Lecture 12: The Cavity Method: introduction to a few basic concepts.

1. Intro: application of the Bethe free energy to counting solutions of K-SAT.

Recall that for K-SAT the partition function is

$$Z = \sum_{\vec{x}} \prod_{a} f_a(x_{\partial a})$$

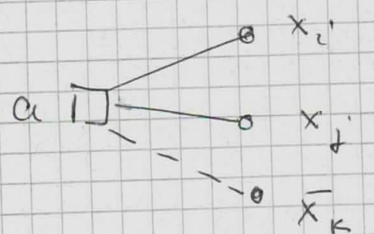

$$f_a(x_{\partial a}) = \mathbb{1}(x_i \vee x_j \vee \bar{x}_k).$$

Check degrees $= K$ (say $K=3$).

Node degrees are assymptotically Poisson(αK).

$$\alpha = \frac{\text{number of clauses}}{\text{number of variables}}.$$

The partition function counts the number of solutions of a given instance. One expects in the SAT phase an exponential number of solutions so we define (in SAT phase)

$$f = \lim_{N \to +\infty} \frac{1}{N} \mathbb{E}(\ln Z).$$

This is an average "entropy" where $\mathbb{E}$ is over the graph ensemble. We can try to compute this quantity thanks to the (Bethe) replica symmetric formula (as in coding). This is the object of homework 11. One finds

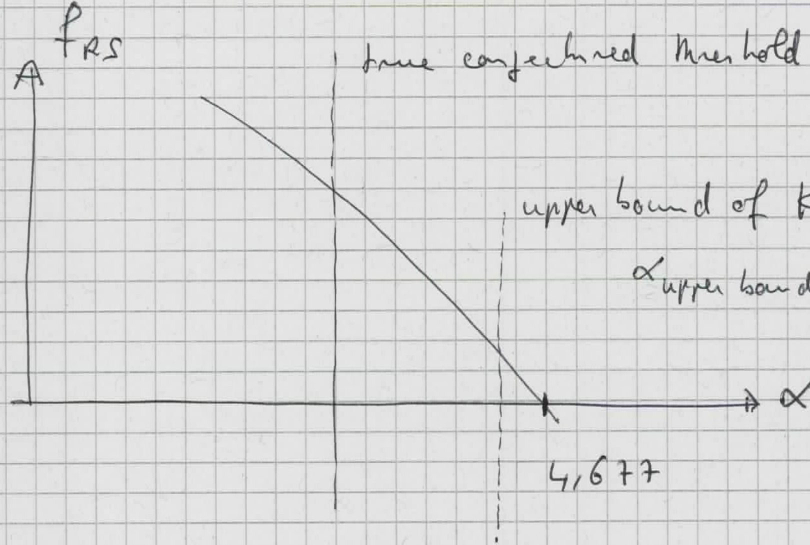

Because of the upper bound to the true threshold we see that the RS prediction is wrong at least for α "large". In fact as we will see later it is believe that the RS predicton is exact for $\alpha < \alpha_{BI} \simeq 3,86$.

Experience with other spin glass problems, and principaly with 20 years of spin-glass research, suggests that the RS formula fails because as α increases the system enters into a phase where many "pure" or "extremal" Gibbs states coexist. In order to understand this statement we have to come back to a few fundamental notions of statistical mechanics.

But before, let us see what is the general picture of the phase diagram for K-SAT that will emerge from the theory developped in this lecture and the next one.

