that $\Sigma(\phi)$ is 'regular enough' without entering details.

[27] As we are only interested in the leading exponential behavior, the details of the definitions of quasi-solutions become irrelevant, as long as (for instance) the fraction of violated BP equations vanishes in the large N limit.

Among Bethe measures, a special role is played by the ones that have short range correlations (are *extremal*). We already mentioned this point in Ch. 12, and shall discuss the relevant notion of correlation decay in Ch. 22. We denote the set of extremal measures as E.

Assumption 2 *The 'canonical' measure μ defined as in Eq. (19.1) can be written as a convex combination of extremal Bethe measures*

$$\mu(\underline{x}) = \sum_{n\in\mathsf{E}} w_n\, \mu^n(\underline{x}) \;, \tag{19.6}$$

with weights related to the Bethe free-entropies $w_n = e^{\mathbb{F}_n}/\Xi$, $\Xi \equiv \sum_{n\in\mathsf{E}} e^{\mathbb{F}_n}$.

Note that Assumption 1 characterizes the number of (approximate) BP fixed points, while Assumption 2 expresses the measure $\mu(\,\cdot\,)$ in terms of extremal Bethe measures. While each such measure gives rise to a BP fixed point by the arguments in the previous Section, it is not clear that the reciprocal holds. The next assumption implies that this is the case, to the leading exponential order.

Assumption 3 *To leading exponential order, the number of extremal Bethe measures equals the number of quasi-solutions of BP equation: the number of extremal Bethe measures with free-entropy $\approx N\phi$ is also given by* $\exp\{N\Sigma(\phi)\}$.

19.2 The 1RSB cavity equations

Within the three assumptions described above, the complexity function $\Sigma(\phi)$ provides basic information on how the measure μ decomposes into Bethe measures. Since the number of extremal Bethe measures with a given free entropy density is exponential in the system size, it is natural to treat them within a statistical physics formalism. BP messages of the original problem will be the new variables and Bethe measures will be the new configurations. This is what 1RSB is about.

We introduce the auxiliary statistical physics problem through the definition of a canonical distribution over extremal Bethe measures: we assign to measure $n \in \mathsf{E}$, the probability $w_n(\mathtt{x}) = e^{\mathtt{x}\mathbb{F}_n}/\Xi(\mathtt{x})$. Here $\mathtt{x}$ plays the role of an inverse temperature (and is often called the **Parisi 1RSB parameter**) [28]. The partition function of this generalized problem is

$$\Xi(\mathtt{x}) = \sum_{n\in\mathsf{E}} e^{\mathtt{x}\mathbb{F}_n} \doteq \int e^{N[\mathtt{x}\phi+\Sigma(\phi)]}\, \mathrm{d}\phi \,. \tag{19.7}$$

According to Assumption 2 above, extremal Bethe measures contribute to μ through a weight $w_n = e^{\mathbb{F}_n}/\Xi$. Therefore the original problem is described by the choice $\mathtt{x} = 1$. But varying $\mathtt{x}$ will allow us to recover the full complexity function $\Sigma(\phi)$.

[28] It turns out that the present approach is equivalent the cloning method discussed in Chapter 12, where $\mathtt{x}$ is the number of clones.

If $\Xi(\mathtt{x}) \doteq e^{N\mathfrak{F}(\mathtt{x})}$, a saddle point evaluation of the integral in (19.7) gives Σ as the Legendre transform of $\mathfrak{F}$:

$$\mathfrak{F}(\mathtt{x}) = \mathtt{x}\phi + \Sigma(\phi)\,, \qquad \frac{\partial \Sigma}{\partial \phi} = -\mathtt{x}\ . \tag{19.8}$$

19.2.1 *Counting BP fixed points*

In order to actually estimate $\Xi(\mathtt{x})$, we need to consider the distribution induced by $w_n(\mathtt{x})$ on the messages $\underline{\nu} = \{\nu_{i\to a}, \widehat{\nu}_{a\to i}\}$, that we shall denote by $\mathsf{P}_x(\underline{\nu})$. The fundamental observation is that this distribution can be written as a graphical model, whose variables are BP messages. A first family of function nodes enforces the BP equations, and a second one implements the weight $e^{\mathtt{x}\mathbb{F}(\underline{\nu})}$. Furthermore, it turns out that the topology of the factor graph in this **auxiliary graphical model** is very close to that of the original factor graph. This suggests to use the BP approximation in this auxiliary model in order to estimate $\Sigma(\phi)$.

The 1RSB approach can be therefore summarized in one sentence:

Introduce a Boltzmann distribution over Bethe measures, write it in the form of a graphical model, and use BP to study this model.

This program is straightforward, but one must be careful not to confuse the two models (the original one and the auxiliary one), and their messages. Let us first simplify the notations of the original messages. The two types of messages entering the BP equations of the original problem will be denoted by $\widehat{\nu}_{a\to i} = \widehat{\mathtt{m}}_{ai}$ and $\nu_{i\to a} = \mathtt{m}_{ia}$; we will denote by $\underline{\mathtt{m}}$ the set of all the $\mathtt{m}_{ia}$ and by $\underline{\widehat{\mathtt{m}}}$ the set of all the $\widehat{\mathtt{m}}_{ai}$. Each of these $2|\mathcal{E}|$ messages is a normalized probability distribution over the alphabet $\mathcal{X}$. With these notations, the original BP equations read:

$$\mathtt{m}_{ia}(x_i) \cong \prod_{b\in\partial i\setminus a} \widehat{\mathtt{m}}_{bi}(x_i)\,, \qquad \widehat{\mathtt{m}}_{ai}(x_i) \cong \sum_{\{x_j\}_{j\in\partial a\setminus i}} \psi_a(\underline{x}_{\partial a}) \prod_{j\in\partial a\setminus i} \mathtt{m}_{ja}(x_j)\,. \tag{19.9}$$

Hereafter we shall write them in the compact form:

$$\mathtt{m}_{ia} = \mathsf{f}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}\right)\,, \qquad \widehat{\mathtt{m}}_{ai} = \hat{\mathsf{f}}_a\left(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}\right)\,. \tag{19.10}$$

Each message set $(\underline{\mathtt{m}}, \underline{\widehat{\mathtt{m}}})$ is given a weight proportional to $e^{\mathtt{x}\mathbb{F}(\underline{\mathtt{m}},\underline{\widehat{\mathtt{m}}})}$, where the free-entropy $\mathbb{F}(\underline{\mathtt{m}}, \underline{\widehat{\mathtt{m}}})$ is written in terms of BP messages

$$\mathbb{F}(\underline{\mathtt{m}}, \underline{\widehat{\mathtt{m}}}) = \sum_{a\in F} \mathbb{F}_a\left(\{\mathtt{m}_{ja}\}_{j\in\partial a}\right) + \sum_{i\in V} \mathbb{F}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i}\right) - \sum_{(ia)\in E} \mathbb{F}_{ia}\left(\mathtt{m}_{ia}, \widehat{\mathtt{m}}_{ai}\right)\ . \tag{19.11}$$

The functions $\mathbb{F}_a, \mathbb{F}_i, \mathbb{F}_{ia}$ have been obtained in (14.28). Let us copy them here for convenience:

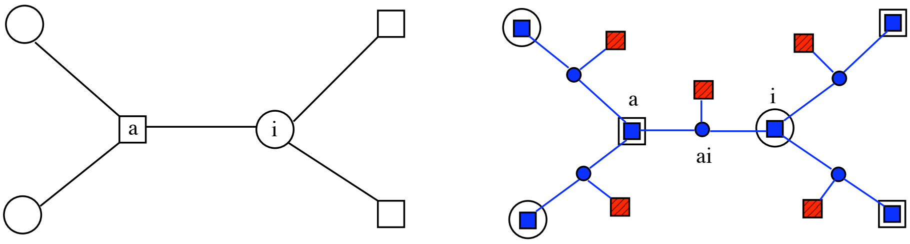

FIG. 19.2. A part of the original factor graph (left) and the corresponding auxiliary factor graph (right)

$$\mathbb{F}_a(\{\mathtt{m}_{ja}\}_{j\in\partial a}) = \log\left[\sum_{\underline{x}_{\partial a}} \psi_a(\underline{x}_{\partial a}) \prod_{j\in\partial a} \mathtt{m}_{ja}(x_j)\right] \,,$$

$$\mathbb{F}_i(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i}) = \log\left[\sum_{x_i} \prod_{b\in\partial i} \widehat{\mathtt{m}}_{bi}(x_i)\right] \,, \tag{19.12}$$

$$\mathbb{F}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) = \log\left[\sum_{x_i} \mathtt{m}_{ia}(x_i)\widehat{\mathtt{m}}_{ai}(x_i)\right] \,. \tag{19.13}$$

We now consider the $2|\mathcal{E}|$ messages $\underline{\mathtt{m}}$ and $\underline{\widehat{\mathtt{m}}}$ as variables in our auxiliary graphical model. The distribution induced my $w_n(\mathtt{x})$ on such messages takes the form

$$\mathsf{P}_{\mathtt{x}}(\underline{\mathtt{m}},\underline{\widehat{\mathtt{m}}}) = \frac{1}{\Xi(x)} \prod_{a\in F} \Psi_a(\{\mathtt{m}_{ja},\widehat{\mathtt{m}}_{ja}\}_{j\in\partial a}) \prod_{i\in V} \Psi_i(\{\mathtt{m}_{ib},\widehat{\mathtt{m}}_{ib}\}_{b\in\partial i}) \prod_{(ia)\in E} \Psi_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ia})\,, \tag{19.14}$$

where we introduced the compatibility functions:

$$\Psi_a = \prod_{i\in\partial a} \mathbb{I}\left(\widehat{\mathtt{m}}_{ai} = \hat{\mathsf{f}}_a\left(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}\right)\right) \; e^{\mathtt{x}\mathbb{F}_a(\{\mathtt{m}_{ja}\}_{j\in\partial a})}\,, \tag{19.15}$$

$$\Psi_i = \prod_{a\in\partial i} \mathbb{I}\left(\mathtt{m}_{ia} = \mathsf{f}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}\right)\right) \; e^{\mathtt{x}\mathbb{F}_i(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i})}\,, \tag{19.16}$$

$$\Psi_{ia} = e^{-\mathtt{x}\mathbb{F}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})}\,. \tag{19.17}$$

The corresponding factor graph is depicted in Fig. 19.2 and can described as follows:

- *For each edge* (i,a) *of the original factor graph*, introduce a variable node in the auxiliary factor graph. The associated variable is the pair $(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})$. Furthermore, introduce a function node connected to this variable, contributing to the weight through a factor $\Psi_{ia} = e^{-\mathtt{x}\mathbb{F}_{ai}}$.
- *For each function node* a *of the original graph* introduce a function node in the auxiliary graph and connect it to all the variable nodes corresponding

to edges (i,a), $i \in \partial a$. The compatibility function Ψ_a associated to this function node has two roles: (i) It enforces the $|\partial a|$ BP equations expressing the variables $\{\widehat{\mathtt{m}}_{ai}\}_{i\in\partial a}$ in terms of the $\{\mathtt{m}_{ia}\}_{i\in\partial a}$, cf. Eq. (19.9); (ii) It contributes to the weight through a factor $e^{\mathtt{x}\mathbb{F}_a}$.

- *For each variable node i of the original graph*, introduce a function node in the auxiliary graph, and connect it to all variable nodes corresponding to edges (i,a), $a \in \partial i$. The compatibility function Ψ_i has two roles: (i) It enforces the $|\partial i|$ BP equations expressing the variables $\{\mathtt{m}_{ib}\}_{b\in\partial i}$ in terms of $\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i}$, cf. Eq. (19.9); (ii) It contributes to the weight through a factor $e^{\mathtt{x}\mathbb{F}_i}$.

Note that we were a bit sloppy in Eqs. (19.15) to (19.17). The messages $\mathtt{m}_{ia}$, $\widehat{\mathtt{m}}_{ai}$ are in general continuous, and indicator functions should therefore be replaced by delta functions. This might pose in turn some definition problem (what is the reference measure on the messages? can we hope for *exact* solutions of BP equations?). One should consider the above as a shorthand for the following procedure. First discretize the messages (and BP equations) in such a way that they can take a finite number q of values. Compute the complexity by letting $N \to \infty$ at fixed q, and take the limit $q \to \infty$ at the end. It is easy to define several alternative, and equally reasonable, limiting procedures. We expect all of them to yield the same result. In practice, the ambiguities in Eqs. (19.15) to (19.17) are solved on a case by case basis.

19.2.2 *Message passing on the auxiliary model*

The problem of counting the number of Bethe measures (more precisely, computing the complexity function $\Sigma(\phi)$) has been reduced to the one of estimating the partition function $\Xi(\mathtt{x})$ of the auxiliary graphical model (19.14). Since we are interested in the case of locally tree-like factor graphs G, the auxiliary factor graph is locally tree-like as well. We can therefore apply BP to estimate its free-entropy density $\mathfrak{F}(\mathtt{x}) = \lim_N N^{-1} \log \Xi(\mathtt{x})$. This will give us the complexity through the Legendre transform of Eq. (19.8).

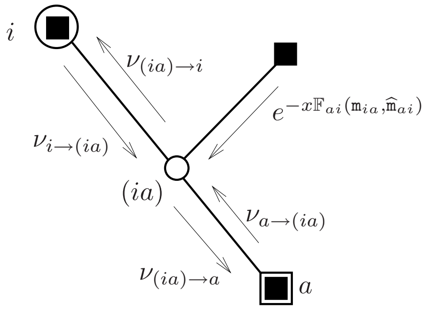

FIG. 19.3. Messages in the auxiliary graphical model.

In the following we denote by $i \in V$ and $a \in F$ a generic variable and function node in the graph G, and by $(ia) \in E$ an edge in G. By extension, we denote in the same way the corresponding nodes in the auxiliary graph. The messages appearing in the BP analysis of the auxiliary model can be classified as follows, cf. Fig. 19.3:

→ From the variable node (ia) are issued two messages: $\nu_{(ia)\to a}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})$ and $\nu_{(ia)\to i}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})$

→ From the function node a are issued $|\partial a|$ messages to nodes $i \in \partial a$, $\widehat{\nu}_{a\to(ai)}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})$

→ From the function node i are issued $|\partial i|$ messages to nodes $a \in \partial i$, $\widehat{\nu}_{i\to(ai)}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})$

→ From the degree-one function node connected to the variable node (ia) is issued a message towards this variable. This message is simply $e^{-\mathtt{x}\mathbb{F}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})}$.

The BP equations on the variable node (ia) take a simple form:

$$
\begin{aligned}
\nu_{(ia)\to a}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) &\cong \widehat{\nu}_{i\to(ia)}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})\; e^{-\mathtt{x}\mathbb{F}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})}\,,\\
\nu_{(ia)\to i}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) &\cong \widehat{\nu}_{a\to(ia)}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})\; e^{-\mathtt{x}\mathbb{F}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})}\,.
\end{aligned}
\tag{19.18}
$$

We can use these equations to eliminate messages $\widehat{\nu}_{i\to(ia)}$, $\widehat{\nu}_{a\to(ia)}$ in favor of $\nu_{(ia)\to a}$, $\nu_{(ia)\to i}$. In order to emphasize this choice (and to simplify notations) we define:

$$
Q_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) \equiv \nu_{(ia)\to a}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})\,,\qquad \widehat{Q}_{ai}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) \equiv \nu_{(ia)\to i}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})\,. \tag{19.19}
$$

We can now write the remaining BP equations of the auxiliary graphical model in terms of $Q_{ia}(\,\cdot\,,\,\cdot\,)$, $\widehat{Q}_{ai}(\,\cdot\,,\,\cdot\,)$. The BP equation associated to the function node corresponding to $i \in V$ reads:

$$
\begin{aligned}
Q_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) \cong &\sum_{\{\mathtt{m}_{ib},\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}} \left[\prod_{c\in\partial i} \mathbb{I}\left(\mathtt{m}_{ic} = \mathsf{f}_i(\{\widehat{\mathtt{m}}_{di}\}_{d\in\partial i\setminus c})\right)\right]\\
&\exp\left\{\mathtt{x}\left[\mathbb{F}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i}\right) - \mathbb{F}_{ai}\left(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}\right)\right]\right\} \prod_{b\in\partial i\setminus a} \widehat{Q}_{bi}(\mathtt{m}_{ib},\widehat{\mathtt{m}}_{bi})\,,
\end{aligned}
\tag{19.20}
$$

and the one associated to the function node corresponding to $a \in F$ is:

$$
\widehat{Q}\;_{ai}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) \cong \sum_{\{\mathtt{m}_{ja},\widehat{\mathtt{m}}_{aj}\}_{j\in\partial a\setminus i}} \left[\prod_{j\in\partial a} \mathbb{I}\left(\widehat{\mathtt{m}}_{aj} = \hat{\mathsf{f}}_a(\{\mathtt{m}_{ka}\}_{k\in\partial a\setminus j})\right)\right] \tag{19.21}
$$

$$
\exp\Big\{\mathtt{x}\Big[\mathbb{F}_a\Big(\{\mathtt{m}_{ja}\}_{j\in\partial a}\Big\} - \mathbb{F}_{ai}\left(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}\right)\Big]\Big) \prod_{j\in\partial a\setminus i} Q_{ja}(\mathtt{m}_{ja},\widehat{\mathtt{m}}_{aj})\,. \tag{19.22}
$$

Equations (19.20), (19.22) can be further simplified, using the following lemma.

Lemma 19.3 *Assume* $\sum_{x_i} \mathtt{m}_{ia}(x_i)\widehat{\mathtt{m}}_{ai}(x_i) > 0$*. Under the condition* $\mathtt{m}_{ia} = \mathsf{f}_i(\{\widehat{\mathtt{m}}_{di}\}_{d\in\partial i\setminus a})$ *(in particular if the indicator functions in Eq. (19.20) evaluate to* 1*), the difference* $\mathbb{F}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i}\right) - \mathbb{F}_{ai}\left(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}\right)$ *can be expressed in terms of* $\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}$*. Explicitly, we have*

$$e^{\mathbb{F}_i-\mathbb{F}_{ia}} = z_{ia}(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}) \equiv \sum_{x_i}\prod_{b\in\partial i\setminus a}\widehat{\mathtt{m}}_{bi}(x_i)\,. \tag{19.23}$$

Analogously, under the condition $\widehat{\mathtt{m}}_{ai} = \hat{\mathsf{f}}_a(\{\mathtt{m}_{ka}\}_{k\in\partial a\setminus i})$ *(in particular if the indicator functions in Eq. (19.22) evaluate to* 1*) the difference* $\mathbb{F}_a\left(\{\mathtt{m}_{ja}\}_{j\in\partial a}\right) - \mathbb{F}_{ai}\left(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}\right)$ *depends only on* $\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}$*. Explicitly:*

$$e^{\mathbb{F}_a-\mathbb{F}_{ia}} = \hat{z}_{ai}(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}) \equiv \sum_{\underline{x}_{\partial a}}\psi_a(\underline{x}_{\partial a})\prod_{j\in\partial a\setminus i}\mathtt{m}_{ja}(x_j)\,. \tag{19.24}$$

Proof: Let us first consider Eq. (19.23). From the definition (14.28), it follows that

$$e^{\mathbb{F}_i-\mathbb{F}_{ia}} = \frac{\sum_{x_i}\prod_{b\in\partial i}\widehat{\mathtt{m}}_{bi}(x_i)}{\sum_{x_i}\mathtt{m}_{ia}(x_i)\widehat{\mathtt{m}}_{ai}(x_i)}\,. \tag{19.25}$$

Substituting $\mathtt{m}_{ia} = f_i(\{\widehat{\mathtt{m}}_{ci}\}_{c\in\partial i\setminus a})$ in the denominator, and keeping track of the normalization constant, we get

$$\sum_{x_i}\mathtt{m}_{ia}(x_i)\widehat{\mathtt{m}}_{ai}(x_i) = \frac{\sum_{x_i}\prod_{b\in\partial i}\widehat{\mathtt{m}}_{bi}(x_i)}{\sum_{x_i}\prod_{b\in\partial i\setminus a}\widehat{\mathtt{m}}_{ai}(x_i)}\,, \tag{19.26}$$

which implies Eq. (19.23).

The derivation of Eq. (19.24) is completely analogous and left as an exercise for the reader. □

Notice that the functions $z_{ia}(\,\cdot\,)$, $\hat{z}_{ai}(\,\cdot\,)$ appearing in Eqs. (19.23), (19.24) are in fact the normalization constants hidden by the $\cong$ notation in Eqs. (19.9).

Because of this lemma, we can seek a solution of Eqs. (19.20), (19.22) with Q_{ia} depending only on $\mathtt{m}_{ia}$, and $\widehat{Q}_{ai}$ depends only on $\widehat{\mathtt{m}}_{ai}$. Hereafter we shall focus on this case, and, with an abuse of notation, we shall write:

$$Q_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) = Q_{ia}(\mathtt{m}_{ia})\,,\qquad \widehat{Q}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}) = \widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai})\,. \tag{19.27}$$

The BP equations for the auxiliary graphical model (19.20), (19.22) then become:

$$Q_{ia}(\mathtt{m}_{ia}) \cong \sum_{\{\widehat{\mathtt{m}}_{bi}\}}\mathbb{I}\left(\mathtt{m}_{ia} = g_i(\{\widehat{\mathtt{m}}_{bi}\})\right)\,\left[z_{ia}(\{\widehat{\mathtt{m}}_{bi}\})\right]^{\mathtt{x}}\prod_{b\in\partial i\setminus a}\widehat{Q}_{bi}(\widehat{\mathtt{m}}_{bi})\,, \tag{19.28}$$

$$\widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai}) \cong \sum_{\{\mathtt{m}_{ja}\}}\mathbb{I}\left(\widehat{\mathtt{m}}_{ai} = f_a(\{\mathtt{m}_{ja}\})\right)\,\left[\hat{z}_{ai}(\{\mathtt{m}_{ja}\})\right]^{\mathtt{x}}\prod_{j\in\partial a\setminus i}Q_{ja}(\mathtt{m}_{ja})\,, \tag{19.29}$$

where $\{\widehat{\mathtt{m}}_{bi}\}$ is a shorthand for $\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}$ and $\{\mathtt{m}_{ja}\}$ a shorthand for $\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}$. The expressions for $z_{ia}(\{\widehat{\mathtt{m}}_{bi}\})$ and $\hat{z}_{ai}(\{\mathtt{m}_{ja}\})$ are given in Eqs. (19.23), (19.24).

Equations (19.28), (19.29) are the **1RSB cavity equations**. As we did in the ordinary BP equations, we can consider them as an update rule for a message passing algorithm. This will be further discussed in the next sections. One sometimes uses the notation $Q_{i\to a}(\,\cdot\,)$, $\widehat{Q}_{a\to i}(\,\cdot\,)$, to emphasize the fact that 1RSB messages are associated to *directed* edges.

Notice that our derivation was based on the assumption that $\sum_{x_i} \mathtt{m}_{ia}(x_i)\widehat{\mathtt{m}}_{ai}(x_i) > 0$. This condition holds if, for instance, the compatibility functions of the original model are bounded away from 0. Under this condition, we have shown that:

Proposition 19.4 *If the 1RSB cavity equations (19.28), (19.29) have a solution $\widehat{Q}, Q$, this corresponds to a solution to the BP equations of the auxiliary graphical model. Reciprocally, if the BP equations of the auxiliary graphical model admit a solution satisfying the condition (19.27), then the resulting messages must be a solution of the 1RSB cavity equations.*

Assumption (19.27) -which is suggestive of a form of “causality”- cannot be further justified within the present approach, but alternative derivations of the 1RSB equations confirm its validity.

19.2.3 *Computing the complexity*

We now compute the free-entropy of the auxiliary graphical model within the BP approximation. We expect the result of this procedure to be asymptotically exact for a wide class of locally tree like graphs, thus yielding the correct free-entropy density $\mathfrak{F}(\mathtt{x}) = \lim_N N^{-1}\log \Xi(\mathtt{x})$.

Assume $\{Q_{ia}, \widehat{Q}_{ai}\}$ to be a solution (or a quasi-solution) of the 1RSB cavity equations (19.28), (19.29). We use the general form (14.27) of Bethe free-entropy, but take into account the degree one factor nodes using the simplified expression derived in Exercise 14.6. The various contributions to the free-entropy are:

→ Contribution from the function node a (here $\{\mathtt{m}_{ia}\}$ is a shorthand for $\{\mathtt{m}_{ia}\}_{i\in\partial a}$):

$$\mathbb{F}_a^{\rm RSB} = \log\left\{\sum_{\{\mathtt{m}_{ia}\}} e^{\mathtt{x}\mathbb{F}_a(\{\mathtt{m}_{ia}\})} \prod_{i\in\partial a} Q_{ia}(\mathtt{m}_{ia})\right\}\,. \tag{19.30}$$

→ Contribution from the function node i ($\{\widehat{\mathtt{m}}_{ai}\}$ is a shorthand for $\{\widehat{\mathtt{m}}_{ai}\}_{a\in\partial i}$):

$$\mathbb{F}_i^{\rm RSB} = \log\left\{\sum_{\{\widehat{\mathtt{m}}_{ai}\}} e^{\mathtt{x}\mathbb{F}_i(\{\widehat{\mathtt{m}}_{ai}\})} \prod_{a\in\partial i} \widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai})\right\}\,. \tag{19.31}$$

→ Contribution from the variable node (ia):

$$\mathbb{F}_{ia}^{\rm RSB} = \log\left\{\sum_{\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}} e^{\mathtt{x}\mathbb{F}_{ia}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})} Q_{ia}(\mathtt{m}_{ia})\widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai})\right\}\,. \tag{19.32}$$

→ The contributions from the two edges $a-(ai)$ and $i-(ai)$ are both equal to $-\mathbb{F}^{\rm RSB}_{ia}$

The Bethe free-entropy of the auxiliary graphical model is equal to:

$$\mathbb{F}^{\rm RSB}(\{Q,\widehat{Q}\}) = \sum_{a\in F}\mathbb{F}^{\rm RSB}_a + \sum_{i\in V}\mathbb{F}^{\rm RSB}_i - \sum_{(ia)\in E}\mathbb{F}^{\rm RSB}_{ia}\,. \tag{19.33}$$

19.2.4 *Summary*

The 1RSB cavity equations (19.28), (19.29) are BP equations for the auxiliary graphical model defined in (19.14). They relate $2|\mathcal{E}|$ messages $\{Q_{ia}(\mathtt{m}_{ia}),\ \widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai})\}$. Each such message is a probability distribution of ordinary BP messages, respectively $\mathtt{m}_{ia}(x_i)$ and $\widehat{\mathtt{m}}_{ai}(x_i)$. These elementary messages are in turn probability distributions on variables $x_i\in\mathcal{X}$.

Given a solution (or an approximate solution) $\{Q_{ia},\ \widehat{Q}_{ai}\}$, one can estimate the free-entropy density of the auxiliary model as

$$\log\Xi(\mathtt{x}) = \mathbb{F}^{\rm RSB}(\{Q,\widehat{Q}\}) + \mathsf{err}_N\,. \tag{19.34}$$

where $\mathbb{F}^{\rm RSB}(\{Q,\widehat{Q}\})$ is given by Eq. (19.33). For a large class of locally tree-like models we expect the BP approximation to be asymptotically exact on the auxiliary model. This means that the error term err_N is $o(N)$.

For such models, the free-entropy density is given by its 1RSB cavity expression $\mathfrak{F}(\mathtt{x}) = \mathrm{f}^{\rm RSB}(\mathtt{x}) \equiv \lim_{N\to\infty}\mathbb{F}^{\rm RSB}(\{Q,\widehat{Q}\})/N$. The complexity $\Sigma(\phi)$ is then computed through the Legendre transform (19.8).

19.2.5 *Random graphical models and density evolution*

Let us consider the case where G is a random graphical model as defined in Sec. 14.6.1. The factor graph is distributed according to one of the ensembles $\mathbb{G}_N(K,\alpha)$ or $\mathbb{D}_N(\Lambda,P)$. Function nodes are taken from a finite list $\{\psi^{(k)}(x_1,\dots,x_k;\widehat{J})\}$ indexed by a label $\widehat{J}$ with distribution $P^{(k)}_{\widehat{J}}$. Each factor $\psi_a(\,\cdot\,)$ is taken equal to $\psi^{(k)}(\cdots;\widehat{J}_a)$ independently with the same distribution. We also introduce explicitly a degree-one factor $\psi_i(x_i)$ connected to each variable node $i\in V$. This are also drawn independently from a list of possible factors $\{\psi(x;J)\}$, indexed by a label J with distribution P_J.

For a random graphical model, the measure $\mu(\,\cdot\,)$ becomes random, and so does its decomposition in extremal Bethe states, in particular the probabilities $\{w_n\}$, and the message sets $\{\nu^n_{i\to a},\widehat{\nu}^n_{a\to i}\}$. In particular, the 1RSB messages $\{Q_{ia},\widehat{Q}_{ai}\}$ become random. It is important to keep in mind the 'two levels' of randomness. Given an edge (ia), the message $\nu^n_{i\to a}$ is random if the Bethe state n is drawn from the distribution w_n. The resulting distribution $Q_{ia}(\mathtt{m})$ becomes a random variable when the graphical model is itself random.

The distributions of $Q_{ia}(\mathtt{m})$, $\widehat{Q}_{ai}(\widehat{\mathtt{m}})$ can then be studied through the density evolution method of Sec. 14.6.2. Let us assume an i.i.d. initialization $Q^{(0)}_{ia}(\,\cdot\,)\overset{\rm d}{=}$

$Q^{(0)}(\,\cdot\,)$ (respectively $\widehat{Q}^{(0)}_{ai}(\,\cdot\,) \overset{\rm d}{=} \widehat{Q}^{(0)}(\,\cdot\,)$), and denote by $Q^{(t)}_{ia}(\,\cdot\,)$, $\widehat{Q}^{(t)}_{ai}(\,\cdot\,)$ the 1RSB messages along edge (ia) after t parallel updates using the 1RSB equations (19.28), (19.29). If (ia) is a uniformly random edge then, as $N \to \infty$, $Q^{(t)}_{ia}(\,\cdot\,)$ converges in distribution[29] to $Q^{(t)}(\,\cdot\,)$ (respectively $\widehat{Q}^{(t)}_{ia}(\,\cdot\,)$ converges in distribution to $\widehat{Q}^{(t)}(\,\cdot\,)$). The distributions $Q^{(t)}(\,\cdot\,)$ and $\widehat{Q}^{(t)}(\,\cdot\,)$ are themselves random variables that satisfy the equations:

$$Q^{(t+1)}(\mathtt{m}) \overset{\rm d}{\cong} \sum_{\{\widehat{\mathtt{m}}_b\}} \mathbb{I}\left(\mathtt{m} = \mathtt{f}(\{\widehat{\mathtt{m}}_b\}; J)\right)\; z(\{\widehat{\mathtt{m}}_b\}; J)^{\mathtt{x}} \prod_{b=1}^{l-1} \widehat{Q}^{(t)}_b(\widehat{\mathtt{m}}_b)\,, \tag{19.35}$$

$$\widehat{Q}^{(t)}(\widehat{\mathtt{m}}) \overset{\rm d}{\cong} \sum_{\{\mathtt{m}_j\}} \mathbb{I}\left(\widehat{\mathtt{m}} = \hat{\mathtt{f}}(\{\mathtt{m}_j\}; \widehat{J})\right)\; \hat{z}(\{\mathtt{m}_j\}; \widehat{J})^{\mathtt{x}} \prod_{j=1}^{k-1} Q^{(t)}_j(\mathtt{m}_j)\,, \tag{19.36}$$

where k and l are distributed according to the edge perspective degree profiles ρ_k and λ_l, the $\left\{\widehat{Q}^{(t)}_b\right\}$ are $k-1$ independent copies of $\widehat{Q}^{(t)}(\,\cdot\,)$, and $\left\{Q^{(t)}_j\right\}$ are $l-1$ independent copies of $Q^{(t)}(\,\cdot\,)$. The functions z and $\hat{z}$ are given by:

$$z(\{\widehat{\mathtt{m}}_b\}; J) = \sum_x \psi(x, J) \prod_{b=1}^{l-1} \widehat{\mathtt{m}}_b(x)$$

$$\hat{z}(\{\mathtt{m}_j\}; \widehat{J}) = \sum_{x_1,\cdots,x_k} \psi^{(k)}(x_1, \cdots, x_k; \widehat{J}) \prod_{j=1}^{k-1} \mathtt{m}_j(x_j) \tag{19.37}$$

Within the 1RSB cavity method, the actual distribution of $Q_{i\to a}$ is assumed to coincide with one of the fixed points of the above density evolution equations. As for the RS case, one hopes that, on large enough instances, the message passing algorithm will converge to messages distributed according to this fixed point equation (meaning that there is no problem in exchanging the limits $t \to \infty$ and $N \to \infty$). This can be checked numerically.

For random graphical models, the 1RSB free-entropy density converges to a finite limit $\mathrm{f}^{\rm RSB}(\mathtt{x})$. This can be expressed in terms of the distributions of Q, $\widehat{Q}$. by taking expectation of Eqs. (19.30) to (19.32), and assuming that 1RSB messages incoming at the same node are i.i.d.. As in (14.77) the result takes the form:

$$\mathrm{f}^{\rm RSB} = \mathrm{f}^{\rm RSB}_{\rm v} + n_{\rm f}\mathrm{f}^{\rm RSB}_{\rm f} - n_{\rm e}\mathrm{f}^{\rm RSB}_{\rm e}\ . \tag{19.38}$$

Here $n_{\rm f}$ is the average number of function nodes per variable (equal to $\Lambda'(1)/P'(1)$ for a graphical model in the $\mathbb{D}_N(\Lambda, P)$ ensemble, and to α for a graphical model in the $\mathbb{G}_N(K, \alpha)$ ensemble) and $n_{\rm e}$ is the number of edges per variable (equal to

[29] We shall not discuss the measure-theoretic subtleties related to this statement. Let us just mention that weak topology is understood on the space of messages $Q^{(t)}$.

$\Lambda'(1)$ and to $K\alpha$ in these two ensembles). The contribution from variable nodes $\mathrm{f}_{\mathrm{v}}^{\mathrm{RSB}}$, function nodes $\mathrm{f}_{\mathrm{f}}^{\mathrm{RSB}}$, and edges $\mathrm{f}_{\mathrm{e}}^{\mathrm{RSB}}$ are:

$$
\begin{aligned}
\mathrm{f}_{\mathrm{v}}^{\mathrm{RSB}} &= \mathbb{E}_{l,J,\{\widehat{Q}\}} \log \left\{ \sum_{\{\widehat{\mathtt{m}}_1,\dots,\widehat{\mathtt{m}}_l\}} \widehat{Q}_1(\widehat{\mathtt{m}}_1)\dots\widehat{Q}_l(\widehat{\mathtt{m}}_l) \left[\sum_{x\in\mathcal{X}} \widehat{\mathtt{m}}_1(x)\dots\widehat{\mathtt{m}}_l(x)\psi(x;J) \right]^{\mathrm{x}} \right\}, \\
\mathrm{f}_{\mathrm{f}}^{\mathrm{RSB}} &= \mathbb{E}_{k,\widehat{J},\{Q\}} \log \left\{ \sum_{\{\mathtt{m}_1,\dots,\mathtt{m}_k\}} Q_1(\mathtt{m}_1)\dots Q_k(\mathtt{m}_k) \right. \\
&\qquad \left. \left[\sum_{x_1,\dots,x_k\in\mathcal{X}} \mathtt{m}_1(x_1)\dots\mathtt{m}_k(x_k)\psi^{(k)}(x_1,\dots,x_k;\widehat{J}) \right]^{\mathrm{x}} \right\}, \\
\mathrm{f}_{\mathrm{e}}^{\mathrm{RSB}} &= \mathbb{E}_{\widehat{Q},Q} \log \left\{ \sum_{\widehat{\mathtt{m}},\mathtt{m}} \widehat{Q}(\widehat{\mathtt{m}})Q(\mathtt{m}) \left[\sum_{x\in\mathcal{X}} \widehat{\mathtt{m}}(x)\mathtt{m}(x) \right]^{\mathrm{x}} \right\}. \qquad (19.39)
\end{aligned}
$$

19.2.6 *Numerical implementation*

Needless to say, it is extremely challenging to find a fixed point of the density evolution equations (19.35), (19.36), and thus determine the distributions of $Q, \widehat{Q}$. A simple numerical approach consists in generalizing the population dynamics algorithm described in the context of the RS cavity method, cf. Sec. 14.6.3.

There are two important issues related to such a generalization:

(*i*) We seek the distribution of $Q(\,\cdot\,)$ (and $\widehat{Q}(\,\cdot\,)$), which is itself a distribution of messages. If we approximate $Q(\,\cdot\,)$ by a sample (a 'population'), we will thus need two level of populations. In other words we will seek a population $\{\mathtt{m}_r^s\}$ with NM items. For each $r \in \{1,\dots,N\}$, the set of messages $\{\mathtt{m}_r^s\}$, $s \in \{1,\dots,M\}$ represents a distribution $Q_r(\,\cdot\,)$ (ideally, it would be an i.i.d. sample from this distribution). At the next level, the population $\{Q_r(\,\cdot\,)\}$, $r \in \{1,\cdots,N\}$ represents the distribution of $Q(\,\cdot\,)$ (ideally, an i.i.d. sample).

Analogously, for function-to-variable messages, we will use a population $\{\widehat{\mathtt{m}}_r^s\}$, with $r \in \{1,\dots,N\}$ and $s \in \{1,\dots,M\}$.

(*ii*) The re-weighting factors $z(\{\widehat{\mathtt{m}}_b\};J)^{\mathrm{x}}$ and $\hat{z}(\{\mathtt{m}_j\};\widehat{J})^{\mathrm{x}}$ appearing in Eqs. (19.35) and (19.36) do not have any analog in the RS context. How can one take such factors into account when $Q(\,\cdot\,)$, $\widehat{Q}(\,\cdot\,)$ are represented as populations? One possibility is to generate an intermediate weighted population, and than sample from it with a probability proportional to the weight.

This procedure is summarized in the following pseudocode.

```
1RSB Population dynamics (Model ensemble, Sizes $N, M$, Iterations $T$)
Initialize $\{\mathtt{m}_r^s\}$;
for $t = 1, \dots, T$:
    for $r = 1, \dots, N$:
        Draw an integer $k$ with distribution $\rho$;
        Draw $i(1), \dots, i(k-1)$ uniformly in $\{1, \dots, N\}$;
        Draw $\widehat{J}$ with distribution $P_{\widehat{J}}^{(k)}$;
        for $s = 1, \dots, M$:
            Draw $s(1), \dots, s(k-1)$ uniformly in $\{1, \dots, M\}$;
            Compute $\widehat{\mathtt{m}}_{\mathtt{temp}}^s = \hat{\mathtt{f}}(\mathtt{m}_{i(1)}^{s(1)}, \cdots, \mathtt{m}_{i(k-1)}^{s(k-1)}; \widehat{J})$
            Compute $W^s = \hat{z}(\mathtt{m}_{i(1)}^{s(1)}, \cdots, \mathtt{m}_{i(k-1)}^{s(k-1)}; \widehat{J})^{\mathtt{x}}$
        end;
        Generate the new population
                $\{\widehat{\mathtt{m}}_r^s\}_{s \in [M]} = \text{Reweight}(\{\widehat{\mathtt{m}}_{\mathtt{temp}}^s, W^s\}_{s \in [M]})$
    end;
    for $r = 1, \dots, N$:
        Draw an integer $l$ with distribution $\lambda$;
        Draw $i(1), \dots, i(l-1)$ uniformly in $\{1, \dots, N\}$;
        Draw $J$ with distribution $P$;
        for $s = 1, \dots, M$:
            Draw $s(1), \dots, s(l-1)$ uniformly in $\{1, \dots, M\}$;
            Compute $\mathtt{m}_{\mathtt{temp}}^s = \mathtt{f}(\widehat{\mathtt{m}}_{i(1)}^{s(1)}, \cdots, \widehat{\mathtt{m}}_{i(l-1)}^{s(k-1)}; J)$
            Compute $W^s = z(\widehat{\mathtt{m}}_{i(1)}^{s(1)}, \cdots, \widehat{\mathtt{m}}_{i(l-1)}^{s(l-1)}; J)^{\mathtt{x}}$
        end;
        Generate the new population
                $\{\mathtt{m}_r^s\}_{s \in [M]} = \text{Reweight}(\{\mathtt{m}_{\mathtt{temp}}^s, W^s\}_{s \in [M]})$
end;
return $\{\widehat{\mathtt{m}}_r^s\}$ and $\{\mathtt{m}_r^s\}$.
```

The re-weighting procedure is given by:

```
Reweight (Population of messages/weights $\{(\mathtt{m}_{\mathtt{temp}}^s, W^s)\}_{s \in [M]}$)
for $s = 1, \dots, M$, set $p^s \equiv W^s / \sum_{s'} W^{s'}$;
for $s = 1, \dots, M$:
    Draw $i \in \{1, \dots, M\}$ with distribution $p^s$;
    Set $\mathtt{m}_{\mathtt{new}}^s = \mathtt{m}_{\mathtt{temp}}^i$;
end;
return $\{\mathtt{m}_{\mathtt{new}}^s\}_{s \in [M]}$.
```

In the large N, M limit, the populations generated by this algorithm should converge to i.i.d. samples distributed as $Q^{(T)}(\,\cdot\,)$, $\widehat{Q}^{(T)}(\,\cdot\,)$, cf. Eq. (19.35), (19.36).

By letting T grow they should represent accurately the fixed points of density evolution, although the caveats expressed in the RS case should be repeated here.

Among the other quantities, the populations generated by this algorithm allow to estimate the 1RSB free-entropy density (19.38). Suppose we have generated a population of messages $\{\widehat{\mathtt{m}}_r^s(\,\cdot\,)\}$, whereby each message is a probability distribution on $\mathcal{X}$. The corresponding estimate of $\mathrm{f}_{\mathrm{v}}^{\mathrm{RSB}}$ is:

$$\widehat{\mathrm{f}}_{\mathrm{v}}^{\mathrm{RSB}} = \mathbb{E}_{l,J}\frac{1}{N^l}\sum_{r(1)\dots r(l)=1}^{N}\log\left\{\frac{1}{M^l}\sum_{s(1),\dots,s(l)=1}^{M}\left[\sum_{x\in\mathcal{X}}\widehat{\mathtt{m}}_{r(1)}^{s(1)}(x)\cdots\widehat{\mathtt{m}}_{r(l)}^{s(l)}(x)\;\psi(x;J)\right]^{\mathrm{x}}\right\}.$$

Similar expressions are easily written for $\mathrm{f}_{\mathrm{f}}^{\mathrm{RSB}}$ and $\mathrm{f}_{\mathrm{e}}^{\mathrm{RSB}}$. Their (approximate) evaluation can be accelerated considerably by summing over a random subset of the l-uples $r(1),\dots,r(l)$ and $s(1),\dots,s(l)$. Further, as in the RS case, it is beneficial to average over iterations (equivalently, over T) in order to reduce statistical errors at small computational cost.

19.3 A first application: XORSAT

Let us apply the 1RSB cavity method to XORSAT. This approach was already introduced in Sec. 18.6, but we want to show how it follows as a special case of the formalism developed in the previous sections. Our objective is to exemplify the general ideas on a well understood problem, and to build basic intuition that will be useful in more complicated applications.

As in Ch. 18 we consider the distribution over $\underline{x} = (x_1,\dots,x_N) \in \{\mathtt{0},\mathtt{1}\}^N$ specified by

$$\mu(\underline{x}) = \frac{1}{Z}\prod_{a=1}^{M}\mathbb{I}\left(x_{i_1(a)}\oplus\cdots\oplus x_{i_k(a)} = b_a\right)\,. \tag{19.40}$$

As usual $\oplus$ denotes sum modulo 2 and, for each $a\in\{1,\cdots,M\}$, $\partial a = \{i_1(a),\dots,i_K(a)\}$ is a subset of $\{1,\cdot,N\}$, and $b_a\in\{\mathtt{0},\mathtt{1}\}$. Random K-XORSAT formulae are generated by choosing both the index set $\{i_1(a),\dots,i_K(a)\}$ and the right hand side b_a uniformly at random.

19.3.1 *BP equations*

The BP equations read:

$$\mathtt{m}_{ia}(x_i) = \frac{1}{z_{ia}}\prod_{b\in\partial i\setminus a}\widehat{\mathtt{m}}_{bi}(x_i)\,, \tag{19.41}$$

$$\widehat{\mathtt{m}}_{ai}(x_i) = \frac{1}{\hat{z}_{ai}}\sum_{\underline{x}_{\partial a\setminus i}}\mathbb{I}\left(x_{i_1(a)}\oplus\cdots\oplus x_{i_K(a)} = b_a\right)\prod_{j\in\partial a\setminus i}\mathtt{m}_{ja}(x_j)\,. \tag{19.42}$$

As in Sec. 18.6, we shall assume that messages can take only three values, which we denote by the shorthands: $\mathtt{m}_{ia} = \mathtt{0}$ if $(\mathtt{m}_{ia}(\mathtt{0}) = 1,\ \mathtt{m}_{ia}(\mathtt{1}) = 0)$; $\mathtt{m}_{ia} = \mathtt{1}$ if $(\mathtt{m}_{ia}(\mathtt{0}) = 0,\ \mathtt{m}_{ia}(\mathtt{1}) = 1)$; $\mathtt{m}_{ia} = *$ if $(\mathtt{m}_{ia}(\mathtt{0}) = \mathtt{m}_{ia}(\mathtt{1}) = 1/2)$.

Consider the first BP equation (19.41), and denote by $n_{\mathtt{0}}$, $n_{\mathtt{1}}$, n_* the number of messages of type $\mathtt{0}$, $\mathtt{1}$, $*$ in the set of incoming messages $\{\widehat{\mathtt{m}}_{bi}\}$, $b \in \partial i \backslash a$. Then Eq. (19.41) can be rewritten as:

$$\mathtt{m}_{ia} = \begin{cases} \mathtt{0} & \text{if } n_{\mathtt{0}} > 0,\ n_{\mathtt{1}} = 0, \\ \mathtt{1} & \text{if } n_{\mathtt{0}} = 0,\ n_{\mathtt{1}} > 0, \\ * & \text{if } n_{\mathtt{0}} = 0,\ n_{\mathtt{1}} = 0, \\ ? & \text{if } n_{\mathtt{0}} > 0,\ n_{\mathtt{1}} > 0, \end{cases} \qquad z_{ia} = \begin{cases} 2^{-n_*} & \text{if } n_{\mathtt{0}} > 0,\ n_{\mathtt{1}} = 0, \\ 2^{-n_*} & \text{if } n_{\mathtt{0}} = 0,\ n_{\mathtt{1}} > 0, \\ 2^{1-n_*} & \text{if } n_{\mathtt{0}} = 0,\ n_{\mathtt{1}} = 0, \\ 0 & \text{if } n_{\mathtt{0}} > 0,\ n_{\mathtt{1}} > 0. \end{cases} \tag{19.43}$$

The computation of the normalization constant z_{ia} will be useful in the 1RSB analysis. Notice that, if $n_{\mathtt{0}} > 0$ and $n_{\mathtt{1}} > 0$, a contradiction arises at node i and therefore $\mathtt{m}_{ia}$ is not defined. However we will see that, because in this case $z_{ia} = 0$, this situation does not create any problem within 1RSB.

In the second BP equation (19.42) denote by $\widehat{n}_{\mathtt{0}}$ (respectively, $\widehat{n}_{\mathtt{1}}$, $\widehat{n}_*$) the number of messages of type $\mathtt{0}$ (resp. $\mathtt{1}$, $*$) among $\{\mathtt{m}_{ja}\}$, $j \in \partial a \backslash i$. Then we get

$$\widehat{\mathtt{m}}_{ai} = \begin{cases} \mathtt{0} & \text{if } n_* = 0, \text{ and } n_{\mathtt{1}} \text{ has the same parity as } b_a, \\ \mathtt{1} & \text{if } n_* = 0, \text{ and } n_{\mathtt{1}} \text{ has not the same parity as } b_a, \\ * & \text{if } n_* > 0. \end{cases} \tag{19.44}$$

In all three cases $\hat{z}_{ai} = 1$.

In Sec. 18.6 we studied the equations (19.41), (19.42) above and deduced that, for typical random instances with $\alpha = M/N < \alpha_{\mathrm{d}}(K)$, they have a unique solution, with $\mathtt{m}_{ia} = \widehat{\mathtt{m}}_{ai} = *$ on each edge.

Exercise 19.2 Evaluate the Bethe free-entropy on this solution, and show that it yields the free-entropy density $\mathrm{f}^{\mathrm{RS}} = (1-\alpha)\log 2$.

19.3.2 *The 1RSB cavity equations*

We now assume that the BP equations (19.43), (19.44) have many solutions, and apply the 1RSB cavity method to study their statistics.

The 1RSB messages Q_{ia}, $\widehat{Q}_{ai}$ are distributions over $\{\mathtt{0}, \mathtt{1}, *\}$. A little effort shows that Eq. (19.28) yields

$$Q_{ia}(\mathtt{0}) = \frac{1}{Z_{ia}} \left\{ \prod_{b \in \partial i \backslash a} \left(\widehat{Q}_{bi}(\mathtt{0}) + 2^{-\mathtt{x}} \widehat{Q}_{bi}(*) \right) - \prod_{b \in \partial i \backslash a} \left(2^{-\mathtt{x}} \widehat{Q}_{bi}(*) \right) \right\}, \tag{19.45}$$

$$Q_{ia}(\mathtt{1}) = \frac{1}{Z_{ia}} \left\{ \prod_{b \in \partial i \backslash a} \left(\widehat{Q}_{bi}(\mathtt{1}) + 2^{-\mathtt{x}} \widehat{Q}_{bi}(*) \right) - \prod_{b \in \partial i \backslash a} \left(2^{-\mathtt{x}} \widehat{Q}_{bi}(*) \right) \right\}, \tag{19.46}$$

$$Q_{ia}(*) = \frac{1}{Z_{ia}}\, 2^{\mathtt{x}} \prod_{b \in \partial i \backslash a} 2^{-\mathtt{x}} \widehat{Q}_{bi}(*) \,. \tag{19.47}$$

For instance, Eq. (19.45) follows from the first line of Eq. (19.43): $\mathtt{m}_{ia} = \mathtt{0}$ if all the incoming messages are $\widehat{\mathtt{m}}_{bi} \in \{*, \mathtt{0}\}$ (first term), unless they are all equal to $*$ (subtracted term). The re-weighting $z_{ia}^{\mathtt{x}} = 2^{-\mathtt{x}n_*}$ decomposes into factors associated to the incoming $*$ messages.

The second group of 1RSB equations, Eq. (19.29), takes the form:

$$\widehat{Q}_{ai}(\mathtt{0}) = \frac{1}{2}\left\{\prod_{j\in\partial a\setminus i}(Q_{ja}(\mathtt{0})+Q_{ja}(\mathtt{1})) + s(b_a)\prod_{j\in\partial a\setminus i}(Q_{ja}(\mathtt{0})-Q_{ja}(\mathtt{1}))\right\}, \tag{19.48}$$

$$\widehat{Q}_{ai}(\mathtt{1}) = \frac{1}{2}\left\{\prod_{j\in\partial a\setminus i}(Q_{ja}(\mathtt{0})+Q_{ja}(\mathtt{1})) - s(b_a)\prod_{j\in\partial a\setminus i}(Q_{ja}(\mathtt{0})-Q_{ja}(\mathtt{1}))\right\}, \tag{19.49}$$

$$\widehat{Q}_{ai}(*) = 1 - \prod_{j\in\partial a\setminus i}(Q_{ja}(\mathtt{0})+Q_{ja}(\mathtt{1}))\,, \tag{19.50}$$

where $s(b_a) = +1$ if $b_a = \mathtt{0}$, and $s(b_a) = -1$ otherwise.

Notice that, if one takes $\mathtt{x} = 0$, the two sets of equations coincide with those obtained in Sec. 18.6, see Eq. (18.35) (the homogeneous linear system, $b_a = \mathtt{0}$, was considered there). As in that section, we look for solutions such that the messages $Q_{ia}(\,\cdot\,)$ (respectively $\widehat{Q}_{ai}(\,\cdot\,)$) take two possible values: either $Q_{ia}(\mathtt{0}) = Q_{ia}(\mathtt{1}) = 1/2$, or $Q_{ia}(*) = 1$. This assumption is consistent with the 1RSB cavity equations (19.45) and (19.50). Under this assumption, the $\mathtt{x}$ dependency drops from these equations and we recover the analysis in Sec. 18.6. In particular, we can repeat the density evolution analysis discussed there. If we denote by Q_* the probability that a randomly chosen edge carries the 1RSB message $Q_{ia}(\mathtt{0}) = Q_{ia}(\mathtt{1}) = 1/2$, then the fixed point equation of density evolution reads:

$$Q_* = 1 - \exp\{-k\alpha Q_*^{k-1}\}\,. \tag{19.51}$$

For $\alpha < \alpha_{\rm d}(K)$ this equation admits the only solution $Q_* = 0$, implying $Q_{ia}(*) = 1$ with high probability. This indicates (once more) that the only solution of the BP equations in this regime is $\mathtt{m}_{ia} = *$ for all $(i,a) \in E$.

For $\alpha > \alpha_{\rm d}$ a couple of non-trivial solutions (with $Q_* > 0$) appear, indicating the existence of a large number of BP fixed points (and hence, Bethe measures). Stability under density evolution suggest to select the largest one. It will also be useful in the following to introduce the probability

$$\widehat{Q}_* = Q_*^{k-1} \tag{19.52}$$

that a uniformly random edge carries a message $\widehat{Q}_{ai}(\mathtt{0}) = \widehat{Q}_{ai}(\mathtt{1}) = 1/2$.

19.3.3 *Complexity*

We can now compute the Bethe free-entropy (19.33) of the auxiliary graphical model. The complexity will be computed through the Legendre transform of the 1RSB free-entropy, see Eq. (19.8).

Let us start by computing the contribution $\mathbb{F}_a^{\rm RSB}$ defined in Eq. (19.30). Consider the weight

$$e^{\mathbb{F}_a(\{\mathtt{m}_{ia}\})} = \sum_{\underline{x}_{\partial a}} \mathbb{I}(x_{i_1(a)} \oplus \cdots \oplus x_{i_K(a)} = b_a) \prod_{i\in\partial a} \mathtt{m}_{ia}(x_i)\,. \tag{19.53}$$

Let $\widehat{n}_{\mathtt{0}}$ (respectively, $\widehat{n}_{\mathtt{1}}$, $\widehat{n}_*$) denote the number of variable nodes $i \in \partial a$ such that $\mathtt{m}_{ia} = \mathtt{0}$ (resp. $\mathtt{1}$, $*$) for $i \in \partial a$. Then we get

$$e^{\mathbb{F}_a(\{\mathtt{m}_{ia}\})} = \begin{cases} 1/2 & \text{if } \widehat{n}_* > 0, \\ 1 & \text{if } \widehat{n}_* = 0 \text{ and } \widehat{n}_{\mathtt{1}} \text{ has the same parity as } b_a, \\ 0 & \text{if } \widehat{n}_* = 0 \text{ and } \widehat{n}_{\mathtt{1}} \text{ has not the same parity as } b_a, \end{cases} \tag{19.54}$$

Taking the expectation of $e^{\mathtt{x}\mathbb{F}_a(\{\mathtt{m}_{ia}\})}$ with respect to $\{\mathtt{m}_{ia}\}$ distributed independently according to $Q_{ia}(\,\cdot\,)$, and assuming $Q_{ia}(\mathtt{0}) = Q_{ia}(\mathtt{1})$ (which is the case in our solution), we get

$$\mathbb{F}_a^{\rm RSB} = \log\left\{ \frac{1}{2}\prod_{i\in\partial a}(1 - Q_{ia}(*)) + \frac{1}{2^{\mathtt{x}}}\left[1 - \prod_{i\in\partial a}(1-Q_{ia}(*))\right]\right\}\,. \tag{19.55}$$

The first term corresponds to the case $\widehat{n}_* = 0$ (the factor 1/2 being the probability that the parity of $\widehat{n}_{\mathtt{1}}$ is b_a), and the second to $\widehat{n}_* > 0$. Within our solution either $Q_{ia}(*) = 0$ or $Q_{ia}(*) = 1$. Therefore only one of the above terms survives: either $Q_{ia}(*) = 0$ for all $i \in \partial a$, yielding $\mathbb{F}_a^{\rm RSB} = -\log 2$, or $Q_{ia}(*) = 1$ for some $i \in \partial a$, implying $\mathbb{F}_a^{\rm RSB} = -\mathtt{x}\log 2$.

Until now we considered a generic K-XORSAT instance. For random instances, we can take the expectation with respect to $Q_{ia}(*)$ independently distributed as in the density evolution fixed point. The first case, namely $Q_{ia}(*) = 0$ for all $i \in \partial a$ (and thus $\mathbb{F}_a^{\rm RSB} = -\log 2$), occurs with probability Q_*^k. The second, i.e. $Q_{ia}(*) = 1$ for some $i \in \partial a$ (and $\mathbb{F}_a^{\rm RSB} = -\mathtt{x}\log 2$), occurs with probability $1 - Q_*^k$. Altogether we obtain:

$$\mathbb{E}\{\mathbb{F}_a^{\rm RSB}\} = -\left[Q_*^k + \mathtt{x}(1 - Q_*^k)\right]\,\log 2 + o_N(1)\,. \tag{19.56}$$

Assuming the messages $Q_{ia}(\,\cdot\,)$ to be short-range correlated, $\sum_{a\in F}\mathbb{F}_a^{\rm RSB}$ will concentrate around its expectation. We then have, with high probability,

$$\frac{1}{N}\sum_{a\in F}\mathbb{F}_a^{\rm RSB} = -\alpha\left[Q_*^k + x(1-Q_*^k)\right]\,\log 2 + o_N(1)\,. \tag{19.57}$$

The contributions from variable node and edge terms can be computed along similar lines. We will just sketch these computations, and invite the reader to work out the details.

Consider the contribution $\mathbb{F}_i^{\rm RSB}$, $i \in V$, defined in (19.31). Assume that $\widehat{Q}_{ai}(*) = 1$ (respectively, $\widehat{Q}_{ai}(\mathtt{0}) = \widehat{Q}_{ai}(\mathtt{1}) = 1/2$) for n_* (resp. $n_{\mathtt{0}}$) of the neighboring function nodes $a \in \partial i$. Then $\mathbb{F}_i^{\rm RSB} = -(n_*\mathtt{x} + n_{\mathtt{0}} - 1)\log 2$ if $n_{\mathtt{0}} \geq 1$, and

$\mathbb{F}_i^{\rm RSB} = -(n_* - 1)\mathtt{x}\log 2$ otherwise. Averaging these expressions over n_0 (a Poisson distributed random variable with mean $k\alpha\widehat{Q}_*$) and n_* (Poisson with mean $k\alpha(1-\widehat{Q}_*)$) we obtain:

$$\frac{1}{N}\sum_{i\in V}\mathbb{F}_i^{\rm RSB} = -\left\{\left[k\alpha\widehat{Q}_* - 1 + e^{-k\alpha\widehat{Q}_*}\right] + \left[k\alpha(1-\widehat{Q}_*) - e^{-k\alpha\widehat{Q}_*}\right]\mathtt{x}\right\}\log 2 + o_N(1)\, . \tag{19.58}$$

Let us finally consider the edge contribution $\mathbb{F}_{(ia)}^{\rm RSB}$ defined in (19.32). If $Q_{ia}(\mathtt{0}) = Q_{ia}(\mathtt{1}) = 1/2$ and $\widehat{Q}_{ai}(\mathtt{0}) = \widehat{Q}_{ai}(\mathtt{1}) = 1/2$, then either $e^{\mathbb{F}_{ai}} = 1$ or $e^{\mathbb{F}_{ia}} = 0$, each with probability 1/2. As a consequence $\mathbb{F}_{(ia)}^{\rm RSB} = -\log 2$. If either $Q_{ia}(*) = 1$ or $\widehat{Q}_{ai}(*) = 1$ (or both), $e^{\mathbb{F}_{ia}^{\rm RSB}} = 1/2$ with probability 1, and therefore $\mathbb{F}_{(ia)}^{\rm RSB} = -\mathtt{x}\log 2$. Altogether we obtain, with high probability

$$\frac{1}{N}\sum_{(ia)\in E}\mathbb{F}_{(ia)}^{\rm RSB} = -k\alpha\left\{Q_*\widehat{Q}_* + (1-Q_*\widehat{Q}_*)\mathtt{x}\right\}\log 2 + o_N(1). \tag{19.59}$$

The free-entropy (19.33) of the auxiliary graphical model is obtained by collecting the various terms. We obtain $\mathbb{F}^{\rm RSB}(\mathtt{x}) = N\mathrm{f}^{\rm RSB}(\mathtt{x}) + o(N)$ where $\mathrm{f}^{\rm RSB}(\mathtt{x}) = [\Sigma_{\rm tot} + \mathtt{x}\,\phi_{\rm typ}]\log 2$ and

$$\Sigma_{\rm tot} = k\alpha Q_*\widehat{Q}_* - k\alpha\widehat{Q}_* - \alpha Q_*^k + 1 - e^{-k\alpha\widehat{Q}_*}\, , \tag{19.60}$$

$$\phi_{\rm typ} = -k\alpha Q_*\widehat{Q}_* + k\alpha\widehat{Q}_* + \alpha Q_*^k - \alpha + e^{-k\alpha\widehat{Q}_*}\, . \tag{19.61}$$

Here Q_* is the largest solution of Eq. (19.51) and $\widehat{Q}_* = Q_*^{k-1}$, a condition that has a pleasing interpretation as shown in the exercise below.

Exercise 19.3 Consider the function $\Sigma_{\rm tot}(Q_*, \widehat{Q}_*)$ defined in Eq. (19.60). Show that the stationary points of this function coincide with the solutions of Eq. (19.51) and $\widehat{Q}_* = Q_*^{k-1}$.

Because of the linear dependence on x, the Legendre transform (19.8) is straightforward

$$\Sigma(\phi) = \begin{cases} \Sigma_{\rm tot} & \text{if } \phi = \phi_{\rm typ}, \\ -\infty & \text{otherwise.} \end{cases} \tag{19.62}$$

This means that there are $2^{N\Sigma_{\rm tot}}$ Bethe measures which all have the same entropy $N\phi_{\rm typ}\log 2$. Furthermore, $\Sigma_{\rm tot} + \phi_{\rm typ} = 1 - \alpha$, confirming that the total entropy is $(1-\alpha)\log 2$. This identity can be also written in the form

$$\frac{1}{2^{N(1-\alpha)}} = \frac{1}{2^{N\Sigma_{\rm tot}}} \times \frac{1}{2^{N\phi_{\rm typ}}}\, , \tag{19.63}$$

which is nothing but the decomposition (19.6) in extremal Bethe measures. Indeed, if $\underline{x}$ is a solution of the linear system, $\mu(\underline{x}) = 1/2^{N(1-\alpha)}$, $w_n \approx 1/2^{N\Sigma_{\rm tot}}$, and (assuming the μ^n to have disjoint supports) $\mu^n(\underline{x}) \approx 1/2^{N\phi_{\rm typ}}$ for the state n which contains $\underline{x}$.

Note that the value of Σ that we find here coincides with the result that we obtained in Sec. 18.5 for the logarithm of the number of clusters in random XORSAT formulae. This provides an independent check of our assumptions, and in particular it shows that the number of clusters is, to leading order, the same as the number of Bethe measures. In particular, the SAT-UNSAT transition occurs at the value of α where the complexity $\Sigma_{\rm tot}$ vanishes. At this value each cluster still contains a large number, $2^{N(1-\alpha_s)}$, of configurations.

Exercise 19.4 Repeat this 1RSB cavity analysis for a linear Boolean system described by a factor graph from the ensemble $\mathbb{D}_N(\Lambda, P)$ (This means a random system of linear equations, whereby the fraction of equations involving k variables is P_k, and the fraction of variables which appear in exactly ℓ equations is Λ_ℓ):

(a) Show that Q_* and $\widehat{Q}_*$ satisfy:

$$\widehat{Q}_* = \rho(Q_*) \quad ; \quad Q_* = 1 - \lambda(1 - \widehat{Q}_*) \ , \tag{19.64}$$

where λ and ρ are the edge perspective degree profiles.

(b) Show that the complexity is given by

$$\Sigma_{\rm tot} = 1 - \frac{\Lambda'(1)}{P'(1)}\, P(Q_*) - \Lambda(1 - \widehat{Q}_*) - \Lambda'(1)(1 - Q_*)\widehat{Q}_* \tag{19.65}$$

and the internal entropy of the clusters is $\phi_{\rm typ} = 1 - \Lambda'(1)/P'(1) - \Sigma_{\rm tot}$.

(c) In the case where all variables have degree strictly larger than 1 (so that $\lambda(0) = 0$), argue that the relevant solution is $Q_* = \widehat{Q}_* = 1$, $\Sigma_{\rm tot} = 1 - \Lambda'(1)/P'(1)$, $\phi_{\rm typ} = 0$. What is the interpretation of this result in terms of the core structure discussed in Sec. 18.3?

19.4 The special value $\mathtt{x} = 1$

Let us return to the general formalism. The $\mathtt{x} = 1$ case plays a special role, in that the weights $\{w_n(\mathtt{x})\}$ of various Bethe measures in the auxiliary model, coincide with the ones appearing in the decomposition (19.6). This fact manifests itself in some remarkable properties of the 1RSB formalism.

19.4.1 *Back to BP*

Consider the general 1RSB cavity equations (19.28), (19.29). Using the explicit form of the re-weighting factors $e^{\mathbb{F}_i - \mathbb{F}_{ia}}$ and $e^{\mathbb{F}_a - \mathbb{F}_{ia}}$ provided in Eqs. (19.23), (19.24), they can be written, for $\mathtt{x} = 1$, as:

$$Q_{ia}(\mathtt{m}_{ia}) \cong \sum_{x_i} \sum_{\{\widehat{\mathtt{m}}_{bi}\}} \mathbb{I}\left(\mathtt{m}_{ia} = g_i(\{\widehat{\mathtt{m}}_{bi}\})\right) \prod_{b\in\partial i\setminus a} \widehat{Q}_{bi}(\widehat{\mathtt{m}}_{bi})\, \widehat{\mathtt{m}}_{bi}(x_i)\,, \tag{19.66}$$

$$\widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai}) \cong \sum_{\underline{x}_{\partial a}} \psi_a(\underline{x}_{\partial a}) \sum_{\{\mathtt{m}_{ja}\}} \mathbb{I}\left(\widehat{\mathtt{m}}_{ai} = f_a(\{\mathtt{m}_{ja}\})\right) \prod_{j\in\partial a\setminus i} Q_{ja}(\mathtt{m}_{ja})\, \mathtt{m}_{ja}(x_j)\,. \tag{19.67}$$

Let us introduce the messages obtained by taking the averages of the 1RSB ones $\{Q_{ia}, \widehat{Q}_{ai}\}$:

$$\nu^{\rm av}_{i\to a}(x_i) \equiv \sum_{\mathtt{m}_{ia}} Q_{ia}(\mathtt{m}_{ia})\, \mathtt{m}_{ia}(x_i)\,, \quad \widehat{\nu}^{\rm av}_{a\to i}(x_i) \equiv \sum_{\widehat{\mathtt{m}}_{ai}} \widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai})\, \widehat{\mathtt{m}}_{ai}(x_i)\,.$$

The interpretation of these quantities is straightforward. Given an extremal Bethe measure sampled according to the distribution w_n, let $\nu^n_{i\to a}(\,\cdot\,)$ (respectively $\widehat{\nu}^n_{a\to i}(\,\cdot\,)$) be the corresponding message along the directed edge $i \to a$ (resp. $a \to i$). Its expectation, with respect to the random choice of the measure, is $\nu^{\rm av}_{i\to a}(\,\cdot\,)$ (respectively $\widehat{\nu}^{\rm av}_{a\to i}(\,\cdot\,)$).

Using the expressions (19.9), one finds that Eqs. (19.66), (19.67) imply

$$\nu^{\rm av}_{i\to a}(x_i) \cong \prod_{b\in\partial i\setminus a} \widehat{\nu}^{\rm av}_{b\to i}(x_i)\,, \tag{19.68}$$

$$\widehat{\nu}^{\rm av}_{a\to i}(x_i) \cong \sum_{\{x_j\}_{j\in\partial a\setminus i}} \psi_a(\underline{x}_{\partial a}) \prod_{j\in\partial a\setminus i} \nu^{\rm av}_{j\to a}(x_j)\,, \tag{19.69}$$

which are nothing but the ordinary BP equations. This suggests that, even if $\mu(\,\cdot\,)$ decomposes into an exponential number of extremal Bethe measures $\mu^n(\,\cdot\,)$, it is itself a (non-extremal) Bethe measure. In particular, there exists a quasi-solution of BP equations associated with it, that allows to compute its marginals.

The reader might be disappointed by these remarks. Why insisting on the 1RSB cavity approach if, when the 'correct' weights are used, one recovers the much simpler BP equations? There are at least two answers:

1. The 1RSB approach provides a much more refined picture: decomposition in extremal Bethe states, long range correlations, complexity. This is useful and interesting *per se.*
2. In the cases of a static (s1RSB) phase, it turns out that the region $\mathtt{x} = 1$ corresponds to an 'unphysical' solution of the 1RSB cavity equations, and that (asymptotically) correct marginals are instead obtained by letting $\mathtt{x} = \mathtt{x}_*$, for some $\mathtt{x}_* \in [0, 1)$. In such cases it is mandatory to resort to the full 1RSB formalism (see Sec. 19.6 below).

19.4.2 *A simpler recursion*

As we stressed above, controlling (either numerically or analytically) the 1RSB distributional recursions (19.35), (19.36) is a difficult task. In the case $\mathtt{x} = 1$, they simplify considerably and lend themselves to a much more accurate numerical study. This remark can be very useful in practice.

As in Sec. 19.2.5, we consider a random graphical model. We shall also assume a 'local uniformity condition.' More precisely, the original model $\mu(\,\cdot\,)$ is a Bethe measure for the message set $\nu^{\rm av}_{i\to a}(x_i) = 1/q$ and $\widehat{\nu}^{\rm av}_{a\to i}(x_i) = 1/q$, where $q = |\mathcal{X}|$ is the size of the alphabet. While such a local uniformity condition is not necessary, it considerably simplify the derivation below. The reader can find a more general treatment in the literature.

Consider Eqs. (19.35) and (19.36) at $\mathtt{x} = 1$. The normalization constants can be easily computed using the uniformity condition. We can then average over the structure of the graph, and the function node distribution: let us denote by $Q_{\rm av}$ and $\widehat{Q}_{\rm av}$ the averaged distributions. They satisfy the following equations:

$$Q^{(t+1)}_{\rm av}(\mathtt{m}) = \mathbb{E}\left\{ q^{l-2} \sum_{\{\widehat{\mathtt{m}}_b\}} \mathbb{I}\left(\mathtt{m} = \mathtt{f}(\{\widehat{\mathtt{m}}_b\}; J)\right)\ z(\{\widehat{\mathtt{m}}_b\}) \prod_{b=1}^{l-1} \widehat{Q}^{(t)}_{\rm av}(\widehat{\mathtt{m}}_b) \right\}\,, \tag{19.70}$$

$$\widehat{Q}^{(t)}_{\rm av}(\widehat{\mathtt{m}}) = \mathbb{E}\left\{ \frac{q^{k-2}}{\overline{\psi}_k} \sum_{\{\mathtt{m}_j\}} \mathbb{I}\left(\widehat{\mathtt{m}} = \hat{\mathtt{f}}(\{\mathtt{m}_j\}; \widehat{J})\right)\ \hat{z}(\{\mathtt{m}_j\}; \widehat{J}) \prod_{j=1}^{k-1} Q^{(t)}_{\rm av}(\mathtt{m}_j) \right\}\,, \tag{19.71}$$

where expectations are taken over $l, k, J, \widehat{J}$, distributed according to the random graphical model. Here $\overline{\psi}_k = \sum_{x_1,\dots,x_{k-1}} \psi(x_1,\dots,x_{k-1},x;\widehat{J})$ can be shown to be independent of x (this is necessary for the uniformity condition to hold).

Equations (19.70) and (19.71) are considerably simpler that the original distributional equations (19.35), (19.36) in that $Q^{(t)}_{\rm av}(\,\cdot\,)$, $\widehat{Q}^{(t)}_{\rm av}(\,\cdot\,)$ are non-random. On the other hand, they still involve a reweighting factor that is difficult to handle. It turns out that this reweighting can be eliminated by introducing a new couple of distributions for each $x \in \mathcal{X}$:

$$\widehat{R}^{(t)}_x(m) \equiv q\, m(x)\, \widehat{Q}^{(t)}_{\rm av}(m)\,, \qquad R^{(t)}_x(m) = q\, m(x)\, Q^{(t)}_{\rm av}(m)\,. \tag{19.72}$$

One can show that Eqs. (19.70), (19.71) imply the following recursions for $R^{(t)}_x$, $\widehat{R}^{(t)}_x$,

$$R^{(t+1)}_x(\mathtt{m}) = \mathbb{E}\left\{ \sum_{\{\widehat{\mathtt{m}}_b\}} \mathbb{I}\left(\mathtt{m} = g(\{\widehat{\mathtt{m}}_b\}; J)\right) \prod_{b=1}^{l-1} \widehat{R}^{(t)}_x(\widehat{\mathtt{m}}_b) \right\}\,, \tag{19.73}$$

$$\widehat{R}^{(t)}_x(\widehat{\mathtt{m}}) = \mathbb{E}\left\{ \sum_{\{x_j\}} \pi(\{x_j\}|x;\widehat{J}) \sum_{\{\mathtt{m}_j\}} \mathbb{I}\left(\widehat{\mathtt{m}} = f(\{\mathtt{m}_j\}; \widehat{J})\right) \prod_{j=1}^{k-1} R^{(t)}_{x_j}(\mathtt{m}_j) \right\} \tag{19.74}$$

Here $\mathbb{E}$ denotes expectation with respect to $l, \widehat{J}, k, J$ and, for any x, $\widehat{J}$, the distribution $\pi(\{x_j\}|x;\widehat{J})$ is defined by

$$\pi(x_1,\dots,x_{k-1}|x;\widehat{J}) = \frac{\psi(x_1,\dots,x_{k-1},x;\widehat{J})}{\sum_{y_1,\dots,y_{k-1}} \psi(y_1,\dots,y_{k-1},x;\widehat{J})}\,. \tag{19.75}$$

Exercise 19.5 Prove formulas (19.73) and (19.74). It might be useful to recall the following explicit expressions for the reweighting factors z and $\hat{z}$:

$$z(\{\widehat{\mathtt{m}}_b\})\, \mathtt{m}(x) = \prod_{b=1}^{l-1} \widehat{\mathtt{m}}_b(x)\,, \tag{19.76}$$

$$\hat{z}(\{\mathtt{m}_j\}; \widehat{J})\, \widehat{\mathtt{m}}(x) = \sum_{\{x_i\},x} \psi(x_1, \dots, x_{k-1}, x; \widehat{J}) \prod_{j=1}^{k-1} \mathtt{m}_j(x_j)\,. \tag{19.77}$$

The equations (19.73), (19.74) have a simple operational description. Let $\widehat{J}$ and k be drawn according to their distribution, and, given x, generate $x_1, \dots, x_{k-1}$ according to the kernel $\pi(x_1, \dots, x_k|x; \widehat{J})$. Then draw independent messages $\mathtt{m}_1, \dots, \mathtt{m}_{k-1}$ with distribution (respectively) $R^{(t)}_{x_1}, \dots, R^{(t)}_{x_{k-1}}$. According to Eq. (19.74), $\widehat{\mathtt{m}} = f(\{\mathtt{m}_j\}; \widehat{J})$ has then distribution $\widehat{R}^{(t)}_x$. For Eq. (19.73), draw J and l according to their distribution. Given x, draw $l-1$ i.i.d. messages $\widehat{\mathtt{m}}_1, \dots, \widehat{\mathtt{m}}_{l-1}$ with distribution $\widehat{R}^{(t)}_x$. Them $\mathtt{m} = g(\{\widehat{\mathtt{m}}_b\}; J)$ has distribution $R^{(t+1)}_x$.

We will see in Ch. 22 that this procedure does indeed coincide with the one for computing 'point-to-set correlations' with respect to the measure $\mu(\,\cdot\,)$.

To summarize, for $\mathtt{x} = 1$ we have succeeded in simplifying the 1RSB density evolution equations in two directions: (*i*) The resulting equations do not involve 'distributions of distributions;' (*ii*) We got rid of the reweighting factor. A third crucial simplification is the following:

Theorem 19.5 *The 1RSB equations have a non trivial solution (meaning a solution different from the RS one) if and only if Eqs. (19.73), (19.74), when initialized with $R^{(0)}_x$ being a singleton distribution on $\mathtt{m}(y) = \mathbb{I}(y = x)$, converge as $t \to \infty$, to a non-trivial distribution.*

This theorem resolves (in the case $\mathtt{x} = 1$) the ambiguity on the initial condition of the 1RSB iteration. In other words, if the 1RSB equations admit a non-trivial solution, it can be reached if we iterate the equations starting from the initial condition mentioned in the theorem. We refer the reader to the literature for the proof.

Exercise 19.6 Show that the free-entropy of the auxiliary model $\mathbb{F}^{\rm RSB}(\mathtt{x})$, evaluated at $\mathtt{x} = 1$, coincides with the RS Bethe free-entropy.

Further, its derivative with respect to $\mathtt{x}$ at $\mathtt{x} = 1$ can be expressed in terms of the fixed point distributions $R^{(\infty)}_x$ and $\widehat{R}^{(\infty)}_x$. In particular the complexity and internal free-entropy can be computed from the fixed points of the simplified equations (19.73), (19.74).

The conclusion of this section is that 1RSB calculations at $\mathtt{x} = 1$ are not technically harder that RS ones. In view of the special role played by the value

$\mathtt{x} = 1$ this observation can be exploited in a number of contexts.

19.5 Survey propagation

The 1RSB cavity method can be applied to other message passing algorithms whenever these have many fixed points. A particularly important case is the min-sum algorithm of Sec. 14.3. This approach (both in its RS and 1RSB versions) is sometimes referred to as the **energetic cavity method** because, in physics terms, the min-sum algorithm aims at computing the ground state configuration and its energy. We will call the corresponding 1RSB message passing algorithm $\mathsf{SP}(\mathtt{y})$ (survey propagation at finite $\mathtt{y}$).

19.5.1 *The* $\mathsf{SP}(\mathtt{y})$ *equations*

The formalism follows closely the one used with BP solutions. To emphasize the similarities, let us adopt the same notation for the min-sum messages as for the BP ones. We define

$$\mathtt{m}_{ja}(x_j) \equiv E_{i\to a}(x_i)\,, \quad \widehat{\mathtt{m}}_{ai}(x_i) \equiv \widehat{E}_{a\to i}(x_i)\ , \tag{19.78}$$

and write the min-sum equations (14.40), (14.41) as:

$$\mathtt{m}_{ia} = \mathsf{f}^{\mathrm{e}}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}\right)\ , \qquad \widehat{\mathtt{m}}_{ai} = \hat{\mathsf{f}}^{\mathrm{e}}_a\left(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}\right)\ . \tag{19.79}$$

The functions $\mathsf{f}^{\mathrm{e}}_i$, $\hat{\mathsf{f}}^{\mathrm{e}}_a$ are defined by Eqs. (14.40), (14.41), that we reproduce here:

$$\mathtt{m}_{ia}(x_i) = \sum_{b\in\partial i\setminus a} \widehat{\mathtt{m}}_{bi}(x_i) - u_{ia}\,, \tag{19.80}$$

$$\widehat{\mathtt{m}}_{ai}(x_i) = \min_{\underline{x}_{\partial a\setminus i}}\left[E_a(\underline{x}_{\partial a}) + \sum_{j\in\partial a\setminus i}\mathtt{m}_{ja}(x_j)\right] - \hat{u}_{ai}\ , \tag{19.81}$$

where $u_{ia}, \hat{u}_{ai}$ are normalization constants (independent of x_i) which ensure that $\min_{x_i}\widehat{\mathtt{m}}_{ai}(x_i) = 0$ and $\min_{x_i}\mathtt{m}_{ia}(x_i) = 0$.

To any set of messages $\{\mathtt{m}_{ia}, \widehat{\mathtt{m}}_{ai}\}$, we associate the Bethe energy

$$\mathbb{F}^{\mathrm{e}}(\underline{\mathtt{m}}, \underline{\widehat{\mathtt{m}}}) = \sum_{a\in F}\mathbb{F}^{\mathrm{e}}_a(\{\mathtt{m}_{ia}\}_{i\in\partial a}) + \sum_{i\in V}\mathbb{F}^{\mathrm{e}}_i(\{\widehat{\mathtt{m}}_{ai}\}_{a\in\partial i}) - \sum_{(ia)\in E}\mathbb{F}^{\mathrm{e}}_{ia}(\mathtt{m}_{ia}, \widehat{\mathtt{m}}_{ai})\,, \tag{19.82}$$

where the various terms are (see Eq. (14.45)):

$$\begin{aligned}
\mathbb{F}^{\mathrm{e}}_a &= \min_{\underline{x}_{\partial a}}\Big[E_a(\underline{x}_{\partial a}) + \sum_{j\in\partial a}\mathtt{m}_{ia}(x_i)\Big]\,, \qquad \mathbb{F}^{\mathrm{e}}_i = \min_{x_i}\Big[\sum_{a\in\partial i}\widehat{\mathtt{m}}_{ai}(x_i)\Big]\,,\\
\mathbb{F}^{\mathrm{e}}_{ia} &= \min_{x_i}\big[\mathtt{m}_{ia}(x_i) + \widehat{\mathtt{m}}_{ai}(x_i)\big]\,.
\end{aligned} \tag{19.83}$$

Having set up the message passing algorithm and the associated energy functional, we can repeat the program developed in the previous Sections. In particular, in analogy with Assumption 1, we have the following

Assumption 4 *There exist exponentially many quasi-solutions* $\{\underline{\mathtt{m}}^n\}$ *of min-sum equations. The number of such solutions with Bethe energy* $\mathbb{F}^{\rm e}(\underline{\mathtt{m}}^n) \approx N\epsilon$ *is (to leading exponential order)* $\exp\{N\Sigma^{\rm e}(\epsilon)\}$*, where* $\Sigma^{\rm e}(\epsilon)$ *is the* ***energetic complexity function****.*

In order to estimate $\Sigma^{\rm e}(\epsilon)$, we introduce an auxiliary graphical model, whose variables are the min-sum messages $\{\mathtt{m}_{ia}, \widehat{\mathtt{m}}_{ai}\}$. These are forced to satisfy (within some accuracy) the min-sum equations (19.80), (19.81). Each solution is given a weight $e^{-\mathtt{y}\mathbb{F}^{\rm e}(\underline{\mathtt{m}},\underline{\widehat{\mathtt{m}}})}$, with $\mathtt{y} \in \mathbb{R}$. The corresponding distribution is:

$$\mathsf{P}_{\mathtt{y}}(\underline{\mathtt{m}}, \underline{\widehat{\mathtt{m}}}) = \frac{1}{\Xi(\mathtt{y})} \prod_{a\in F} \Psi_a(\{\mathtt{m}_{ja}, \widehat{\mathtt{m}}_{ja}\}_{j\in\partial a}) \prod_{i\in V} \Psi_i(\{\mathtt{m}_{ib}, \widehat{\mathtt{m}}_{ib}\}_{b\in\partial i}) \prod_{(ia)\in E} \Psi_{ia}(\mathtt{m}_{ia}, \widehat{\mathtt{m}}_{ia}) , \tag{19.84}$$

where:

$$\Psi_a = \prod_{i\in\partial a} \mathbb{I}\left(\widehat{\mathtt{m}}_{ai} = \hat{\mathsf{f}}^{\rm e}_a\left(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}\right)\right) \; e^{-\mathtt{y}\mathbb{F}^{\rm e}_a(\{\mathtt{m}_{ja}\}_{j\in\partial a})} , \tag{19.85}$$

$$\Psi_i = \prod_{a\in\partial i} \mathbb{I}\left(\mathtt{m}_{ia} = \mathsf{f}^{\rm e}_i\left(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}\right)\right) \; e^{-\mathtt{y}\mathbb{F}^{\rm e}_i(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i})} , \tag{19.86}$$

$$\Psi_{ia} = e^{\mathtt{y}\mathbb{F}^{\rm e}_{ia}(\mathtt{m}_{ia}, \widehat{\mathtt{m}}_{ai})} . \tag{19.87}$$

Since the auxiliary graphical model is again locally tree-like, we can hope to derive asymptotically exact results through belief propagation. Messages of the auxiliary problem, to be denoted as $Q_{ia}(\,\cdot\,)$, $\widehat{Q}_{ai}(\,\cdot\,)$, are distributions over the min-sum messages. The $\mathsf{SP}(\mathtt{y})$ equations are obtained by further making the independence assumption (19.27).

The reader has certainly noticed that the whole procedure is extremely close to our study in Sec. 19.2.2. We can apply our previous analysis verbatim to derive the $\mathsf{SP}(\mathtt{y})$ update equations. The only step that requires some care is the formulation of the proper analog of Lemma 19.3. This becomes:

Lemma 19.6 *Assume that* $\mathtt{m}_{ia}(x_i) + \widehat{\mathtt{m}}_{ai}(x_i) < \infty$ *for at least one value of* $x_i \in \mathcal{X}$*. If* $\mathtt{m}_{ia} = \mathsf{f}^{\rm e}_i(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a})$*, then*

$$\mathbb{F}^{\rm e}_i - \mathbb{F}^{\rm e}_{ia} = u_{ia}(\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}) \equiv \min_{x_i} \Big\{ \sum_{b\in\partial i\setminus a} \widehat{\mathtt{m}}_{bi}(x_i) \Big\} . \tag{19.88}$$

Analogously, if $\widehat{\mathtt{m}}_{ai} = f^{\rm e}_a(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i})$*, then*

$$\mathbb{F}^{\rm e}_a - \mathbb{F}^{\rm e}_{ia} = \hat{u}_{ai}(\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}) \equiv \min_{\underline{x}_{\partial a}} \Big\{ E_a(\underline{x}_{\partial a}) + \sum_{k\in\partial a\setminus i} \mathtt{m}_{ka}(x_k) \Big\} . \tag{19.89}$$

Using this lemma, the same derivation as in Sec. 19.2.2 leads to

Proposition 19.7 *The* $\mathsf{SP}(\mathtt{y})$ *equations are (with the shorthands* $\{\widehat{\mathtt{m}}_{bi}\}$ *for* $\{\widehat{\mathtt{m}}_{bi}\}_{b\in\partial i\setminus a}$ *and* $\{\mathtt{m}_{ja}\}$ *for* $\{\mathtt{m}_{ja}\}_{j\in\partial a\setminus i}$*):*

$$Q_{ia}(\mathtt{m}_{ia}) \cong \sum_{\{\widehat{\mathtt{m}}_{bi}\}} \mathbb{I}\left(\mathtt{m}_{ia} = g_i^{\rm e}(\{\widehat{\mathtt{m}}_{bi}\})\right) e^{-\mathtt{y} u_{ia}(\{\widehat{\mathtt{m}}_{bi}\})} \prod_{b\in\partial i\setminus a} \widehat{Q}_{bi}(\widehat{\mathtt{m}}_{bi}) \,, \qquad (19.90)$$

$$\widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai}) \cong \sum_{\{\mathtt{m}_{ja}\}} \mathbb{I}\left(\widehat{\mathtt{m}}_{ai} = f_a^{\rm e}(\{\mathtt{m}_{ja}\})\right) e^{-\mathtt{y}\hat{u}_{ai}(\{\mathtt{m}_{ja}\})} \prod_{j\in\partial a\setminus i} Q_{ja}(\mathtt{m}_{ja}) \,. \qquad (19.91)$$

In the following we shall reserve the name **survey propagation** (SP) for the $\mathtt{y} = \infty$ case of these equations.

19.5.2 *Energetic complexity*

The Bethe free-entropy for the auxiliary graphical model is given by

$$\mathbb{F}^{\rm RSB,e}(\{Q,\widehat{Q}\}) = \sum_{a\in F} \mathbb{F}_a^{\rm RSB,e} + \sum_{i\in V} \mathbb{F}_i^{\rm RSB,e} - \sum_{(ia)\in E} \mathbb{F}_{ia}^{\rm RSB,e} \,, \qquad (19.92)$$

and allows to count the number of min-sum fixed points. The various terms are formally identical to the ones in Eqs. (19.30), (19.31) and (19.32), provided $\mathbb{F}_\cdot(\,\cdot\,)$ is replaced everywhere by $-\mathbb{F}_\cdot^{\rm e}(\,\cdot\,)$ and $\mathtt{x}$ by $\mathtt{y}$. We reproduce them here for convenience:

$$\mathbb{F}_a^{\rm RSB,e} = \log\left\{ \sum_{\{\mathtt{m}_{ia}\}} e^{-\mathtt{y}\mathbb{F}_a^{\rm e}(\{\mathtt{m}_{ia}\})} \prod_{i\in\partial a} Q_{ia}(\mathtt{m}_{ia}) \right\} \,, \qquad (19.93)$$

$$\mathbb{F}_i^{\rm RSB,e} = \log\left\{ \sum_{\{\widehat{\mathtt{m}}_{ai}\}} e^{-\mathtt{y}\mathbb{F}_i^{\rm e}(\{\widehat{\mathtt{m}}_{ai}\})} \prod_{a\in\partial i} \widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai}) \right\} \,, \qquad (19.94)$$

$$\mathbb{F}_{ia}^{\rm RSB,e} = \log\left\{ \sum_{\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai}} e^{-\mathtt{y}\mathbb{F}_{ia}^{\rm e}(\mathtt{m}_{ia},\widehat{\mathtt{m}}_{ai})} Q_{ia}(\mathtt{m}_{ia})\widehat{Q}_{ai}(\widehat{\mathtt{m}}_{ai}) \right\} \,. \qquad (19.95)$$

Assuming that the Bethe free-entropy gives the correct free-entropy of the auxiliary model, the energetic complexity function $\Sigma^{\rm e}(\epsilon)$ can be computed from $\mathbb{F}^{\rm RSB,e}(\mathtt{y})$ through the Legendre transform: in the large N limit we expect $\mathbb{F}^{\rm RSB,e}(\{Q,\widehat{Q}\}) = N\mathfrak{F}^{\rm e}(\mathtt{y}) + o(N)$ where

$$\mathfrak{F}^{\rm e}(\{Q,\widehat{Q}\}) = \Sigma^{\rm e}(\epsilon) - \mathtt{y}\epsilon \,, \qquad \frac{\partial\Sigma^{\rm e}}{\partial\epsilon} = \mathtt{y} \,. \qquad (19.96)$$

Finally, the 1RSB population dynamics algorithm can be used to sample -approximately- the $\mathsf{SP}(\mathtt{y})$ messages in random graphical models.

19.5.3 *Constraint satisfaction and binary variables*

In Sec. 14.3.3 we noticed that the min-sum messages simplify significantly when one deals with constraint satisfaction problems. In such problems, the energy function takes the form $E(\underline{x}) = \sum_a E_a(\underline{x}_{\partial a})$, where $E_a(\underline{x}_{\partial a}) = 0$ if constraint a is satisfied by the assignment $\underline{x}$, and $E_a(\underline{x}_{\partial a}) = 1$ otherwise. As discussed in

Sec. 14.3.3 the min-sum equations then admit solutions with $\widehat{\mathtt{m}}_{ai}(x_i) \in \{0,1\}$. Furthermore, one does not need to keep track of the variable-to-function node messages $\mathtt{m}_{ia}(x_i)$, but only of their 'projection' on $\{0,1\}$.