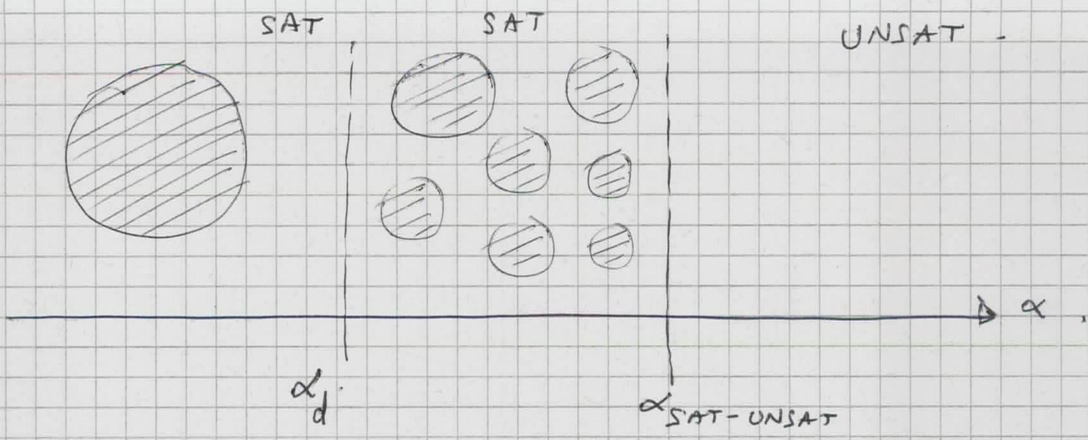

Solution space forms one "cluster". All pairs of solutions can be connected by a path through soln space where one goes from one solution to the next one by flipping only a finite nb of variables.

Solution space becomes fragmented or "clustered". Two solutions in different cluster can be connected only by flipping a macroscopic number $O(N)$ of variables. The number of such clusters is exponential in N.

The cavity method predicts such a picture at finite and zero temperature. At finite temperature clusters correspond to different "extremal measures" capturing all possible limits of the Gibbs distribution in the infinite system size limit.

2. Notion of pure / extremal state.

This goes back to fundamental work of Dobrushin – Lanford and Ruelle (~ in 1965–70). Here we will be content with a heuristic physical picture.

Recall the phase diagram of the Ising model (see Lect 4). (or if you wish that of the CW model).

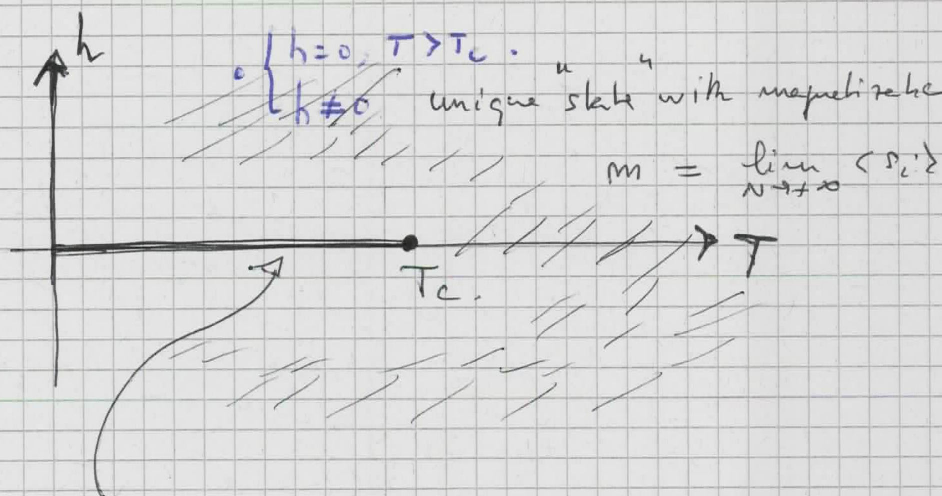

- For $h=0$ & $T<T_c$, Coexistence of two "states" with distinct magnetizations

$$m_{\pm} = \lim_{h\to 0_{\pm}} \lim_{N\to+\infty} \langle s_i\rangle(h).$$

In principle one could compute all possible marginals of $\mu^{(N)}_{Gibbs}(\vec{s})$. Then take the limit $N\to+\infty$. From the marginals one could in principle reconstruct the infinite volume Gibbs state – μ.

For $h\neq 0$ and $h=0$, $T>T_c$: unique inf vol Gibbs state. $\mu_{h,T}$

For $h=0$, $T<T_c$: at least two inf vol Gibbs states. $\mu_{\pm,T}$.

Typical configurations of measure μ_+ have fraction $\frac{1+m_+}{2}$ of ↑ spins & fraction $\frac{1-m_+}{2}$ of ↓ spins.

Typical confs of measure μ_- have fraction $\frac{1+m_-}{2}$ of ↑ spins & fraction $\frac{1-m_-}{2}$ of ↓ spins.