In other words, in constraint satisfaction problems the min-sum messages take $2^{|\mathcal{X}|}-1$ possible values (the all-1 message cannot appear). As a consequence, the $\mathsf{SP}(\mathtt{y})$ messages $\widehat{Q}_{ai}(\,\cdot\,)$ and $Q_{ia}(\,\cdot\,)$ simplify considerably: they are points in the $(2^{|\mathcal{X}|}-1)$-dimensional simplex.

If the min-sum messages are interpreted in terms of warnings, as we did in Sec. 14.3.3, then $\mathsf{SP}(\mathtt{y})$ messages keep track of the warnings' statistics (over pure states). One can use this interpretation to derive directly the $\mathsf{SP}(\mathtt{y})$ update equations without going through the whole 1RSB formalism. Let us illustrate this approach on the important case of binary variables $|\mathcal{X}|=2$.

The min-sum messages $\widehat{\mathtt{m}}$ and $\mathtt{m}$ (once projected) can take three values: $(\widehat{\mathtt{m}}(0),\widehat{\mathtt{m}}(1)) \in \{(0,1),(1,0),(0,0)\}$. We shall denote them respectively as $\mathtt{0}$ (interpreted as a warning: "take value 0"), $\mathtt{1}$ (interpreted as a warning: "take value 1") and $*$ (interpreted as a warning: "you can take any value"). Warning propagation (WP) can be described in words as follows.

Consider the message from variable node i to function node a. This depends on all the messages to i from function nodes $b \in \partial i \setminus a$. Suppose that $\widehat{n}_{\mathtt{0}}$ (respectively, $\widehat{n}_{\mathtt{1}}$, $\widehat{n}_{*}$) of these messages are of type $\mathtt{0}$ (resp. $\mathtt{1}$, $*$) for $i \in \partial a$. If $\widehat{n}_{\mathtt{0}} > \widehat{n}_{\mathtt{1}}$, i sends to a a $\mathtt{0}$ message. If $\widehat{n}_{\mathtt{1}} > \widehat{n}_{\mathtt{0}}$, it sends to a a $\mathtt{1}$ message. If $\widehat{n}_{\mathtt{1}} = \widehat{n}_{\mathtt{0}}$, it send to a a $*$ message. The 'number of contradictions' among the messages that it receives is: $\mathbb{F}^{\rm e}_i - \mathbb{F}^{\rm e}_{ia} = u_{ia} = \min(\widehat{n}_{\mathtt{1}},\widehat{n}_{\mathtt{0}})$.

Now consider the message from function node a to variable node i. It depends on the ones coming from neighboring variables $j \in \partial a \setminus i$. Partition the neighbors into subsets $\mathcal{P}_*, \mathcal{P}_{\mathtt{0}}, \mathcal{P}_{\mathtt{1}}$, whereby $\mathcal{P}_{\mathtt{m}}$ is the set of indices j such that $\mathtt{m}_{ja} = \mathtt{m}$. For each value of $x_i \in \{0,1\}$, the algorithm computes the minimal value of $E_a(\underline{x}_{\partial a})$ such that the variables in $\mathcal{P}_{\mathtt{0}}$ (respectively, $\mathcal{P}_{\mathtt{1}}$) are fixed to 0 (resp. to 1). More explicitly, let us define a function $\Delta_{\mathcal{P}}(x_i)$ as follows:

$$\Delta_{\mathcal{P}}(x_i) = \min_{\{x_j\}_{j\in\mathcal{P}_*}} E_a(x_i, \{x_j\}_{j\in\mathcal{P}_*}, \{x_k = 0\}_{k\in\mathcal{P}_{\mathtt{0}}}, \{x_l = 1\}_{l\in\mathcal{P}_{\mathtt{1}}})\,. \tag{19.97}$$

The following table then gives the outgoing message $\widehat{\mathtt{m}}_{ai}$ and the number of contradictions at a, $\mathbb{F}^{\rm e}_a - \mathbb{F}^{\rm e}_{ai} = \hat{u}_{ai}$ as a function of the values $\Delta_{\mathcal{P}}(0)$ and $\Delta_{\mathcal{P}}(1)$:

$\Delta_{\mathcal{P}}(0)$	$\Delta_{\mathcal{P}}(1)$	$\widehat{\mathtt{m}}_{ai}$	$\hat{u}_{ai}$
0	0	$*$	0
0	1	$\mathtt{0}$	0
1	0	$\mathtt{1}$	0
1	1	$*$	1

Having established the WP update rules, it is immediate to write the $\mathsf{SP}(\mathtt{y})$ equations. Consider a node, and one of its neighbors to which it sends messages. For each possible configuration of incoming warnings on the node, denoted by `input`, we found the rules to compute the outgoing warning $\texttt{output} = \widehat{\texttt{OUT}}(\texttt{input})$ and the number of contradictions $\delta\mathbb{F}^{\rm e}(\texttt{input})$. $\mathsf{SP}(\mathtt{y})$ messages are distributions over $(\mathtt{0},\mathtt{1},*)$: $(Q_{ia}(\mathtt{0}), Q_{ia}(\mathtt{1}), Q_{ia}(*))$ and $(\widehat{Q}_{ai}(\mathtt{0}), \widehat{Q}_{ai}(\mathtt{1}), \widehat{Q}_{ai}(*))$. Notice that these messages are only marginally more complicated than ordinary BP messages. Let $\mathbb{P}(\texttt{input})$ denote the probability of a given input assuming independent warnings with distribution $Q_{ia}(\,\cdot\,)$ (respectively, $\widehat{Q}_{ai}(\,\cdot\,)$). The probability of an outgoing message $\texttt{output} \in \{\mathtt{0},\mathtt{1},*\}$ is then:

$$\mathbb{P}(\texttt{output}) \cong \sum_{\texttt{input}} \mathbb{P}(\texttt{input})\mathbb{I}(\widehat{\texttt{OUT}}(\texttt{input}) = \texttt{output})e^{-\mathtt{y}\delta\mathbb{F}^{\rm e}(\texttt{input})}\,. \tag{19.98}$$

Depending whether the node we are considering is a variable or function node, this probability distribution corresponds to the outgoing message $Q_{ia}(\,\cdot\,)$ or $\widehat{Q}_{ai}(\,\cdot\,)$.

It can be shown that the Bethe energy (19.83) associated with a given fixed point of the WP equations coincides with the total number of contradictions. This is expressed as the number of contradictions at function nodes, plus those at variable nodes, minus the number of edges (i,a) such that the warning in direction $a \to i$ contradicts the one in direction $i \to a$ (the last term avoids double counting). It follows that the Bethe free-entropy of the auxiliary graphical model $\mathbb{F}^{\rm RSB,e}(\mathtt{y})$ weights each WP fixed point depending on its number of contradictions, as it should.

19.5.4 *XORSAT again*

Let us know apply the $\mathsf{SP}(\mathtt{y})$ formalism to random K-XORSAT instances. We let the energy function $E(\underline{x})$ count the number of unsatisfied linear equations:

$$E_a(\underline{x}_{\partial a}) = \begin{cases} 0 & \text{if } x_{i_1(a)} \oplus \cdots \oplus x_{i_K(a)} = b_a, \\ 1 & \text{otherwise.} \end{cases} \tag{19.99}$$

The simplifications discussed in the previous subsection apply to this case. The 1RSB population dynamics algorithm can be used to compute the free-entropy density $\mathfrak{F}^{\rm e}(\mathtt{y})$. Here we limit ourselves to describing the results of this calculation for the case $K = 3$.

Let us stress that the problem we are considering here is different from the one investigated in Section 19.3. While there we were interested in the uniform measure over solutions (thus focusing on the satisfiable regime $\alpha < \alpha_{\rm s}(K)$), here we are estimating the minimum number of unsatisfied constraints (which is most interesting in the unsatisfiable regime $\alpha > \alpha_{\rm s}(K)$).

It is easy to show that the $\mathsf{SP}(\mathtt{y})$ equations always admit a solution in which $Q_{ia}(*) = 1$ for all (i,a), indicating that the min-sum equations have a unique solution. This corresponds to a density evolution fixed point whereby $Q(*) = 1$ with probability 1, yielding $\mathfrak{F}^{\rm e}(\mathtt{y})$ independent of $\mathtt{y}$. For $\mathtt{y}$ smaller than an α-dependent threshold $\mathtt{y}^*(\alpha)$, this is the only solution we find. For larger values of $\mathtt{y}$, the $\mathsf{SP}(\mathtt{y})$ equations have a non-trivial solution. Fig. 19.4 shows the result for the free-entropy density $\mathfrak{F}^{\rm e}(\mathtt{y})$, for three values of α.

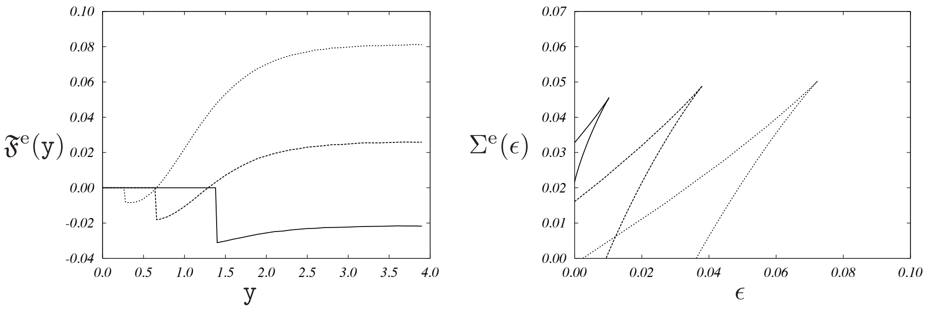

FIG. 19.4. Random 3-XORSAT at $\alpha = 0.87$, 0.97 and 1.07. Recall that, for $K = 3$, $\alpha_{\rm d}(K) \approx 0.818$ and $\alpha_{\rm s}(K) \approx 0.918$. Left frame: Free-entropy density $\mathfrak{F}^{\rm e}(\mathtt{y})$ as a function of $\mathtt{y}$, obtained using the population dynamics algorithm, with $N = 2 \cdot 10^4$ and $t = 5 \cdot 10^3$ (α increases from bottom to top). Right frame: Complexity $\Sigma^{\rm e}(\epsilon)$ as a function of energy density (equal to the number of violated constraints per variable). α increases from left to right.

Above this threshold density evolution converges to a 'non-trivial' 1RSB fixed point. The complexity functions $\Sigma^{\rm e}(\epsilon)$ can be deduced by Legendre transform, cf. Eq. (19.96), which requires differentiating $\mathfrak{F}^{\rm e}(\mathtt{y})$ and plotting $(\epsilon, \Sigma^{\rm e})$ in parametric form. The derivative can be computed numerically in a number of ways:

1. Compute analytically the derivative of $\mathbb{F}^{\rm RSB,e}(\mathtt{y})$ with respect to $\mathtt{y}$. This turns out to be a functional of the fixed point distributions of Q, $\widehat{Q}$, and can therefore be easily evaluated.
2. Fit the numerical results for the function $\mathfrak{F}^{\rm e}(\mathtt{y})$ and differentiate the fitting function
3. Approximate the derivative as difference at nearby values of $\mathtt{y}$.

In the present case we followed the second approach using the parametric form $\mathfrak{F}^{\rm fit}(\mathtt{y}) = a + b\, e^{-\mathtt{y}} + c\, e^{-2\mathtt{y}} + d\, e^{-3\mathtt{y}}$. As shown in Fig. 19.4 the resulting parametric curve $(\epsilon, \Sigma^{\rm e})$ is multiple valued (this is a consequence of the fact that $\mathfrak{F}^{\rm e}(\mathtt{y})$ is not concave). Only the concave part of $\mathfrak{F}^{\rm e}(\mathtt{y})$ is retained as physically meaningful. Indeed the convex branch is 'unstable' (in the sense that further RSB would be needed) and it is not yet understood whether it has any meaning.

For $\alpha \in [\alpha_{\rm d}(K), \alpha_{\rm s}(K)[$, $\Sigma^{\rm e}(\epsilon)$ remains positive down to $\epsilon = 0$. The intercept $\Sigma^{\rm e}(\epsilon = 0)$ coincides with the complexity of clusters of SAT configurations, as computed in Ch. 18 (see Theorem 18.2). For $\alpha > \alpha_{\rm s}(K)$ (UNSAT phase) $\Sigma^{\rm e}(\epsilon)$ vanishes at $\epsilon_{\rm gs}(K, \alpha) > 0$. The energy density $\epsilon_{\rm gs}(K, \alpha)$ is the minimal fraction of violated equations, in a random XORSAT linear system. Notice that $\Sigma^{\rm e}(\epsilon)$ is not defined above a second energy density $\epsilon_{\rm d}(K, \alpha)$. This indicates that we should take $\Sigma^{\rm e}(\epsilon) = -\infty$ there: above $\epsilon_{\rm d}(K, \alpha)$ one recovers a simple problem with a unique Bethe measure.

Figure 19.5 shows the values of $\epsilon_{\rm gs}(K, \alpha)$ and $\epsilon_{\rm d}(K, \alpha)$ as functions of α for $K = 3$ (random 3-XORSAT).

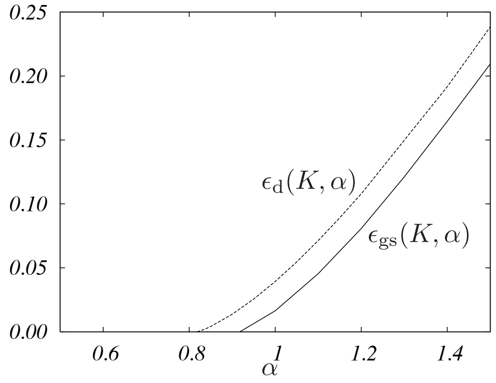

FIG. 19.5. Asymptotic ground state energy (= minimal number of violated constraints) per variable $\epsilon_{\rm gs}(K,\alpha)$ for random K = 3-XORSAT formulae. $\epsilon_{\rm gs}(K,\alpha)$ vanishes for $\alpha < \alpha_{\rm s}(K)$. The dashed line $\epsilon_{\rm d}(K,\alpha)$ is the highest energy density e such that configurations with $E(\underline{x}) < Ne$ are clustered. It vanishes for $\alpha < \alpha_{\rm d}(K)$.

19.6 The nature of 1RSB phases

In the last sections we discussed how to compute the complexity function $\Sigma(\phi)$ (or its 'zero temperature' version, the energetic complexity $\Sigma^{\rm e}(\epsilon)$). Here we want to come back to the problem of determining some qualitative properties of the measure $\mu(\,\cdot\,)$ for random graphical models, on the basis of its decomposition into extremal Bethe measures:

$$\mu(\underline{x}) = \sum_{n\in\mathsf{E}} w_n \mu^n(\underline{x})\,. \tag{19.100}$$

Assumptions 2 and 3 imply that, in this decomposition, we introduce a negligible error if we drop all the states n but the ones with free-entropy $\phi_n \approx \phi_*$, where

$$\phi_* = \text{argmax}\,\{\phi + \Sigma(\phi):\ \Sigma(\phi)\geq 0\}\ . \tag{19.101}$$

In general, $\Sigma(\phi)$ is strictly positive and continuous in an interval $[\phi_{\rm min},\phi_{\rm max}]$ with $\Sigma(\phi_{\rm max})=0$, and

$$\Sigma(\phi) = \mathtt{x}_*(\phi_{\rm max}-\phi) + O((\phi_{\rm max}-\phi)^2)\,, \tag{19.102}$$

for ϕ close to $\phi_{\rm max}$.

It turns out that the decomposition (19.100) has different properties depending on the result of the optimization (19.101). One can distinguish two phases (see Fig. 19.6): d1RSB (dynamic one-step replica symmetry breaking) when the max is achieved in the interior of $[\phi_{\rm min},\phi_{\rm max}]$ and, as a consequence $\Sigma(\phi_*)>0$; s1RSB (static one-step replica symmetry breaking) when the max is achieved at $\phi_*=\phi_{\rm max}$ and therefore $\Sigma(\phi_*)=0$ (this case occurs iff $\mathtt{x}_*\leq 1$).

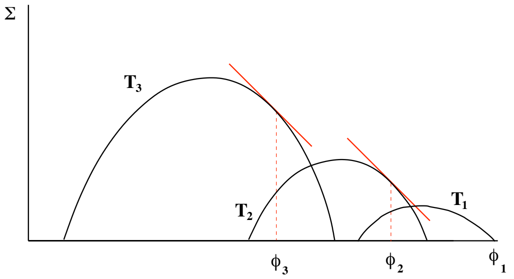

FIG. 19.6. A sketch of the complexity Σ versus free-entropy-density ϕ in a finite-temperature problem with 1RSB phase transition, at three temperatures $T_1 < T_2 < T_3$. A random configuration $\underline{x}$ with distribution $\mu(\underline{x})$ is found with high probability in a cluster of free-entropy-density ϕ_1, ϕ_2, ϕ_3 respectively. T_2 and T_3 are above the condensation transition: ϕ_2, ϕ_3 are the points where $\partial\Sigma/\partial\phi = -1$. T_1 is below the condensation transition: ϕ_1 is the largest value of ϕ where Σ is positive.

19.6.1 *Dynamical 1RSB*

Assume $\Sigma_* = \Sigma(\phi_*) > 0$. Then we can restrict the sum (19.100) to those states n such that $\phi_n \in [\phi_* - \varepsilon, \phi_* + \varepsilon]$, if we allow for an exponentially small error. To the leading exponential order there are $e^{N\Sigma_*}$ such states whose weights are $w_n \in [e^{-N(\Sigma_*+\varepsilon')}, e^{-N(\Sigma_*-\varepsilon')}]$.

Different states are expected to have essentially disjoint support. By this we mean that there exists subsets $\{\Omega_n\}_{n\in\mathsf{E}}$ of the configuration space $\mathcal{X}^N$ such that, for any $m \in \mathsf{E}$

$$\mu^m(\Omega_m) \approx 1\,, \qquad \sum_{n\in\mathsf{E}\backslash m} w_n \mu^n(\Omega_m) \approx 0\,. \tag{19.103}$$

Further, different states are separated by 'large free-energy barriers.' This means that one can choose the above partition in such a way that only an exponentially small (in N) fraction of the probability measure is on its boundaries.

This structure has two important consequences:

Glassy dynamics. Let us consider a local Markov Chain dynamics that satisfies detailed balance with respect to the measure $\mu(\,\cdot\,)$. As an example we can consider the Glauber dynamics introduced in Ch. 4 (in order to avoid trivial reducibility effects, we can assume in this discussion that the compatibility functions $\psi_a(\underline{x}_{\partial a})$ are bounded away from 0).

Imagine initiating the dynamics at time 0 with an equilibrated configuration $\underline{x}(0)$ distributed according to $\mu(\,\cdot\,)$. This is essentially equivalent to picking a state n uniformly at random among the typical ones, and then sampling $\underline{x}(0)$ from $\mu^n(\,\cdot\,)$. Because of the exponentially large barriers, the dynamics will stay confined in Ω_n for an exponentially large time, and equilibrate among states only on larger time scales.

This can be formalized as follows. Denote by $D(\underline{x},\underline{x}')$ the Hamming distance in $\mathcal{X}^N$. Take two i.i.d. configuration with distribution μ and let Nd_0 be the expectation value of their Hamming distance. Analogously take two i.i.d. configuration with distribution μ^n, and let Nd_1 be the expectation value of their Hamming distance. When the state n is chosen randomly with distribution w_n, we expect d_1 not to depend on the state n asymptoticaly for large sizes. Furthermore: $d_1 < d_0$. Then we can consider the (normalized) expected Hamming distance between configurations at time t in Glauber dynamics $d(t) = \langle D(\underline{x}(0),\underline{x}(t))\rangle/N$. For any $\varepsilon < d_0 - d_1$, the correlation time $\tau(\varepsilon) \equiv \inf\{t : \, d(t) \geq d_0 - \varepsilon\}$ is expected to be exponentially large in N

Short-range correlations in finite-dimensional projections. We motivated the 1RSB cavity method with the emergence of long-range correlations due to decomposition of $\mu(\,\cdot\,)$ into many extremal Bethe measures. Surprisingly, such correlations cannot be detected by probing a bounded (when $N \to \infty$) number of variables. More precisely, if $i(1), \dots, i(k) \in \{1, \cdots, N\}$ are uniformly random variable indices, then, in the d1RSB phase:

$$\mathbb{E}|\langle f_1(x_{i(1)})f_2(x_{i(2)})\cdots f_k(x_{i(k)})\rangle - \langle f_1(x_{i(1)})\rangle\langle f_2(x_{i(2)})\rangle\cdots\langle f_k(x_{i(k)})\rangle| \overset{N\to\infty}{\to} 0\,.$$

(Here $\langle\,\cdot\,\rangle$ denote the expectation with respect to the measure μ, and $\mathbb{E}$ the expectation with respect to the graphical model in a random ensemble). This finding can be understood intuitively as follows. If there are long range correlations among subsets of k variables, then it must be true that conditioning on the values of $k-1$ of them changes the marginal distribution of the k-th one. On the other hand, we think that long range correlations arise because far apart variables 'know' that the whole system is in the same state n. But conditioning on a bounded number $(k-1)$ of variables cannot select in any significant way among the $e^{N\Sigma_*}$ relevant states, and thus cannot change the marginal of the k-th one.

An alternative argument makes use of the observation that, if $\underline{x}^{(1)}$ and $\underline{x}^{(2)}$ are two i.i.d. configurations with distribution $\mu(\,\cdot\,)$, then their distance $D(\underline{x}^{(1)},\underline{x}^{(2)})$ concentrates in probability. This is due to the fact that the two configurations will be, with high probability, in different states $n_1 \neq n_2$ (the probability of $n_1 = n_2$ being $e^{-N\Sigma_*}$), whose distance depends weakly on the states couple.

Let us finally notice that the absence of long range correlations among bounded subset of variables is related to the observation that $\mu(\,\cdot\,)$ is itself a Bethe measure (although a non-extremal one) in a d1RSB phase, cf. Sec. 19.4.1. Indeed, each BP equation involves a bounded subset of the variables and can be violated only because of correlations among them.

As we shall discuss in Sec. 22.1.2, long range correlations in a d1RSB phase can be probed through more sophisticated "point-to-set" correlation functions.

19.6.2 *Static 1RSB*

In this case the decomposition (19.100) is dominated by a few states of near-to-maximal free-entropy $\phi_n \approx \phi_{\max}$. If we 'zoom' near the edge by letting $\phi_n =$

$\phi_{\max} + s_n/N$, then the 'free-entropy shifts' s_n form a point process with density $\exp(-\mathtt{x}_* s)$.

The situation is analogous to the one we found in the random energy model for $T < T_{\rm c}$. Indeed it is expected that the weights $\{w_n\}$ converge to the same universal Poisson-Dirichlet process found there, and to depend on the model details only through the parameter $\mathtt{x}_*$ (we have already discussed this universality using replicas in Ch. 8). In particular, if $\underline{x}^{(1)}$ and $\underline{x}^{(2)}$ are two i.i.d. replicas with distribution μ, and n_1, n_2 are the states they belong to, then the probability for them to belong to the same state is

$$\mathbb{E}\left\{\mathbb{P}_\mu(n_1 = n_2)\right\} = \mathbb{E}\left\{\sum_{n\in\mathsf{E}} w_n^2\right\} = 1 - x_*\,. \qquad (19.104)$$

Here $\mathbb{E}$ denote expectation with respect to the graphical model distribution.

As a consequence, the distance $D(\underline{x}^{(1)}, \underline{x}^{(2)})$ between two i.i.d. replicas does not concentrate (the overlap distribution is non-trivial). This in turn can only be true if the two-point correlation function does not vanish at large distances. Long-range correlations of this type make BP break down. The original graphical model $\mu(\,\cdot\,)$ is no longer a Bethe measure: its local marginals cannot be described in terms of a set of messages. The 1RSB description, according to which $\mu(\,\cdot\,)$ is a convex combination of Bethe measures, is unavoidable.

At this point we are left with a puzzle. How to circumvent the argument given in Section 19.4.1 that, if the 'correct' weight $\mathtt{x} = 1$ is used, then the marginals as computed within 1RSB still satisfy BP equations? The conundrum is that, within a $\mathsf{s1RSB}$ phase, the parameter $\mathtt{x} = 1$ is *not* the correct one to be used in the 1RSB cavity equations (although it is the correct one to weight states). In order to explain this, let us first notice that, if the complexity is convex and behaves as in Eq. (19.102) near its edge, with a slope $-\mathtt{x}_* > -1$, then the optimization problem (19.101) has the same result as

$$\phi_* = \operatorname{argmax}\left\{\mathtt{x}\phi + \Sigma(\phi)\,:\; \Sigma(\phi) \ge 0\right\}\,. \qquad (19.105)$$

for any $\mathtt{x} \ge \mathtt{x}_*$. Therefore, in the 1RSB cavity equations we could in principle use any value of $\mathtt{x}$ larger or equal to $\mathtt{x}_*$ (this would select the same states). However, the constraint $\Sigma(\phi) \ge 0$ cannot be enforced locally and does not show up in the cavity equations. If one performs the computation of Σ within the cavity method using a value $\mathtt{x} > \mathtt{x}_*$, then one finds a negative value of Σ which must be rejected (it is believed to be related to the contribution of some exponentially rare instances). Therefore, in order to ensure that one studies the interval of ϕ such that $\Sigma(\phi) \ge 0$, one must *impose* $\mathtt{x} \le \mathtt{x}_*$ in the cavity method. In order to select the states with free-entropy density $\phi_{\max}$, we must thus choose the Parisi parameter that corresponds to $\phi_{\max}$, namely $\mathtt{x} = \mathtt{x}_*$.

19.6.3 *When does 1RSB fail?*

The 1RSB cavity method is a powerful tool, but does not always provide correct answers, even for locally tree-like models, in the large system limit. The main assumption of the 1RSB approach is that, once we pass to the auxiliary graphical model (which 'enumerates' BP fixed points) a simple BP procedure is asymptotically exact. In other words, the auxiliary problem has a simple 'replica symmetric' structure and no glassy phase. This is correct in some cases, such as random XORSAT or SAT close to their SAT-UNSAT threshold, but it may fail in others.

A mechanism leading to a failure of the 1RSB approach is that the auxiliary graphical model is incorrectly described by BP. This may happen because the auxiliary model measure decomposes in many Bethe states. In such a case, one should introduce a second auxiliary model, dealing with the multiplicity of BP fixed points of the first one. This is usually referred to as 'two-step replica symmetry breaking' (2RSB). Obviously one can find situations in which it is necessary to iterate this construction, leading to a R-th level auxiliary graphical model (R-RSB). Continuous (or full) RSB corresponds to the large-R limit.

While such developments are conceptually clear (at least from an heuristic point of view), they are technically challenging. So far limited results have been obtained beyond 1RSB. For a brief survey, we refer to Ch. 22.

Appendix: $\mathsf{SP}(\mathtt{y})$ equations for XORSAT

This appendix provides technical details on the 1RSB treatment of random K-XORSAT, within the 'energetic' formalism. The results of this approach were discussed in Sec. 19.5.4. In particular we will derive the behavior of the auxiliary free-entropy $\mathfrak{F}^{\rm e}(\mathtt{y})$ at large $\mathtt{y}$, and deduce the behavior of the complexity $\Sigma^{\rm e}(\epsilon)$ at small ϵ. This section can be regarded as an exercise in applying the $\mathsf{SP}(\mathtt{y})$ formalism. We shall skip many details and just give the main intermediate results of the computation.

XORSAT is a constraint satisfaction problems with binary variables. We can thus apply the simplified method of Sec. 19.5.3. The projected min-sum messages can take three values: $\mathtt{0}$, $\mathtt{1}$, $*$. Exploiting the symmetry of XORSAT between $\mathtt{0}$ and $\mathtt{1}$, $\mathsf{SP}(\mathtt{y})$ messages can be parametrized by a single number, e.g. by the sum of their weights on $\mathtt{0}$ and $\mathtt{1}$. We will therefore write: $Q_{ia}(\mathtt{0}) = Q_{ia}(\mathtt{1}) = \zeta_{ia}/2$ (thus implying $Q_{ia}(*) = 1-\zeta_{ia}$), and $\widehat{Q}_{ai}(\mathtt{0}) = \widehat{Q}_{ai}(\mathtt{1}) = \eta_{ai}/2$ (whence $\widehat{Q}_{ai}(*) = 1-\eta_{ai}$).

In terms of these variables, the $\mathsf{SP}(\mathtt{y})$ equation at function node a reads:

$$\eta_{ai} = \prod_{j\in\partial a\setminus i} \zeta_{ja} \,. \tag{19.106}$$

The $\mathsf{SP}(\mathtt{y})$ equation at variable node i is a bit more complicated. Let us consider all the $|\partial i|-1$ incoming messages $\widehat{Q}_{bi}$, $b\in\partial i\setminus a$. Each of them is parameterized by a number η_{bi}. We let $\underline{\eta} = \{\eta_{bi},\ b\in\partial i\setminus a\}$ and define the function $B_q(\underline{\eta})$ as follows:

$$B_q(\underline{\eta}) = \sum_{S\subset\{\partial i\setminus a\}} \mathbb{I}(|S|=q) \prod_{b\in\partial i\setminus\{S\cup\{a\}\}} (1-\eta_{bi}) \prod_{c\in S} \eta_{cj} \,. \tag{19.107}$$

Let $A_{q,r}(\underline{\eta}) = B_{q+r}(\underline{\eta})\binom{q+r}{q}2^{-(q+r)}$. After some thought one obtains the update equation:

$$\zeta_{ia} = \frac{2\sum_{q=0}^{|\partial i|-2}\sum_{r=q+1}^{|\partial i|-1} A_{q,r}(\underline{\eta})e^{-\mathtt{y}q}}{\sum_{q=0}^{\lfloor(|\partial i|-1)/2\rfloor} A_{q,q}(\underline{\eta})e^{-\mathtt{y}q} + 2\sum_{q=0}^{|\partial i|-2}\sum_{r=q+1}^{|\partial i|-1} A_{q,r}(\underline{\eta})e^{-\mathtt{y}q}} \tag{19.108}$$

The auxiliary free-entropy $\mathbb{F}^{\mathrm{RSB,e}}(\mathtt{y})$ has the general form (19.92), with the various contributions expressed as follows in terms of the parameters $\{\zeta_{ia}, \eta_{ai}\}$:

$$e^{\mathbb{F}_a^{\mathrm{RSB,e}}} = 1 - \frac{1}{2}(1-e^{-\mathtt{y}})\prod_{i\in\partial a}\zeta_{ia}\,, \qquad e^{\mathbb{F}_{ai}^{\mathrm{RSB,e}}} = 1 - \frac{1}{2}\eta_{ai}\zeta_{ia}(1-e^{-\mathtt{y}})\,,$$

$$e^{\mathbb{F}_i^{\mathrm{RSB,e}}} = \sum_{q=0}^{d_i}\sum_{r=0}^{d_i-q} A_{q,r}\left(\{\eta_{ai}\}_{a\in\partial i}\right)\, e^{-\mathtt{y}\min(q,r)}\,. \tag{19.109}$$

Let us consider random K-XORSAT instances with constraint density α. Equations (19.106), (19.108) get promoted to distributional relations that determine the asymptotic distribution of η and ζ on a randomly chosen edge (i,a). The 1RSB population dynamics algorithm can be used to approximate these distributions. We encourage the reader to implement it, and obtain a numerical estimate of the auxiliary free-entropy density $\mathfrak{F}^{\mathrm{e}}(\mathtt{y})$.

It turns out that, at large $\mathtt{y}$, one can control the distributions of η, ζ analytically, provided their qualitative behavior satisfies the following assumptions (that can be checked numerically):

- With probability t one has $\eta = 0$, and with probability $1-t$, $\eta = 1 - e^{-\mathtt{y}}\hat{\eta}$, where t has a limit in $]0,1[$, and $\hat{\eta}$ converges to a random variable with support on $[0,\infty[$, as $\mathtt{y}\to\infty$.
- With probability s one has $\zeta = 0$, and with probability $1-s$, $\zeta = 1 - e^{-\mathtt{y}}\hat{\zeta}$, where s has a limit in $]0,1[$, and $\hat{\zeta}$ converges to a random variable with support on $[0,\infty[$, as $\mathtt{y}\to\infty$.

Under these assumptions, we shall expand the distributional version of Eqs. (19.106), (19.108) keeping terms up to first order in $e^{-\mathtt{y}}$. We shall use $t, s, \hat{\eta}, \hat{\zeta}$ to denote the limit quantities mentioned above.

It is easy to see that t, s must satisfy the equations $(1-t) = (1-s)^{k-1}$ and $s = e^{-K\alpha(1-t)}$. These are identical to Eqs. (19.51) and (19.52), whence $t = 1-\widehat{Q}_*$ and $s = 1 - Q_*$.

Equation (19.106) leads to the distributional equation:

$$\hat{\eta} \overset{\mathrm{d}}{=} \hat{\zeta}_1 + \cdots + \hat{\zeta}_{K-1}\,, \tag{19.110}$$

where $\hat{\zeta}_1, \dots, \hat{\zeta}_{K-1}$ are $K-1$ i.i.d. copies of the random variable $\hat{\zeta}$.

The update equation (19.108) is more complicated. There are in general l inputs to a variable node, where l is Poisson with mean $K\alpha$. Let us denote by

m the number of incoming messages with $\eta = 0$. The case $m = 0$ yields $\zeta = 0$ and is taken care of in the relation between t and s. If we condition on $m \geq 1$, the distribution of m is

$$\mathbb{P}(m) = \frac{\lambda^m}{m!} e^{-\lambda} \frac{1}{1 - e^{-\lambda}} \, \mathbb{I}(m \geq 1) \, , \tag{19.111}$$

where $\lambda = K\alpha(1-t)$. Conditional on m, Eq. (19.108) simplifies as follows:

- If $m = 1$: $\hat{\zeta} \overset{\rm d}{=} \hat{\eta}$.
- If $m = 2$: $\hat{\zeta} = 1$ identically.
- If $m \geq 3$: $\hat{\zeta} = 0$ identically.

The various contributions to the free-entropy (19.38) are given by:

$$\begin{aligned} \mathrm{f}_{\rm f}^{\rm RSB,e} &= (1-s)^k \left[- \log 2 + e^{-\mathrm{y}} (1 + K \langle \hat{\zeta} \rangle) \right] + o(e^{-\mathrm{y}}) \, , & (19.112) \\ \mathrm{f}_{\rm v}^{\rm RSB,e} &= \frac{\lambda^2}{2} e^{-\lambda} \left[- \log 2 + e^{-\mathrm{y}} (1 + 2\langle \hat{\eta} \rangle) \right] \\ &\quad + \sum_{m=3}^{\infty} \frac{\lambda^m}{m!} e^{-\lambda} \left[(1-m) \log 2 + e^{-\mathrm{y}} m (1 + \langle \hat{\eta} \rangle) \right] + o(e^{-\mathrm{y}}) \, , & (19.113) \\ \mathrm{f}_{\rm e}^{\rm RSB,e} &= (1-t)(1-s) \left[- \log 2 + e^{-\mathrm{y}} (1 + \langle \hat{\eta} \rangle + \langle \hat{\zeta} \rangle) \right] + o(e^{-\mathrm{y}}) \ , & (19.114) \end{aligned}$$

where $\langle \hat{\eta} \rangle$ and $\langle \hat{\zeta} \rangle$ are the expectation values of $\hat{\eta}$ and $\hat{\zeta}$. This gives for the free-entropy density $\mathfrak{F}^{\rm e}(\mathrm{y}) = \mathrm{f}_{\rm f}^{\rm RSB,e} + \alpha \mathrm{f}_{\rm v}^{\rm RSB,e} - K\alpha \, \mathrm{f}_{\rm e}^{\rm RSB,e} = \Sigma_0 + e^{-\mathrm{y}} \epsilon_0 + o(e^{-\mathrm{y}})$, with:

$$\Sigma_0 = \left[1 - \frac{\lambda}{k} - e^{-\lambda} \left(1 + \frac{k-1}{k} \lambda \right) \right] \log 2 \, , \tag{19.115}$$

$$\epsilon_0 = \frac{\lambda}{k} \left[1 - e^{-\lambda} \left(1 + \frac{k}{2} \lambda \right) \right] \, . \tag{19.116}$$

Taking the Legendre transform, cf. Eq. (19.96), we obtain the following behavior of the energetic complexity as $\epsilon \to 0$:

$$\Sigma^{\rm e}(\epsilon) = \Sigma_0 + \epsilon \log \frac{\epsilon_0 e}{\epsilon} + o(\epsilon) \ , \tag{19.117}$$

This shows in particular that the ground state energy density is proportional to $(\alpha - \alpha_s)/|\log(\alpha - \alpha_s)|$ close to the SAT-UNSAT transition (when $0 < \alpha - \alpha_s \ll 1$).

Exercise 19.7 In the other extreme, show that at large α one gets $\epsilon_{\rm gs}(K, \alpha) = \alpha/2 + \sqrt{2\alpha} \epsilon_*(K) + o(\sqrt{\alpha})$, where the positive constant $\epsilon_*(K)$ is the absolute value of the ground state energy of the fully connected K-spin model studied in Sec. 8.2. This indicates that there is no interesting intermediate asymptotic regime between the $M = \Theta(N)$ (discussed in the present chapter) and $M = \Theta(N^{K-1})$ (discussed with the replica method in Ch. 8)

Notes

The cavity method originated as an alternative to the replica approach in the study of the Sherrington-Kirkatrick model (Mézard, Parisi and Virasoro, 1985*b*). The 1RSB cavity method for locally tree-like factor graphs was developed in the context of spin glasses in (Mézard and Parisi, 2001). Its application to zero temperature problems (counting solutions of the min-sum equations), was also first described in the spin glass context in (Mézard and Parisi, 2003). The presentation in this chapter differs in its scope from those work, which were more focused in computing averages over random instances. For a rigorous treatment of the notion of Bethe measure, we refer to (Dembo and Montanari, 2008*b*).

The idea that the 1RSB cavity method is in fact equivalent to applying BP on an auxiliary model appeared in several paper treating the cases of coloring and satisfiability with $\mathtt{y} = 0$ (Parisi, 2002; Braunstein and Zecchina, 2004; Maneva, Mossel and Wainwright, 2005). The treatment presented here generalizes these works, with the important difference that the variables of our auxiliary model are messages rather than node quantities.

The analysis of the $\mathtt{x} = 1$ case is strictly related to the problem of reconstruction on a tree. This has been studied in (Mézard and Montanari, 2006), where the reader will find the proof of Theorem 19.5 and the expression of the free-entropy of exercise 19.6.

The $\mathsf{SP}(\mathtt{y})$ equations for one single instance have been written first in the context of the K-satisfiability problem in (Mézard and Zecchina, 2002), see also (Mézard, Parisi and Zecchina, 2003). The direct derivation of $\mathsf{SP}(\mathtt{y})$ equations in binary variable problems, shown in Sec. 19.5.3, was done originally for satisfiability in (Braunstein, Mézard and Zecchina, 2005), see also (Braunstein and Zecchina, 2004) and (Maneva, Mossel and Wainwright, 2005). The application of the 1RSB cavity method to the random XORSAT problem, and its comparison to the exact results, was done in (Mézard, Ricci-Tersenghi and Zecchina, 2003).

An alternative to the cavity approach followed throughout this book is provided by the replica method of Ch. 8. As we saw, it was first invented in order to treat fully connected models (i.e. models on complete graphs), cf. (Mézard, Parisi and Virasoro, 1987), and subsequently developed in the context of sparse random graphs (Mézard and Parisi, 1985; Dominicis and Mottishaw, 1987; Mottishaw and Dominicis, 1987; Wong and Sherrington, 1988; Goldschmidt and Lai, 1990). The technique was further improved in the paper (Monasson, 1998), that offers a very lucid presentation of the method.

20

RANDOM K-SATISFIABILITY

This chapter applies the cavity method to the random K-satisfiability problem. We will study both the phase diagram (in particular, we will determine the SAT-UNSAT threshold $\alpha_{\rm s}(K)$) and the algorithmic applications of message passing. The whole chapter is based on heuristic derivations: it turns out that the rigorization of the whole approach is still in its infancy. Neither the conjectured phase diagram, nor the efficiency of message passing algorithms have been yet confirmed rigorously. But the computed value of $\alpha_{\rm s}(K)$ is conjectured to be exact, and the low-complexity message passing algorithms that we will describe turn out to be particularly efficient in finding solutions.

We will start in Sec. 20.1 by writing the BP equations, following the approach exposed in Ch. 14. The statistical analysis of such equations provides a first (replica symmetric) estimate of $\alpha_{\rm s}(K)$. This however turns out to be incorrect. The reason of this failure is traced back to the incorrectness of the replica symmetric assumption close to the SAT-UNSAT transition. The system undergoes a 'structural' phase transition at a clause density smaller than $\alpha_{\rm s}(K)$. Nevertheless, BP empirically converges in a wide range of clause densities, and it can be used to find SAT assignments on large instances provided the clause density α is not too close to $\alpha_{\rm s}(K)$. The key idea is to use BP as a heuristic guide in a sequential decimation procedure.

In Sec. 20.2 we apply the 1RSB cavity method developed in Ch. 19. The statistical analysis of the 1RSB equations gives the values for $\alpha_{\rm s}(K)$ summarized in Table 20.2.4. From the algorithmic point of view, one can use SP instead of BP as a guide in the decimation procedure. We shall explain and study numerically the corresponding 'survey-guided decimation' algorithm, which is presently the most efficient algorithm to find SAT assignments in large random satisfiable instances with a clause density close to the threshold $\alpha_{\rm s}(K)$.

This chapter focuses on K-SAT with $K \geq 3$. The $K = 2$ problem is quite different: satisfiability can be proved in polynomial time, the SAT-UNSAT phase transition is driven by a very different mechanism, and the threshold is known to be $\alpha_{\rm s}(2) = 1$. It turns out that a (more subtle) qualitative difference also distinguishes $K = 3$ from $K \geq 4$. In order to illustrate this point, we will use both 3-SAT and 4-SAT as running examples.

Coloring random graphs turns out to be very similar to random K-satisfiability. Section 20.4 presents a few highlights in the study of random graph colorings. In particular, we emphasize how the techniques used for K-satisfiability are successful in this case as well.

20.1 Belief Propagation and the replica symmetric analysis

We already studied some aspects of random K-SAT in Ch. 10, where we derived in particular some rigorous bounds on the SAT/UNSAT threshold $\alpha_{\mathrm{s}}(K)$. Here we will study the problem using message passing approaches. Let us start by summarizing our notations.

An instance of the K-satisfiability problem is defined by M clauses (indexed by $a, b \cdots \in \{1, \dots, M\}$) over N Boolean variables $x_1, \dots, x_N$ taking values in $\{0, 1\}$. We denote by ∂a the set of variables in clause a, and by ∂i the set of clauses in which variable x_i appears. Further, for each $i \in \partial a$, we introduce the number J_{ai} which takes value 1 if x_i appears negated in clause a, and takes value 0 if the variable appears unnegated.

It will be convenient to distinguish elements of ∂a according to the values of J_{ai}. We let $\partial_0 a \equiv \{i \in \partial a \text{ s.t. } J_{ai} = 0\}$ and $\partial_1 a = \{i \in \partial a \text{ s.t. } J_{ai} = 1\}$. Similarly we denote by $\partial_0 i$ and $\partial_1 i$ the neighborhoods of i: $\partial_0 i = \{a \in \partial i \text{ s.t. } J_{ai} = 0\}$ and $\partial_1 i = \{a \in \partial i \text{ s.t. } J_{ai} = 1\}$.

As usual, the indicator function over clause a being satisfied is denoted by $\psi_a(\,\cdot\,)$: $\psi_a(\underline{x}_{\partial a}) = 1$ if clause a is satisfied by the assignment $\underline{x}$ and $\psi_a(\underline{x}_{\partial a}) = 0$ if it is not. Given a SAT instance, we begin by studying the uniform measure over SAT assignments:

$$\mu(\underline{x}) = \frac{1}{Z} \prod_{a=1}^{M} \psi_a(\underline{x}_{\partial a})\,. \tag{20.1}$$

We will represent this distribution with a factor graph, as in Fig. 10.1, and in this graph we draw dashed edges when $J_{ai} = 1$, and full edges when $J_{ai} = 0$.

20.1.1 *The BP equations*

The BP equations for a general model of the form (20.1) have already been written in Chapter 14. Here we want to rewrite them in a more compact form, that is convenient both for analysis and implementation. They are best expressed using the following notation. Consider a variable node i connected to factor node a and partition its neighborhood as $\partial i = \{a\} \cup \mathcal{S}_{ia} \cup \mathcal{U}_{ia}$, where (see Fig. 20.1):

$$\begin{aligned} &\text{if } J_{ai} = 0 \text{ then} \quad \mathcal{S}_{ia} = \partial_0 i \setminus \{a\},\ \mathcal{U}_{ia} = \partial_1 i\,, \\ &\text{if } J_{ai} = 1 \text{ then} \quad \mathcal{S}_{ia} = \partial_1 i \setminus \{a\},\ \mathcal{U}_{ai} = \partial_0 i\,. \end{aligned} \tag{20.2}$$

Since the variables x_i's are binary, the BP messages at any time $\nu_{i\to a}(\,\cdot\,)$, $\widehat{\nu}_{a\to i}(\,\cdot\,)$, can be parameterized by a single real number. We fix the parameterization by letting $\zeta_{ia} \equiv \nu_{i\to a}(x_i = J_{ai})$ (which obviously implies $\nu_{i\to a}(x_i = 1 - J_{ai}) = 1 - \zeta_{ia}$), and $\hat{\zeta}_{ai} \equiv \widehat{\nu}_{a\to i}(x_i = J_{ai})$ (yielding $\widehat{\nu}_{a\to i}(x_i = 1 - J_{ai}) = 1 - \hat{\zeta}_{ai}$).

A straightforward calculation allows to express the BP equations (here in fixed point form) in terms of these variables:

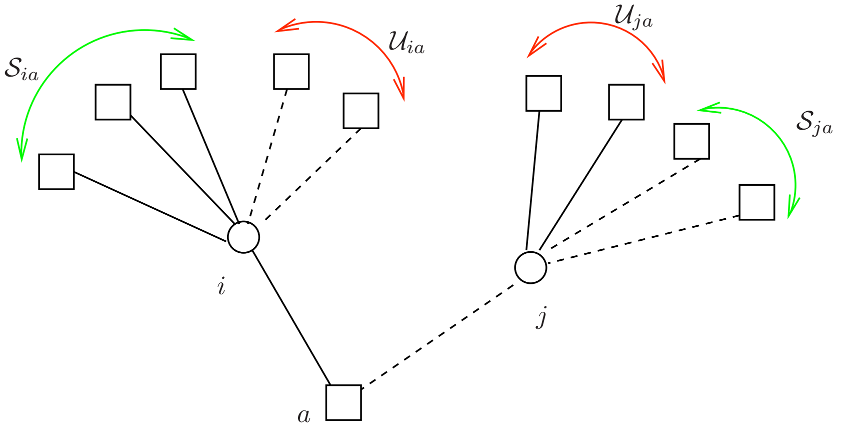

FIG. 20.1. The set $\mathcal{S}_{ia}$ contains all checks b in $\partial i \setminus a$ such that $J_{bi} = J_{ai}$, the set $\mathcal{U}_{ia}$ contains all checks b in $\partial i \setminus a$ such that $J_{bi} = 1 - J_{ai}$

$$\zeta_{ia} = \frac{\left[\prod_{b\in\mathcal{S}_{ia}} \hat{\zeta}_{bi}\right] \left[\prod_{b\in\mathcal{U}_{ia}} (1-\hat{\zeta}_{bi})\right]}{\left[\prod_{b\in\mathcal{S}_{ia}} \hat{\zeta}_{bi}\right] \left[\prod_{b\in\mathcal{U}_{ia}} (1-\hat{\zeta}_{bi})\right] + \left[\prod_{b\in\mathcal{U}_{ia}} \hat{\zeta}_{bi}\right] \left[\prod_{b\in\mathcal{S}_{ia}} (1-\hat{\zeta}_{bi})\right]},$$
$$\hat{\zeta}_{ai} = \frac{1-\prod_{j\in\partial a\setminus i}\zeta_{ja}}{2-\prod_{j\in\partial a\setminus i}\zeta_{ja}}, \tag{20.3}$$

with the convention that a product over zero term is equal to 1. Notice that evaluating the right hand side takes (respectively) $O(|\partial i|)$ and $O(|\partial a|)$ operations. This should be contrasted with the general implementation of the BP equations, cf. Ch. 14 , that requires $O(|\partial i|)$ operations at variable nodes but $O(2^{|\partial a|})$ at function nodes.

The Bethe free-entropy takes the usual form, cf. Eq. (14.27), $\mathbb{F} = \sum_{a\in F} \mathbb{F}_a + \sum_{i\in V} \mathbb{F}_i - \sum_{(ia)\in E} \mathbb{F}_{ia}$. The various contributions can be expressed in terms of the parameters ζ_{ia}, $\hat{\zeta}_{ai}$ as follows

$$\mathbb{F}_a = \log\left[1-\prod_{i\in\partial a}\zeta_{ia}\right]; \; \mathbb{F}_i = \log\left[\prod_{a\in\partial_0 i}\hat{\zeta}_{ai}\prod_{b\in\partial_1 i}(1-\hat{\zeta}_{bi}) + \prod_{a\in\partial_0 i}(1-\hat{\zeta}_{ai})\prod_{b\in\partial_1 i}\hat{\zeta}_{bi}\right];$$
$$\mathbb{F}_{ai} = \log\left[(1-\zeta_{ia})(1-\hat{\zeta}_{ai}) + \zeta_{ia}\hat{\zeta}_{ai}\right] . \tag{20.4}$$

Given the messages, the BP estimate for the marginal on site i is:

$$\nu_i(x_i) \cong \prod_{a\in\partial i} \widehat{\nu}_{a\to i}(x_i)\,. \tag{20.5}$$

20.1.2 *Statistical analysis*

Let us now consider a random K-sat formula, i.e. a uniformly random formula with N variables and $M = N\alpha$ clauses. The resulting factor graph will be distributed according to the $\mathbb{G}_N(K, M)$ ensemble. Given a variable index i, the numbers $|\partial_0 i|$, $|\partial_1 i|$ of variables in which x_i appears directed or negated, converge to independent Poisson random variables of mean $K\alpha/2$.

If (i, a) is a uniformly random edge in the graph, the corresponding fixed point messages ζ_{ia}, $\hat{\zeta}_{ai}$ are random variables (we assume here that an 'approximate' fixed point exists). Within the RS assumption, they converge in distribution, as $N \to \infty$, to random variables ζ, $\hat{\zeta}$ whose distribution satisfy the RS distributional equations

$$\hat{\zeta} \overset{\mathrm{d}}{=} \frac{1 - \zeta_1 \dots \zeta_{K-1}}{2 - \zeta_1 \dots \zeta_{K-1}} \, , \tag{20.6}$$

$$\zeta \overset{\mathrm{d}}{=} \frac{\hat{\zeta}_1 \dots \hat{\zeta}_p (1 - \hat{\zeta}_{p+1}) \dots (1 - \hat{\zeta}_{p+q})}{\hat{\zeta}_1 \dots \hat{\zeta}_p (1 - \hat{\zeta}_{p+1}) \dots (1 - \hat{\zeta}_{p+q}) + (1 - \hat{\zeta}_1) \dots (1 - \hat{\zeta}_p) \hat{\zeta}_{p+1} \dots \hat{\zeta}_{p+q}} \, . \tag{20.7}$$

Here p and q are two i.i.d. Poisson random variables with mean $K\alpha/2$ (corresponding to the sizes of $\mathcal{S}$ and $\mathcal{U}$), $\zeta_1, \dots, \zeta_{K-1}$ are i.i.d. copies of ζ, and $\hat{\zeta}_1, \dots, \hat{\zeta}_{p+q}$ are i.i.d. copies $\hat{\zeta}$.

The distributions of ζ and $\hat{\zeta}$ can be approximated using the population dynamics algorithm. The resulting samples can then be used to estimate the free-entropy density, as outlined in the exercise below.

Exercise 20.1 Argue that, within the RS assumptions, the large N limit of the Bethe free-entropy density is given by $\lim_{N\to\infty} \mathbb{F}/N = \mathrm{f}^{\mathrm{RS}} = \mathrm{f}^{\mathrm{RS}}_{\mathrm{v}} + \alpha \mathrm{f}^{\mathrm{RS}}_{\mathrm{c}} - K\alpha \mathrm{f}^{\mathrm{RS}}_{\mathrm{e}}$, where:

$$\begin{aligned}
\mathrm{f}^{\mathrm{RS}}_{\mathrm{v}} &= \mathbb{E} \log \left[\prod_{a=1}^{p} \hat{\zeta}_a \prod_{a=p+1}^{p+q} (1 - \hat{\zeta}_a) + \prod_{a=1}^{p} (1 - \hat{\zeta}_a) \prod_{a=p+1}^{p+q} \hat{\zeta}_a \right] \, , \\
\mathrm{f}^{\mathrm{RS}}_{\mathrm{c}} &= \mathbb{E} \log \left[1 - \zeta_1 \cdots \zeta_{K-1} \right] \, , \\
\mathrm{f}^{\mathrm{RS}}_{\mathrm{e}} &= \mathbb{E} \log \left[(1 - \zeta_1)(1 - \hat{\zeta}_1) + \zeta_1 \hat{\zeta}_1 \right] \, .
\end{aligned} \tag{20.8}$$

Here $\mathbb{E}$ denotes the expectation with respect to: $\zeta_1, \dots, \zeta_K$ which are i.i.d. copies of ζ; $\hat{\zeta}_1, \dots, \hat{\zeta}_{p+q}$ which are i.i.d. copies of $\hat{\zeta}$; p and q which are i.i.d. Poisson random variables with mean $K\alpha/2$.

Fig. 20.2 shows an example of the entropy density found within this approach for 3-SAT. For each value of α in a mesh, we used a population of size 10^4, and ran the algorithm for $3 \cdot 10^3$ iterations. Messages are initialized uniformly in $]0, 1[$, and the first 10^3 iterations are not used for computing the free-entropy.

The predicted entropy density is strictly positive and decreasing in α for $\alpha \leq \alpha_*(K)$, with $\alpha_*(3) \approx 4.6773$. Above $\alpha_*(K)$ the RS distributional equations do not seem to admit any solution with $\zeta, \hat{\zeta} \in [0, 1]$. This is revealed numerically by the fact that the denominator of Eq. (20.7) vanishes during the population updates. Since one finds a RS entropy density which is positive for all $\alpha < \alpha_*(K)$,

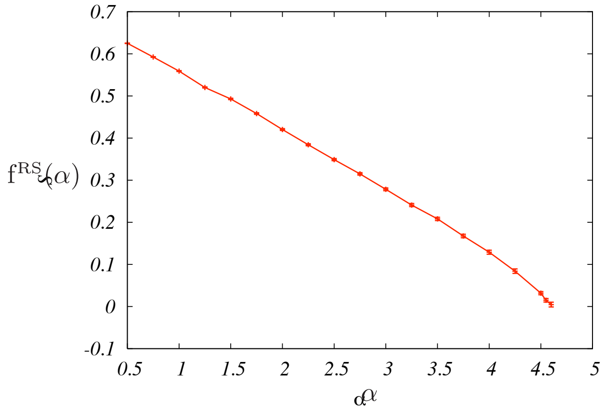

Fig. 20.2. RS prediction for the asymptotic entropy density of random 3-SAT formulae, plotted versus the clause density α for 3-SAT. The result is expected to be correct for $\alpha \leq \alpha_{\rm c}(3) = \alpha_{\rm d}(3) \approx 3.86$.

the value $\alpha_*(K)$ is the RS prediction for the SAT-UNSAT threshold. It turns out that $\alpha_*(K)$ can be computed without population dynamics, as outlined by the exercise below.

Exercise 20.2 How to compute $\alpha_*(K)$? The idea is that above this value of the clause density any solution of the RS distributional equations has $\hat{\zeta} = 0$ with positive probability. In this case the denominator of Eq. (20.7) vanishes with positive probability, leading to a contradiction.

We start by regularizing Eq. (20.7) with a small parameter ϵ: Each $\hat{\zeta}_i$ is replaced by $\max(\hat{\zeta}_i, \epsilon)$. Let us denote by x the probability that $\hat{\zeta}$ is of order ϵ, and by y the probability that ζ is of order $1 - \epsilon$. Consider the limit $\epsilon \to 0$.

(*a*) Show that $x = y^{K-1}$

(*b*) Show that $1 - 2y = e^{-K\alpha x} I_0(K\alpha x)$, where $I_0(z)$ is the Bessel function with Taylor expansion $I_0(t) = \sum_{p=0}^{\infty} \frac{1}{p!^2} \left(\frac{t}{2}\right)^{2p}$.

[Hint: Suppose that there are p' variables among $\hat{\zeta}_1 \dots \hat{\zeta}_p$, and q' among $\hat{\zeta}_{p+1} \dots \hat{\zeta}_{p+q}$, that are of order ϵ. Show that this update equation gives $\zeta = O(\epsilon)$ if $p' > q'$, $\zeta = 1 - O(\epsilon)$ if $p' < q'$, and $\zeta = O(1)$ when $p' = q'$.

(*c*) Let $\alpha_*(K)$ the largest clause density such that the two equations derived in (*a*) and (*b*) admit the unique solution $x = y = 0$. Show that, for $\alpha \geq \alpha_*(K)$ a new solution appears with $x, y > 0$.

(*d*) By solving numerically the above equations show that $\alpha_*(3) \approx 4.6673$ and $\alpha_*(4) \approx 11.83$.

Unhappily this RS computation is incorrect at α large enough, and, as a consequence, the prediction for the SAT-UNSAT phase transition is wrong as

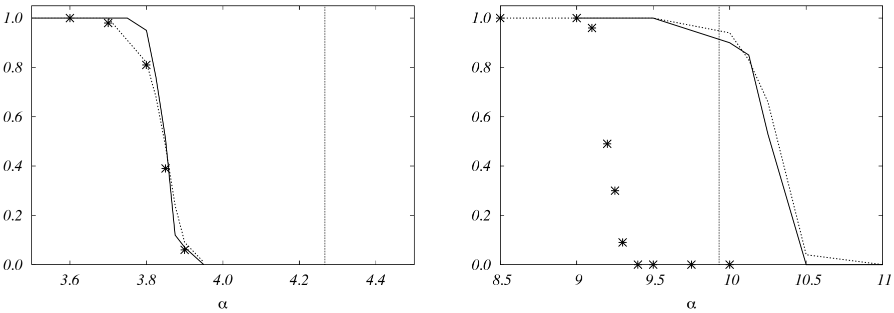

FIG. 20.3. Empirical probability that BP converges to a fixed point, plotted versus the clause density α, for 3-SAT (left plot) and 4-SAT (right plot). The statistics is over 100 instances, with $N = 5 \cdot 10^3$ variables (dashed curve) and $N = 10^4$ variables (full curve). There is an indication of a phase transition occurring for $\alpha_{\rm BP} \approx 3.85$ ($K = 3$) and $\alpha_{\rm BP} \approx 10.3$ ($K = 4$.)

Data points show the empirical probability that BP-guided decimation finds a SAT assignment, computed over 100 instances with $N = 5 \cdot 10^3$. The vertical lines correspond to the SAT-UNSAT threshold.

well. In particular, it contradicts the upper bound $\alpha_{\rm UB,2}(K)$, found in Ch. 10 (for instance, in the two cases $K = 3, 4$, one has $\alpha_{\rm UB,2}(3) \approx 4.66603 < \alpha_*(3)$, and $\alpha_{\rm UB,2}(4) \approx 10.2246 < \alpha_*(4)$). The largest α such that the RS entropy density is correct is nothing but the condensation transition $\alpha_{\rm c}(K)$. We will further discuss this phase transition below and in Ch. 22.

There is another way to realize that something is wrong with the RS assumption close to the SAT-UNSAT phase transition. The idea is to look at the BP iteration.

20.1.3 *BP-Guided Decimation*

The simplest experiment consists in iterating the BP equations (20.3) on a randomly generated K-SAT instance. We start from uniformly random messages, and choose the following convergence criterion defined in terms of a small number δ: The iteration is halted at the first time $t_*(\delta)$ such that no message has changed by more than δ over the last iteration.

Fixing a large time $t_{\max}$, one can estimate the probability of convergence within $t_{\max}$ iterations by repeating the same experiment many times. Fig.20.3 shows this probability for $\delta = 10^{-2}$ and $t_{\max} = 10^3$, plotted versus α. The probability curves show a sharp decrease around a critical value of α, $\alpha_{\rm BP}$ which is robust to variations of $t_{\max}$ and δ. This numerical result is indicative of a threshold behavior: The typical convergence time $t_*(\delta)$ stays finite (or grows moderately) with N when $\alpha < \alpha_{\rm BP}$. Above $\alpha_{\rm BP}$, BP fails to converge in a time $t_{\max}$ on a typical random instance.

When it converges, BP can be used in order to find a SAT assignment, using it as an heuristic guide for a sequential decimation procedure. Each time the value of a new variable has to be fixed, BP is iterated until the convergence criterion, with parameter δ, is met (alternatively, one may be more bold and use the BP messages after a time $t_{\max}$ even when they have not converged). Then one uses the BP messages in order to decide: (*i*) Which variable to fix; (*ii*) Which value should the variable take.

In the present implementation these decisions are taken on the basis of a simple statistics: the variables bias. Given the BP estimate $\nu_i(\,\cdot\,)$ of the marginal of x_i, we define the bias as $\pi_i \equiv \nu_i(0) - \nu_i(1)$.

```
BP-Guided Decimation (SAT formula $\mathcal{F}$, Accuracy $\epsilon$, Iterations $t_{\max}$)
For all $n \in \{1, \dots, N\}$:
    Call BP($\mathcal{F}$,$\epsilon$, $t_{\max}$);
    If BP does not converge, return 'NOT found' and exit;
    For each variable node $j$, compute the bias $\pi_j$;
    Find a variable $i(n)$ with the largest absolute bias $|\pi_{i(n)}|$;
    If $\pi_{i(n)} \geq 0$, fix $x_{i(n)}$ to $x^*_{i(n)} = 0$;
    Otherwise, fix $x_{i(n)}$ to $x^*_{i(n)} = 1$;
    Replace $\mathcal{F}$ by the formula obtained after this reduction
End-For;
Return the assignment $\underline{x}^*$
```

A pseudocode for BP was given in Sec. 14.2. Let us emphasize that the same decimation procedure could be used not only with BP, but with other types of guidance, as soon as we have some way to estimate the marginals.

The empirical success probability of the BP-Guided decimation on random formulae are shown in Fig. 20.3 (estimated from 100 instances of size $N = 5 \cdot 10^4$) for several values of α. The qualitative difference between 3-SAT and 4-SAT emerges clearly from this data. In 3-SAT, the decimation procedure returns a SAT assignment about every time it converges, i.e. with probability close to one for $\alpha \lesssim 3.85$. In 4-SAT, BP converges most of the times if $\alpha \lesssim 10.3$. This value is larger than the conjectured SAT-UNSAT threshold $\alpha_{\rm s}(4) \approx 9.931$ (and also larger than the best rigorous upper bound $\alpha_{\rm UB,2}(4) \approx 10.2246$.) On the other hand, the BP guided decimation finds SAT assignments only when $\alpha \lesssim 9.25$. It is believed that the cases $K \geq 5$ behave as $K = 4$.

20.1.4 *On the validity of the RS analysis*

These experiments suggest that something is not correct in the RS assumptions for α large enough. The precise mechanism by which they are incorrect depends however on the value of K. For $K = 3$, the BP fixed point become unstable, and this leads to errors in decimations. In fact, the local stability of the BP fixed point can be computed along the lines of Sec. 17.4.2. The result is that it become unstable at $\alpha_{\rm st}(3) \approx 3.86$. On the contrary, for $K \geq 4$ the fixed point remains stable but does not correspond to the correct marginals. Local stability is not a good enough test in this case.

Correspondingly, one can define two type of thresholds:

(*i*) A stability threshold $\alpha_{\rm st}(K)$ beyond which BP does not have a locally stable fixed point.

(*ii*) A 1RSB condensation threshold $\alpha_{\rm c}(K)$ beyond which there is no BP fixed point giving a correct estimate of the local marginals and free-entropy.

One should clearly have $\alpha_{\rm c}(K) \leq \alpha_{\rm st}(K)$. Our study suggests that $\alpha_{\rm c}(3) = \alpha_{\rm st}(3) \simeq 3.86$ while, for $K \geq 4$, one has a strict inequality $\alpha_{\rm c}(K) < \alpha_{\rm st}(K)$.

The reason for the failure of BP is the decomposition of the measure (20.1) in many pure states. This happens at a third critical value $\alpha_{\rm d}(K) \leq \alpha_{\rm c}(K)$, referred to as the dynamical transition, in accordance with our discussion of spin glasses in Sec. 12.3: $\alpha_{\rm d}(K)$ is the critical clause density above which Glauber dynamics will become inefficient. If $\alpha_{\rm d}(K) < \alpha < \alpha_{\rm c}(K)$, one expects, as we discussed in Sec. 19.4.1, that there exist many pure states, and many quasi-solutions to BP equations among which one will give the correct marginals.

At this point the reader might well be discouraged. This is understandable: we started seeking one threshold (the SAT-UNSAT transition $\alpha_{\rm s}(K)$) and rapidly ended up defining a number of other thresholds, $\alpha_{\rm d}(K) \leq \alpha_{\rm c}(K) \leq \alpha_{\rm st}(K) \leq \alpha_{\rm s}(K)$ to describe a zoology of exotic phenomena. It turns out that, while the understanding of the proliferation of pure states is necessary to get the correct value of $\alpha_{\rm s}(K)$, one does not need a detailed description of the clusters, which is a challenging task. Luckily, there exists a *shortcut*, through the use of the energetic cavity method. It turns out that the sketchy description of clusters that we get from this method, as if looking at them *from far*, is enough to determine $\alpha_{\rm s}$. Even more than that. The sketch will be a pretty useful and interesting one. In Sec. 20.3, we will discuss a more detailed picture obtained through the full-fledged 1RSB cavity method applied to the model (20.1).

20.2 Survey propagation and the 1RSB phase

The use of the energetic 1RSB cavity method can be motivated in two ways. From a first point of view, we are changing problem. Instead of computing marginals of the distribution (20.1), we consider the problem of minimizing the energy function

$$E(\underline{x}) = \sum_{a=1}^{M} E_a(\underline{x}_{\partial a}) \,. \tag{20.9}$$

Here $E_a(\underline{x}_{\partial a}) = 0$ if clause a is satisfied by the assignment $\underline{x}$, and $E_a(\underline{x}_{\partial a}) = 1$ otherwise. The SAT-UNSAT threshold $\alpha_{\rm s}(K)$ is thus identified as the critical value above which the ground state energy $\min E(\underline{x})$ vanishes.

With the cavity method we shall estimate the ground state energy density, and find that it vanishes below some threshold. This is then identified as $\alpha_{\rm s}(K)$. This identification amounts to assuming that, for generic large random K-SAT problems, there is no interval of α where the ground state energy is positive but sub-linear in N. This assumption is reasonable, but of course it does not hold in more general situations. If, for instance, we added to a random K-SAT formula a small unsatisfiable sub-formula (including $o(N)$ variables), our approach would not detect the change, while the formula would be always unsatisfiable.

For $\alpha < \alpha_{\rm s}(K)$ the cavity method provides a rough picture of zero energy pure states. This brings us to the second way of motivating this 'sketch.' We saw that describing a pure (Bethe) state in a locally tree-like graph amounts to assigning a set of cavity messages, i.e. of marginal distributions for the variables. The simplified description of the energetic 1RSB method only distinguishes between marginals that are concentrated on a single value, and marginals that are not. The concentrated marginals are described exactly, while the other ones are just summarized by a single statement, "not concentrated".

20.2.1 *The* SP(y) *equations*

The satisfiability problem involves only hard constraints and binary variables. We can thus use the simplified SP(y) equations of Sec. 19.5.3. The messages are triples: $(Q_{ia}(\texttt{0}), Q_{ia}(\texttt{1}), Q_{ia}(*))$ for variable-to-function messages, and $(\widehat{Q}_{ai}(\texttt{0}), \widehat{Q}_{ai}(\texttt{1}), \widehat{Q}_{ai}(*))$ for function-to-variable messages.

In the case of K-satisfiability, these can be further simplified. The basic observation is that, if $J_{ai} = 0$ then $\widehat{Q}_{ai}(\texttt{1}) = 0$, and if $J_{ai} = 1$ then $\widehat{Q}_{ai}(\texttt{0}) = 0$. This can be shown either starting from the general formalism in Sec. 19.5.3, or reconsidering the interpretation of warning propagation messages. Recall that a "0" message means that the constraint a 'forces' variable x_i to take value 0 in order to minimize the system's energy. In K-SAT this can happen only if $J_{ai} = 0$, because $x_i = 0$ is then the value that satisfies the clause a. With this remark in mind, the function-to-variable node message can be parameterized by a single real number. We will choose it to be $\widehat{Q}_{ai}(\texttt{0})$ if $J_{ai} = 0$, and $\widehat{Q}_{ai}(\texttt{1})$ if $J_{ai} = 1$, and we shall denote it as $\widehat{Q}_{ai}$. This number $\widehat{Q}_{ai}$ is the probability that there is a warning sent from a to i which forces the value of variable x_i.

Analogously, it is convenient to adopt a parameterization of the variable-to-function message $Q_{ia}(\texttt{m})$ which takes into account the value of J_{ai}. Precisely, recall that Q_{ia} is supported on three types of messages: $\texttt{m}(0) = 0, \texttt{m}(1) > 0$, or $\texttt{m}(0) = \texttt{m}(1) = 0$, or $\texttt{m}(0) > 0, \texttt{m}(1) = 0$. Let us denote the corresponding weights as $Q_{ia}(\texttt{0})$, $Q_{ia}(*)$, $Q_{ia}(\texttt{1})$. If $J_{ai} = 0$, we then define $Q^{\rm S}_{ia} \equiv Q_{ia}(\texttt{0})$, $Q^{*}_{ia} \equiv Q_{ia}(*)$ and $Q^{\rm U}_{ia} \equiv Q_{ia}(\texttt{1})$. Vice-versa, if $J_{ai} = 1$, we let $Q^{\rm S}_{ia} \equiv Q_{ia}(\texttt{1})$, $Q^{*}_{ia} \equiv Q_{ia}(*)$ and $Q^{\rm U}_{ia} \equiv Q_{ia}(\texttt{0})$.

Below we summarize these notations with the corresponding interpretations. We emphasize that 'probability' refers here to the random choice of a pure state, cf. Sec. 19.1.

$Q^{\rm S}_{ia}$: probability that x_i is forced by the clauses $b \in \partial i \setminus a$ to satisfy a,
$Q^{\rm U}_{ia}$: probability that x_i is forced by the clauses $b \in \partial i \setminus a$ to violate a,
Q^{*}_{ia}: probability that x_i is not forced by the clauses $b \in \partial i \setminus a$.

$\widehat{Q}_{ai}$: probability that x_i is forced by clause a to satisfy it.

The 1RSB cavity equations have been written in Sec. 19.5.3.

Exercise 20.3 Write explicitly the 1RSB equations in terms of the messages $Q^{\rm S}, Q^{\rm U}, Q^*, \widehat{Q}$ applying the procedure of Sec. 19.5.3.

Alternatively, they can be guessed having in mind the above interpretation. Clause a forces variable x_i to satisfy it if and only if all the other variables entering clause a are forced (by some other clause) not to satisfy a. This means:

$$\widehat{Q}_{ai} = \prod_{j \in \partial a \setminus i} Q^{\rm U}_{ja}\,. \tag{20.10}$$

Consider on the other hand variable node i, and assume for definiteness that $J_{ia} = 0$ (the opposite case gives rise to identical equations). Remember that, in this case, $\mathcal{S}_{ia}$ denotes the subset of clauses $b \neq a$ in which $J_{ib} = 0$, and $\mathcal{U}_{ia}$ the subset in which $J_{ib} = 1$. Assume that the clauses in $\Omega^{\rm S} \subseteq \mathcal{S}_{ia}$, and $\Omega^{\rm U} \subseteq \mathcal{U}_{ia}$ force x_i to satisfy them. Then x_i is forced to satisfy or violate a depending whether $|\Omega^{\rm S}| > |\Omega^{\rm U}|$ or $|\Omega^{\rm S}| < |\Omega^{\rm U}|$. Finally, x_i is not forced if $|\Omega^{\rm S}| = |\Omega^{\rm U}|$. The energy shift is equal to the number of 'forcing' clauses in $\partial i \setminus a$ that are violated when x_i is chosen to satisfy the largest number of them, namely $\min(|\Omega^{\rm U}|, |\Omega^{\rm S}|)$. We thus get the equations

$$Q^{\rm U}_{ia} \cong \sum_{|\Omega^{\rm U}|>|\Omega^{\rm S}|} e^{-y|\Omega^{\rm S}|} \prod_{b \in \Omega^{\rm U}\cup\Omega^{\rm S}} \widehat{Q}_{bi} \prod_{b \notin \Omega^{\rm U}\cup\Omega^{\rm S}} (1-\widehat{Q}_{bi})\,, \tag{20.11}$$

$$Q^{\rm S}_{ia} \cong \sum_{|\Omega^{\rm S}|>|\Omega^{\rm U}|} e^{-y|\Omega^{\rm U}|} \prod_{b \in \Omega^{\rm U}\cup\Omega^{\rm S}} \widehat{Q}_{bi} \prod_{b \notin \Omega^{\rm U}\cup\Omega^{\rm S}} (1-\widehat{Q}_{bi})\,, \tag{20.12}$$

$$Q^{*}_{ia} \cong \sum_{|\Omega^{\rm U}|=|\Omega^{\rm S}|} e^{-y|\Omega^{\rm U}|} \prod_{b \in \Omega^{\rm U}\cup\Omega^{\rm S}} \widehat{Q}_{bi} \prod_{b \notin \Omega^{\rm U}\cup\Omega^{\rm S}} (1-\widehat{Q}_{bi})\,. \tag{20.13}$$

The overall normalization is fixed by the condition $Q^{\rm U}_{ia} + Q^{*}_{ia} + Q^{\rm S}_{ia} = 1$.

As usual, Eqs (20.10-20.13) can be understood either as defining a mapping from the space of messages $\{\widehat{Q}_{ai}, Q_{ia}\}$ onto itself or as a set of fixed point conditions. In both cases they are referred to as the $\mathsf{SP(y)}$ equations for the satisfiability problem. From the computational point of view, these equations involve a sum over $2^{|\partial i|-1}$ terms. This is often too much if we want to iterate the $\mathsf{SP(y)}$ equations on large K-SAT formulae: the average degree of a variable node in a random K-SAT formula with clause density α is $K\alpha$. Further, in the most interesting regime –close to the SAT-UNSAT threshold– $\alpha = \Theta(2^K)$, and the sum is over $2^{\Theta(K2^K)}$ terms, which becomes rapidly unpractical. It is thus important to notice that the sums can be computed efficiently by interpreting them as convolutions.

Exercise 20.4 Consider a sequence of independent Bernoulli random variables $X_1, \dots, X_n, \dots$, with means (respectively) $\eta_1, \dots, \eta_n, \dots$. Let $W_n(m)$ be the probability that the sum $\sum_{b=1}^{n} X_b$ is equal to m.

(a) Show that these probabilities satisfy the recursion

$$W_n(m) = \eta_n W_{n-1}(m-1) + (1-\eta_n) W_{n-1}(m)\,,$$

for $m \in \{0, \dots, n\}$. Argue that these identities can be used together with the initial condition $W_0(m) = \mathbb{I}(m=0)$, to compute $W_n(m)$ in $O(n^2)$ operations.

(b) How can one compute the right hand sides of Eqs. (20.11-20.13) in $O(|\partial i|^2)$ operations?

20.2.2 *The free-entropy* $\mathbb{F}^{\mathrm{RSB,e}}$

Within the 1RSB energetic cavity method, the free-entropy $\mathbb{F}^{\mathrm{RSB,e}}(\{Q, \widehat{Q}\})$ provides detailed information on the minimal energy of (Bethe) pure states. These pure states are nothing but metastable minima of the energy function (i.e. minima whose energy cannot be decreased with a bounded number of spin flips).

The 1RSB free-entropy is expressed in terms of a set of messages $\{Q_{ia}, \widehat{Q}_{ai}\}$ that provide a (quasi-)solution of the $\mathsf{SP}(\mathtt{y})$ equations (20.10-20.13). Following the general theory in Sec. 19.5.2, it can be written in the form

$$\mathbb{F}^{\mathrm{RSB,e}}(\{Q, \widehat{Q}\}) = \sum_{a\in C} \mathbb{F}_a^{\mathrm{RSB,e}} + \sum_{i\in V} \mathbb{F}_i^{\mathrm{RSB,e}} - \sum_{(i,a)\in E} \mathbb{F}_{ia}^{\mathrm{RSB,e}}\,. \tag{20.14}$$

Equation (19.95) yields

$$e^{\mathbb{F}_{ia}^{\mathrm{RSB,e}}} = 1 - (1 - e^{-\mathtt{y}})\widehat{Q}_{ai} Q_{ia}^{\mathrm{U}}\,. \tag{20.15}$$

The contribution $\mathbb{F}_a^{\mathrm{RSB,e}}$ defined in (19.93) can be computed as follows. The reweighting $\mathbb{F}_a^{\mathrm{e}}(\{\mathtt{m}_{ia}\})$ is always equal to 0, except for the case where all the variables in clause a receive a warning requesting that they point in the "wrong direction", namely the direction which does not satisfy the clause. Therefore:

$$e^{\mathbb{F}_a^{\mathrm{RSB,e}}} = 1 - (1 - e^{-\mathtt{y}}) \prod_{i\in\partial a} Q_{ia}^{\mathrm{U}}\,.$$

Finally, the contribution $\mathbb{F}_i^{\mathrm{RSB,e}}$ defined in (19.94) depends on the messages sent from check nodes $b \in \partial i$. Let us denote by $\Omega^{\mathrm{S}} \subseteq \partial_0 i$ the subset of check nodes

$b \in \partial_0 i$ such that clause b forces x_i to satisfy it. Similarly, defined as $\Omega^{\rm U} \subseteq \partial_1 i$ the subset of $\partial_1 i$ such that clause b forces x_i to satisfy it. We then have:

$$e^{\mathbb{F}_i^{\rm RSB,e}} = \sum_{\Omega^{\rm U},\Omega^{\rm S}} e^{-{\tt y}\min(\Omega^{\rm S},\Omega^{\rm U})} \left[\prod_{b\in\Omega^{\rm U}\cup\Omega^{\rm S}} \widehat{Q}_{bi}\right] \left[\prod_{b\notin\Omega^{\rm U}\cup\Omega^{\rm S}} (1-\widehat{Q}_{bi})\right] . \tag{20.16}$$

Exercise 20.5 Show that, for any $i \in \partial a$, $\mathbb{F}_{ia}^{\rm RSB,e} = \mathbb{F}_a^{\rm RSB,e}$.

20.2.3 *Large* ${\tt y}$ *limit: the SP equations*

Consider now the case of satisfiable instances. A crucial problem is then to characterize satisfying assignments and to find them efficiently. This amounts to focusing on zero energy assignments, which are selected by taking the ${\tt y} \to \infty$ limit within the energetic cavity method.

We can take the limit ${\tt y} \to \infty$ in the ${\sf SP}({\tt y})$ equations (20.11-20.13). This yields

$$\widehat{Q}_{ai} = \prod_{j\in\partial a\setminus i} Q^{\rm U}_{ja}\,, \tag{20.17}$$

$$Q^{\rm U}_{ja} \cong \prod_{b\in\mathcal{S}_{ja}} (1-\widehat{Q}_{bj}) \left[1-\prod_{b\in\mathcal{U}_{ja}}(1-\widehat{Q}_{bj})\right] , \tag{20.18}$$

$$Q^{\rm S}_{ja} \cong \prod_{b\in\mathcal{U}_{ja}} (1-\widehat{Q}_{bj}) \left[1-\prod_{b\in\mathcal{S}_{ja}}(1-\widehat{Q}_{bj})\right] , \tag{20.19}$$

$$Q^{*}_{ja} \cong \prod_{b\in\partial j\setminus a} (1-\widehat{Q}_{bj})\,, \tag{20.20}$$

where the normalization is always fixed by the condition $Q^{\rm U}_{ja}+Q^{\rm S}_{ja}+Q^{*}_{ja}=1$.

The ${\tt y}=\infty$ equations have a simple interpretation. Consider a variable x_j appearing in clause a, and assume it receives a warning from clause $b \neq a$ independently with probability $\widehat{Q}_{bj}$. Then $\prod_{b\in\mathcal{S}_{ja}}(1-\widehat{Q}_{bj})$ is the probability that variable j receives no warning forcing it in the direction which satisfies clause a. The product $\prod_{b\in\mathcal{U}_{ja}}(1-\widehat{Q}_{bj})$ is the probability that variable j receives no warning forcing it in the direction which violates clause a. Therefore $Q^{\rm U}_{ja}$ is the probability that variable j receives at least one warning forcing it in the direction which violates clause a, *conditional to the fact that there are no contradictions in the warnings received by j from clauses $b \neq a$*. Analogous interpretations hold for $Q^{\rm S}_{ja}$ and Q^{*}_{ja}. Finally, $\widehat{Q}_{ai}$ is the probability that all variables in $\partial a \setminus i$ are forced in the direction violating clause a, under the same *condition of no contradiction.*

Notice that the ${\tt y}=\infty$ equations are a relatively simple modification of the BP equations in (20.3). However, the interpretation of the messages is very different in the two cases.

Finally the free-entropy in the $\mathtt{y} = \infty$ limit is obtained as

$$\mathbb{F}^{\mathrm{RSB,e}} = \sum_{a \in C} \mathbb{F}_a^{\mathrm{RSB,e}} + \sum_{i \in V} \mathbb{F}_i^{\mathrm{RSB,e}} - \sum_{(i,a) \in E} \mathbb{F}_{ia}^{\mathrm{RSB,e}}, \tag{20.21}$$

where

$$\mathbb{F}_{ia}^{\mathrm{RSB,e}} = \log\left\{1 - Q_{ia}^{\mathrm{U}} \widehat{Q}_{ai}\right\}, \tag{20.22}$$

$$\mathbb{F}_i^{\mathrm{RSB,e}} = \log\left\{\prod_{b \in \partial_0 i} (1 - \widehat{Q}_{bi}) + \prod_{b \in \partial_1 i} (1 - \widehat{Q}_{bi}) - \prod_{b \in \partial i} (1 - \widehat{Q}_{bi})\right\}, \tag{20.23}$$

$$\mathbb{F}_a^{\mathrm{RSB,e}} = \log\left\{1 - \prod_{j \in \partial a} Q_{ja}^{\mathrm{U}}\right\}. \tag{20.24}$$

Exercise 20.6 Show that, if the SP messages satisfy the fixed point equations (20.17) to (20.20), the free-entropy can be rewritten as $\mathbb{F}^{\mathrm{RSB,e}} = \sum_i \mathbb{F}_i^{\mathrm{RSB,e}} + \sum_a (1 - |\partial a|) \mathbb{F}_a^{\mathrm{RSB,e}}$.

20.2.4 *The SAT-UNSAT threshold*

The $\mathsf{SP}(\mathtt{y})$ equations (20.10-20.13) always admit a 'no warning' fixed point corresponding to $\widehat{Q}_{ai} = 0$, and $Q_{ia}^{\mathrm{S}} = Q_{ia}^{\mathrm{U}} = 0$, $Q_{ia}^{*} = 1$ for each $(i,a) \in E$. Other fixed points can be explored numerically by iterating the equations on large random formulae.

Within the cavity approach, the distribution of the message associated to a uniformly random edge (i,a) satisfies a distributional equation. As explained in Sec. 19.2.5, this distributional equation is obtained by promoting $\widehat{Q}_{ai}$, $(Q_{ia}^{\mathrm{U}}, Q_{ia}^{\mathrm{S}}, Q_{ia}^{*})$ to random variables and reading Eqs. (20.10-20.13) as equalities in distribution. The distribution can then be studied by the population dynamics of Sec. 19.2.6. It obviously admits a no-warning (or 'replica symmetric') fixed point, with $\widehat{Q} = 0$, $(Q^{\mathrm{U}}, Q^{\mathrm{S}}, Q^{*}) = (0, 0, 1)$ identically, but (as we will see) in some cases one also finds a different, 'non-trivial' fixed point distribution.

Given a fixed point, the 1RSB free-entropy density $\mathfrak{F}^{\mathrm{e}}(\mathtt{y})$ is estimated by taking the expectation of Eq. (20.14) (both with respect to degrees and fields) and dividing by N. When evaluated on the no-warning fixed point, the free-entropy density $\mathfrak{F}^{\mathrm{e}}(\mathtt{y})$ vanishes. This means that the number of clusters of SAT assignments is sub-exponential, so that the corresponding complexity density vanishes. To a first approximation, this solution corresponds to low-energy assignments forming a single cluster. Note that the energetic cavity method counts the number of clusters of SAT assignments, and not the number of SAT assignments itself (which is actually exponentially large).

Figure 20.4 shows the outcome of a population dynamics computation. We plot the free-entropy density $\mathfrak{F}^{\mathrm{e}}(\mathtt{y})$ as a function of $\mathtt{y}$ for random 3-SAT, at a few values of the clause density α. These plots are obtained initializing the

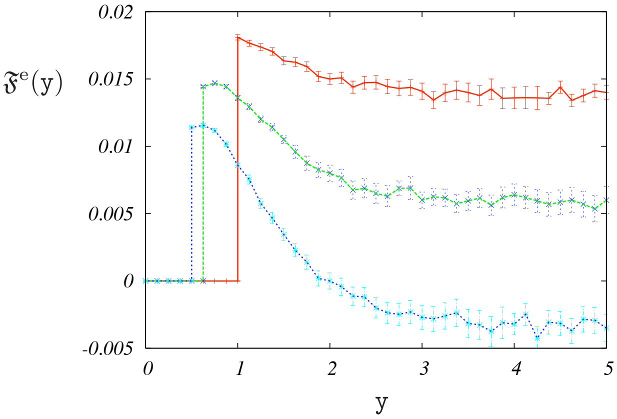

FIG. 20.4. 1RSB free-entropy density for 3-SAT, computed from the population dynamics analysis of the SP equation, at $\alpha = 4.1,\ 4.2,\ 4.3$ (from top to bottom). For each α, y, a population of size 12000 has been iterated $12 \cdot 10^6$ times. The resulting $\mathfrak{F}^{\rm e}$ has been computed by averaging over the last $8 \cdot 10^6$ iterations.

population dynamics recursion with i.i.d. messages $\{\widehat{Q}_i\}$ uniformly random in $[0, 1]$. For $\alpha < \alpha_{\rm d,SP} \simeq 3.93$, the iteration converges to the 'no-warning' fixed point where all the messages $\widehat{Q}$ are equal to 0.

For $\alpha > \alpha_{\rm d,SP}$, and when y is larger than a critical value $\mathrm{y}_{\rm d}(\alpha)$ the iteration converges to a non-trivial fixed point. This second solution has a non-vanishing value of the free-entropy density $\mathfrak{F}^{\rm e}(\mathrm{y})$. The energetic complexity $\Sigma^{\rm e}(\epsilon)$ is obtained from $\mathfrak{F}^{\rm e}(\mathrm{y})$ via the Legendre transform (19.96).

In practice, the Legendre transform is computed by fitting the population dynamics data, and then transforming the fitting curve. Good results are obtained with a fit of the form $\mathfrak{F}^{\rm e}_{\rm fit}(\mathrm{y}) = \sum_{r=0}^{r_*} \psi_r\, e^{-r\mathrm{y}}$ with r_* between 2 and 4. The resulting curves $\Sigma^{\rm e}(\epsilon)$ (or more precisely their concave branches[30]) are shown in Fig. 20.5.

Exercise 20.7 Show that $\Sigma^{\rm e}(\epsilon = 0) = \lim_{\mathrm{y}\to\infty} \mathfrak{F}^{\rm e}(\mathrm{y})$

The energetic complexity $\Sigma^{\rm e}(\epsilon)$ is the exponential growth rate number of (quasi-)solutions of the min-sum equations with energy density u. As can be seen in Fig. 20.5, for $\alpha = 4.1$ or 4.2 (and in general, in an interval above $\alpha_{\rm d}(3)$) one finds $\Sigma^{\rm e}(\epsilon = 0) > 0$. The interpretation is that there exist exponentially many solutions of the min-sum equations with zero energy density.

On the contrary when $\alpha = 4.3$ the curve starts at a positive ϵ or, equivalently the 1RSB complexity curve has $\Sigma^{\rm e}(\epsilon = 0) < 0$. Of course, the typical number

[30] $\Sigma^{\rm e}(\epsilon)$ has a second, convex branch which joins the concave part at the maximal value of ϵ; the precise meaning of this second branch is not known.

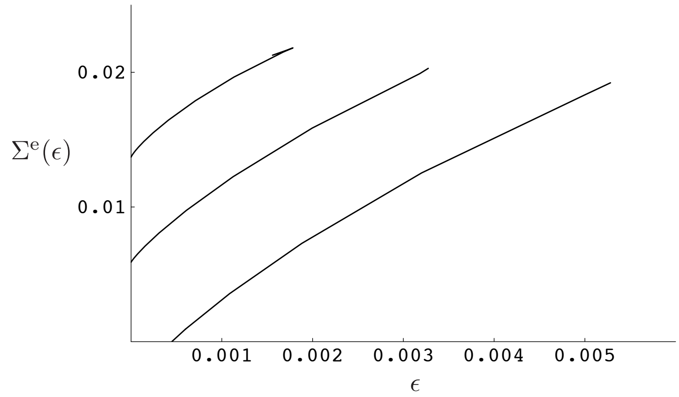

FIG. 20.5. Energetic complexity density Σ^{e} plotted versus energy density ϵ, for the 3-SAT problem at $\alpha = 4.1,\ 4.2,\ 4.3$ (from top to bottom). These curves have been obtained as the Legendre transform of the free-entropy fits of Fig. 20.4.

of min-sum solutions cannot decrease exponentially. The result $\Sigma^{\mathrm{e}}(\epsilon = 0) < 0$ is interpreted as a consequence of the fact that a typical random formula does not admit any (approximate) solution of the min-sum equations with energy density $\epsilon = 0$. Given the correspondence between min-sum fixed points and clusters of low-energy assignments, this in turns implies that a typical random formula does not have any SAT assignment.

From Fig. 20.5 one expects that the SAT-UNSAT transition lies between $\alpha = 4.2$ and $\alpha = 4.3$. A more precise estimate can be obtained by plotting $\mathfrak{F}^{\mathrm{e}}(\mathrm{y} \to \infty)$ versus α, and locating the value of α where it vanishes. For 3-SAT one obtains the SAT-UNSAT threshold estimate $\alpha_{\mathrm{s}}(3) = 4.26675 \pm 0.00015$. The predictions of this method for $\alpha_{\mathrm{s}}(K)$ are shown in the Table 20.2.4. In practice, reliable estimates can be obtained with population dynamics only for $K \leq 7$. The reason is that $\alpha_{\mathrm{s}}(K)$ increases exponentially with K, and the size of the population needed in order to achieve a given precision should increase accordingly (the average number of independent messages entering the distributional equations is $K\alpha$).

For large K, one can formally expand the distributional equations, which yields a series for $\alpha_{\mathrm{s}}(K)$ in powers of 2^{-K}. The first two terms (seven terms have been computed) of this expansion are:

$$\alpha_{\mathrm{s}}(K) = 2^K \log 2 - \frac{1}{2}(1 + \log 2) + O(2^{-K} K^2) \tag{20.25}$$

20.2.5 *SP-Guided Decimation*

The analysis in the last few pages provides a refined description of the set of solutions of random formulae. This knowledge can be exploited to efficiently

K	3	4	5	6	7	8	9	10
$\alpha_{\rm s}(K)$	4.2667	9.931	21.117	43.37	87.79	176.5	354.0	708.9

Table 20.1 *Predictions of the 1RSB cavity method for the SAT-UNSAT threshold of random K satisfiability*

find some solutions, much in the same way as we used belief propagation in Sec. 20.1.3. The basic strategy is again to use the information provided by the SP messages as a clever heuristic in a decimation procedure.

The first step consists in finding an approximate solution of the $\mathsf{SP}(\mathtt{y})$ equations (20.10-20.13), or of their simplified $\mathtt{y} = \infty$ version (20.17-20.20), on a given instance of the problem. To be definite, we shall focus on the latter case, since $\mathtt{y} = \infty$ selects zero energy states. We can seek solutions of the SP equations by iteration, exactly as we would do with BP. We initialize SP messages, generally as i.i.d. random variable with some common distribution, and then update them according to Eqs. (20.17-20.20). Updates can be implemented, for instance, in parallel, until a convergence criterion has been met.

Figure 20.6 shows the empirical probability that the iteration converges before $t_{\rm max} = 1000$ iterations on random formulae as a function of the clause density α. As a convergence criterion we required that the maximal difference between any two subsequent values of a message is smaller than $\delta = 10^{-2}$. Messages were initialized by drawing, for each edge, $\widehat{Q}_{ai} \in [0, 1]$ independently and uniformly at random. It is clear that SP has better convergence properties than BP for $K = 3$, and indeed it converges even for α larger than the SAT-UNSAT threshold.

The numerics suggests the existence of two thresholds $\alpha_{\rm d,SP}(K)$, $\alpha_{\rm u,SP}(K)$ characterizing the convergence behavior as follows (all the statements below should be interpreted as holding with high probability in the large N limit):

For $\alpha < \alpha_{\rm d,SP}$: the iteration converges to the trivial fixed point defined by $\widehat{Q}_{ai} = 0$ for all edges $(i, a) \in G$.

For $\alpha_{\rm d,SP} < \alpha < \alpha_{\rm u,SP}$: the iteration converges to a 'non-trivial' fixed point.

For $\alpha_{\rm u,SP} < \alpha$: the iteration does not converge.

In the interval $\alpha_{\rm d,SP}(K) < \alpha < \alpha_{\rm U,SP}(K)$ it is expected that an exponential number of fixed points exist but most of them will be degenerate and correspond to 'disguised' WP fixed points. In particular $\widehat{Q}_{ai} = 0$ or 1 for all the edges (i, a). On the other hand, the fixed point actually reached by iteration is stable with respect to changes in the initialization. This suggest the existence of a unique non-degenerate fixed point. The threshold $\alpha_{\rm d,SP}(K)$ is conjectured to be the same as defined for the distributional equation in the previous section, this is why we used the same name. In particular $\alpha_{\rm d,SP}(K = 3) \approx 3.93$ and $\alpha_{\rm d,SP}(K = 4) \approx 8.30$. One further obtains $\alpha_{\rm u,SP}(K = 3) \approx 4.36$ and $\alpha_{\rm u,SP}(K = 4) \approx 9.7$.

SP can be used in a decimation procedure . After iterating the SP equations until convergence, one computes the following **SP marginal** for each variable

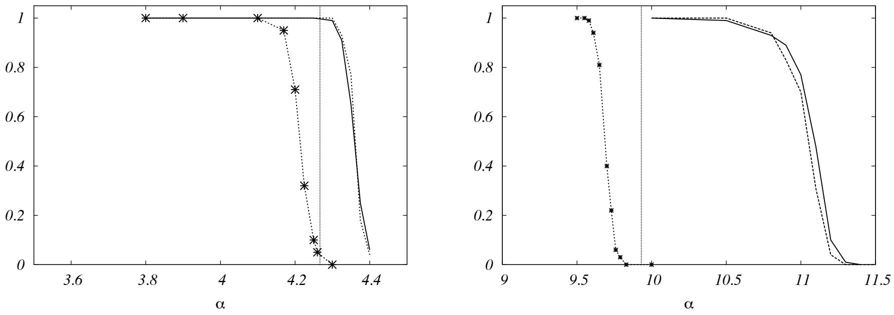

FIG. 20.6. Empirical convergence probability of SP (initialized from uniformly random messages) plotted versus the clause density α for 3-SAT (left), and 4-SAT (right). The average is over 100 instances, with $N = 5 \cdot 10^3$ (solid line) and $N = 10^4$ variables (dashed line). Data points show the empirical probability that SP-guided decimation finds a SAT assignment, computed over 100 instances with $N = 5 \cdot 10^3$. The vertical lines are the predicted SAT-UNSAT thresholds.

$i \in \{1, \dots, N\}$

$$\begin{aligned}
w_i(1) &\cong \prod_{a\in\partial_0 i}(1-\widehat{Q}_{ai})\;\left[1-\prod_{a\in\partial_1 i}(1-\widehat{Q}_{ai})\right]\,,\\
w_i(0) &\cong \prod_{a\in\partial_1 i}(1-\widehat{Q}_{ai})\;\left[1-\prod_{a\in\partial_0 i}(1-\widehat{Q}_{ai})\right]\,,\\
w_i(*) &\cong \prod_{a\in\partial i}(1-\widehat{Q}_{ai})\,, \qquad (20.26)
\end{aligned}$$

with the normalization condition $w_i(1)+w_i(0)+w_i(*) = 1$. The interpretations of these SP marginals is the following: $w_i(1)$ (resp. $w_i(0)$) is the probability that the variable i receives a warning forcing it to take the value $x_i = 1$ (resp. $x_i = 0$), *conditioned* to the fact that it does not receive contradictory warnings. The variable bias is then defined as $\pi_i \equiv w_i(0) - w_i(1)$. The variable with the largest absolute bias is selected and fixed according to the bias sign. This procedure is then iterated as with BP-guided decimation.

It typically happens that, after fixing some fraction of the variables with this method, the SP iteration on the reduced instance converges to the trivial fixed point $\widehat{Q}_{ai} = 0$. According to our interpretation, this means that the resulting problem is described by a unique Bethe measure, and SAT assignments are no longer clustered. In fact, in agreement with this interpretation, one finds that, typically, simple algorithms are able to solve the reduced problem. A possible approach is to run BP guided decimation. An even simpler alternative is to apply a simple local search algorithms, like Walksat or simulated annealing.

The pseudocode for this algorithm is as follows.

```
SP-Guided Decimation (Formula $\mathcal{F}$, SP parameter $\epsilon$, $t_{\rm max}$,
                                   WalkSAT parameters $f$, $p$)
Set $U = \emptyset$;
Repeat until FAIL or $U = V$:
    Call SP$(\mathcal{F}, \epsilon, t_{\rm max})$. If it does not converge, FAIL;
    For each $i \in V \setminus U$ compute the bias $\pi_i$;
    Let $j \in V \setminus U$ have the largest value of $|\pi_i|$;
    If $|\pi_j| \leq 2K\epsilon$ call WalkSAT$(\mathcal{F}, f, p)$;
      Else fix $x_j$ according to the sign of $\pi_j$,
         and define $\mathcal{F}$ as the new formula obtained after fixing $x_j$;
End-Repeat;
Return the current assignment;
```

```
SP (Formula $\mathcal{F}$, Accuracy $\epsilon$, Iterations $t_{\rm max}$ )
Initialize SP messages to i.i.d. random variables;
For $t \in \{0, \dots, t_{\rm max}\}$
    For each $(i, a) \in E$
      Compute the new value of $\widehat{Q}_{ai}$ using Eq. (20.10)
    For each $(i, a) \in E$
      Compute the new value of $Q_{ai}$ using Eqs. (20.11-20.13)
    Let $\Delta$ be the maximum difference with previous iteration;
    If $\Delta < \epsilon$ return current messages;
End-For;
Return 'Not Converged';
```

The WalkSAT pseudocode was given in Sec. 10.2.3.

In Fig. 20.6 we plot the empirical success probability of SP-Guided Decimation for random 3-SAT and 4-SAT formulae as a function of the clause density α. A careful study suggests that the algorithm finds a satisfying assignment with high probability when $\alpha \lesssim 4.252$ (for $K = 3$) and $\alpha \lesssim 9.6$ (for $K = 4$). These values are slightly smaller than the conjectured locations of the SAT-UNSAT threshold $\alpha_{\rm s}(3) \approx 4.2667$ and $\alpha_{\rm s}(4) \approx 9.931$.

Apart from the SP routine (that builds upon the statistical mechanics insight) the above algorithm is quite naive and could be improved in a number of directions. One possibility is to allow the algorithm to backtrack, i.e. to release some variables that had been fixed at a previous stage of the decimation. Further, we did not use at any step the information provided by the free-entropy $\mathfrak{F}^{\rm e}(\mathrm{y} = \infty)$ that can be computed at little extra cost. Since this gives an estimate of the logarithm of the number solutions clusters, it can also be reasonable to make choices that maximize the value of $\mathfrak{F}^{\rm e}$ in the resulting formula.

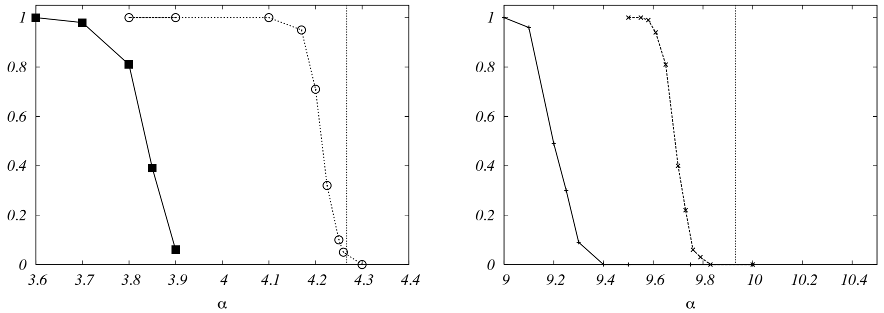

FIG. 20.7. Performance of BP-inspired decimation and SP-inspired decimation on 3-SAT (left plot) and 4-SAT (right plot) problems. Probability of finding a SAT assignment versus clause density, averaged over 100 instances with $N = 5 \cdot 10^3$ variables. The SP based algorithm (dotted line) performs better than the BP based one (full line). The vertical lines are the SAT-UNSAT thresholds.

As can be deduced from Fig. 20.7, SP-Guided Decimation outperforms BP-Guided Decimation. Empirically this algorithm, or small variations of it, provide the most efficient procedure for solving large random K-SAT formulae close to the SAT-UNSAT threshold. Furthermore, it has extremely low complexity. Each SP iteration requires $O(N)$ operations, which yields $O(Nt_{\max})$ operations per SP call. In the implementation outlined above this implies a $O(N^2 t_{\max})$ complexity. This can however be reduced to $O(Nt_{\max})$ by noticing that fixing a single variable does not affect the SP messages significantly. As a consequence, SP can be called every $N\delta$ decimation steps for some small δ. Finally, the number of iterations required for convergence seem to grow very slowly with N, if it does at all. One should probably think of $t_{\max}$ as a big constant or $t_{\max} = O(\log N)$

In order to get a better understanding of how SP-guided decimation works, it is useful to monitor the evolution of the energetic complexity curve $\Sigma^{\rm e}(\epsilon)$ while decimating. When SP iteration has converged on a given instance, one can use (20.21) to compute the free-entropy, and by a Legendre transform the curve $\Sigma^{\rm e}(\epsilon)$.

In Fig. 20.8 we consider a run of SP-Guided Decimation on one random 3-SAT formula with $N = 10^4$ at $\alpha = 4.2$. the complexity curve of the residual formula ($N\Sigma^{\rm e}(\epsilon)$ versus the number of violated clauses $N\epsilon$) is plotted every 1000 decimation steps. One notices two main effects: (1) The zero-energy complexity $N\Sigma^{\rm e}(0)$ decreases, showing that some clusters of solutions are lost along the decimation; (2) The number of violated clauses in the most numerous metastable clusters, the so-called 'threshold energy', decreases as well[31], implying that the

[31] Because of the instability of the 1RSB solution at large energies (see Chapter 22), the threshold energies obtained within the 1RSB approach are not exact. However one expects the actual behavior to be quantitatively close to the 1RSB description.

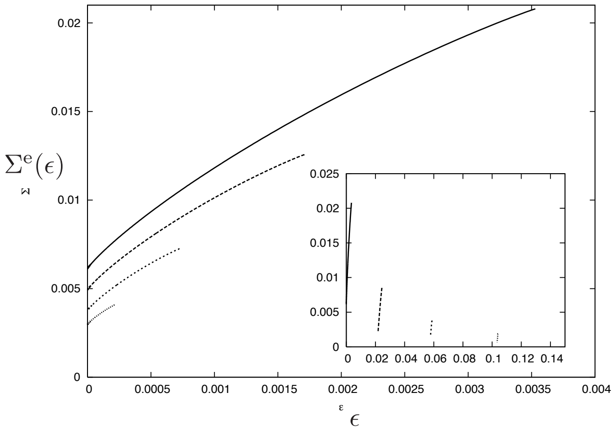

FIG. 20.8. Decimation process: The complexity versus energy density ($1/N$ times the number of violated clauses) measured on a single instance of random 3-SAT with $N = 10000$ and $\alpha = 4.2$ (top curve), and on the decimated instances obtained after fixing $1000, 2000, 3000$ variables with the survey inspired decimation procedure (from top to bottom). For comparison, the inset shows the same complexity versus total energy after fixing to arbitrary values $1000, 2000, 3000$ randomly chosen variables

problem becomes simpler: the true solutions are less and less hidden among metastable minima.

The important point is that the effect (2) is much more pronounced than (1). After fixing about half of the variables, the threshold energy vanishes. SP converges to the trivial fixed point, the resulting instance becomes 'simple,' and is solved easily by Walksat.

20.3 Some ideas on the full phase diagram

20.3.1 *Entropy of clusters*

The energetic 1RSB cavity method has given two important results: on one hand, a method to locate the SAT-UNSAT transition threshold α_{s}, which is conjectured to be exact, on the other, a powerful message passing algorithm: SP. These results were obtained at a cost: we completely forgot about the size of the clusters of SAT assignments, their 'internal entropy'.

In order to get a finer understanding of geometry of the set of solutions in the SAT phase, we need to get back to the uniform measure over SAT assignments of (20.1), and use the 1RSB method of Sec. 19.2. Our task is in principle straightforward: we need to estimate the 1RSB free entropy $\mathfrak{F}(\mathtt{x})$, and perform the Legendre transform (19.8) in order to get the complexity function $\Sigma(\phi)$. Recall that $\Sigma(\phi)$ is the exponential growth rate of the number of clusters with free-entropy $N\phi$ (in the present case, since we restrict to SAT configurations, the free-entropy of a cluster is equal to its entropy).

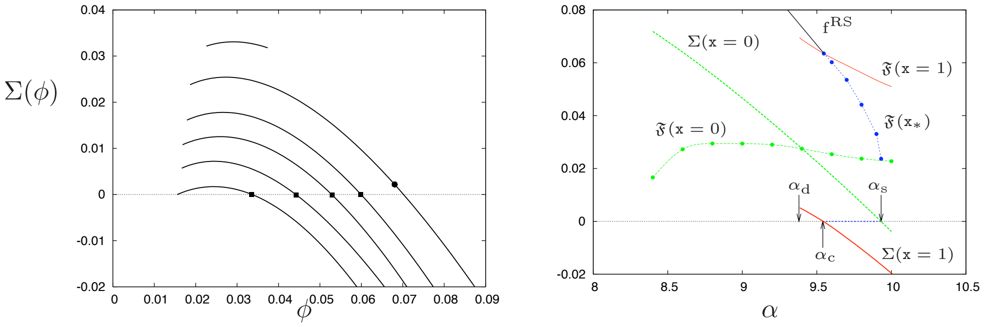

FIG. 20.9. 1RSB analysis of random 4-SAT. Left: Complexity versus internal entropy density of clusters, for $\alpha = 9.3, 9.45, 9.6, 9.7, 9.8, 9.9$ (from top to bottom). When sampling uniformly SAT configurations, one finds either configurations in an exponentially large number of clusters (dot on the curve $\alpha = 9.45$, which is the point where $\mathrm{d}\Sigma/\mathrm{d}\phi = -1$), or a condensed phase where the measure is dominated by a few clusters (squares on the curves with $\alpha \geq 9.6$). Right: Complexity $\Sigma(\mathtt{x})$ and free-entropy density $\mathfrak{F}(\mathtt{x})$ at a few key values of $\mathtt{x}$: $\mathtt{x} = 0$ corresponds to the maximum of $\Sigma(\phi)$, $\mathtt{x} = 1$ to the point with $\mathrm{d}\Sigma/\mathrm{d}\phi = -1$, and $\mathtt{x} = \mathtt{x}_*$ to $\Sigma(\phi) = 0$. The dynamical transition is at $\alpha_{\mathrm{d}} \approx 9.38$, the condensation transition at $\alpha_{\mathrm{c}} \approx 9.547$, and the SAT-UNSAT transition at $\alpha_{\mathrm{s}} \approx 9.931$.

This is a rather demanding task from the numerical point of view. Let us understand why: each BP message is parameterized by one real number in $[0, 1]$, as we saw in (20.3). A 1RSB message characterizes the distribution of this number, so it is a pdf on $[0, 1]$. One such distribution is associated to each directed edge of the factor graph. For the study of the phase diagram, one needs to perform a statistical analysis of the 1RSB messages. Within the population dynamics approach this means that we must use a (large) population of distribution functions. For each value of $\mathtt{x}$, the algorithm must be run for a large enough number of iterations to estimate $\mathfrak{F}(\mathtt{x})$. This is at the limit of what can be done numerically. Fortunately it can be complemented by two simpler computations: the SP approach which gives the results corresponding to $\mathtt{x} = 0$, and the study of the $\mathtt{x} = 1$ case using the simplification described in Sec. 19.4.

20.3.2 *The condensation transition for $K \geq 4$*

We shall not provide any technical detail of these computations, but focus on the main results using $K = 4$-SAT as a running example. As shown by Fig. 20.9, this system displays the full scenario of phase transitions explained in Sec. 19.6. Upon increasing the clause density α, one finds first a RS phase for $\alpha < \alpha_{\mathrm{d}}$, then a d1RSB phase with exponentially many relevant states for $\alpha_{\mathrm{d}} < \alpha < \alpha_{\mathrm{c}}$, then a s1RSB phase with condensation of the measure on a few states, for $\alpha_{\mathrm{c}} < \alpha < \alpha_{\mathrm{s}}$. The system becomes UNSAT for $\alpha > \alpha_{\mathrm{s}}$.

Fig. 20.9 shows the evolution of the complexity versus internal entropy density of the clusters when α increases (note that increasing α plays the same role as decreasing the temperature in the general scenario sketched in Fig. 19.6). For a given α, almost all clusters have an internal entropy density ϕ_0 corresponding to the maximum of $\Sigma(\phi)$. The complexity at the maximum, $\Sigma(\phi_0) = \mathfrak{F}(\mathtt{x} = 0)$, is equal to the complexity at zero energy density that we found with the energetic 1RSB cavity method. When sampling SAT configurations uniformly, almost all of them are found in clusters of internal entropy density ϕ_1 such that $\Sigma(\phi) + \phi$ is maximum, conditioned to the fact that $\Sigma(\phi) \geq 0$. In the d1RSB phase one has $\Sigma(\phi_1) > 0$, in the s1RSB one has $\Sigma(\phi_1) = 0$. The condensation point $\alpha_{\rm c}$ can therefore be found through a direct (and more precise) study at $\mathtt{x} = 1$. Indeed it is identified as the value of clause density such that the two equations: $\Sigma(\phi) = 0$, $\mathrm{d}\Sigma/\mathrm{d}\phi = -1$ admit a solution.

Exercise 20.8 Using the Legendre transform 19.8, show that this condensation point $\alpha_{\rm c}$ is the one where the 1RSB free-entropy function $\mathfrak{F}(\mathtt{x})$ satisfies $\mathfrak{F}(1) - \mathfrak{F}'(1) = 0$ (where $'$ means derivative with respect to $\mathtt{x}$). As we saw in Sec. 19.4, the value of $\mathfrak{F}(1)$ is equal to the RS free-entropy. As for the value of the internal entropy $\mathfrak{F}'(1)$, it can also be obtained explicitly from the $\mathtt{x} = 1$ formalism. Writing down the full $\mathtt{x} = 1$ formalism for random satisfiability, including this computation of $\mathfrak{F}'(1)$, is an interesting (non-trivial) exercise.

The dynamical transition point $\alpha_{\rm d}$ is defined as the smallest value of α such that there exists a non-trivial solution to the 1RSB equation at $\mathtt{x} = 1$ (in practice it is best studied using the point-to-set correlation which will be described in Ch. 22). Notice from Fig. 20.9 that there can exist clusters of SAT assignments even at $\alpha < \alpha_{\rm d}$: for $\alpha = 4.3$, there exists a branch of $\Sigma(\phi)$, around the point ϕ_0 where it is maximum, but this branch disappears, if one increases ϕ, before one can find a point where $\mathrm{d}\Sigma/\mathrm{d}\phi = -1$. The interpretation of this regime is that an exponentially small fraction of the solutions are grouped in well separated clusters. The vast majority of the solutions belongs instead to a single, well connected 'replica symmetric' cluster. As we saw in the energetic cavity method, the first occurrence of the clusters around ϕ_0 occurs at the value $\alpha_{\rm d,SP}$ which is around 8.3 for 4-SAT.

The same scenario has been found in the studies of random K-SAT with $K = 5, 6$, and it is expected to hold for all $K \geq 4$. The situation is somewhat different at $K = 3$, as the condensation point $\alpha_{\rm c}$ coincides with $\alpha_{\rm d}$: the 1RSB phase is always condensed. Table 20.3.2 summarizes the values of the thresholds.

20.4 An exercise: coloring random graphs

Recall that a proper q-coloring of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is an assignment of colors $\{1, \dots, q\}$ to the vertices of q in such a way that no edge has the two adjacent vertices of the same color. Hereafter we shall refer to a proper q-coloring as to a 'coloring' of $\mathcal{G}$. Colorings of a random graph can be studied following the approach just described for satisfiability, and reveal a strikingly similar behavior.

K	$\alpha_{\rm d}$	$\alpha_{\rm c}$	$\alpha_{\rm s}$
3	3.86	3.86	4.2667
4	9.38	9.547	9.931
5	19.16	20.80	21.117
6	36.53	43.08	43.37

Table 20.2 *Predictions of the 1RSB cavity method for the non-trivial SP, dynamical, condensation, and SAT-UNSAT threshold of random K-satisfiability*

Here we shall just present some key steps of this analysis: this section can be seen as a long exercise in applying the cavity method. We shall focus on the case of random regular graphs, which is technically simpler. In particular, many results can be derived without resorting to a numerical resolution of the cavity equations. The reader is encouraged to work out the many details which are left aside.

We shall adopt the following description of the problem: to each vertex $i \in \mathcal{V}$ of a graph $G = (\mathcal{V}, \mathcal{E})$, associate a variable $x_i \in \{1, \cdots, q\}$. The energy of a color assignment $\underline{x} = \{x_1, \cdots, x_N\}$ is given by the number of edges whose vertices have the same color:

$$E(\underline{x}) = \sum_{(ij)\in\mathcal{E}} \mathbb{I}(x_i = x_j) \ . \tag{20.27}$$

If the graph is colorable, one is also interested in the uniform measure over proper colorings:

$$\mu(\underline{x}) = \frac{1}{Z}\, \mathbb{I}(E(\underline{x}) = 0) = \frac{1}{Z} \prod_{(ij)\in\mathcal{E}} \mathbb{I}(x_i \neq x_j) \ , \tag{20.28}$$

where Z is the number of proper colorings of $\mathcal{G}$. The factor graph associated with $\mu(\,\cdot\,)$ is easily constructed. Associate one variable node to each vertex of $i \in \mathcal{G}$, one function node to each edge $(ij) \in \mathcal{C}$, and connect this function it to the variable nodes corresponding to i and j. The probability distribution $\mu(\underline{x})$ is therefore a pairwise graphical model.

We will assume that $\mathcal{G}$ is a random regular graphs of degree c. Equivalently, the corresponding factor graph is distributed according to the $\mathbb{D}_N(\Lambda, P)$ ensemble, with $\Lambda(x) = x^c$ and $P(x) = x^2$. The important technical simplification is that, for any fixed r, the radius-r neighborhood around a random a vertex i is with high probability a tree of degree c, i.e. it is non-random. In other words, the neighborhood of most of the nodes is the same.

Let us start with the RS analysis of the graphical model (20.28). As we saw in Sec. 14.2.5, we can get rid of function-to-variable node messages, and work with variable-to-function messages $\nu_{i\to j}(x_i)$. The BP equations read

$$\nu_{i\to j}(x) \cong \prod_{k\in\partial i\setminus j} (1 - \nu_{k\to i}(x)) \ . \tag{20.29}$$

Because of the graph regularity, there exists solutions of these equations such that messages take the same value on all edges. In particular, Eq. (20.29) admits the solution $\nu_{i\to j}(\,\cdot\,) = \nu_{\rm unif}(\,\cdot\,)$, where $\nu_{\rm unif}(\,\cdot\,)$ is the uniform messages: $\nu_{\rm unif}(x) = 1/q$ for $x \in \{1,\dots,q\}$. The corresponding free-entropy density (equal here to the entropy density) is

$$\mathrm{f}^{\rm RS} = \log q + \frac{c}{2}\log\left(1-\frac{1}{q}\right) \, . \tag{20.30}$$

It can be shown that this coincides with the 'annealed' estimate $N^{-1}\log \mathbb{E}Z$. It decreases with the degree c of the graph and becomes negative for c larger than $c_{\rm UB}(q) \equiv 2\log q/\log(q/(q-1))$, similarly to what we saw in Fig. 20.2. Markov inequality implies that, with high probability, a random c-regular graph does not admit a proper q-coloring for $c > c_{\rm UB}(q)$. Further, the RS solution is surely incorrect for $c > c_{\rm UB}(q)$.

The stability analysis of this solution shows that the spin glass susceptibility diverges as $c \uparrow c_{\rm st}(q)$, with $c_{\rm st}(q) = q^2 - 2q + 2$. For $q \geq 4$, $c_{\rm st}(q) > c_{\rm UB}(q)$.

In order to correct the above inconsistencies, one has to resort to the energetic 1RSB approach. Let us focus onto $\mathtt{y} \to \infty$ limit (equivalently, on the zero energy limit). In this limit one obtains the SPequations. This can be written in terms of messages $Q_{i\to j}(\,\cdot\,)$ that have the following interpretation

$$Q_{i\to j}(x) = \text{probability that, in absence of } (i,j),\ x_i \text{ is forced to value } x,$$
$$Q_{i\to j}(*) = \text{probability that, in absence of } (i,j),\ x_i \text{ is not forced.}$$

Recall that 'probability' is interpreted here with respect to a random Bethe state.

An SP equation express the message $Q_{i\to j}(\,\cdot\,)$ in terms of the $c-1$ incoming messages $Q_{k\to i}(\,\cdot\,)$ with $k \in \partial i \setminus j$. To keep notations simple, we fix an edge $i \to j$ and denote it by 0, while we use $1\,\dots,c-1$ to label the edges $k \to i$ with $k \in \partial i \setminus j$. Then, for any x in $\{1,\cdots,q\}$, one has:

$$Q_0(x) = \frac{\sum_{(x_1\dots x_{c-1})\in\mathcal{N}(x)}\ Q_1(r_1)Q_2(x_2)\cdots Q_{c-1}(x_{c-1})}{\sum_{(x_1\dots x_{c-1})\in\mathcal{D}}\ Q_1(r_1)Q_2(x_2)\cdots Q_{c-1}(x_{c-1})} \, . \tag{20.31}$$

where:

- $\mathcal{D}$ is the set of tuples $(x_1,\cdots,x_{c-1}) \in \{*,1,\cdots,q\}^n$ such that there exist $z \in \{1,\cdots,q\}$ with $z \neq x_1,\dots,x_{c-1}$. According to the interpretation above, this means that there is no contradiction among the warmings to i.
- $\mathcal{N}(x)$ is the set of tuples $(x_1,\cdots,x_{c-1}) \in \mathcal{D}$ such that, for any $z \neq x$ there exists $k \in \{1,\dots,c-1\}$ such that $x_k = z$. In other words, x is the only color for vertex i that is compatible with the warnings.

$Q_0(*)$ is determined by the normalization condition $Q_0(*) + \sum_x Q_0(x) = 1$.

On a random regular graph of degree c, these equations admit a solution with $Q_{i\to j}(\,\cdot\,) = Q(\,\cdot\,)$ independent of the edge (i,j). Furthermore, if we assume

q	$c_{\rm d,SP}$	$c_{\rm d}$	$c_{\rm c}$	$c_{\rm s}$
3	5	5	6	6
4	9	9	10	10
5	13	14	14	15
6	17	18	19	20
10	36	39	43	44

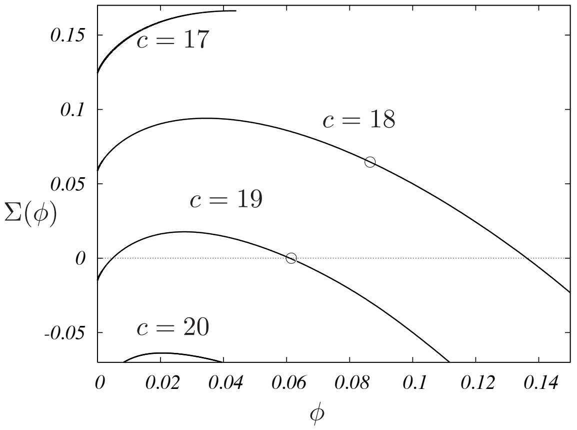

FIG. 20.10. Results of the 1RSB analysis of proper q-colorings of random regular graphs. The table gives the thresholds: appearance of non-trivial SP solutions $c_{\rm d,SP}$, dynamical $c_{\rm d}$, condensation $c_{\rm c}$, colorable/uncolorable $c_{\rm s}$. The figure shows the clusters complexity as a function of their internal entropy density. Here $q = 6$ and the graph degrees are $c = 17$ (RS), $c = 18$ (d1RSB), $c = 19$ (s1RSB) and $c = 20$ (uncolorable). The circles denote the points of slope -1 on the complexity curves.

this solution to be symmetric under permutation of colors, the corresponding message can by parameterized by a single number $a \in [0, 1/q]$:

$$\begin{aligned} &Q(x) = a \text{ for } x \in \{1, \cdots, q\}\ , \\ &Q(*) = 1 - qa\ . \end{aligned} \tag{20.32}$$

Plugging this Ansatz in Eq. (20.31), we get:

$$a = \frac{\sum_{r=0}^{q-1} (-1)^r \binom{q-1}{r} (1-(r+1)a)^{c-1}}{\sum_{r=0}^{q-1} (-1)^r \binom{q}{r+1} (1-(r+1)a)^{c-1}}\,. \tag{20.33}$$

The complexity $\Sigma^{\rm e}(\epsilon = 0)$ yielding the exponential growth rate of the number of clusters of proper colorings, is given by $\Sigma^{\rm e}(e = 0) = \lim_{\mathsf{y}\to\infty} \mathfrak{F}^{\rm e}(\mathsf{y})$. One finds:

$$\Sigma^{\rm e}(\epsilon = 0; c, q) = \log\left(\sum_{r=0}^{q-1} (-1)^r \binom{q}{r+1} (1-(r+1)a)^c\right) - \frac{c}{2}\log(1-qa^2)\,. \tag{20.34}$$

Given the number of colors q, one can study what happens when the degree c grows (which amounts to increasing the density of constraints). The situation is very similar to the one found in satisfiability. For $c \geq c_{\rm d,SP}(q)$, there exists a pair of non-trivial solution to Eq.(20.33) with $a > 0$. The complexity $\Sigma^{\rm e}(e = 0)$ can be computed from (20.34) (evaluated on the largest solution a of Eq. (20.33)), and is decrasing in c. It becomes negative for $c \geq c_{\rm s}(q)$. The degree $c_{\rm s}(q)$ is thus the 1RSB prediction for the SAT-UNSAT threshold.

When $c < c_{\rm s}(q)$, the uniform measure over valid colorings can be studied, and in particular one can characterize the distribution of entropy of clusters. Fig. 20.10 shows the complexity as function of internal entropy density of clusters. The similarity to Fig. 20.9 is obvious. One can define two particularly relevant thresholds: $c_{\rm d}$ is the smallest degree such that the 1RSB equations at $\mathtt{x} = 1$ have a non-trivial solution, and $c_{\rm c}$ is the smallest degree such that the uniform measure over proper colorings is 'condensed'. The table in Fig. 20.10 gives some examples of these thresholds. An asymptotic analysis for large q shows that:

$$c_{\rm d,SP} = q(\log q + \log\log q + 1 - \log 2 + o(1)) \tag{20.35}$$

$$c_{\rm d} = q(\log q + \log\log q + O(1)) \tag{20.36}$$

$$c_{\rm c} = 2q\log q - \log q - 2\log 2 + o(1) \tag{20.37}$$

$$c_{\rm s} = 2q\log q - \log q - 1 + o(1) \tag{20.38}$$

These predictions can be rephrased into a statement on the **chromatic number**, i.e. the minimal number of colors needed to color a graph. Because of the heuristic nature of the approach, we formulate it as a conjecture:

Conjecture 20.1 *With high probability, the chromatic number of a random regular graph with N vertices and degree $c \geq 4$ is equal to $\chi_{\rm chrom}(c)$, where*

$$\chi_{\rm chrom}(c) = \max\{q : \;\; \Sigma^{\rm e}(\epsilon = 0; c, q) > 0\} \; . \tag{20.39}$$

Here $\Sigma^{\rm e}(\epsilon = 0; c, q)$ is given by Eq. (20.34) with a the largest solution of (20.33) in the interval $[0, 1/q]$.

Using the numbers in table 20.10, this conjecture predicts for instance that $\chi_{\rm chrom}(c) = 3$ for $c = 4, 5$, $\chi_{\rm chrom}(c) = 4$ for $c = 6, 7, 8, 9$, and $\chi_{\rm chrom}(c) = 5$ for $10 \leq c \leq 14$.

On the side of rigorous results, a clever use of the first and second moment methods allows to prove the following result:

Theorem 20.2 *With high probability, the chromatic number of a random regular graph with N vertices and degree c is either k or $k+1$ or $k+2$, where k is the smallest integer such that $c < 2k\log k$. Furthermore, if $c > (2k-1)\log k$, then with high probability the chromatic number is either k or $k+1$.*

One can check explicitely that the results of the 1RSB cavity conjecture agree with this theorem, that proves the correct leading behavior at large c.

While this presentation was focused on random regular graphs, a large class of random graph ensembles can be analyzed along the same lines.

Notes

Random K-satisfiability was first analyzed using the replica symmetric cavity method in (Monasson and Zecchina, 1996; Monasson and Zecchina, 1996). The resulting equations are equivalent to a density evolution analysis of belief propagation. BP was used as an algorithm for finding SAT assignments in (Pumphrey, 2001). This study concluded that BP is ineffective in solving satisfiability problems, mainly because it assigned variables in a one-shot fashion, unlike in decimation.

The 1RSB cavity method was applied to random satisfiability in (Mézard, Parisi and Zecchina, 2003; Mézard and Zecchina, 2002), where the value of $\alpha_{\rm c}$ was computed for 3-SAT. This approach was applied to larger K in (Mertens, Mézard and Zecchina, 2006), which also derived the large K asymptotics. The SPY and SP equations for satisfiability were first written in (Mézard and Zecchina, 2002), where SP-inspired decimation was introduced (Fig. 20.8 is borrowed from this paper). A more algorithmic presentation of SP was then developed in (Braunstein, Mézard and Zecchina, 2005), together with an optimized source code for SP and decimation (Braunstein, Mézard and Zecchina, 2004). The idea of backtracking was suggested in (Parisi, 2003), but its performances have not been systematically studied yet.

The condensation phenomenon was discussed in (Krzakala, Montanari, Ricci-Tersenghi, Semerjian and Zdeborova, 2007), in relation with studies of the entropic complexity in colouring (Mézard, Palassini and Rivoire, 2005*b*; Krzakala and Zdeborova, 2007) and in satisfiability (Montanari, Ricci-Tersenghi and Semerjian, 2008).

The analysis in this chapter is heuristic, and is waiting for a rigorous proof. Let us point out that one important aspect of the whole scenario has been established rigorously for $K \geq 8$: it has been shown that in some range of clause density below $\alpha_{\rm s}(K)$, the SAT assignments are grouped into exponentially many clusters, well separated from each other (Mézard, Mora and Zecchina, 2005*a*; Achlioptas and Ricci-Tersenghi, 2006; Daudé, Mézard, Mora and Zecchina, 2008). This result can be obtained by a study of 'x-satisfiability' problem, that requires to determine whether a formula has *two* SAT assignments differing in xN variables. Bounds on the x-satisfiability threshold can be obtained through the first and second moment methods.

The coloring problem has been first studied with the energetic 1RSB cavity method by (Mulet, Pagnani, Weigt and Zecchina, 2002; Braunstein, Mulet, Pagnani, Weigt and Zecchina, 2003): these papers contain the derivation of the SAT/UNSAT threshold and the SP equations. A detailed study of the entropy of clusters, and the computation of the other thresholds, has carried out in (Krzakala and Zdeborova, 2007). These papers also study the case of Erdös Rényi graphs. Theorem 20.2 was proven in (Achlioptas and Moore, 2004), and its analogue for Erdös Rényi graphs in (Achlioptas and Naor, 2005).

21

GLASSY STATES IN CODING THEORY

In Ch. 15 we studied the problem of decoding random LDPC codes, and found two phase transitions, that characterize the code performances in the large blocklength limit. Consider, for instance, communication over a binary symmetric channel with crossover probability p. Under belief propagation decoding, the bit error rate vanishes in the large blocklength limit below a first threshold $p_{\rm d}$ and remains strictly positive for $p > p_{\rm d}$. On the other hand, the minimal bit error rate achievable with the same ensemble (i.e. the bit error rate under symbol MAP decoding) vanishes up to a larger noise level $p_{\rm c}$ and is bounded away from 0 for $p > p_{\rm c}$.

In principle, one should expect each decoding algorithm to have a different threshold. This suggests not to attach too much importance to the BP threshold $p_{\rm d}$. On the contrary, we will see in this chapter that $p_{\rm d}$ is, in some sense, a 'universal' characteristics of the code ensemble: above $p_{\rm d}$, the decoding problem is plagued by an exponential number of metastable states (Bethe measures). In other words the phase transition which takes place at $p_{\rm d}$ is not only algorithmic, it is a *structural* phase transition. This transition turns out to be a dynamical 1RSB glass transition and this suggests that $p_{\rm d}$ is the largest possible threshold for a large class of local decoding algorithms.

We have already seen in the last section of Ch. 15 that the two thresholds $p_{\rm d}$ and $p_{\rm c}$ are closely related and can both be computed formally within the RS cavity method, i.e. in terms of the density evolution fixed point. The analysis below will provide a detailed explanation of this connection in terms of the glass transition studied in Ch.19.

In the next section we start by a numerical investigation of the role of metastable states in decoding. Sec. 21.2 considers the particularly instructive case of the binary erasure channel, where the glassy states can be analyzed relatively easily using the energetic 1RSB cavity method. The analysis of general memoryless channels is described in Sec. 21.3. Finally, Sec. 21.4 draws the connection between metastable states, which are a main object of study in this chapter, and trapping sets (subgraphs of the original factor graph that are often regarded as responsible for coding failures).

21.1 Local search algorithms and metastable states

The codewords of an LDPC code are solutions of a constraint satisfaction problem. The variables are the bits of a word $\underline{x} = (x_1, x_2, \dots, x_N)$, with $x_i \in \{0, 1\}$, and the constraints are the parity check equations, i.e. a set of linear equations mod 2. This is analogous to the XORSAT problem considered in Ch. 18, although the ensembles of linear systems used in coding are different.

An important difference with XORSAT is that we are looking for a *specific* solution of the linear system, namely the transmitted codeword. The received message $\underline{y}$ gives us a hint of where to look for this solution. For notational simplicity, we shall assume that the output alphabet $\mathcal{Y}$ is discrete, and the channel is a binary input memoryless output symmetric (BMS- see Ch. 15) channel with transition probability[32] $\mathcal{Q}(y|x)$. The probability that $\underline{x}$ is the transmitted codeword, given the received message $\underline{y}$, is given by the usual formula (15.1) $\mathbb{P}(\underline{x}|\underline{y}) = \mu_y(\underline{x})$ where:

$$\mu_y(\underline{x}) \cong \prod_{i=1}^{N} \mathcal{Q}(y_i|x_i) \prod_{a=1}^{M} \mathbb{I}(x_{i_1^a} \oplus \cdots \oplus x_{i_{k(a)}^a} = 0)\,. \tag{21.1}$$

It is natural to associate an optimization problem to the code. Define the energy $E(\underline{x})$ of a word $\underline{x}$ (also called a 'configuration') as *twice* the number of parity check equations violated by $\underline{x}$ (the factor 2 is introduced for future simplifications). Codewords coincide with the global minima of this energy function, with zero energy.

We already know that decoding consist in computing marginals of the distribution $\mu_y(\underline{x})$ (symbol MAP decoding), or finding its argmax (word MAP decoding). In the following we shall discuss two closely related problems: (i) optimizing the energy function $E(\underline{x})$ within a subset of the configuration space defined by the received word and the channel properties; (ii) sampling from a 'tilted' Boltzmann distribution associated to $E(\underline{x})$.

21.1.1 *Decoding through constrained optimization*

Let us start by considering the word-MAP decoding problem. We shall exploit our knowledge of the BMS channel. Conditional on the received word $\underline{y} = (y_1, y_2, \ldots, y_N)$, the log-likelihood for $\underline{x}$ to be the channel input is:

$$L_{\underline{y}}(\underline{x}) = \sum_{i=1}^{N} \log \mathcal{Q}(y_i|x_i)\,. \tag{21.2}$$

We shall later use the knowledge that the input word was a codeword, but $L_{\underline{y}}(\underline{x})$ is well defined for any $\underline{x} \in \{\mathtt{0},\mathtt{1}\}^N$, regardless of whether it is a codeword or not, so let us first characterize its properties.

Assume without loss of generality that the codeword $\underline{\mathtt{0}}$ had been transmitted. By the law of large numbers, for large N the log-likelihood of this codeword is close to $-Nh$, where h is the channel entropy: $h = -\sum_y \mathcal{Q}(y|\mathtt{0}) \log \mathcal{Q}(y|\mathtt{0})$. The probability of an order-N deviation away from this value is exponentially small

[32] Throughout this chapter we adopt a different notation for the channel transition probability than in the rest of the book, in order to avoid confusion with 1RSB messages.

in N. This suggests to look for the transmitted codeword among those $\underline{x}$ such that $L_{\underline{y}}(\underline{x})$ is close to h.

The corresponding '**typical pairs**' decoding strategy goes as follows: Given the channel output $\underline{y}$, look for a codeword $\underline{x} \in \mathfrak{C}$, such that $L_{\underline{y}}(\underline{x}) \geq -N(h+\delta)$. We shall refer to this condition as the 'distance constraint'. For instance, in the case of the BSC channel, it amounts to constraining the Hamming distance between the codeword $\underline{x}$ and the received codeword $\underline{y}$ to be small enough. If exactly one codeword satisfies the distance constraint, return it. If there is no such codeword, or if there are several of them, declare an error. Here $\delta > 0$ is a parameter of the algorithm, which should be thought of as going to 0 *after* $N \to \infty$.

Exercise 21.1 Show that the block error probability of typical pairs decoding is independent of the transmitted codeword.
[Hint: use the linear structure of LDPC codes, and the symmetry property of the BMS channel.]

Exercise 21.2 This exercise aims at convincing the reader that typical pairs decoding is 'essentially' equivalent to maximum likelihood (ML) decoding.

(a) Show that the probability that no codeword exists with $L_{\underline{y}}(\underline{x}) \in [-N(h+\delta), -N(h-\delta)]$ is exponentially small in N.
[Hint: apply Sanov Theorem, cf. Sec. 4.2, to the type of the received codeword.]

(b) Upper bound the probability that ML succeeds and typical pairs decoding fails in terms of the probability that there exists an incorrect codeword $\underline{x}$ with $L_{\underline{y}}(\underline{x}) \geq -N(h+\delta)$, but no incorrect codeword $L_{\underline{y}}(\underline{x}) \geq -N(h-\delta)$.

(c) Estimate the last probability for Shannon's random code ensemble. Show in particular that it is exponentially small for all noise levels strictly smaller than the MAP threshold and δ small enough.

Since codewords are global minima of the energy function $E(\underline{x})$ we can rephrase typical pairs decoding as an optimization problem:

$$\text{Minimize} \quad E(\underline{x}) \qquad \text{subject to} \quad L_{\underline{y}}(\underline{x}) \geq -N(h+\delta)\,. \tag{21.3}$$

Neglecting exponentially rare events, we know that there always exists at least one solution with cost $E(\underline{x}) = 0$, corresponding to the transmitted codeword. Therefore, typical pairs decoding is successful if and only if the minimum is non-degenerate. This happens with high probability for $p < p_{\rm c}$. On the contrary, for $p > p_{\rm c}$, the optimization admits other minima with zero cost (incorrect codewords). We already explored this phenomenon in chapters 11 and 15, and we shall discuss it further below. For $p > p_{\rm c}$ there exists an exponential number of codewords whose likelihood is larger or equal to the likelihood of the transmitted one.

Similarly to what we have seen in other optimization problems (such as MAX-XORSAT or MAX-SAT), generically there exists an intermediate regime $p_{\rm d} < p < p_{\rm c}$, which is characterized by an exponentially large number of metastable states. For these values of p, the global minimum of $E(\underline{x})$ is still the transmitted codeword, but is 'hidden' by the proliferation of deep local minima. Remarkably, the threshold for the appearence of an exponential number of metastable states coincides with the BP threshold $p_{\rm d}$. Thus, for $p \in]p_{\rm d}, p_{\rm c}[$ MAP decoding would be successful, but message passing decoding fails. In fact no practical algorithm which succeeds in this regime is known. A cartoon of this geometrical picture is presented in Fig. 21.1.

At this point, the reader might be puzzled by the observation that finding configurations with $E(\underline{x}) = 0$ is *per se* a polynomial task. Indeed it amounts to solving a linear system modulo 2, and can be done by Gauss elimination. However, the problem (21.3) involves the condition $L_{\underline{y}}(\underline{x}) \geq -N(h+\delta)$ which is *not* a linear constraint modulo 2. If one resorts to local-search based decoding algorithms, the proliferation of metastable states for $p > p_{\rm d}$ can block the algorithms. We shall discuss this phenomenon on two local search strategies: Δ-local search and simulated annealing.

21.1.2 Δ *local-search decoding*

A simple local search algorithm consists in starting from a word $\underline{x}(0)$ such that $L_{\underline{y}}(\underline{x}(0)) \geq -N(h+\delta)$ and then recursively constructing $\underline{x}(t+1)$ by optimizing the energy function within a radius Δ neighborhood around $\underline{x}(t)$:

```
Δ LOCAL SEARCH (channel output $\underline{y}$, search size $\Delta$, likelihood resolution $\delta$)
1: Find $\underline{x}(0)$ such that $L_{\underline{y}}(\underline{x}(0)) \geq -N(h+\delta)$ ;
2: for $t = 0, \dots t_{\max} - 1$:
3:     Choose a uniformly random connected set $U \subset \{1, \dots, N\}$
       of variable nodes in the factor graph with $|U| = \Delta$;
4:     Find the configuration $\underline{x}'$ that minimizes the energy subject
       to $x'_j = x_j$ for all $j \notin U$;
5:     If $L_{\underline{y}}(\underline{x}') \geq -N(h+\delta)$, set $\underline{x}(t+1) = \underline{x}'$;
       otherwise, set $\underline{x}(t+1) = \underline{x}(t)$;
6: end;
7: return $\underline{x}(t_{\max})$.
```

(Recall that a set of variable nodes U is 'connected' if, for any $i, j \in U$, there exists a path in the factor graph connecting i to j, such that all variable nodes along the path are in U as well.)

Exercise 21.3 A possible implementation of step 1 consists in setting $x_i(0) = \arg\max_x \mathcal{Q}(y_i|x)$. Show that this choice meets the likelihood constraint.

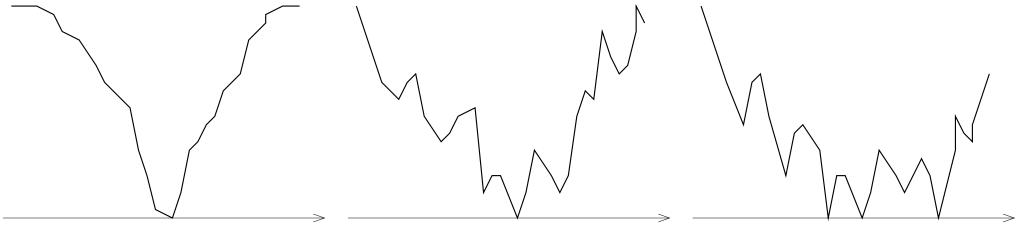

FIG. 21.1. Three possible cartoon landscapes for the energy function $E(\underline{x})$ (the number of violated checks), plotted in the space of all configurations $\underline{x}$ with $L_{\underline{y}}(\underline{x}) \geq N(h-\delta)$. On the left: the energy as a unique global minimum with $E(\underline{x}) = 0$ (the transmitted codeword) and no (deep) local minima. Center: many deep local minima appear although the global minimum remains non-degenerate. Right: More than one codeword is compatible with the likelihood constraint, and the global minimum $E(\underline{x}) = 0$ becomes degenerate.

If the factor graph has bounded degree (which is the case with LDPC ensembles), and Δ is bounded as well, each execution of the cycle above implies a bounded number of operations. As a consequence if we let $t_{\max} = O(N)$, the algorithm has linear complexity. A computationally heavier variant consists in choosing U at step 3 greedily. This means going over all such subsets and then taking the one that maximizes the decrease in energy $|E(\underline{x}(t+1)) - E(\underline{x}(t))|$.

Obviously the energy $E(\underline{x}(t))$ of the configuration produced after t iterations is a non-increasing function of t. If it vanishes at some time $t \leq t_{\max}$, then the algorithm implements a typical pairs decoder. Ideally, one would like a characterization of the noise levels and code ensembles such that $E(\underline{x}(t_{\max})) = 0$ with high probability.

The case $\Delta = 1$ was analyzed in Ch. 11, under the name of 'bit-flipping' algorithm, for communicating over the channel BSC(p). We saw that there exists a threshold noise level p_1 such that, if $p < p_1$ the algorithm returns with high probability the transmitted codeword. It is reasonable to think that the algorithm will be unsuccessful with high probability for $p > p_1$.

Analogously, one can define thresholds p_Δ for each value of Δ. Determining these thresholds analytically is an extremely challenging problem.

One line of approach could consist in first studying **Δ-stable configurations**. We say that a configuration $\underline{x}$ is Δ-stable if, for any configuration $\underline{x}'$ such that $L_{\underline{y}}(\underline{x}') \geq -N(h+\delta)$ and $d(\underline{x}, \underline{x}') \leq \Delta$, $E(\underline{x}') \geq E(\underline{x})$.

Exercise 21.4 Show that, if no Δ-stable configurations exists, then the greedy version of the algorithm will find a codeword after at most M steps (M being the number or parity checks).

While this exercise hints at a connection between the energy landscape and the difficulty of decoding, one should be aware that the problem of determining p_Δ cannot be reduced to determining whether Δ-stable states exist or to estimate their number. The algorithm indeed fails if, after a number t of iterations, the distribution of $\underline{x}(t)$ is (mostly) supported in the basin of attraction of Δ-stable states. The key difficulty is of course to characterize the distribution of $\underline{x}(t)$.

21.1.3 *Decoding through simulated annealing*

A more detailed understanding of the role of metastable configurations in the decoding problem can be obtained through the analysis of the MCMC decoding procedure that we discussed in Sec. 13.2.1. We thus soften the parity check constraints through the introduction of an inverse temperature $\beta = 1/T$ (this should not be confused with the temperature introduced in Ch. 6, which instead multiplied the codewords log-likelihood). Given the received word $\underline{y}$, we define the following distribution over the transmitted message $\underline{x}$, cf. Eq. (13.10):

$$\mu_{y,\beta}(\underline{x}) \equiv \frac{1}{Z(\beta)} \exp\{-\beta E(\underline{x})\} \prod_{i=1}^{N} \mathcal{Q}(y_i|x_i) \,. \tag{21.4}$$

This is the 'tilted Boltzmann form' that we alluded to before. In the low-temperature limit it reduces to the familiar a posteriori distribution which we would like to sample: $\mu_{y,\beta=\infty}(\underline{x})$ is supported on the codewords, and gives to each of them a weight proportional to its likelihood. At infinite temperature, $\beta = 0$, the distribution factorizes over the bits x_i. More precisely, under $\mu_{y,\beta=0}(\underline{x})$, the bits x_i are independent random variables with marginal $\mathcal{Q}(y_i|x_i)/(\mathcal{Q}(y_i|\mathtt{0}) + \mathcal{Q}(y_i|\mathtt{1}))$. Sampling from this measure is very easy.

For $\beta \in]0, \infty[$, $\mu_{y,\beta}(\,\cdot\,)$ can be regarded as a distribution of possible channel inputs for a code with 'soft' parity check constraints. Notice that, unlike the $\beta = \infty$ case, it depends in general on the actual parity check matrix and not just on the codebook $\mathfrak{C}$. This is actually a good feature of the tilted measure: performances of practical algorithms do indeed depend upon the parity check matrix representation of $\mathfrak{C}$. It is therefore necessary to take it into account.

We shall sample from $\mu_{y,\beta}(\,\cdot\,)$ using Glauber dynamics, cf. Sec. 13.2.1. We have already seen in that section that decoding through sampling at a fixed β fails above a certain noise level. Let us now try to improve on it using a simulated annealing procedure in which β is increased gradually according to an annealing schedule $\beta(t)$, with $\beta(0) = 0$. This decoder uses as input the received word $\underline{y}$, the annealing schedule, and some maximal numbers of iterations $t_{\max}$, n:

```
SIMULATED ANNEALING DECODER ( $\underline{y}$, $\{\beta(t)\}$, $t_{\max}$, $n$ )
Generate $\underline{x}_*(0)$ form $\mu_{y,0}(\,\cdot\,)$;
for $t = 0, \dots t_{\max} - 1$:
    Set $\underline{x}(0;t) = \underline{x}_*(t-1)$;
    Let $\underline{x}(j;t)$, $j \in \{1, \dots, n\}$ be the configurations produced by
        $n$ successive Glauber updates at $\beta = \beta(t)$;
    Set $\underline{x}_*(t) = \underline{x}(n;t)$;
end
return $\underline{x}(t_{\max})$.
```

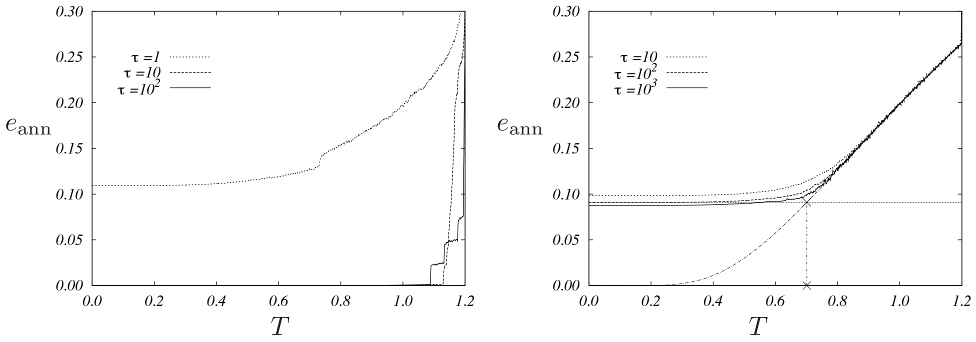

FIG. 21.2. Decoding random codes from the $(5, 6)$ LDPC ensemble through simulated annealing. Here we consider blocklength $N = 12000$ and transmission over the BSC(p) with $p = 0.12$ (left) and 0.25 (right). The system is annealed through $t_{\rm max} = 1200$ temperature values equally spaced between $T = 1.2$ and $T = 0$. At each temperature $n = N\tau$ updates are executed. Statistical errors are comparable with the size of jumps along the curves.

Its algorithmic complexity is proportional to the total number of Glauber updates $nt_{\rm max}$. If we want the algorithm to be efficient, this should grow linearly or slightly super-linearly with N. The intuition is that the first (small β) steps allow the Markov chain to equilibrate across the configuration space while, as β gets larger, the sample concentrates onto (or near to) codewords. Hopefully at each stage $\underline{x}_*(t)$ will be approximately distributed according to $\mu_{y,\beta(t)}(\,\cdot\,)$.

Figure 21.2 shows the result obtained by the simulated annealing decoder, using random LDPC codes from the $(5, 6)$ regular ensemble, used over the binary symmetric channel at crossover probabilities $p = 0.12$ and 0.25 (for this ensemble, $p_{\rm d} \approx 0.139$ and $p_{\rm c} \approx 0.264$). The annealing schedule is linear in the temperature, namely $\beta(t) = 1/T(t)$ with

$$T(t) = T(0) - \big\{T(0) - T(t_{\rm max})\big\} \left(\frac{t}{t_{\rm max}}\right) , \qquad (21.5)$$

with $T(0) = 1.2$ and $T(t_{\rm max}) = 0$. The performance of decoding can be evaluated through the number of violated checks in the final configuration, which is half $E(\underline{x}(t_{\rm max}))$. The figure shows the energy density averaged over 10 repetitions of the decoding experiment (each time with a new code randomly chosen from the ensemble), $e(t) = \frac{1}{N}\langle E(\underline{x}(t))\rangle$, versus the temperature $T(t)$. As the number of updates performed at each temperature increases, the number of violated checks per variable seems to converge to a well defined limiting value, that depends on t only through the corresponding temperature

$$\frac{1}{N}\langle E(\underline{x}(t))\rangle \to e_{\rm ann}(\beta(t))\, . \qquad (21.6)$$

Further, $E(\underline{x}(t))/N$ seems to concentrate around its mean as $N \to \infty$.

At small p, the curve $e_{\rm ann}(\beta)$ quickly converges to 0 as $\beta \to \infty$: a codeword (the transmitted one) is found efficiently. In fact, already at $\beta = 1$, the numerical result for $e_{\rm ann}(\beta)$ is indistinguishable from 0. We expect that $e_{\rm ann}(\beta)$ coincides within numerical accuracy with the theoretical prediction for the equilibrium average

$$e_{\rm eq}(\beta) \equiv \frac{1}{N} \lim_{N\to\infty} \langle E(\underline{x}) \rangle_\beta \, . \tag{21.7}$$

This agrees with the above observations since $e_{\rm eq}(\beta) = O(e^{-10\beta})$ (the lowest excitation over the ground state amounts to flipping a single bit, its energy is equal to 10). The numerics thus suggest that $\underline{x}(t_{\rm max})$ is indeed approximately distributed according to $\mu_{y,\beta(t)}(\,\cdot\,)$.

At large p, $e_{\rm ann}(\beta)$ has instead a non-vanishing $\beta \to \infty$ limit: the annealing algorithm does not find any codeword. The returned word $\underline{x}_*(t_{\rm max})$ typically violates $\Theta(N)$ parity checks. On the other hand, in the equilibrated system at $\beta = \infty$, the energy vanishes by construction (we know that the transmitted codeword satisfies all checks). Therefore the simulation has fallen out of equilibrium at some finite β, thus yielding a distribution of $\underline{x}(t_{\rm max})$ which is very different from $\mu_{y,\beta=\infty}(\,\cdot\,)$. The data in Fig. 21.2 shows that the energy varies very slowly at low temperatures, which confirms the fact that the system is out of equilibrium.

We shall argue below that this slowing down is in fact due to a dynamical glass phase transition occuring at a well defined temperature $T_{\rm d} = 1/\beta_{\rm d}$. Below this temperature, $\underline{x}(t_{\rm max})$ gets trapped with high probability into a pure state corresponding to a deep local minimum of $E(\underline{x})$ with positive energy, and never reaches a global minimum of the energy (i.e. a codeword).

This is related to the 'energy landscape' picture discussed in the previous section. Indeed, the success of the simulated annealing decoder for $p \le p_{\rm d}$ can be understood as follows. At small noise the 'tilting' factor $\prod_i \mathcal{Q}(y_i|x_i)$ effectively selects a portion of the configuration space around the transmitted codeword (more or less like the likelihood constraint above) and this portion is small enough that there is no metastable state inside it. An interesting aspect of simulated annealing decoding is that it can be analyzed on the basis of a purely static calculation. Indeed for any $\beta \le \beta_{\rm d}$, the system is still in equilibrium and its distribution is simply given by Eq. (21.4). Its study, and the determination of $\beta_{\rm d}$, will be the object of the next sections.

Before moving to this analysis, let us make a last remark about simulated annealing: for any finite β, the MCMC algorithm is able to equilibrate if it is iterated a large number of times (a direct consequence of the fact that Glauber dynamics is irreducible and aperiodic). This raises a paradox, as it seems to imply that the annealing energy always coincide with the equilibrium one, and the system never falls out of equilibrium during the annealing process. The conundrum is that, in the previous discussion we tacitly assumed that the number of Monte Carlo steps cannot grow exponentially with the system size. To be more precise, one can for instance define the annealing energy as

$$e_{\rm ann}(\beta) \equiv \lim_{t_{\max}\to\infty} \lim_{N\to\infty} \frac{1}{N} \left\langle E_N(\underline{x}(t_\beta = \lfloor (1-\beta(0)/\beta) t_{\max} \rfloor)) \right\rangle , \qquad (21.8)$$

where we assumed $\beta(t_{\max}) = \infty$ The important point is that the limit $N \to \infty$ is taken before $t_{\max} \to \infty$: in such a case simulated annealing can be trapped in metastable states.

21.2 The binary erasure channel

If communication takes place over the binary erasure channel BEC(ϵ), the analysis of metastable states can be carried out in details by adopting the point of view of constrained optimization introduced in Sec. 21.1.1.

Suppose that the all zero codeword $\underline{x}_* = (\mathtt{0}, \cdots, \mathtt{0})$ has been sent, and let Let $\underline{y} \in \{\mathtt{0}, *\}^N$ be the channel output. We shall denote by $U = U(\underline{y})$ the set of erased bits. The log-likelihood for the word $\underline{x}$ to be the input can take two possible values: $L_{\underline{y}}(\underline{x}) = |U| \log \epsilon$ if $x_i = \mathtt{0}$ for all $i \notin U$, and $L_{\underline{y}}(\underline{x}) = -\infty$ otherwise. Of course the input codeword belongs to the first set: $L_{\underline{y}}(\underline{x}_*) = |U| \log \epsilon$. The strategy of Sec. 21.1.1 reduces therefore to minimizing $E(\underline{x})$ (i.e. minimizing the number of violated parity checks) among all configurations $\underline{x}$ such that $x_i = \mathtt{0}$ on all the non-erased positions.

When the noise ϵ is smaller than the MAP threshold, there is a unique minimum with energy 0, namely the transmitted codeword $\underline{x}_*$. Our aim is to study the possible existence of metastable states, using the energetic cavity method of Sec. 19.5. This problem is closely related to XORSAT, whose analysis was presented analysis in Ch. 18 and Ch. 19: Once all the non-erased bits have been fixed to $x_i = \mathtt{0}$, decoding amounts to solving a homogeneous system of linear equations among the remaining bits. If one uses a code from the LDPC$_N(\Lambda, P)$ ensemble, the degree profiles of the remaining nodes are $\Lambda(x), R(x)$, where the probability of a check node to have degree k, R_k, is given in terms of the original P_k by:

$$R_k = \sum_{k'=k}^{k_{\max}} P_{k'} \binom{k'}{k} \epsilon^k (1-\epsilon)^{k'-k} , \qquad (21.9)$$

and the corresponding edge perspective degree profile is given as usual by $r_k = kR_k / \sum_p pR_p$.

Exercise 21.5 Show that $r(u) = \sum_k r_k u^{k-1} = \rho(1 - \epsilon(1-u))$.

Assuming as usual that the number of metastable states - solutions of min-sum equations- of energy Ne grows like $\exp(N\Sigma^{\rm e}(e))$, we will use the 1RSB energetic cavity method to compute the energetic complexity $\Sigma^{\rm e}(e)$. This can be done using the $\mathsf{SP(y)}$ equations on the original factor graph. As our problem involves only hard constraints and binary variables, we can use the simplified formalism of Sec.19.5.3. Each min-sum message can take three possible values, $\mathtt{0}$ (the meaning of which is "take value 0"), $\mathtt{1}$ ("take value 1") and $*$ ("you can take any value"). The $\mathsf{SP}(\mathtt{y})$ messages are distributions on these three values or, equivalently, normalized triplets.

21.2.1 *The energetic 1RSB equations*

Let us now turn to the statistical analysis of these messages. We denote by $Q = (Q_{\mathtt{0}}, Q_{\mathtt{1}}, Q_*)$ the messages from variable to check, and $\widehat{Q}$ the messages from check to variables. We first notice that, if a bit is not erased, then it sends a sure $\mathtt{0}$ message $Q = (1, 0, 0)$ to all its neighboring checks. This means that the distribution of Q has a mass at least $1-\epsilon$ on sure $\mathtt{0}$ messages. We can write:

$$Q = \begin{cases} (1,0,0) & \text{with probability } (1-\epsilon)\ , \\ \widetilde{Q} & \text{with probability } \epsilon\ . \end{cases} \tag{21.10}$$

The distributional equations of $\widetilde{Q}$ and $\widehat{Q}$ can then be obtained exactly as in Secs. 19.5 and 19.6.3.

Exercise 21.6 Show that the distributions of $\widetilde{Q}$ and $\widehat{Q}$ satisfy the equations:

$$\widetilde{Q}_\sigma \overset{\mathrm{d}}{=} \mathsf{F}_{l,\sigma}(\widehat{Q}^1, \cdots, \widehat{Q}^{l-1}) \tag{21.11}$$

$$\begin{pmatrix} \widehat{Q}_{\mathtt{0}} \\ \widehat{Q}_{\mathtt{1}} \\ \widehat{Q}_* \end{pmatrix} \overset{\mathrm{d}}{=} \begin{pmatrix} \frac{1}{2}\prod_{i=1}^{k-1}(\widetilde{Q}_{\mathtt{0}}^i + \widetilde{Q}_{\mathtt{1}}^i) + \frac{1}{2}\prod_{i=1}^{k-1}(\widetilde{Q}_{\mathtt{0}}^i - \widetilde{Q}_{\mathtt{1}}^i) \\ \frac{1}{2}\prod_{i=1}^{k-1}(\widetilde{Q}_{\mathtt{0}}^i + \widetilde{Q}_{\mathtt{1}}^i) - \frac{1}{2}\prod_{i=1}^{k-1}(\widetilde{Q}_{\mathtt{0}}^i - \widetilde{Q}_{\mathtt{1}}^i) \\ 1 - \prod_{i=1}^{k-1}(1 - \widetilde{Q}_{*,i}) \end{pmatrix} \tag{21.12}$$

where we defined, for $\sigma \in \{\mathtt{0}, \mathtt{1}, *\}$

$$\mathsf{F}_{l,\sigma}(\widehat{Q}^1, \ldots, \widehat{Q}^{l-1}) \equiv \frac{\mathsf{Z}_{l,\sigma}(\{\widehat{Q}^a\})}{\mathsf{Z}_{l,\mathtt{0}}(\{\widehat{Q}^a\}) + \mathsf{Z}_{l,\mathtt{1}}(\{\widehat{Q}^a\}) + \mathsf{Z}_{l,*}(\{\widehat{Q}^a\})} \tag{21.13}$$

$$\mathsf{Z}_{l,\sigma}(\{\widehat{Q}^a\}) \equiv \sum_{\Omega_{\mathtt{0}},\Omega_{\mathtt{1}},\Omega_*}^{(\sigma)} e^{-y\min(|\Omega_{\mathtt{0}}|,|\Omega_{\mathtt{1}}|)} \prod_{a\in\Omega_{\mathtt{0}}} \widehat{Q}_{\mathtt{0}}^a \prod_{a\in\Omega_{\mathtt{1}}} \widehat{Q}_{\mathtt{1}}^a \prod_{a\in\Omega_*} \widehat{Q}_*^a\ . \tag{21.14}$$

Here we denoted by $\sum_{\Omega_{\mathtt{0}},\Omega_{\mathtt{1}},\Omega_*}^{(\sigma)}$ the sum over partitions of $\{1, \cdots, l-1\} = \Omega_{\mathtt{0}} \cup \Omega_{\mathtt{1}} \cup \Omega_*$ such that $|\Omega_{\mathtt{0}}| > |\Omega_{\mathtt{1}}|$ (for the case $\sigma = \mathtt{0}$), $|\Omega_{\mathtt{0}}| = |\Omega_{\mathtt{1}}|$ (for $\sigma = *$), or $|\Omega_{\mathtt{0}}| < |\Omega_{\mathtt{1}}|$ (for $\sigma = \mathtt{1}$). Furthermore, k, l, are random integers, with distributions respectively r_k and λ_l, the $\{\widetilde{Q}^i\}$ are $l-1$ i.i.d. copies of $\widetilde{Q}$, and $\{\widehat{Q}^a\}$ are $k-1$ i.i.d. copies of $\widehat{Q}$.

Given a solution of the 1RSB equations, one can compute the Bethe free-entropy density $\mathbb{F}^{\mathrm{RSB,e}}(Q, \widehat{Q})$ of the auxiliary problem. Within the 1RSB cavity method we estimate the free-entropy density of the auxiliary model using Bethe approximation as: $\mathfrak{F}^{\mathrm{e}}(\mathtt{y}) = \frac{1}{N}\mathbb{F}^{\mathrm{RSB,e}}(Q, \widehat{Q})$. This gives access to the energetic complexity function $\Sigma^{\mathrm{e}}(e)$ through the Legendre transform $\mathfrak{F}^{\mathrm{e}}(\mathtt{y}) = \Sigma^{\mathrm{e}}(e) - \mathtt{y}\, e$. Within the 1RSB cavity method we estimate the latter using Bethe approximation: $\mathfrak{F}^{\mathrm{e}}(\mathtt{y}) = \mathrm{f}^{\mathrm{RSB,e}}(\mathtt{y})$.

Exercise 21.7 Computation of the free-entropy. Using Eq. (19.92) show that the Bethe free-entropy of the auxiliary graphical model is $N\mathrm{f}^{\mathrm{RSB,e}} + o(N)$, where:

$$\mathrm{f}^{\mathrm{RSB,e}} = -\Lambda'(1)\epsilon\,\mathbb{E}\,\log z_{\mathrm{e}}(\widetilde{Q},\widehat{Q}) + \epsilon\,\mathbb{E}\,\log z_{\mathrm{v}}(\{\widehat{Q}^a\};l) + \\ + \frac{\Lambda'(1)}{P'(1)}\mathbb{E}\,\log z_{\mathrm{f}}(\{\widetilde{Q}^i\};k)\,. \tag{21.15}$$

Here expectations are taken over l (with distribution Λ_l), k (with distribution R_k defined in (21.9)), $\widetilde{Q},\widehat{Q}$ as well as their i.i.d. copies $\widetilde{Q}^i,\widehat{Q}^a$. The contributions of edges (z_{e}), variable (z_{v}) and function nodes (z_{f}) take the form:

$$z_{\mathrm{e}}(\widetilde{Q},\widehat{Q}) = 1 + (e^{-\mathsf{y}}-1)\Big(\widetilde{Q}_{\mathsf{0}}\widehat{Q}_{\mathsf{1}} + \widetilde{Q}_{\mathsf{1}}\widehat{Q}_{\mathsf{0}}\Big)\,, \tag{21.16}$$

$$z_{\mathrm{v}}(\{\widehat{Q}^i\};l) = \sum_{\Omega_0,\Omega_1,\Omega_*}\prod_{b\in\Omega_0}\widehat{Q}^b_{\mathsf{0}}\prod_{b\in\Omega_1}\widehat{Q}^b_{\mathsf{1}}\prod_{b\in\Omega_*}\widehat{Q}^b_{*}\;e^{-\mathsf{y}\,\min(|\Omega_0|,|\Omega_1|)}\,, \tag{21.17}$$

$$z_{\mathrm{f}}(\{\widetilde{Q}^i\};k) = 1 + \frac{1}{2}(e^{-\mathsf{y}}-1)\left\{\prod_{i=1}^{k}(\widetilde{Q}^i_{\mathsf{0}}+\widetilde{Q}^i_{\mathsf{1}}) - \prod_{i=1}^{k}(\widetilde{Q}^i_{\mathsf{0}}-\widetilde{Q}^i_{\mathsf{1}})\right\}\,, \tag{21.18}$$

where the sum in the second equation runs over the partitions $\Omega_0\cup\Omega_1\cup\Omega_* = [l]$.

21.2.2 *BP threshold and onset of metastability*

A complete study of the distributional equations (21.11), (21.12) is a rather challenging task. On the other hand they can be solved approximately through population dynamics. It turns out that the distribution obtained numerically shows different symmetry properties depending on the value of ϵ. Let us define a distribution $\widetilde{Q}$ (or $\widehat{Q}$) to be 'symmetric' if $\widetilde{Q}_{\mathsf{0}} = \widetilde{Q}_{\mathsf{1}}$, and 'positive' if $\widetilde{Q}_{\mathsf{0}} > \widetilde{Q}_{\mathsf{1}}$. We know from the BP decoding analysis that directed edges in the graph can be distinguished in two classes: those that eventually carry a message $\mathsf{0}$ under BP decoding, and those that instead carry a message $*$ even after a BP fixed point has been reached. It is natural to think that edges of the first class correspond to a positive 1RSB message $\widetilde{Q}$ (i.e., even among metastable states the corresponding bits are biased to be $\mathsf{0}$), while edges of the second class correspond instead to a symmetric message $\widetilde{Q}$.

This motivates the following hypothesis concerning the distributions of $\widetilde{Q}$ and $\widehat{Q}$. We assume that there exist weights $\xi,\hat{\xi}\in[0,1]$ and random distributions b, $\hat{\mathsf{b}}$, c, $\hat{\mathsf{c}}$, such that: b, $\hat{\mathsf{b}}$ are symmetric, c, $\hat{\mathsf{c}}$ are positive, and

$$\widetilde{Q} \overset{\mathrm{d}}{=} \begin{cases} \mathsf{b} & \text{with probability } \xi \\ \mathsf{c} & \text{with probability } 1-\xi, \end{cases} \tag{21.19}$$

$$\widehat{Q} \overset{\mathrm{d}}{=} \begin{cases} \hat{\mathsf{b}} & \text{with probability } \hat{\xi}, \\ \hat{\mathsf{c}} & \text{with probability } 1-\hat{\xi}. \end{cases} \tag{21.20}$$

In other words ξ (respectively $\hat{\xi}$) denotes the probability that Q (resp. $\widehat{Q}$) is symmetric.

Equation (21.11) shows that, in order for $\widetilde{Q}$ to be symmetric, all the input $\widehat{Q}^i$ must be symmetric. On the other hand, Eq. (21.12) implies that $\widehat{Q}$ is symmetric if at least one of the input $\widetilde{Q}^a$ must be symmetric. Using the result of Exercise 21.5, we thus find that our Ansatz is consistent only if the weights $\xi, \hat{\xi}$ satisfy the equations:

$$\xi = \lambda(\hat{\xi}) \qquad \hat{\xi} = 1 - \rho(1 - \epsilon\xi)\,, \tag{21.21}$$

If we define $z \equiv \epsilon\xi$, $\hat{z} \equiv \hat{\xi}$, these coincide with the density evolution fixed point conditions for BP, cf. Eqs. (15.34). This is not surprising in view of the physical discussion which lead us to introduce Ansatz (21.19), (21.20): ξ corresponds to the fraction of edges that remain erased at the BP fixed point. On the other hand, we will see that this observation implies that BP stops to converge to the correct fixed point at the same threshold noise $\epsilon_{\rm d}$ where metastable states start to appear.

For $\epsilon \le \epsilon_{\rm d}$, Eqs. (21.21) admit the unique solution $\xi = \hat{\xi} = 0$, corresponding to the fact that BP decoding recovers the full transmitted message. As a consequence we can take $Q(\,\cdot\,) \overset{\rm d}{=} \mathsf{c}(\,\cdot\,)$, $\widehat{Q}(\,\cdot\,) \overset{\rm d}{=} \hat{\mathsf{c}}(\,\cdot\,)$ to have almost surely positive mean. In fact it is not hard to check that a consistent solution of Eqs. (21.11), (21.12) is obtained by taking

$$\widehat{Q} = \widetilde{Q} = (1, 0, 0) \qquad \text{almost surely.} \tag{21.22}$$

Since the cavity fields do not fluctuate from state to state (their distribution is almost surely a point mass), the structure of this solution indicates that no metastable state is present for $\epsilon \le \epsilon_{\rm d}$. This is confirmed by the fact that the free entropy density of this solution $\mathfrak{F}^{\rm e}(\mathtt{y})$ vanishes for all $\mathtt{y}$.

Above a certain noise threshold, for $\epsilon > \epsilon_{\rm d}$, Eq. (21.21) still possesses the solution $\xi = \hat{\xi} = 0$, but a new solution with $\xi, \hat{\xi} > 0$ appears as well. We have discussed this new solution in the density evolution analysis of BP decoding: it is associated with the fact that the BP iterations have a fixed point in which a finite fraction of the bits remains undetermined. Numerical calculations show that that, for $\epsilon > \epsilon_{\rm d}$, the iteration of Eqs. (21.11), (21.12) converges to a non-trivial distribution. In particular $\widetilde{Q}$ (resp. $\widehat{Q}$) is found to be symmetric with probability $\xi > 0$ (resp $\hat{\xi} > 0$), where the values of $\xi, \hat{\xi}$ are the non-trivial solution of (21.21). The free-entropy of the auxiliary model $\mathfrak{F}^{\rm e}(\mathtt{y})$, can be computed using (21.15). Its Legendre transform is the energetic complexity curve $\Sigma^{\rm e}(e)$.

Figure 21.3 shows the typical outcome of such a calculation for LDPC ensembles, when $\epsilon_{\rm d} < \epsilon < \epsilon_{\rm c}$. In this whole regime, there exists a zero energy word, the transmitted (all $\mathtt{0}$) codeword. This is described by the solution $\xi = \hat{\xi} = 0$. On top of this, the non-trivial solution gives a complexity curve $\Sigma^{\rm e}(e)$ which is positive in an interval of energy densities $(e_{\rm c}, e_{\rm d})$. A positive complexity means

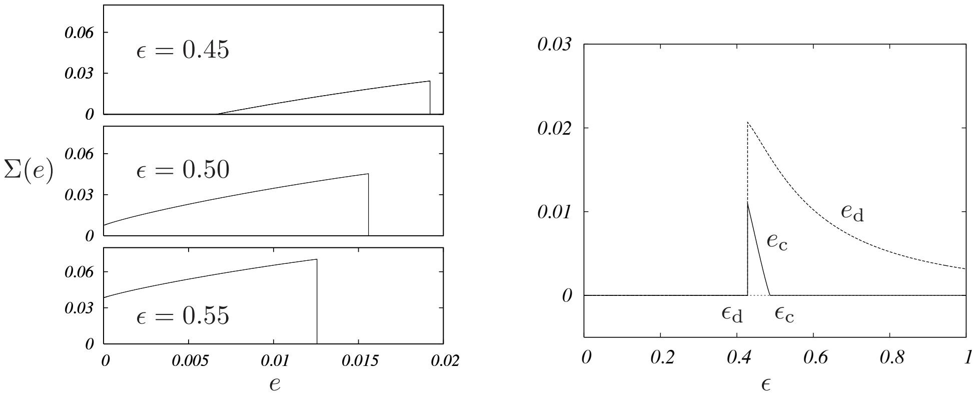

FIG. 21.3. Metastable states for random elements of the $(3, 6)$ regular ensemble used over the BEC(ϵ) (for this ensemble $\epsilon_{\rm d} \approx 0.4294$ and $\epsilon_{\rm c} \approx 0.4882$). Left frame: complexity as a function of the energy density for three values of the channel parameter above $\epsilon_{\rm d}$. Right frame: the maximum and minimum energy density $e_{\rm d}$ and $e_{\rm c}$ of metastable states as a function of the erasure probability.

that an exponential number of metastable states is present. But since $e_{\rm c} > 0$, these metastable states violate a finite fraction of the parity checks.

As ϵ increases both $e_{\rm d}$ and $e_{\rm c}$ decrease. At $\epsilon_{\rm c}$, $e_{\rm c}$ vanishes continuously and $e_{\rm c} = 0$, $e_{\rm d} > 0$ for all $\epsilon \geq \epsilon_{\rm c}$. In other words, at noise levels larger than $\epsilon_{\rm c}$ there appears an exponential number of zero energy 'metastable' states. These are codewords, that are indeed separated by energy barriers with height $\Theta(N)$. Consistently with this interpretation $\Sigma(e = 0) = {\rm f}^{\rm RS}_{h,u}$ where ${\rm f}^{\rm RS}_{h,u}$ is the RS free-entropy density (15.48) estimated on the non-trivial fixed point of density evolution.

The notion of metastable states thus allows to compute the BP and MAP thresholds within a unified framework. The BP threshold is the noise level where an exponential number of metastable states appears. This shows that this threshold is not only associated with a specific decoding algorithm, but it also has a structural, geometric meaning. On the other hand the MAP threshold coincides with the noise level where the energy of the lowest-lying metastable states vanishes.

Figure 21.4 shows the results of some numerical experiments with the simulated annealing algorithm of Sec. 21.1.3. Below the BP threshold, and for a slow enough annealing schedule the algorithm succeeds in finding a codeword (a zero energy state) in linear time. Above the threshold, even at the slowest annealing rate we could not find a codeword. Furthermore, the residual energy density at zero temperature is close to $e_{\rm d}$, suggesting that the optimization procedure is indeed trapped among the highest metastable states. This suggestion is further confirmed by Fig. 21.5 which compares the ϵ dependence of $e_{\rm d}$ with the residual energy under simulated annealing. Once again, there is rough agreement between

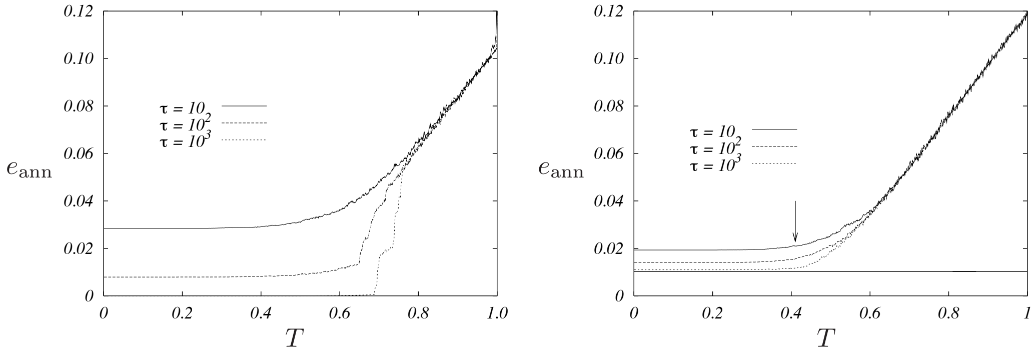

FIG. 21.4. Decoding random codes from the $(3, 6)$ regular ensemble used over the BEC(ϵ). In both cases $N = 10^4$, and the annealing schedule consists of $t_{\max}10^3$ equidistant temperatures in $T = 1/\beta \in [0, 1]$. At each value of the temperature $n = N\tau$ Monte Carlo updates are performed. On the left $\epsilon = 0.4 < \epsilon_{\rm d}$. On the right $\epsilon = 0.6 > \epsilon_{\rm d}$; the horizontal line corresponds to the energy density of the highest metastable states $e_{\rm d}(\epsilon = 0.6)$.

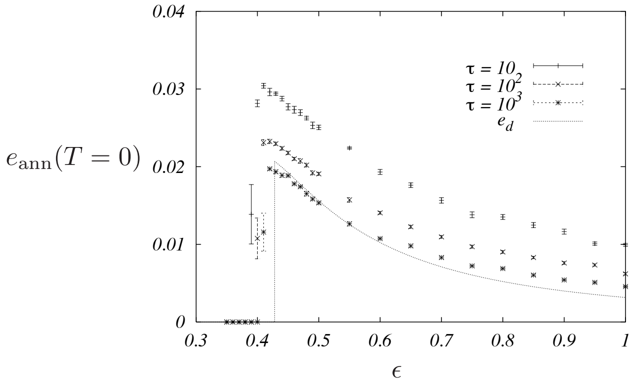

FIG. 21.5. Decoding random codes from the $(3, 6)$ regular ensemble used over the BEC(ϵ). Here we plot the minimum energy density achieved through simulated annealing versus the channel parameter. The continuous line is the energy of the highest lying metastable states. Size and annealing schedule as in Fig. 21.4.

the two (let us stress that one should not expect perfect agreement between the residual energy in Fig. 21.5 and $e_{\rm d}$: the former does indeed depend on the whole dynamical annealing process).

21.3 General binary memoryless symmetric channels

One would like to generalize to other channel models the above analysis of metastable states in the constrained optimization formulation of decoding. In general the computation is technically more intricate than for the BEC. The reason is that in general channels, the distance condition $L_{\underline{y}}(\underline{x}) \geq -N(h+\delta)$ cannot be written in terms of 'local' binary constraints. As a consequence, one cannot use the simplified approach of Sec. 19.5.3 and the general 1RSB formalism is required.

We shall follow this line of approach, but rather than pushing it to the point of determining the full complexity function, we will only determine whether the model (21.4) undergoes a dynamical phase transition as β increases from 0 to ∞, and locate the critical point $\beta_{\rm d}(p)$ (here p denotes the channel parameter). This is indeed the most important piece of information for our purposes. If a dynamical phase transition occurs at some $\beta_{\rm d} < \infty$, then for $\beta > \beta_{\rm d}$ the measure (21.4) decomposes into an exponential number of metastable pure states. As β crosses $\beta_{\rm d}$ the system is trapped in one of these and falls out of equilibrium. Upon further cooling (increase of β) the energy density of the annealed system remains higher than the equilibrium one and does not vanish as $\beta \to \infty$. This analysis allows to determine the noise threshold of the simulated annealing decoder, as the largest noise level p such that there is no finite $\beta_{\rm d}$.

In the following we first write the general 1RSB equations at finite β, and present some results obtained by solving them numerically. Finally we give a heuristic argument showing that $\beta_{\rm d}(p)$ goes to infinity exactly for $p \downarrow p_{\rm d}$.

21.3.1 *The 1RSB cavity approach*

We shall apply the 1RSB cavity approach of Ch. 19 to the decoding problem. Given a code and the received message $\underline{y}$, we want to study the probability distribution $\mu_{y,\beta}(\underline{x})$ defined in Eq. (21.4), and understand whether it decomposes in exponentially many extremal Bethe measures. The BP equations are simple generalizations of those written in Ch. 15 for the case $\beta = \infty$. In terms of the log-likelihoods

$$h_{i\to a} = \frac{1}{2}\log\frac{\nu_{i\to a}(0)}{\nu_{i\to a}(1)}\,, \qquad u_{a\to i} = \frac{1}{2}\log\frac{\widehat{\nu}_{a\to i}(0)}{\widehat{\nu}_{a\to i}(1)}$$

$$B_i = \frac{1}{2}\log\frac{\mathcal{Q}(y_i|0)}{\mathcal{Q}(y_i|1)} \equiv B(y_i)\,, \tag{21.23}$$

they read:

$$h_{i\to a} = B_i + \sum_{b\in\partial i\setminus a} u_{b\to i} \equiv \mathsf{f}_i(\{u_{b\to i}\})\,, \tag{21.24}$$

$$u_{a\to i} = \operatorname{atanh}\Big\{\tanh\beta \prod_{j\in\partial a\setminus i}\tanh h_{j\to a}\Big\} \equiv \hat{\mathsf{f}}_a(\{h_{j\to a}\})\,. \tag{21.25}$$

The corresponding Bethe free-entropy is given by (unlike in Ch. 15, here we use natural logarithms)

$$\mathbb{F}(\underline{u},\underline{h}) = -\sum_{(ia)\in E} \log\left[\sum_{x_i} \widehat{\nu}_{u_{a\to i}}(x_i)\nu_{h_{i\to a}}(x_i)\right] + \sum_{i=1}^{N} \log\left[\sum_{x_i} \mathcal{Q}(y_i|x_i) \prod_{a\in\partial i} \widehat{\nu}_{u_{a\to i}}(x_i)\right]$$
$$+\sum_{a=1}^{M} \log\left[\sum_{\underline{x}_{\partial a}} \exp(-\beta E_a(\underline{x}_{\partial a})) \prod_{i\in\partial a} \nu_{h_{i\to a}}(x_i)\right] . \qquad (21.26)$$

As in (15.44), we shall introduce a “shifted” free-entropy density ϕ defined as

$$\phi = \frac{1}{N}\,\mathbb{F}(\underline{u},\underline{h}) - \sum_{y} \mathcal{Q}(y|0)\log\mathcal{Q}(y|0)\,, \qquad (21.27)$$

Recall that the 1RSB cavity approach assumes that, to leading exponential order, the number $\mathcal{N}(\phi)$ of Bethe measures with a shifted free-entropy density equal to ϕ is equal to the number of quasi-solutions of Eqs. (21.24), (21.25). We shall write as usual $\mathcal{N}(\phi) \doteq \exp(N\Sigma(\phi))$, and our aim is to compute the complexity $\Sigma(\phi)$, using as in Ch. 19 an auxiliary graphical model which counts the number of solutions of BP equations, weighted by a factor $\exp(N\mathtt{x}\phi)$. If the free-entropy of the auxiliary model is $\mathfrak{F}(\mathtt{x}) = \lim_{N\to\infty}\mathbb{F}^{\rm RSB}(\mathtt{x})/N$, then $\Sigma(\phi)$ is given by the Legendre transform $\mathfrak{F}(\mathtt{x}) = \mathtt{x}\phi + \Sigma(\phi)$, $\partial\Sigma/\partial\phi = -\mathtt{x}$.

For a given code and received $\underline{y}$, the basic objects involved in the 1RSB approach are the distributions of the fields $h_{i\to a}$ and $u_{b\to j}$ denoted respectively as Q_{ia} and $\widehat{Q}_{bj}$. They satisfy the following 1RSB equations:

$$Q_{ia}(h_{i\to a}) \cong \int \delta\left(h_{i\to a} = \mathsf{f}_i(\{u_{b\to i}\})\right)\,(z_{ia})^{\mathtt{x}} \prod_{b\in\partial i\setminus a} {\rm d}\widehat{Q}_{bi}(u_{b\to i})\,, \qquad (21.28)$$

$$\widehat{Q}_{ai}(u_{a\to i}) \cong \int \delta\left(u_{a\to i} = \hat{\mathsf{f}}_a(\{h_{j\to a}\})\right)(\hat{z}_{ai})^{\mathtt{x}} \prod_{j\in\partial a\setminus i} {\rm d}Q_{ja}(h_{j\to a})\,. \qquad (21.29)$$

Exercise 21.8 Show that the factors z_{ia} and $\hat{z}_{ai}$ in these equations, defined in (19.23), (19.24), are given by:

$$z_{ia}(\{u_{b\to i}\},B_i) = \frac{2\cosh(B_i + \sum_{b\in\partial i\setminus a} u_{b\to i})}{\prod_{b\in\partial i\setminus a}(2\cosh(u_{b\to i}))}\,, \qquad (21.30)$$
$$\hat{z}_{ai}(\{h_{j\to a}\}) = 1 + e^{-2\beta}\,. \qquad (21.31)$$

Although in this case $\hat{z}_{ai}$ is a constant and can be absorbed in the normalization, we shall keep it explicitly in the following.

We now turn to the statistical analysis of these equations. Picking up a uniformly random edge in the Tanner graph of a code from the LDPC$_N(\Lambda,P)$ ensemble, the densities $\widehat{Q}$ and Q become themselves random objects which satisfy the distributional equations:

$$Q(h) \stackrel{\rm d}{=} \frac{1}{Z} \int z(\{u_a\};B(y))^{\tt x}\, \delta\Big(h - {\sf f}_{l-1}(\{u_a\};B(y))\Big) \prod_{a=1}^{l-1} {\rm d}\widehat{Q}_a(u_a)\,, \quad (21.32)$$

$$\widehat{Q}(u) \stackrel{\rm d}{=} \frac{1}{\widehat{Z}} \int \hat{z}(\{h_i\})^{\tt x}\, \delta\left(u - \hat{\sf f}_{k-1}(\{h_i\})\right) \prod_{i=1}^{k-1} {\rm d}Q_i(h_i)\,. \quad (21.33)$$

where k, l, y are random variables, $\{\widehat{Q}_a\}$ are $l-1$ i.i.d. copies of $\widehat{Q}$, and $\{Q_i\}$ are $k-1$ i.i.d. copies of Q. Further, l is drawn from the edge perspective variable degree profile λ, k is drawn from the edge perspective check degree profile ρ, and y is drawn from $\mathcal{Q}(\,\cdot\,|{\tt 0})$, the distribution of channel output upon input ${\tt 0}$. The functions $\hat{\sf f}_{k-1}(\{h_i\}) = {\rm atanh}(\tanh\beta \prod_{i=1}^{k-1}\tanh(h_i))$, and ${\sf f}_{l-1}(\{u_a\};B) = B - \sum_{a=1}^{l-1} u_a$ are defined analogously to Eqs. (21.24), (21.25). The functions $z(\,\cdot\,)$ and $\hat{z}(\,\cdot\,)$ are given similarly by the expressions in (21.30), (21.31).

The 1RSB free-entropy density (i.e. the entropy density of the auxiliary model) is estimated as $\mathfrak{F}({\tt x}) = {\rm f}^{\rm RSB}(Q,\widehat{Q})$ where ${\rm f}^{\rm RSB}(Q,\widehat{Q})$ is the expected free-entropy density and Q and $\widehat{Q}$ are distributed according to the 'correct' solution of the distributional equations Eqs. (21.32), (21.33).

$${\rm f}^{\rm RSB}(Q,\widehat{Q}) = -\Lambda'(1)\,\mathbb{E}\log z_{\rm e}(Q,\widehat{Q}) + \mathbb{E}\log z_{\rm v}(\{\widehat{Q}_a\};l,y) + \frac{\Lambda'(1)}{P'(1)}\,\mathbb{E}\log z_{\rm f}(\{Q_i\};k)\,.$$

Here the expectation is taken with respect to k i.i.d. copies of $\widehat{Q}$ and l i.i.d. copies of Q, and with respect to $k \stackrel{\rm d}{=} P.$, $l \stackrel{\rm d}{=} \Lambda.$ and $y \stackrel{\rm d}{=} \mathcal{Q}(\,\cdot\,|{\tt 0})$. Finally, $z_{\rm e}, z_{\rm v}, z_{\rm f}$ read:

$$z_{\rm e}(Q,\widehat{Q}) = \int {\rm d}Q(h)\,{\rm d}\widehat{Q}(u)\Big[\sum_{x={\tt 0}}^{\tt 1} \nu_h(x)\nu_u(x)\Big]^{\tt x}\,, \quad (21.34)$$

$$z_{\rm v}(\{\widehat{Q}_a\};l,y) = \int \prod_{a=1}^{l} {\rm d}\widehat{Q}_a(u_a)\Big[\sum_{x={\tt 0}}^{\tt 1} \frac{\mathcal{Q}(y|x)}{\mathcal{Q}(y|{\tt 0})} \prod_{a=1}^{l} \nu_{u_a}(x)\Big]^{\tt x}\,, \quad (21.35)$$

$$z_{\rm f}(\{Q_i\};k) = \int \prod_{i=1}^{l} {\rm d}Q_i(h_i)\Big[\sum_{\{x_1,\cdots,x_k\}} \prod_{i=1}^{k} \nu_{h_i}(x_i) \left(\mathbb{I}\Big(\sum_i x_i = {\rm even}\Big) + e^{-2\beta}\,\mathbb{I}\Big(\sum_i x_i = {\rm odd}\Big)\right)\Big]^{\tt x}. \quad (21.36)$$

A considerable amount of information is contained in the 1RSB free-energy density $\mathfrak{F}({\tt x})$. For instance, one could deduce from it the energetic complexity by taking the appropriate $\beta \to \infty$ limit. Here we shall not attempt at developing a full solution of the 1RSB distributional equations, but use them to detect the occurrence of a dynamical phase transition.

21.3.2 *Dynamical phase transition*

The location of the dynamical phase transition location $\beta_{\rm d}(p)$ is determined as the smallest value of β such that the distributional equations (21.32), (21.33)

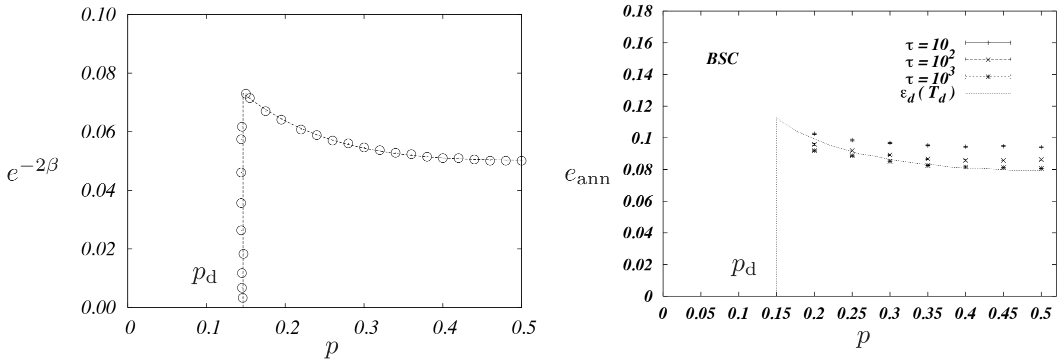

FIG. 21.6. Left: Dynamic phase transition for random codes from the $(5, 6)$ ensemble used over the BSC(p) (circles are obtained through sampled density evolution; the dashed line is a guide for the eye). Right: residual energy density after simulated annealing, as measured in numerical simulations. The dashed line gives the equilibrium energy at the dynamical transition temperature $T_{\rm d}$.

have a non-trivial solution at $\mathtt{x} = 1$. For $\beta > \beta_{\rm d}(p)$, the distribution (21.4) decomposes into an exponential number of pure states. As a consequence, we expect simulated annealing to fall out of equilibrium when $\beta_{\rm d}(p)$ is crossed.

In Fig. 21.6 left frame, we show the result of applying such a technique to the $(5, 6)$ regular ensemble used for communication over the BSC(p). At small p, no dynamic phase transition is revealed through this procedure at any positive temperature. Above a critical value of the noise level p, the behavior changes dramatically and a phase transition is encountered at a critical point $\beta_{\rm d}(p)$ that decreases monotonically for larger p. By changing both β and p, one can identify a phase transition line that separates the ergodic and non-ergodic phases. Remarkably, the noise level at which a finite $\beta_{\rm d}$ appears is numerically indistinguishable from $p_{\rm d} \approx 0.145$.

Does the occurrence of a dynamical phase transition for $p \gtrsim p_{\rm d}$ indeed influence the behavior of the simulated annealing decoder? Some numerical confirmation was already presented in Fig. 21.2. Further support in favor of this thesis is provided by Fig. 21.6, right frame, which plots the residual energy density of the configuration produced by the decoder as $\beta \to \infty$. Above $p_{\rm d}$ this becomes strictly positive and only slowly dependent on the cooling rate. It is compared with the equilibrium value of the internal energy at $\beta_{\rm d}(p)$. This would be the correct prediction if the system didn't decrease any more its energy after it falls out of equilibrium at $\beta_{\rm d}(p)$. Although we do not expect this to be strictly true, the resulting curve provides a good first estimate.

21.3.3 *Metastable states and BP threshold*

One crucial element of this picture can be confirmed analytically, for a generic BMS channel family ordered by physical degradation with respect to p: At zero temperature, the dynamical transition, signaling the proliferation of metastable Bethe states, occurs exactly at the decoding threshold $p_{\rm d}$. More precisely, the argument below proves that at $\beta = \infty$ there cannot exist any non-trivial $\mathtt{x} = 1$ solution of Eqs. (21.32), (21.33) for $p < p_{\rm d}$, while there exists one for $p > p_{\rm d}$. We expect that, for most channel families, the same situation should hold for β large enough (and dependent on p), but this has not been proven yet.

Let us consider the 1RSB equations (21.32), (21.33) in the case $\beta = \infty$. Assuming that the degree profiles are such that $l \geq 2$ and $k \geq 2$ (a reasonable requirement for useful code ensembles), it is clear that they have a special 'no-error' solution associated with the sent codeword in which $Q(h) = \delta_\infty(h)$ and $\widehat{Q}(u) = \delta_\infty(h)$ almost surely. It is a simple exercise to check that the (shifted) free-entropy density of this solution is equal to 0.

The important question is whether there exist other solutions beyond the 'no-error' one. We can make use of the simplification occuring at $\mathtt{x} = 1$. As we saw in Sec. 19.4.1, the expectation values of the messages, $\nu^{\rm av}_{i\to a}(x_i) \equiv \sum_{\nu_{ia}} Q_{ia}(\nu_{ia})\nu_{ia}(x_i)$ and $\widehat{\nu}^{\rm av}_{a\to i}(x_i) \equiv \sum_{\widehat{\mathtt{m}}_{ai}} \widehat{Q}_{ai}(\widehat{\nu}_{ai})\widehat{\nu}_{ai}(x_i)$ satisfy the BP equations.

Let us first study the case $p < p_{\rm d}$. We have seen in Ch. 15 that there is a unique solution of BP equations: the no-error solution. This shows that in this low noise regime, there cannot exist any non-trivial 1RSB solution. We conclude that there is no glass phase in the regime $p < p_{\rm d}$

We now turn to the case $p > p_{\rm d}$ (always with $\beta = \infty$), and use the analysis of BP presented in Ch. 15. That analysis revealed that, when $p > p_{\rm d}$, the density evolution of BP messages admits at least one 'replica symmetric' fixed point distinct from the no-error one.

We shall now use this replica symmetric fixed point in order to construct a non-trivial 1RSB solution. The basic intuition behind this construction is that each Bethe measure consists of a single configuration, well separated from other ones. Indeed, each Bethe measure can be identified with a zero-energy configuration, i.e. with a codeword. If this is true, then, with respect to each of these Bethe measures the local distribution of a variable is deterministic, either a unit mass on `0` or a unit mass on `1`. Therefore we seek a solution where the distribution of Q and $\widehat{Q}$ is supported on functions of the form:

$$Q(h) = \frac{1}{2}(1 + \tanh \tilde{h})\, \delta_{+\infty}(h) + \frac{1}{2}(1 - \tanh \tilde{h})\, \delta_{-\infty}(h)\,, \qquad (21.37)$$

$$\widehat{Q}(u) = \frac{1}{2}(1 + \tanh \tilde{u})\, \delta_{+\infty}(u) + \frac{1}{2}(1 + \tanh \tilde{u})\, \delta_{-\infty}(u)\ , \qquad (21.38)$$

where $\tilde{h}$ and $\tilde{u}$ are random variables.

Exercise 21.9 Show that this Ansatz solves Eqs. (21.32), (21.33) at $\beta = \infty$ if and only if the distributions of $\tilde{h}$, $\tilde{u}$ satisfy:

$$\tilde{h} \overset{\mathrm{d}}{=} B(y) + \sum_{a=1}^{l-1} \tilde{u}\,, \qquad \tilde{u} \overset{\mathrm{d}}{=} \operatorname{atanh}\Big[\prod_{i=1}^{k-1} \tanh \tilde{h}_i\Big]\,. \tag{21.39}$$

It is easy to check that the random variables $\tilde{h}$ and $\tilde{u}$ satisfy the same equations as the fixed point of density evolution for BP (see Eq. (15.11)). We conclude that, for $p > p_{\rm d}$ and $\mathtt{x} = 1$, a solution to the 1RSB equations is given by the Ansatz (21.37), (21.38), if $\tilde{h}$, $\tilde{u}$ are drawn from the fixed point distributions of Eq. (15.11).

It turns out that a similar solution is easily found for any value of $\mathtt{x} > 0$, provided $\beta = \infty$. The only place where $\mathtt{x}$ plays a role is in the reweighting factor of Eq. (21.35): when $\mathtt{x} \neq 1$, the only modification in the distributional equations (21.39) is that $B(y)$ should be multiplied by $\mathtt{x}$. Therefore one can obtain the 1RSB solution for any $\mathtt{x} > 0$ if one knows the solution to the RS cavity equations (i.e. the fixed point of the density evolution for BP) in a slightly modified problem in which $B(y)$ is changed to $\mathtt{x}B(y)$. Technically this is equivalent to studying the modified measure

$$\mu_y(\underline{x}) \cong \prod_{a=1}^{M} \mathbb{I}(x_{i_1^a} \oplus \cdots \oplus x_{i_{k(a)}^a} = 0) \prod_{i=1}^{N} \mathcal{Q}(y_i|x_i)^{\mathtt{x}}\,, \tag{21.40}$$

within the RS approach of Ch. 15 (such a modified measure was already introduced in Ch. 6).

Let us assume that we have found a non-trivial fixed point for this auxiliary problem, characterized by the distributions $\mathsf{a}_{\rm RS}^{(\mathtt{x})}(h)$, and $\hat{\mathsf{a}}_{\rm RS}^{(\mathtt{x})}(u)$, and call $\mathrm{f}^{\rm RS}(\mathtt{x})$ the corresponding value of the free-entropy density defined in (15.45). The 1RSB equations with reweighting parameter $\mathtt{x}$ have a solution of the type (21.37), (21.38), provided $\tilde{h}$ is distributed according to $\mathsf{a}_{\rm RS}^{(\mathtt{x})}(\,\cdot\,)$, and $\tilde{u}$ is distributed according to $\hat{\mathsf{a}}_{\rm RS}^{(\mathtt{x})}(\,\cdot\,)$. The 1RSB free-entropy density $\mathfrak{F}(\mathtt{x}) = \mathbb{E}\,\mathbb{F}^{\rm RSB}(\mathtt{x})/N$ is simply given by:

$$\mathfrak{F}(\mathtt{x}) = \mathrm{f}^{\rm RS}(\mathtt{x})\,. \tag{21.41}$$

Therefore the problem of computing $\mathfrak{F}(\mathtt{x})$, and its Legendre transform the complexity $\Sigma(\phi)$, reduce to a replica symmetric computation. This is a simple generalization of the problem Ch. 15, whereby the decoding measure is modified by raising it to the power $\mathtt{x}$, as in Eq. (21.40). Notice however that the interpretation is now different. In particular $\mathtt{x}$ has to be properly chosen in order to focus on dominant pure states.

The problem can be easily studied numerical using the population dynamics algorithm. Fig. 21.7 shows an example of the complexity $\Sigma(\phi)$ for a BSC channel.

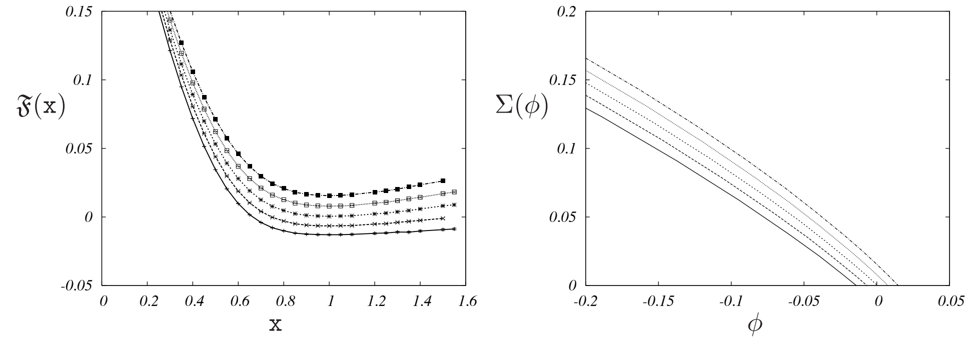

FIG. 21.7. Left: The free-entropy of the auxiliary model, $\mathfrak{F}(\mathtt{x})$, as a function of the weight parameter $\mathtt{x}$, for a $(3, 6)$ code on the BSC channel (recall that $p_{\rm d} \approx 0.084$ and $p_{\rm c} \approx 0.101$ in this case). From bottom to top: $p = 0.090, 0.095, 0.100, 0.105, 0.110$. Right: The complexity $\Sigma(\phi)$ plotted versus the shifted free-entropy density ϕ. From left to right: $p = 0.090, 0.095, 0.100, 0.105, 0.110$.

The regime $p_{\rm d} < p < p_{\rm c}$ is characterized by the existence of a band of metastable states with negative shifted free-entropy $\phi \leq \phi_0 < 0$. They are in principle irrelevant when compared to the 'no-error' solution which has $\phi = 0$, confirming that MAP decoding will return the transmitted codeword. In fact they are even unphysical: ϕ is nothing but the conditional entropy density of the transmitted codeword given the received message. As a consequence it must be non-negative. However the solution extends to $\beta < \infty$, where it makes perfect sense (it describes non-codeword metastable configurations), thus solving the puzzle.

The appearance of metastable states coincides with the noise threshold above which BP decoding fails. When $p > p_{\rm c}$ the top end of the band ϕ_0 becomes positive: the 'glassy' states dominate the measure and MAP decoding fails.

21.4 Metastable states and near-codewords

In a nutshell, the failure of BP decoding for $p > p_{\rm d}$ can be traced back to configurations (words) $\underline{x}$ that: (i) Are deep local minima of the energy function $E(\underline{x})$ (that counts the number of violated parity checks); (ii) Have a significant weight under the measure $\prod_i Q(y_{|}x_i)$.

Typically, such configurations are not codewords, although they can be very close to codeword from the energy point of view. An interesting qualitative analogy can be drawn between this analysis, and various notions that have been introduced to characterize the so-called **error floor**.

Let us start by describing the error floor problem. We saw that for $p < p_{\rm d}$ the bit error rate under BP decoding vanishes when the blocklength $N \to \infty$. Unhappily, the blocklength cannot be taken arbitrarily large because of two types of practical considerations. First, coding a block of N bits simultaneously implies a communication delay proportional to N. Second, any hardware implementation

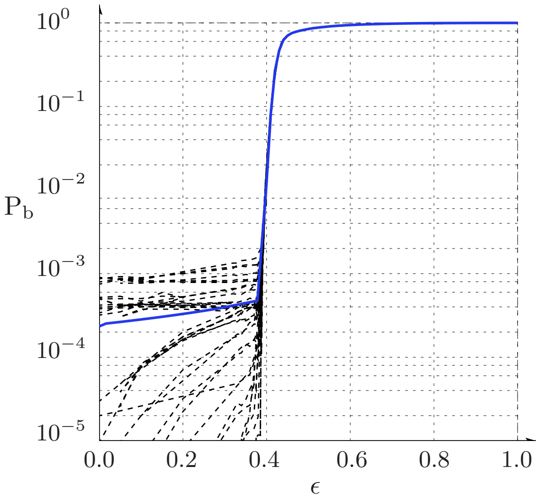

FIG. 21.8. Bit error probability for 40 random elements of the $(3, 6)$ regular ensemble with $N = 2500$ used over the BEC(ϵ). The continuous curve corresponds to the average error probability.

of BP decoding becomes increasingly difficult as N get larger. Depending on the application, one can be forced to consider a maximum blocklength between 10^3 and 10^5.

This brings up the problem of characterizing the bit error rate at moderate blocklength. Figure 21.8 shows the outcomes of numerical simulations for random elements of the $(3, 6)$ ensemble used over the erasure channel. One can clearly distinguish two regimes: a rapid decrease of the error probability in the 'waterfall region' $\epsilon \lesssim \epsilon_{\rm d} \approx 0.429$ (in physics terms, the 'critical regime'); a flattening at lower noise values, in the 'error floor'. It is interesting to note that the error floor level is small but highly dependent (in relative terms) on the graph realization.

We know that the error floor should vanish when taking codes with larger and larger blocklength, but we would like a prediction of its value *given* the graph G. With the notable exception of the erasure channel, this problem is largely open. However several heuristics have been developed. The basic intuition is that the error floor is due to small subgraphs of the Tanner graph that are prone to error. If U is the set of variable nodes in such a subgraph, we can associate to it a configuration $\underline{x}$ that takes value $\mathtt{1}$ on U and 0 otherwise (throughout our analysis we are assuming that the codeword $\underline{\mathtt{0}}$ has been transmitted). This $\underline{x}$ needs not to be a codeword but it is in some sense 'close' to it.

Once a class $\mathcal{F}$ of such subgraphs is identified, the error probability is estimated by assuming that any type of error is unlikely, and errors on different subsets are roughly independent:

$$\mathrm{P_B}(G) \approx \sum_{U \in \mathcal{F}} \mathbb{P}\{\text{BP decoder fails on } U\} \ . \tag{21.42}$$

If the subset U are small, each of the terms on the right hand side can be evaluated efficiently via importance sampling.

It is interesting to have a look at some definitions of the class of subgraphs $\mathcal{F}$ that have been introduced in the literature. In each case the subgraph is characterized by two integers (w, e) that describe how dangerous/close to codewords they are (small w or e corresponding to dangerous subgraphs). In practice one restricts the sum in Eq. (21.42) to small w, e.

Trapping sets. (or **near codewords**) A trapping set is a subgraph including the variable nodes in U, all the adjacent check nodes and the edges that connect them. It is a (w, e) near-codeword if the number of variable nodes is $|U| = w$ and the number of check nodes of odd degree is e.

In our framework a trapping set is simply a configuration $\underline{x}$ with weight (number of non-zero entries) equal to w and energy $E(\underline{x}) = 2e$. Notice that hardly any restriction is imposed on trapping sets. Special constraints are sometimes added depending on the channel model, and on the decoding algorithm (if not BP).

Adsorbing sets. A (w, e) adsorbing set is a (w, e) trapping set that satisfies two further requirements: (i) Each variable node is adjacent to more check nodes of even degree (with respect to the subgraph) than of odd degree; (ii) It does not contain a (w', e) adsorbing set with $w' < w$.

The first condition implies that the corresponding configuration $\underline{x}$ is a local minimum of $E(\underline{x})$ stable with respect to 1 flip.

The connection between small weak subgraphs and error probability is still somewhat vague. The 'energy landscape' $E(\underline{x})$ might provide some hints towards bridging this gap.

Notes

This chapter is largely based on the analysis of metastable states in (Montanari, 2001*b*), (Montanari, 2001*a*) and (Franz, Leone, Montanari and Ricci-Tersenghi, 2002). One step replica symmetry breaking was also investigated in (Migliorini and Saad, 2006). The approach was extended to asymmetric channels in (Neri, Skantzos and Bollé, 2008).

Typical pairs decoding presented here is slightly different from the original procedure of (Aji, Jin, Khandekar, MacKay and McEliece, 2001).

Stopping sets were introduced in (Di, Proietti, Richardson, Telatar and Urbanke, 2002), and inspired much of the subsequent research on error floors. The idea that small subgraphs of the Tanner graph are responsible for error floors was first convincingly demonstrated for general channel models in (MacKay and Postol, 2003) and (Richardson, 2003). Absorbing sets are defined in (Dolecek, Zhang, Anantharam and Nikolić, 2007).

After its invention, simulated annealing was the object of a significant amount of work within operations research and probability. A review can be found in

(Aarts, Korst and van Laarhoven, 2003). A detailed comparison between 1RSB analysis and simulated annealing experiments for models on sparse graphs is presented in (Montanari and Ricci-Tersenghi, 2004).

22

AN ONGOING STORY

This book describes a unified approach to a number of important problems in information theory, physics and computer science. We have presented a consistent set of methods to address these problems, but the field is far from being fully understood, and there remain many open challenges. This chapter provides a synthetic description of some of these challenges, as well as a survey of recent progress. Our ambition is to set an agenda for the newly developed field that we have been describing. We will distinguish roughly three types of directions.

The first one, to be discussed in Sec. 22.1, is the main challenge. It aims at a better qualitative understanding of models on sparse random graphs. At the core of the cavity method lies the postulate that such systems can have only a limited number of 'behaviors' (phases). Each phase corresponds to a different pattern of replica symmetry breaking (replica symmetric -RS, one-step replica symmetry breaking -1RSB, etc. . .). In turn they also have a description in terms of pure states decomposition, as well as in terms of long range correlations. Understanding the fundamental reasons and conditions for the universality of these phases, as well as the equivalence among their characterizations would be extremely important.

The second direction, described in Sec. 22.2, concerns the development of the cavity formalism itself. We have mainly focused on systems in which either the RS or 1RSB cavity method is expected to be asymptotically exact in the large size limit. This expectation is in part based on some internal consistency checks of the 1RSB approach. An important one consists in verifying that the 1RSB 'solution' is stable with respect to small perturbations. Whenever this test is passed, physicists feel confident enough that the cavity method provides exact conjectures (thresholds, minimum cost per variable, etc. . .). If the test is not passed, higher order RSB is thought to be needed. The situation is much less satisfactory in this case, and the cavity method poses some technical problems even at the heuristic level.

Section 22.3 lists a number of fascinating questions that arise in the connexion between the existence of glassy phase transitions and algorithmic slowdown. These are particularly important in view of the applications in computer science and information theory: sparse graphical models can be useful for a number of practically relevant tasks, as the example of LDPC codes in channel coding has shown. There is some empirical evidence that phase transitions have an impact on algorithms behavior and efficiency. Physicists hope that this impact can be understood (to some extent) in a unified way, and is ultimately related to the geometric structure of the set of solutions, and to correlation properties of the measure. While some general arguments in favour of this statement have been put forward, the actual understanding is still very poor.

22.1 Gibbs measures and long-range correlations

At an abstract level, the cavity method explored in the last few chapters relies on a (yet unproven) *structural theorem*. Consider a generic graphical model, a probability distribution on N variables, $\underline{x}$, taking values in a discrete space $\mathcal{X}^N$:

$$\mu(\underline{x}) = \frac{1}{Z} \prod_{a \in F} \psi_a(\underline{x}_{\partial a})\,. \tag{22.1}$$

The cavity method postulates that, for large classes of models taken from some appropriate ensembles, the model is qualitatively described in the large N limit by one out of a small number of generic scenarios, or phases. The postulated qualitative features of such phases are then cleverly used to derive quantitative predictions (e.g. phase transition locations.)

Needless to say, we are not able to state precisely, let alone to prove, such a structural theorem in this generality. The complete set of necessary hypotheses is unknown. However we discussed several examples, from XORSAT to diluted spin glasses or error correcting codes. In principle, it is not necessary that the factor graph be locally tree-like, but in practice locally tree-like models are the ones that we can control most effectively. Such a structure implies that when one digs a cavity in the graph, the variables on the boundary of the cavity are far apart. This leads to a simple structure of their correlation in the large system limit, and hence to the possibility of writing asymptotically exact recursion equations.

Here we do not want to discuss in more details the hypotheses. It would certainly be a significant achievement to prove such a structural theorem even in a restricted setting (say, for the uniform measure over solutions of random K-SAT formulae). We want instead to convey some important features of the phases postulated within the cavity approach. In particular there is a key aspect that we want to stress. Each of the various phases mentioned can be characterized from two, complementary, points of view:

1. In terms of decomposition of the distribution $\mu(\,\cdot\,)$ into ‘lumps’ or ‘clusters’. Below we shall propose a precise definition of the lumps, and they will be called **pure states**.
2. In terms of correlations among far apart variables on the factor graph. We shall introduce two notions of *correlation decay* that differ in a rather subtle way but correspond to different phases.

These two characterizations are in turn related to the various aspects of the cavity method.

22.1.1 *On the definition of pure states*

The notion of pure state is a crucial one in rigorous statistical mechanics. Unfortunately, standard definitions are tailored to translation-invariant models on infinite graphs. The graphical models that we have in mind are sparse random graphs (in this class we include labeled random graphs, whereby the labels specify the nature of function nodes), and standard approaches don't apply to them. In particular, we need a concrete definition that is meaningful for finite graphs.

Consider a sequence of *finite* graphical models $\{\mu_N(\,\cdot\,)\}$, indexed by the number of variable nodes N. A **pure state decomposition** is defined by assigning, for each N, a partition of the configuration space $\mathcal{X}^N$ into $\mathcal{N}_N$ subsets $\Omega_{1,N},\dots,\Omega_{\mathcal{N}_N,N}$:

$$\mathcal{X}^N = \Omega_{1,N}\cup\cdots\cup\Omega_{\mathcal{N}_N,N}\,. \tag{22.2}$$

The pure state decomposition must meet the following conditions:

1. The measure of each subset in the partition is bounded away from 1:
$$\max\{\mu_N(\Omega_{1,N}),\dots,\mu_N(\Omega_{\mathcal{N},N})\}\le 1-\delta\,. \tag{22.3}$$
2. The subsets are separated by 'bottlenecks.' More precisely, for $\Omega\subseteq\mathcal{X}^N$, define its ϵ-boundary as
$$\partial_\epsilon\Omega\equiv\{x\in\mathcal{X}^N\,:\,1\le d(x,\Omega)\le N\epsilon\}\,. \tag{22.4}$$
where $d(x,\Omega)$ is the minimum Hamming distance between x and any configuration $x'\in\Omega$. Then we require
$$\lim_{N\to\infty}\max_r\frac{\mu_N(\partial_\epsilon\Omega_{r,N})}{\mu_N(\Omega_{r,N})}=0\,, \tag{22.5}$$
for some $\epsilon>0$. Notice that the measure of $\partial_\epsilon\Omega_{r,N}$ can be small for two reasons, either because $\Omega_{r,N}$ is small itself (and therefore has a small boundary) or because the boundary of $\Omega_{r,N}$ is much smaller than its interior. Only the last situation corresponds to a true bottleneck, as is enforced by the denominator $\mu_N(\Omega_{r,N})$ in (22.5).
3. The conditional measure on the subset $\Omega_{r,N}$, defined by
$$\mu_N^r(\underline{x})\equiv\frac{1}{\mu_N(\Omega_{r,N})}\,\mu_N(\underline{x})\mathbb{I}(\underline{x}\in\Omega_{r,N}) \tag{22.6}$$
cannot be further decomposed according to the two conditions above.

Given such a partition, the distribution $\mu_N(\,\cdot\,)$ can be written as a convex combination of distributions with disjoint support

$$\mu_N(\,\cdot\,)=\sum_{r=1}^{\mathcal{N}_N}w_r\,\mu_N^r(\,\cdot\,)\,,\qquad w_r\equiv\mu_N(\Omega_{r,N})\,. \tag{22.7}$$

Notice that this decomposition is not necessarily unique, as shown by the example below. Non-uniqueness is due to the fact that sets of configurations of $\mathcal{X}^N$ with negligeable weight can be attributed to one state or another. On the other hand, the conditional measures $\mu_N^r(\,\cdot\,)$ should depend weakly on the precise choice of decomposition.

Example 22.1 Consider the ferromagnetic Ising model on a random regular graph of degree $(k+1)$. The Boltzmann distribution reads

$$\mu_N(\underline{x}) = \frac{1}{Z_N(\beta)} \exp\left\{\beta \sum_{(i,j)\in E} x_i x_j\right\}, \tag{22.8}$$

with $x_i \in \mathcal{X} = \{+1, -1\}$. To avoid irrelevant complications, let's assume that N is odd. Following the discussion of Sec. 17.3, we expect this distribution to admit a non-trivial pure state decomposition for $k \tanh \beta > 1$, with partition $\Omega_+ \cup \Omega_- = \mathcal{X}^N$. Here Ω_+ (respectively Ω_-) is the set of configurations for which $\sum_i x_i$ is positive (negative). With respect to this decomposition $w_+ = w_- = 1/2$.

Of course an (asymptotically) equivalent decomposition is obtained by letting Ω_+ be the set of configurations with $\sum_i x_i \geq C$ for some fixed C.

It is useful to recall that the condition (22.5) implies that any 'local' Markov dynamics that satisfies detailed balance with respect to $\mu_N(\,\cdot\,)$ is slow. More precisely, assume that

$$\frac{\mu_N(\partial_\epsilon \Omega_{r,N})}{\mu_N(\Omega_{r,N})} \leq \exp\{-\Delta(N)\}\,. \tag{22.9}$$

Then any Markov dynamics that satisfies detailed balance with respect to μ_N and flips at most $N\epsilon$ variables at each step, has relaxation time larger than $C\exp\{\Delta(N)\}$ (where C is an N-independent constant that depends on the details of the model). Moreover, if the dynamics is initialized in $\underline{x} \in \Omega_{r,N}$, it will take a time of order $C\exp\{\Delta(N)\}$ to get at distance $N\epsilon$ from $\Omega_{r,N}$.

In many cases based on random factor graph ensembles, we expect Eq. (22.9) to hold with a $\Delta(N)$ which is linear in N. In fact in the definition of pure state decomposition we might ask a bound of the form (22.9) to hold, for some function $\Delta(N)$ (e.g. $\Delta(N) = N^\psi$, with some appropriately chosen ψ). This implies that pure states are stable on time scales shorter than $\exp\{\Delta(N)\}$.

22.1.2 *Notions of correlation decay*

The above discussion on relaxation times brings up a second key concept: **correlation decay**. According to an important piece of wisdom in statistical mechanics, physical systems that have only short-range correlations should relax rapidly to their equilibrium distribution. The hand-waving reason is that, if different degrees of freedom (particles, spins, etc) are independent, then the system relaxes on microscopic time scales (namely the relaxation time of a single particle, spin, etc). If they are not independent, but correlations are short ranged, they can be coarse grained in such a way that they become nearly independent, Roughly speaking, this means that one can construct 'collective' variables from blocks of original variables. Such conditional variables take $|\mathcal{X}|^B$ values, where

B is the block size, and are nearly independent under the original (Boltzmann) distribution.

As we are interested in models on non-Euclidean graphs, the definition of correlation decay must be precised. We will introduce two distinct types of criteria. Although they may look similar at first sight, it turns out that they are not, and each of them will characterize a distinct generic phase.

The simplest approach, widely used in physics, consists in considering two-points correlation functions. Averaging them over the two positions defines a susceptibility. For instance, in the case of Ising spins $x_i \in \mathcal{X} = \{1, -1\}$, we have already discussed the spin glass susceptibility

$$\chi^{\rm SG} = \frac{1}{N} \sum_{i,j\in V} \left(\langle x_i x_j\rangle - \langle x_i\rangle\langle x_j\rangle\right)^2 \,, \tag{22.10}$$

where $\langle\,\cdot\,\rangle$ denotes the expectation value with respect to μ. When $\chi^{\rm SG}$ is bounded as $N \to \infty$, this is an indication of short range correlations. Through the fluctuation dissipation theorem (cf. Sec. 2.3), this is equivalent to stability with respect to local perturbations. Let us recall the mechanism of this equivalence. Imagine a perturbation of the model (22.16) that acts on a single variable x_i. Stability requires that the effect of such a perturbation on the expectation of a global observable $\sum_j f(x_j)$ should be bounded. The change in the marginal at node j due to a perturbation at i, is proportional to the covariance $\langle x_i x_j\rangle - \langle x_i\rangle\langle x_j\rangle$. As in Sec. 12.3.2, the average effect of the perturbation at i on the variables x_j, $j \neq i$ often vanishes (more precisely $\lim_{N\to\infty} \frac{1}{N}\sum_{j\in V}\big(\langle x_i x_j\rangle - \langle x_i\rangle\langle x_j\rangle\big) = 0$) because terms related to different vertices j cancel. The *typical* effect of the perturbation is captured by the spin glass-susceptibility.

Generalizing this definition to arbitrary alphabets is easy. We need to use a measure of how much the joint distribution $\mu_{ij}(\,\cdot\,,\,\cdot\,)$ of x_i and x_j is different from the product of the marginals $\mu_i(\,\cdot\,)$ times $\mu_j(\,\cdot\,)$. One such measure is provided by the variation distance:

$$||\mu_{ij}(\,\cdot\,,\,\cdot\,) - \mu_i(\,\cdot\,)\mu_j(\,\cdot\,)|| \equiv \frac{1}{2}\sum_{x_i,x_j} |\mu_{ij}(x_i,x_j) - \mu_i(x_i)\mu_j(x_j)|\,. \tag{22.11}$$

We then define the two-points correlation by averaging this distance over the vertices i, j

$$\chi^{(2)} \equiv \frac{1}{N}\sum_{i,j\in V} ||\mu_{ij}(\,\cdot\,,\,\cdot\,) - \mu_i(\,\cdot\,)\mu_j(\,\cdot\,)||\,. \tag{22.12}$$

Exercise 22.1 Consider again the case of Ising variables, $\mathcal{X} = \{+1, -1\}$. Show that $\chi^{\rm SG} = o(N)$ if and only if $\chi^{(2)} = o(N)$.

[Hint: Let $C_{ij} \equiv \langle x_i x_j\rangle - \langle x_i\rangle\langle x_j\rangle$. Show that $C_{ij} = 2||\mu_{ij}(\,\cdot\,,\,\cdot\,) - \mu_i(\,\cdot\,)\mu_j(\,\cdot\,)||$. Then use $\chi^{\rm SG} = N\mathbb{E}\{C_{ij}^2\}$, $\chi^{(2)} = N\mathbb{E}\{|C_{ij}|\}/2$, the expectation $\mathbb{E}$ being over uniformly random $i, j \in V$.]

Of course one can define l-points correlations in an analogous manner:

$$\chi^{(l)} \equiv \frac{1}{N^{l-1}} \sum_{i(1),\dots,i(l)\in V} ||\mu_{i(1)\dots i(l)}(\,\cdots) - \mu_{i(1)}(\,\cdot\,)\cdots\mu_{i(l)}(\,\cdot\,)|| \,. \tag{22.13}$$

The l-points correlation $\chi^{(l)}$ has a useful interpretation in terms of a thought experiment. Suppose you are given an N-dimensional distribution $\mu(\underline{x})$ and have access to the marginal $\mu_{i(1)}(\,\cdot\,)$ at a uniformly random variable node $i(1)$. You want to test how stable is this marginal with respect to small perturbations. Perturbations affect $l-1$ randomly chosen variable nodes $i(2)$,..., $i(l)$ changing $\mu(\underline{x})$ into $\mu'(\underline{x}) \cong \mu(\underline{x})(1+\delta_2(x_{i(2)}))\cdots(1+\delta_l(x_{i(l)}))$. The effect of the resulting perturbation on $\mu_{i(1)}$, to the first order in the product $\delta_2\cdots\delta_l$, is bounded in expectation by $\chi^{(l)}$ (this is again a version of the fluctuation dissipation theorem).

Definition 22.2. (First type of correlation decay) *The graphical model given by $\mu(\,\cdot\,)$ is said to be **stable to small perturbations** if, for all finite l, $\chi^{(l)}/N \to 0$ as $N \to \infty$.*

In practice in sufficiently homogeneous (mean field) models, this type of stability is equivalent to the one found using only $l = 2$.

Let us now introduce another type of criterion for correlation decay. Again we look at a variable node i, but now we want to check how strongly x_i is correlated with *all the* 'far apart' variables. Of course we must define what 'far apart' means. Fix an integer ℓ and define $\mathsf{B}(i,\ell)$ as the ball of radius ℓ centered at i, and $\overline{\mathsf{B}}(i,\ell)$ its complement, i.e. the subset of variable nodes j such that $d(i,j) \geq \ell$. We then want to estimate the correlation between x_i and $\underline{x}_{\overline{\mathsf{B}}(i,\ell)} = \{x_j : j \in \overline{\mathsf{B}}(i,\ell)\}$. This amounts to measuring the distance between the joint distribution $\mu_{i,\overline{\mathsf{B}}(,\ell)}(\,\cdot\,,\,\cdot\,)$ and the product of the marginals $\mu_i(\,\cdot\,)\mu_{\overline{\mathsf{B}}(,\ell)}(\,\cdot\,)$. If we use the total variation distance defined in (22.11) we obtain the following **point-to-set correlation function**

$$G_i(\ell) \equiv ||\mu_{i,\overline{\mathsf{B}}(i,\ell)}(\,\cdot\,,\,\cdot\,) - \mu_i(\,\cdot\,)\mu_{\overline{\mathsf{B}}(i,\ell)}(\,\cdot\,)|| \,. \tag{22.14}$$

The function $G_i(\ell)$ can be interpreted according to two distinct but equally suggestive thought experiments. The first one comes from the theory of structural glasses (it is meant to elucidate the kind of long range correlations arising in a fragile glass). Imagine to draw a reference configuration $\underline{x}^*$ from the distribution $\mu(\,\cdot\,)$. Now generate a second configuration $\underline{x}$ as follows: variables outside the ball, with $i \in \overline{\mathsf{B}}(i,\ell)$, are forced to the reference configuration: $x_i = x_i^*$. Variables at distance smaller than ℓ (denoted by $\underline{x}_{\mathsf{B}(i,\ell)}$) are instead drawn from the conditional distribution $\mu(\underline{x}_{\mathsf{B}(i,\ell)}|\underline{x}^*_{\overline{\mathsf{B}}(i,\ell)})$. If the model $\mu(\,\cdot\,)$ has some form of *rigidity* (long range correlations), then x_i should be close to x_i^*. The correlation $G_i(\ell)$ measures how much the distributions of x_i and x_i^* differ.

The second experiment is closely related to the first one, but has the flavour of a statistics (or computer science) question. Someone draws the configuration

$\underline{x}^*$ as above from the distribution $\mu(\,\cdot\,)$. She then reveals to you the values of far apart variables in the reference configuration, i.e. the values x_j^* for all $j \in \overline{\mathsf{B}}(i,\ell)$. She asks you to *reconstruct* the value of x_i^*, or to guess it as well as you can. The correlation function $G_i(\ell)$ measures how likely you are to guess correctly (assuming unbounded computational power), compared to the case in which no-variable has been revealed to you.

This discussion suggests the following definition:

Definition 22.3. (Second type of correlation decay) *The graphical model $\mu(\,\cdot\,)$ is said to satisfy the **non-reconstructibility** (or **extremality**) condition if for all i's, $G_i(\ell) \to 0$ as $\ell \to \infty$. (More precisely, we require that there exists a function $\delta(\ell)$, with $\lim_{\ell\to\infty} \delta(\ell) = 0$, such that $G_i(\ell) \le \delta(\ell)$ for all i and N). In the opposite case, i.e. if $G_i(\ell)$ remains bounded away from zero at large distance, the model is said **reconstructible**.*

22.1.3 *Generic scenarios*

We shall now describe the correlation decay properties and the pure state decomposition for the three main phases that we have encountered in the previous chapters: RS, dynamical 1RSB, and static 1RSB. When dealing with models on locally tree-like random graphs, each of these phases can also be studied using the appropriate cavity approach, as we shall recall.

Here we focus on phases that appear ‘generically’. This means that we exclude: (*i*) Critical points, that are obtained by fine-tuning some parameters of the model; (*ii*) Multiplicities due to global symmetries, like for instance in the zero-field ferromagnetic Ising model. Of course there also exist other types of generic phases, such as higher order RSB phases that will be discussed in the next section, and maybe some more that have not been explored yet.

Replica symmetric. In this phase there exists no non-trivial decomposition into pure states of the form (22.7). In other words $\mathcal{N}_N = 1$ with high probability.

Correlations decay according to both criteria: the model is stable to small perturbations and it satisfies the non-reconstructibility condition. Therefore it is short-range correlated in the strongest sense.

Finally, the replica symmetric cavity method of Ch. 14 yields asymptotically exact predictions.

Dynamical 1RSB. In this phase the measure $\mu(\,\cdot\,)$ admits a non trivial decomposition of the form (22.7) into an exponential number of pure states: $\mathcal{N}_N = e^{N\Sigma+o(N)}$ with high probability for some $\Sigma > 0$. Furthermore, most of the measure is carried by states of equal size. More precisely, for any $\delta > 0$, all but an exponentially small fraction of the measure is comprised in states $\Omega_{r,N}$ such that

$$-\Sigma - \delta \le \frac{1}{N}\log \mu(\Omega_{r,N}) \le -\Sigma + \delta\,. \qquad (22.15)$$

From the correlation point of view, this phase is stable to small perturbations, but it is reconstructible. In other words, a finite number of probes would fail to

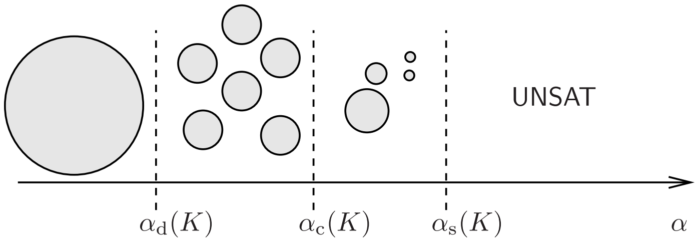

FIG. 22.1. A pictorial view of the different phases in K-SAT with $K \geq 4$, depending on the number of clauses per variable α. Form left to right: replica symmetric, dynamical 1RSB, static 1RSB and UNSAT.

reveal long range correlations. But long range correlations of the point-to-set type are instead present, and they are revealed, for instance, by a slowdown of reversible Markov dynamics.

The glass order parameter overlap distribution $P(q)$ is trivial in this phase (as implied by (12.31)), but its glassy nature can be found through the ϵ-coupling method of Sec. 12.3.4.

The model is solved exactly (in the sense of determining its asymptotic free-energy density) within the 1RSB cavity method. The thermodynamically dominant states, i.e. those satisfying (22.15), correspond to the 1RSB parameter $\mathtt{x} = 1$.

Static 1RSB. This is the 'genuine' 1RSB phase analogous to the low temperature phase of the random energy model. The model admits a non-trivial pure states decomposition with wildly varying weights. For any $\delta > 1$, a fraction $1 - \delta$ of the measure is comprised in the $k(N, \delta)$ pure states with largest weight. The number $k(N, \delta)$ converges, when $N \to \infty$, to a *finite* random variable (taking integer values). If we order the weights according to their magnitude $w^{(1)} \geq w^{(2)} \geq w^{(3)} \geq \cdots$, they converge to a Poisson-Dirichlet process, cf. Ch. 8.

This phase is not stable to small perturbation, and it is reconstructible: It has long range correlations according to both criteria. The asymptotic overlap distribution function $P(q)$ has two delta-function peaks, as in Fig.12.3.

Again, it is solved exactly within the 1RSB cavity method.

These three phases are present in a variety of models, and are often separated by phase transitions. The 'clustering' or 'dynamical' phase transition separates the RS and dynamical 1RSB phases, while a condensation phase transition separates the dynamical 1RSB from the static 1RSB phase. Fig. 22.1.3 describes the organization of various phases in random K-SAT with $K \geq 4$, as we discussed in Sec. 20.3. For $\alpha < \alpha_{\rm d}(K)$ the model is RS; for $\alpha_{\rm d}(K) < \alpha < \alpha_{\rm c}(K)$, it is dynamically 1RSB; for $\alpha_{\rm c}(K) < \alpha < \alpha_{\rm s}(K)$, it is statically 1RSB, for $\alpha_{\rm s}(K) < \alpha$ it is UNSAT. Fig. 22.1.3 shows the point-to-set correlation function in random 4-SAT. It clearly develops long-range correlations at $\alpha \geq \alpha_{\rm d} \approx 9.38$. Notice the

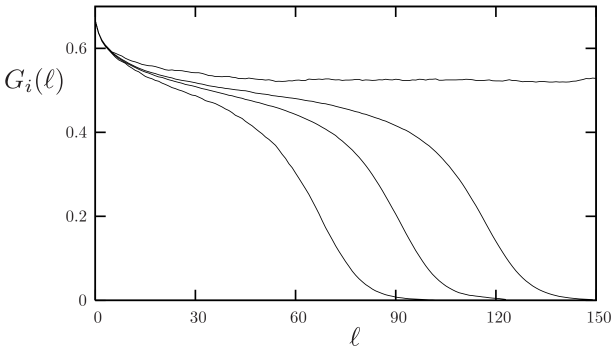

FIG. 22.2. The point-to-set correlation function defined in (22.14) is plotted versus distance for random 4-satisfiability, at clause densities $\alpha = 9.30$, 9.33, 9.35 and 9.40 (from bottom to top).

peculiar development of correlations through a plateau whose width increases with α, and diverges at $\alpha_{\rm d}$. This is typical of the dynamical 1RSB transition.

22.2 Higher levels of replica symmetry breaking

For some of the models studied in this book the RS, or the 1RSB cavity method are thought to yield asymptotically exact predictions. However, in general higher orders of RSB are necessary. We shall sketch how to construct these higher order solutions hierarchically in locally tree-like graphical models. In particular, understanding the structure of the 2RSB solution allows to derive a 'stability criterion' for the 1RSB approach. It is on the basis of this criterion that, for instance, our derivation of the SAT-UNSAT threshold in Ch. 20 is conjectured to give an exact result.

22.2.1 *The high-level picture*

Let us first briefly summarize the RS/1RSB approach. Consider an ensemble of graphical models defined through the distribution (22.1) with a locally tree-like factor graph structure. Within the RS cavity method, the local marginals of $\mu(\,\cdot\,)$ are accurately described in terms of the message sets $\{\nu_{i\to a}\}$, $\{\widehat{\nu}_{a\to i}\}$. Given a small (tree-like) subgraph induced by the vertex set $U \subset V$, the effect of the rest of the graph $G \setminus G_U$ on U is described by a factorized measure on the boundary of U.

One-step replica symmetry breaking relaxes this assumption, by allowing for long-range correlations, with a peculiar structure. Namely, the probability distribution $\mu(\,\cdot\,)$ is assumed to decompose into the convex combination of Bethe measures $\mu_r(\,\cdot\,)$. Within each 'state' r, the local marginals of the measure restricted to this state are well described in terms of a set of messages $\{\nu^r_{i\to a}\}$ (by 'well described' we mean that the description becomes asymptotically exact at large N). Sampling at random a state r defines a probability distribution

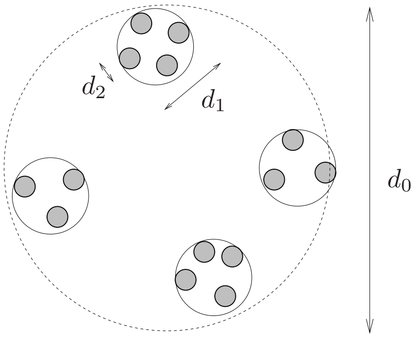

FIG. 22.3. Cartoon of the distribution $\mu(\underline{x})$ for a model described by two-step replica symmetry breaking. The probability mass is concentrated on the gray 'lumps' of radius d_2, which are organized in 'clouds' of radius $d_1 > d_2$. The dashed circle corresponds to the typical distance d_0 between clouds.

$\mathsf{P}(\{\nu\}, \{\widehat{\nu}\})$ over messages. This distribution is then found to be described by an 'auxiliary' graphical model which is easily deduced from the original one. In particular the auxiliary factor graph inherits the structure of the original one, and therefore it is again locally tree-like. 1RSB amounts to using the RS cavity method to study of this auxiliary graphical model over messages.

In some cases 1RSB is expected to be asymptotically exact in the thermodynamic limit. However, this is not always the case: it may fail because the measure $\mathsf{P}(\{\nu\}, \{\widehat{\nu}\})$ decomposes into multiple pure states. Higher-order RSB is used to study this type of situation by iterating the above construction.

More precisely, the two-step replica symmetry breaking (2RSB) method starts from the 'auxiliary' distribution $\mathsf{P}(\{\nu\}, \{\widehat{\nu}\})$. Instead of studying it with the RS method as we did so far, we use instead the 1RSB method to study $\mathsf{P}(\{\nu\}, \{\widehat{\nu}\})$ (introducing therefore an auxiliary auxiliary model, that is studied by the RS method).

The 2RSB Ansatz admits a hand-waving interpretation in terms of the qualitative features of the original model $\mu(\,\cdot\,)$. Reconsider again 1RSB. The interpretation was that $\mu(\,\cdot\,)$ is the convex combination of 'pure states' $\mu^r(\,\cdot\,)$, each forming a well separated lump in configuration space. Within 2RSB, lumps have a hierarchical organization, i.e. they are grouped into 'clouds'. Each lump is addressed by giving a 'cloud index' r_1, and, within the cloud, a 'lump index' r_2. The measure thus decomposes as

$$\mu(\underline{x}) = \sum_{r_1 \in S_1,\, r_2 \in S_2(r_1)} w_{r_1,r_2}\, \mu^{r_1,r_2}(\underline{x})\,. \tag{22.16}$$

Here $S_2(r_1)$ is the set of indices of the lumps inside cloud r_1. A pictorial sketch of this interpretation is shown in Fig. 22.2.1.

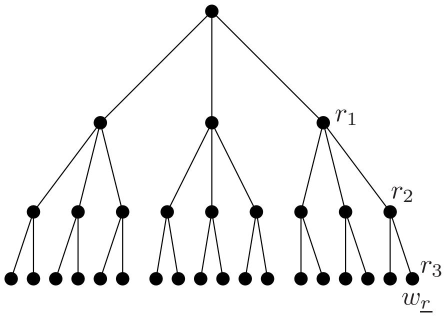

FIG. 22.4. Hierarchical structure of the distribution $\mu(\underline{x})$ within k-step replica symmetry breaking. Here $k = 3$.

Even the most forgiving reader should be puzzled by all this. For instance, what is the difference between $\mathcal{N}_1$ clouds, each involving $\mathcal{N}_1$ lumps, and just $\mathcal{N}_1\mathcal{N}_2$ lumps? In order to distinguish between these two cases one can look at a properly defined distance, say the Hamming distance divided by N, between two i.i.d. configurations drawn with distribution $\mu(\,\cdot\,)$ (in physics jargon, two replicas). If one conditions on the two configurations to belong to the same lump, to different lumps within the same cloud, or to different clouds, the normalized distances concentrate around three values, respectively d_2, d_1, d_0, with $d_2 < d_1 < d_0$. As in the case of 1RSB, one could in principle distinguish dynamic and static 2RSB phases depending on the number of relevant clouds and lumps within clouds. For instance in the most studied case of static 2RSB, these numbers are subexponential. As a consequence, the asymptotic distribution of the distance between two replicas has non-zero weight on each of the three values d_0, d_1, d_2 (in other words, the overlap distribution $P(q)$ is the combination of three delta functions).

Of course this whole construction can be bootstrapped further, by having clouds grouped into larger structures etc... Within k-RSB, the probability distribution $\mu(\,\cdot\,)$ is a convex combination of 'states' $\mu^{\underline{r}}(\,\cdot\,)$ where $\underline{r} = (r_1, r_2, \dots, r_k)$ indexes the leaves of a k-generations tree. The indices r_1, r_2, $\dots, r_k$ correspond to the nodes encountered along the path between the root and the leaf. This translates into a hierarchy of auxiliary graphical models. By allowing k to be arbitrarily large, this hierarchy is expected to determine the asymptotic properties of a large class of models. In particular one can use it to compute the free-entropy per variable $\phi \equiv \lim_{N\to\infty} N^{-1} \log Z_N$.

The resulting description of $\mu(\underline{x})$ has a natural ultrametric structure, as discussed in Ch. 8 and recalled in Fig. 22.4. This structure is captured by the generalized random energy model (GREM), a simple model that generalizes the REM discussed in Chapter 5. While presenting the solution of the GREM would take us too far, it is instructive to give its definition.

Example 22.4 The GREM is a simple model for the probability distribution $\mu(\,\cdot\,)$, within k-step RSB. Its definition involves one parameter $N \in \mathbb{N}$ that corresponds to the system size, and several others (to be denoted as $\{a_0, a_1, \dots, a_{k-1}\}$, $\{d_0, d_2, \dots, d_{k-1}\}$ and $\{\Sigma_0, \Sigma_1, \dots, \Sigma_{k-1}\}$) that are thought to be fixed as $N \to \infty$. States are associated with the leaves of a k-generations tree. Each leaf is indexed by the path $\underline{r} = (r_0, \dots, r_{k-1})$ that connects it to the root, cf. Fig. 22.4.

The GREM does not describe the structure of each state $\mu_{\underline{r}}(\,\cdot\,)$ (that can be thought as supported on a single configuration). It only describes the distribution of distances between the states, and the distribution of the weights $w_{\underline{r}}$ appearing in the decomposition (22.16).

A node at level i has $\exp\{N\Sigma_i\}$ offsprings. The total number of states is therefore $\exp\{N(\Sigma_0 + \cdots + \Sigma_{k-1})\}$. Two random configurations drawn from states $\underline{r}$ and $\underline{s}$ have distance $d_{i(\underline{r},\underline{s})}$, where $i(\underline{r},\underline{s})$ is the largest integer i such that $r_i = s_i$. Finally, the weight of state $\underline{r}$ has the form

$$w_{\underline{r}} = \frac{1}{Z} \exp\{-\beta(E^{(0)}_{r_0} + \cdots + E^{(k-1)}_{r_{k-1}})\}\,, \tag{22.17}$$

where $E^{(i)}_r$ are independent normal random variables with mean 0 and variance Na_i. The interested reader is invited to derive the thermodynamic properties of the GREM, for instance the free-energy as a function of the temperature.

22.2.2 *What does 2RSB look like?*

Higher order RSB has been studied in some detail in many ‘fully connected’ models such as the p-spin Ising model considered in Chapter 8. On the contrary, if one considers models on sparse graphs as we do here, any cavity calculation beyond 1RSB is technically very challenging. In order to understand why, it is interesting to have a superficial look at how a 2RSB cavity calculation would be formally set up without any attempt at justifying it.

For the sake of simplicity we shall consider a model of the form (22.1) with pairwise interactions. Therefore all the factor nodes have degree 2, and BP algorithms can be simplified by using only one type of messages passed along the edges of an ordinary graph, cf. Sec. 14.2.5. Consider a variable node $0 \in V$ of degree $(l+1)$, and denote l of its neighbors by $\{1, \dots, l\}$. We let $\nu_1, \dots, \nu_l$ be the messages from (respectively) $1, \dots, l$, and ν_0 the message from 0 to its $(l+1)$-th neighbor.

As we saw in Sec. 14.2.5, the RS cavity equation (i.e. the BP fixed point equation) at node 0 reads

$$\nu_0(x_0) = \frac{1}{z\{\nu_i\}} \prod_{i=1}^{k} \sum_{x_i} \psi_{0i}(x_0, x_i)\nu_i(x_i)\,, \tag{22.18}$$

where $z\{\nu_i\}$ is determined by the normalization condition of $\nu_0(\,\cdot\,)$. In order to lighten the notation, it is convenient to introduce a function f_0 that, evaluated on l messages $\nu_1, \dots, \nu_l$ returns the message ν_0 as above. We will therefore write Eq. (22.18) in shorthand form as $\nu_0 = \mathsf{f}_0\{\nu_i\}$. Each ν_i is a point in the $(|\mathcal{X}|-1)$-dimensional simplex.

The 1RSB cavity equations are obtained from Eq. (22.18) by promoting the messages ν_i to random variables with distribution $Q_i(\,\cdot\,)$, cf. Ch. 19. The equations depend on the 1RSB parameter (a real number), that we denote here as $\mathtt{x}_1$. Adopting a continuous notation for the messages distributions, we get

$$Q_0(\nu_0) = \frac{1}{Z\{Q_i\}} \int z\{\nu_i\}^{\mathtt{x}_1}\, \delta(\nu_0 - \mathsf{f}_0\{\nu_i\}) \prod_{i=1}^{l} \mathrm{d}Q_i(\nu_i)\,, \tag{22.19}$$

Analogously to the replica-symmetric case, Eq. (22.18), we shall write $Q_0 = \mathsf{F}_0\{Q_i\}$ as a shorthand for this equation. The function F_0 takes as argument l distributions $Q_1, \dots, Q_l$ and evaluates a new distribution Q_0 (each of the Q_i's is a distribution over the $(|\mathcal{X}|-1)$-dimensional simplex).

At this point the formal similarity of Eqs. (22.18) and (22.19) should be clear. The 2RSB cavity equations are obtained by promoting the distributions Q_i to random variables (taking values in the set of distributions over the $|\mathcal{X}|$-dimensional simplex)[33]. Their probability distributions are denoted as $\mathcal{Q}_i$, and the resulting equations depend on one further real parameter $\mathtt{x}_2$. Formally the 2RSB equation can be written as

$$\mathcal{Q}_0(Q_0) = \frac{1}{\mathcal{Z}\{\mathcal{Q}_i\}} \int Z\{Q_i\}^{\mathtt{x}_2/\mathtt{x}_1}\, \delta(Q_0 - \mathsf{F}_0\{Q_i\}) \prod_{i=1}^{l} \mathrm{d}\mathcal{Q}_i(Q_i)\,. \tag{22.20}$$

This equation might look scary, as $\mathcal{Q}_i(\,\cdot\,)$ are distributions over distributions over a compact subset of the reals. It is useful to rewrite it in a mathematically more correct form. This is done by requiring, for any measurable set of distributions $\mathcal{A}$ (see the footnote), the following equality to hold:

$$\mathcal{Q}_0(\mathcal{A}) = \frac{1}{\mathcal{Z}\{\mathcal{Q}_i\}} \int Z\{Q_i\}^{\mathtt{x}_2/\mathtt{x}_1}\, \mathbb{I}(\mathsf{F}_0\{Q_i\} \in \mathcal{A}) \prod_{i=1}^{l} \mathrm{d}\mathcal{Q}_i(Q_i)\,. \tag{22.21}$$

The interpretation of the 2RSB messages $\mathcal{Q}_i$ is obtained by analogy with the 1RSB one. Let α_1 be the index of a particular cloud of states and $Q_i^{\alpha_1}(\,\cdot\,)$ be the distribution of the message ν_i over the lumps in cloud α_1. Then $\mathcal{Q}_i$ is the distribution of $Q_i^{\alpha_1}$ when one picks up a cloud index α_1 randomly (each cloud being sampled with a weight that depends on $\mathtt{x}_1$.)

[33] The mathematically inclined reader might be curious about the precise definition of a probability distribution over the space of distributions. It turns out that given a measure space Ω (in our case the $(|\mathcal{X}|-1)$ dimensional simplex), the set of distribution over Ω can be given a measurable structure that makes 2RSB equations well defined. This is done by using the smallest σ-field under which the mapping $Q \mapsto Q(A)$ is measurable for any $A \subseteq \Omega$ measurable.

In principle Eq. (22.20) can be studied numerically by generalizing the population dynamics approach of Ch. 19. In the present case one can think of two implementations: for one given instance, one can generalize the SP algorithm, but this generalization involves, on each directed edge of the factor graph, a population of populations. If instead one wants to perform a statistical analysis of these messages, seeking a fixed point of the corresponding density evolution, one should use a population of populations of populations! This is obviously challenging from the point of view of computer resources (both memory and time). To the best of our knowledge it has been tried only once, in order to compute the ground state energy of the spin glass on random 5-regular graphs. Because the graph is regular it looks identical at any finite distance from any given point. One can therefore seek a solution such that the $\mathcal{Q}_i$ on all edges are the same, and one is back to the study of populations of populations. The results have been summarized in Table 17.4.5: if one looks at the ground state energy, the 2RSB method provides a small correction of order 10^{-4} to the 1RSB value, and this correction seems to be in agreement with the numerical estimates of the ground state.

22.2.3 *Local stability of the 1RSB phase*

The above discussion of 2RSB will help us to check the stability of the 1RSB phase. The starting point consists in understanding the various ways in which the 2RSB formalism can reduce to the 1RSB one.

The first obvious reduction consists in taking the 2RSB distribution $\mathcal{Q}_i$ to be a Dirac delta at Q_i^*. In other words, for any continuous functional $\mathcal{F}$ on the space of distributions

$$\int \mathcal{F}(Q_i) \;\, \mathrm{d}\mathcal{Q}_i(Q_i) = \mathcal{F}(Q_i^*)\,. \tag{22.22}$$

It is not hard to check that, if $\{Q_i^*\}$ solves the 1RSB equation Eq. (22.19), this choice of $\{\mathcal{Q}_i\}$ solves Eq. (22.20) independently of $\mathtt{x}_2$.

There exists however a second reduction, that corresponds to taking $\mathcal{Q}_i(\,\cdot\,)$ a non-trivial distribution, but supported on Dirac deltas: let us denote by δ_{ν^*} a 1RSB distribution which is a Dirac delta on the message $\nu = \nu^*$. Given a set of messages$\{Q_i^*\}$ that solves the 1RSB equation Eq. (22.19), we construct $\mathcal{Q}_i(\,\cdot\,)$ as a superposition of Dirac deltas over all values of ν^*, each one appearing with a weight $Q_i^*(\nu^*)$. Again this distribution is more precisely defined by its action on a continuous functional $\mathcal{F}(Q)$:

$$\int \mathcal{F}(Q_i) \;\, \mathrm{d}\mathcal{Q}_i(Q_i) = \int \mathcal{F}(\delta_{\nu^*}) \;\, \mathrm{d}Q_i^*(\nu^*)\,. \tag{22.23}$$

Exercise 22.2 Suppose that $\{Q_i^*\}$ solves the analog of the 1RSB equation Eq. (22.19) in which the parameter $\mathtt{x}_1$ has been changed into $\mathtt{x}_2$. Show that $\mathcal{Q}_i$ defined by Eq. (22.23) solves Eq. (22.20) independently of $\mathtt{x}_1$.

[Hint: Show that, when evaluated on Dirac deltas, the normalization Z appearing in (22.19) is related to the normalization z in (22.18) by $Z\{\delta_{\nu_i}\} = (z\{\nu_i\})^{\mathtt{x}_1}$.]

In view of the interpretation of the 2RSB messages $\mathcal{Q}_i$ outlined in the previous section, and cartooned in Fig. 22.2.1, these two reductions correspond to qualitatively different limit situations. In the first case, described by Eq. (22.22), the distribution over clouds becomes degenerate: there is essentially one cloud (by this we mean that the number of clouds is not exponentially large in N: the corresponding complexity vanishes). In the second case, described by Eq. (22.23), it is the distribution within each cloud that trivializes: there is only one cluster (in the same sense as above) in each cloud.

What are the implications of these remarks? Within the 1RSB approach one needs to solve Eq. (22.19) in the space of didtributions over BP messages: let us call this the '1RSB space'. When passing to 2RSB, one seeks a solution of (22.20) within a larger '2RSB space,' namely the space of distributions over distributions over BP messages. Equations (22.22) and (22.23) provide two ways for embedding the 1RSB space inside the 2RSB space.

When one finds a 1RSB solution, one should naturally ask whether there exists a proper 2RSB as well (i.e. a solution outside the 1RSB subspace). If this is not the case, physicists usually conjecture that the 1RSB solution is asymptotically correct (for instance it yields the correct free-energy per spin). This check has been carried out for models on complete graph (e.g. the fully connected p-spin glasses). So far, the difficulty of studying the 2RSB equations have prevented its implementation for sparse factor graph.

Luckily there is a convenient (albeit less ambitious) alternative: check the **local stability of 1RSB** solutions with respect to higher order RSB. Given a 1RSB solution, one looks at it as a point in the 2RSB space according to the two possible embeddings, and one studies the effect of a small perturbation. More precisely, consider the iteration of 2RSB equations (22.20):

$$\mathcal{Q}_{i\to j}^{(t+1)}(Q_0) = \frac{1}{\mathcal{Z}\{\mathcal{Q}_{l\to i}\}} \int Z\{Q_{l\to i}\}^r \, \delta(Q_{i\to j} - \mathsf{F}_i\{Q_{l\to i}\}) \prod_{l\in\partial i\setminus j} \mathrm{d}\mathcal{Q}_{l\to i}^{(t)}(Q_{l\to i}) \, .$$

Given the factor graph G, we initiate this iteration from a point close to the 1RSB solution described by either of the embeddings (22.22) or (22.23) and see if, the iteration converges back to the 1RSB fixed point. This is studied by linearizing the iteration in an appropriate 'perturbation' parameter. If the iteration does not converge to the 1RSB fixed point, the 1RSB solution is said unstable. The instability is named of 'type I' if it occurs when embedding (22.22) is used and named of 'type II' for embedding (22.23).

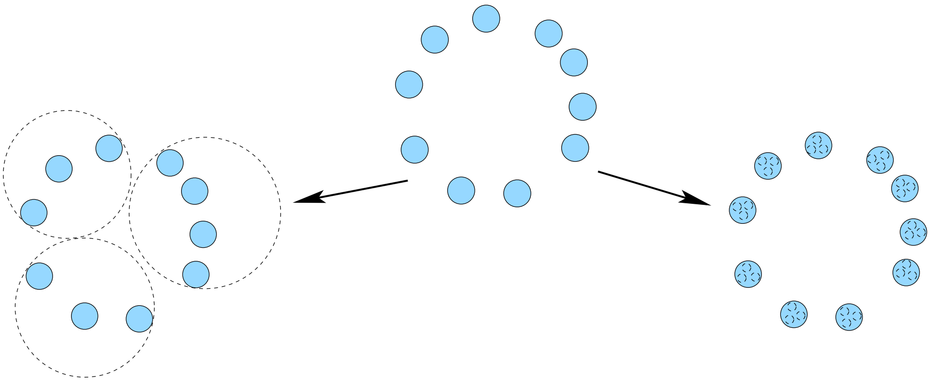

FIG. 22.5. Cartoon of the two types of local instabilities from a 1RSB solution towards 2RSB.

An alternative approach for checking the local stability of a 1RSB solution consists in computing the spin glass susceptibility, which describes the reaction of the model (22.16) to a perturbation that acts on a single variable x_i. As we discussed above, the effect of this perturbation (studied in linear order) remains finite when the spin glass susceptibility $\chi^{(2)}$ is finite. One should therefore compute $\chi^{(2)}$ assuming that the 1RSB solution is correct and check that it is finite. However, the 1RSB picture implies a second condition: each single lump r should also be stable to small perturbations. More precisely, we define $\chi^{\mathrm{SG},r}$ as the spin glass susceptibility with respect to the measure $\mu^r(\,\cdot\,)$ restricted to state r. Denoting by $\langle\,\cdot\,\rangle_r$ the expectation value with respect to μ^r, the 'intra-state' susceptibility, $\chi^{\mathrm{SG,intra}}$, is a weighted average of $\chi^{\mathrm{SG},r}$ over the state r:

$$\chi^{\mathrm{SG,intra}} = \sum_r w_r\, \chi^{\mathrm{SG},r}, \tag{22.24}$$

$$\chi^{\mathrm{SG},r} = \frac{1}{N}\sum_{i,j}\left(\langle x_i x_j\rangle_r - \langle x_i\rangle_r\langle x_j\rangle_r\right)^2 . \tag{22.25}$$

Within the susceptibility approach, the second condition consists in computing $\chi^{\mathrm{SG,intra}}$ with the 1RSB approach and requiring that it stays finite as $N \to \infty$.

It is generally believed that these two approaches to the local stability of the 1RSB phase coincide. Type I stability should be equivalent to $\chi^{(2)}$ being finite; it means that the system is stable with respect to the grouping of states into clusters. Type II stability should be equivalent to $\chi^{\mathrm{SG,intra}}$ being finite; it means that the system is stable towards a splitting of the states into sub-states. A pictorial representation of the nature of the two instabilities in the spirit of Fig. 22.2.1 is shown in Fig. 22.2.3.

The two approaches to stability computations have been developed in several special cases, and are conjectured to coincide in general. Remarkably 1RSB is

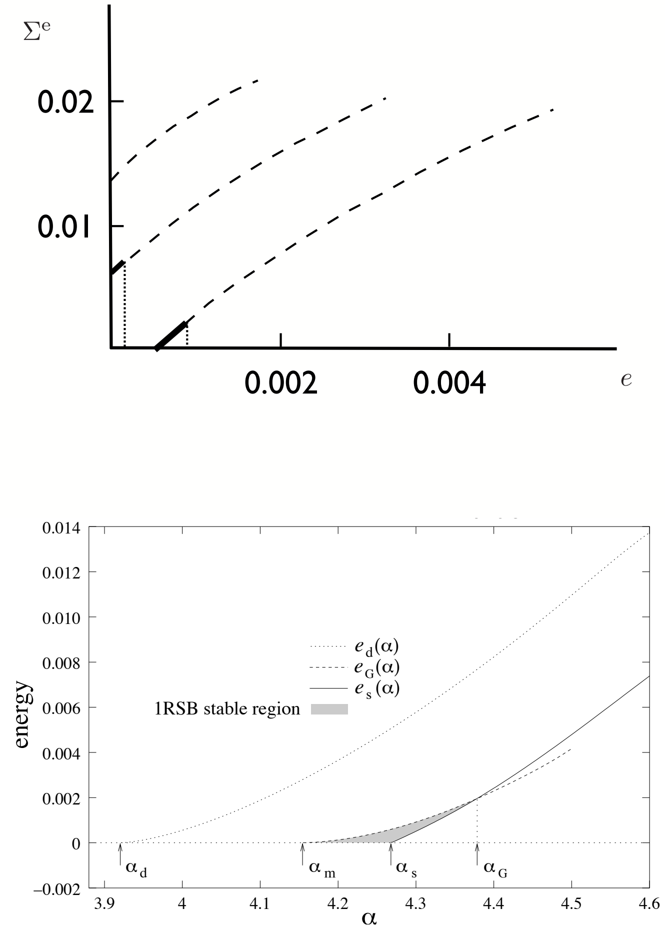

FIG. 22.6. Top: The energetic complexity Σ^e in a random 3-SAT problem, computed within the 1RSB cavity method, is plotted versus the density e of violated clauses, for $\alpha = 4.1$, 4.2, and 4.3 (from top to bottom). The curve reproduces Fig. 20.5, but it now shows the stable and unstable regions. The full thick line, below $e_G(\alpha)$, gives the part of the complexity curve for which the 1RSB computation is locally stable (absent for $\alpha = 4.1 < \alpha_m(3)$, where the full curve is unstable). This is the only part that is computed reliably by 1RSB, the dashed part is unstable. Bottom: In the same random 3-SAT problem, plotted versus the clause density α: the continuous line gives the minimum density of unsatisfied clauses as predicted within 1RSB (this is the value of e where $\Sigma^e(e)$ starts to become positive). The dotted line gives the threshold energy density as predicted within 1RSB (the maximal value of e where $\Sigma^e(e)$ exists). The gray area indicates the region of local stability of the 1RSB stability. The ground state energy density predicted by 1RSB is wrong for $\alpha > \alpha_G$ (although probably very close to the actual value), because in this region there is an instability towards higher order RSB. It is conjectured that the stable region, $\alpha_m < \alpha < \alpha_s$, is in a 1RSB phase: if this conjecture holds the 1RSB prediction α_s for the SAT-UNSAT threshold is correct. For $K = 3$ one has $\alpha_m(3) = 4.153(1)$, $\alpha_s(3) = 4.2667(1)$, $\alpha_G(3) = 4.390(5)$.

unstable in several interesting cases and higher order RSB would be needed to obtain exact predictions.

Stability computations are somewhat involved, and a detailed description is beyond our scope. Nevertheless, we want to give an example of the results that can be obtained through a local stability analysis. Consider random K-SAT formulae, with N variables and $M = N\alpha$ clauses. Let $e_{\rm s}(\alpha)$ denote the minimum number of unsatisfied clauses per variable, in the large system limit. The limit $e_{\rm s}(\alpha)$ can be computed along the lines of Ch. 20 using the 1RSB cavity method: for a given α, one computes the energetic complexity density $\Sigma^{\rm e}(e)$ versus the density of violated clauses e. Then $e_{\rm s}(\alpha)$ is found as the minimal value of u such that $\Sigma^{\rm e}(e) > 0$. It vanishes for $\alpha < \alpha_{\rm s}(K)$ (the SAT-UNSAT threshold) and departs continuously from 0, increasing monotonically for $\alpha > \alpha_{\rm s}(K)$.

The stability computation shows that, for a given α, there is in general an instability of type II which appears above some value $e = e_{\rm G}(\alpha)$: only the part of $\Sigma^{\rm e}(e)$ with $e \le e_{\rm G}(\alpha)$ is in a locally stable 1RSB phase. When $\alpha < \alpha_{\rm m}(K)$, $e_{\rm G}(\alpha) = 0$ and the whole 1RSB computation is unstable. For $\alpha > \alpha_{\rm G}(K)$, $e_{\rm G}(\alpha) < e_{\rm s}(\alpha)$ (the ground state energy density) and again 1RSB is unstable (this implies that the 1RSB prediction for $e_{\rm s}(\alpha)$ is not correct). The conclusion is that the 1RSB calculation is stable only in an interval $]\alpha_{\rm m}(K), \alpha_{\rm G}(K)[$. Figure 22.2.3 summarizes this discussion for 3-SAT. For all values of K, the stable interval $]\alpha_{\rm m}(K), \alpha_{\rm G}(K)[$ contains the SAT-UNSAT threshold $\alpha_{\rm s}(K)$.

The stability check leads to the conjecture that the 1RSB prediction for $\alpha_{\rm s}(K)$ is exact. Let us stress however that stability has been checked only with respect to small perturbations. A much stronger argument would be obtained if one could do the 2RSB computation and show that it has no solution apart from the two 'embedded 1RSB solutions' that we discussed above.

22.2.4 *Open problems within the cavity method*

The main open problem is of course to prove that the 1RSB cavity approach yields correct predictions in some models. This was achieved until now only for a class of models on the complete graph. Here we want to point out a number of open questions that wait for an answer, even at a heuristic level, within the 1RSB cavity method itself.

Distributional equations. Cavity predictions are expressed in terms of fixed point of equations of the form (22.19). When considering models on ensembles of random graphs, this can be read as an equation for the probability distribution of $Q_0(\,\cdot\,)$ (that is taken identical to the one of $Q_1(\,\cdot\,)$,..., $Q_k(\,\cdot\,)$.)

Currently such equations are mostly studied using the population dynamics method of Sec. 14.6.4. The main alternative explored so far has been to formally expand the equations for large degrees. Population dynamics is powerful and versatile. However in many cases, this approach is too coarse, particularly as soon as one wants to study k-RSB with $k \ge 2$. It is intrinsically hampered by statistical errors, that are of the order of the inverse square root of population size. In some models (for instance, in graph ensembles with large but bounded average degree), statistical fluctuations are too large for the population sizes that can be implemented on ordinary PCs (typically $10^7 \div 10^8$ elements). This limits the possibility to distinguish, for instance, 2RSB from 1RSB effects, because high precision is generally required to see the difference. Furthermore, metastability is the crux (and the limit) of the whole population dynamics approach. Therefore it would be interesting to make progress in two directions:

- Analytical tools and generic results on the cavity equations; this could provide important guiding principles for any numerical study.
- New efficient and stable numerical methods.

A step forward has been made by the reconstruction algorithm discussed in Theorem 19.5, but unfortunately it is limited to one value of the rescaling parameter, $\mathtt{x} = 1$.

Local stability. Local stability criteria provide an important guidance in heuristic studies. It would be important to put these results on firmer grounds. Two specific tasks could be, for instance:

- Prove that, if all 1RSB solutions of the cavity equations are locally unstable, then there must exist a 2RSB solution outside the 1RSB subspace.
- Prove that, if a solution of the cavity equations is locally unstable, it does not describe correctly the model.

Occurrence of k-RSB. A number of random graphical models have been studied within the cavity (or replica) method. In most cases, one finds that the system is either RS, or 1RSB, or FRSB. The cases in which a 2RSB phase is found are rare, and they always involve some kind of special construction of the compatibility function (for instance, a fully connected model which is a superposition of two p-spin glass interactions, with $p_1 = 3$ and $p_2 = 16$ displays 2RSB). Therefore one should

- Find a ‘natural’ model for which 2RSB is asymptotically exact, or understand why this is impossible.

Full replica-symmetry breaking. We saw that k-RSB provides, as k increases, a sequence of ‘nested’ schemes that aim at computing various quantities like local marginals, free-entropy density, etc..., in the large system limit. A k-th order scheme includes all the lower l-RSB schemes with $l < k$ as nested subspaces of the set of feasible solutions to the cavity equations. On the other hand, as the number of steps increases, the description of the set of feasible solutions becomes more and more complicated (distributions of distributions of...).

Surprisingly, in the case of fully connected models, there exists a compact description of the space of feasible solutions in the FRSB limit $k \to \infty$. An outstanding problem is to find an analogous description in the case of models on sparse graphs. This would allow to look for the best solution in the k-RSB space for all k.

- Find a description of the space of full replica-symmetry breaking messages for models on sparse graphs.

Variational aspect. It is widely believed that if one finds a consistent solution of the cavity k-RSB equations, the free-energy density computed with this solution, is always a lower bound to the correct free energy density of the model (in particular the k-RSB ground state energy density prediction is a lower bound to the true one). This should hold for a large class of models with a statistical $+1/-1$ symmetry. While this has been proven in some specific cases, one would like to:

- Find a general proof that the free-energy computed with the cavity method is a lower bound to the correct free-energy of the model.

22.3 Phase structure and the behavior of algorithms

A good part of this book has been devoted to the connection between the various phases in random graphical models, and the behavior of algorithms. There exists by now substantial evidence (empirical, heuristic, and, in some cases, rigorous) that such a connection exists. For instance, we have seen on the example of codes in Ch.21 how the appearance of a 1RSB phase, and the corresponding proliferation of metastable states, determines the noise threshold where BP decoding fails. Developing a broader understanding of this connection, and determining the class of algorithms to which it applies, is a very important problem.

We propose here a list of broad research problems, whose advancement will probably help to clarify this issue. We always have in mind a graphical model of the form (22.1), with a locally tree-like factor graph.

Impact of the dynamical transition on Monte Carlo dynamics. Consider the problem of sampling from the distribution (22.1) using a Monte Carlo Markov Chain (MCMC) algorithm. The Markov chain is assumed to flip a sub-linear ($o(N)$) number of variables at each step, and to satisfy detailed balance with respect to the probability distribution $\mu(\,\cdot\,)$.

One expects that, if the system is in a 1RSB phase, the relaxation time of this algorithm will increase rapidly (probably exponentially) with system size. Intuitive arguments in favor of this statement can be obtained from each of the two characterizations of the 1RSB phases introduced in Sec. 22.1. The argument is different whether we start from the pure state decomposition, or from the characterization in terms of correlations. In the first case, the relaxation time is estimated through the time to cross a bottleneck, see also Ch. 13. In the second case, one can define a correlation length ℓ_i^* through the point-to-set correlation function $G_i(\ell)$, cf. Eq. (22.14). In order for the system to relax, information has to travel a distance ℓ_i^*. But if ℓ_i^* diverges with size, so must the relaxation time.

This picture is intuitively satisfying, but it is far from being proved, and should be formulated more precisely. For instance it often happens that in RS phases there exist small isolated metastable states that make the relaxation time (the inverse spectral gap of the MCMC) formally large. But even in such cases, numerical simulations indicate that Glauber dynamics equilibrates rapidly within the RS phase. This observation is probably related to the fact that the initial condition is chosen uniformly random, and that equilibration is only checked on local observables. A number of questions arise:

- Why is metastability irrelevant 'in practice' in a RS phase? Is it because of local measurements? Or because of the uniform initial condition? If the latter is true, what is so special about the uniform initial condition?
- Within a RS phase, can one approximate partition functions efficiently?

Message passing and the estimation of marginals.

For a number of models on sparse random graphs within the RS and (sometimes) dynamical 1RSB phases, message passing methods like belief propagation or survey propagation show, empirically, good performances. More precisely, they return good approximations of local expectation values if initialized from uniform messages.

Current rigorous techniques for analyzing BP often aim at proving that it is accurate regardless of the initialization. As a consequence, results are dominated by the behavior under *worst case* initializations that are not used in practice. As an illustration, consider applying BP to the uniform measure over solutions of a random K-SAT formula. The analysis under worst case initialization allows to prove that BP is accurate only for $\alpha \leq (2\log K)/K[1+o(1)]$. This threshold is embarrassingly small when compared to the dynamical transition point that terminates the RS phase $\alpha_{\rm d}(K) = 2^K \log K/K[1+o(1)]$.

In general we have no good mathematical control of when BP or SP converge or/and give good approximations of marginals. Empirically it seems that SP is able to converge in some regions of 1RSB phases where BP does not. We have no real understanding of this fact beyond the hand-waving argument that 1RSB correctly captures the structure of correlations in these phases.

Here are a number of open questions on these issues:

- Why are BP/SP performances on random instances, with uniformly random initialization, much better than in the worst case? What is special about the uniform initialization? What are the features of random instances that make them easier? Can these features be characterized and checked efficiently?
- Under what conditions do the BP (or the SP) algorithms converge and give good approximations to local marginals? When their naive iteration does not converge, can one systematically either force convergence or use time averages of the messages?
- It seems that, on sparse random graphical models, BP or SP outperforms local MCMC algorithms. In particular these message passing algorithms can have (at least in principle), good performances within the dynamical 1RSB phase. Can one demonstrate this possibility convincingly in some model?

Message passing algorithms and optimization.

If one seeks a solution to a random constraint satisfaction problem using message passing, the main approach so far has been the use of decimation: one first computes all local marginals, then decides, based on this knowledge, how to fix a variable, and then iterate the procedure. In general this procedure converges when the number of constraints per variable is not too large, but it fails above a critical value of this number, which is strictly smaller than the SAT-UNSAT threshold. No one knows how to determine analytically this threshold.

An alternative to decimation is the reinforcement method: instead of fixing a variable based on the knowledge of local marginals, it modifies some local factors applying to each individual variables, based on this same information. So far, optimizing this modification is an art, and its critical threshold cannot be estimated either.

- How to *predict* the performances of BP+ decimation or SP+decimation. For instance, empirically these methods find solutions to random K-SAT formulae with high probability for $\alpha < \alpha_{\mathrm{BP}}(K)$ (or $\alpha < \alpha_{\mathrm{SP}}(K)$), but we have no prediction for these algorithmic thresholds. In what class of problems is SP better than BP?
- Similar questions for BP+reinforcement or SP+reinforcement.
- Find new ways to use the local marginal information found by message passing in order to exhibit solutions.
- In an UNSAT phase, the message passing procedure is able to give an estimate of the minimal number of violated constraints. Is it possible to use this information, and the one contained in the messages, in order to prove unsatisfiability for one given instance?

The above questions focus on sparse random instances. Message passing techniques have been (partially) understood and sharpened for this type of instances. They naturally arise in a large class of applications where the graphical model is random, or pseudo-random, *by design.* The theory of sparse graph codes is a clear example in this direction. In the limit of large block-lengths, random constructions proved to be generally superior to deterministic ones. More recently sparse graph constructions have been proposed for data compression (both lossless and lossy), online network measurements, multi-terminal communications, distributed storage, group testing, etc…

On the other hand, being able to deal with structured graphs would open an even much broader class of applications. When applied to structured problems, message passing algorithms often fail to converge. This is typically the reason why the decimation method may fail, even when the marginals of the original problem are well estimated by message passing: the instance found after fixing many variables is no longer random. Finding appropriate modifications of message passing for structured graphs would therefore be very interesting.

- How to use message passing in order to improve the solution of some general classes of (non-random) constraint satisfaction problems. Can it be coupled efficiently to other general methods (such as MCMC)?

Notes

The present chapter was inevitably elliptic. We will provide a few pointers to recent research without any ambition to be comprehensive.

The connection between correlation lengths and phase transitions is a classical topic in statistical mechanics which has been recently revived by the interest in the glass transition. A good starting point for learning about this subject in the context of glasses is the paper (Bouchaud and Biroli, 2004) which describes the 'freezing' thought experiment in Sec. 22.1.2.

The description of point-to-set correlations in terms of 'reconstruction' problems is taken from (Evans, Kenyon, Peres and Schulman, 2000). This paper studies the reconstruction phase transition for Ising models on trees. Results for a wide class of models on trees are surveyed in (Mossel and Peres, 2003; Mossel, 2004). We also refer to (Gerschenfeld and Montanari, 2007) for the generalization to non-tree graphs. The connection between 'reconstruction' and 'dynamical' 1RSB phase transition was first pointed out in (Mézard and Montanari, 2006). The implications of this phase transition on dynamics were explored in (Berger, Kenyon, Mossel and Peres, 2005; Martinelli, Sinclair and Weitz, 2004; Montanari and Semerjian, 2006*b*). The definition of pure states presented in this chapter as well as the location of the dynamical and condensation phase transitions for random K-SAT and coloring of random graphs are from (Krzakala, Montanari, Ricci-Tersenghi, Semerjian and Zdeborova, 2007).

The GREM has been introduced by (Derrida, 1985) and studied in details in (Derrida and Gardner, 1986). A 2RSB phase in fully connected models has been found by (Crisanti and Leuzzi, 2007). There are very few results about higher order RSB in models on sparse random graphs. For spin glasses, one can use perturbative expansions close to the critical point (Viana and Bray, 1985), or for large degrees (Goldschmidt and Dominicis, 1990). The 2RSB computation of ground state energy for spin glasses mentioned in Sec. 22.2 is from (Montanari, 2003). The method for verifying the local stability of the 1RSB solution in sparse systems was first devised in (Montanari and Ricci-Tersenghi, 2003), and applied to random satisfiability problems in (Montanari, Parisi and Ricci-Tersenghi, 2004). A complete list of stability thresholds, including their asymptotic behavior, for random K-SAT can be found in (Mertens, Mézard and Zecchina, 2006). The interpretation of 1RSB instability in terms of susceptibilities is discussed in (Rivoire, Biroli, Martin and Mézard, 2003).

The fact that the free-energy computed with the cavity (or replica) method is a lower bound to the true one can be proven in some fully connected models using the inequalities of (Guerra, 2003). The same strategy also yields rigorous bounds in some diluted systems (Franz and Leone, 2003; Franz, Leone and Toninelli, 2003; Panchenko and Talagrand, 2004) but it still relies on some details of the structure of the models, and a general proof applicable to all cases is lacking.

The reinforcement algorithm has been introduced and discussed for SAT in (Chavas, Furtlehner, Mézard and Zecchina, 2005).

There exist only scarce results on the algorithmic consequences of the structure of the solution space. Some recent analyses can be found in (Altarelli, Monasson and Zamponi, 2007; Montanari, Ricci-Tersenghi and Semerjian, 2007; Ardelius and Aurell, 2006; Alava, Ardelius, Aurell, Kaski, Krishnamurthy, Orponen and Seitz, 2007). The convergence and correctness of BP for random K-satisfiability at small enough α was proven in (Montanari and Shah, 2007).

This book covered only a small subsets of problems that lie at the intersection between information theory, computer science and statistical physics. It would be difficult to provide an exhaustive list of references on the topics we did not touch: we will limit ourselves to a few 'access points'.

As we mentioned, channel coding is only one of the fundamental problems addressed by information theory. Data compression, in particular in its 'lossy' version, is a key component in many modern technologies, and presents a number of open problems (Ciliberti, Mézard and Zecchina, 2005; Wainwright and Maneva, 2005). Some other statistics problems like group testing are similar in spirit to data compression (Mézard, Tarzia and Toninelli, 2007).

Modern wireless and wireline communication systems are intrisically multiuser systems. Finding optimal coding schemes in a multiuser context is a widely open subject of great practical interest. Even the information theoretic capacity of such systems is unknown. Two fields that benefited from tools or analogies with statistical mechanics are multiuser detection (Tanaka, 2002; Guo and Verdú, 2002) and networking (Kelly, 1991). Always within a communications context, a large effort has been devoted to characterizing large communication networks such as the Internet. A useful review is provided by (Kleinberg, Kumar, Raghavan, Rajagopalan and Tomkins, 1999).

Statistical mechanics concepts have been applied to the analysis of fluctuations in financial markets (Bouchaud and Potters, 2003) or to model interactions among economic agents (Challet, Marsili and Zhang, 2005). Finally, biology presents a number of problems in which randomness, interaction between different components, and robustness play important roles. Stochastic models on networks, and inference algorithms have been studied in a number of contexts, from neural networks (Baldassi, Braunstein, Brunel and Zecchina, 2007; Coolen, Kuehn and Sollich, 2005) to phylogeny (Mossel, 2003), to gene expression (Friedman, Linial, Nachman and Peér, 2000).

A few of these topics, and others, are reviewed in the recent school proceedings (Bouchaud, Mézard and Dalibard, 2007).