These two state are the extremal (pure) states. One can form other states that are convex linear superpositions (mixed states)

$$\mu_W = W\mu_+ + (1-W)\mu_- ; \quad 0<W<1 \quad ..$$

The general theory of Gibbs measures shows that infinite volume Gibbs states form a convex set. Extremal points are called extremal or pure states. Non-trivial convex superpositions are called mixed states.

2.1. Correlations in pure versus mixed states.

For pure states one can show that

$$\langle S_i S_j \rangle_{pure} - \langle S_i \rangle_{pure} \langle S_j \rangle_{pure} \rightarrow 0 \quad \text{as } |i-j| \rightarrow +\infty$$

usually exponentially fast. This is called decay of correlations (or clustering not to be confused with the clusters of paragraph 1). Here is one ~~..~~ two intuitions for this behavior:

- At very high temperatures one would expect that degrees of freedom become independent. Indeed the Gibbs distribution

$e^{-\beta \mathcal{H}(\vec{s})} \sim$ uniform distr as $\beta \to 0$.

At high temperature there is a unique state which is thus pure and

$$\langle s_i s_j \rangle_{pure} - \langle s_i \rangle \langle s_j \rangle_{pure} \sim e^{-|i-j|/\xi(\beta)}$$

As long as there is no phase transition $\xi(\beta)$ the correlation length remains finite.

- Consider now the situation $h = 0$, $T < T_c$ where two pure states coexist. In the limit $T \to 0$ the measure μ_+ is concentrated on ~~all~~ nearly all $\uparrow$ spins thus:

$$\langle s_i s_j \rangle_+ - \langle s_i \rangle_+ \langle s_j \rangle_+ \approx \langle 1 \rangle - \langle 1 \rangle \langle 1 \rangle = 0.$$

Similarly the measure μ_- is concentrated on nearly all $\downarrow$ spins.
Thus $\langle s_i s_j \rangle_- - \langle s_i \rangle_- \langle s_j \rangle_- \approx 1 - 1 = 0.$

<u>For mixed states correlations do not decay.</u> Indeed consider the mixed state

$$\langle - \rangle_{\frac{1}{2}} = \frac{1}{2} \langle - \rangle_+ + \frac{1}{2} \langle - \rangle_- .$$

Then
$$\begin{aligned}
\langle s_i s_j \rangle_{1/2} &= \frac{1}{2} \langle s_i s_j \rangle_+ + \frac{1}{2} \langle s_i s_j \rangle_- \\
&\approx \frac{1}{2} \langle s_i \rangle_+ \langle s_j \rangle_+ + \frac{1}{2} \langle s_i \rangle_- \langle s_j \rangle_- \quad \text{as } |i-j| \to +\infty \\
&\approx \frac{1}{2} m_+^2 + \frac{1}{2} m_-^2 = m_\pm^2 \quad \text{(by symmetry)}.
\end{aligned}$$

On the other hand $\langle S_i \rangle_{1/2} = \frac{1}{2} \langle S_i \rangle_+ + \frac{1}{2} \langle S_i \rangle_- = 0$

$\Rightarrow \quad \langle S_i \rangle_{1/2} \langle S_j \rangle_{1/2} = 0 .$

Thus $\quad \langle S_i S_j \rangle_{1/2} - \langle S_i \rangle_{1/2} \langle S_j \rangle_{1/2} \longrightarrow m^2$

as $|i-j| \to +\infty$.

One says that correlations (in a mixed state) have long range. Long range correlations are a signal that many pure states coexist.

Remark: often pure states are related by a symmetry group. In the Ising model this symmetry is $S_i \to -S_i$ and is obvious. Moreover it is manifest in the Hamiltonian. When the system selects a pure state we say that it breaks a symmetry of the Hamiltonian. Sometimes (even very often) it is not obvious what is the symmetry and this symmetry is not at all manifest in the microscopic Hamiltonian. This is the case with the replica symmetry in spin-glasses. There the symmetry group (if there really is one) is a strange object (not yet defined by mathematicians) that permutes 0 number of [illegible] items.

3. Glassy behavior and complex landscape.

In the Ising like models there are a small number of pure states. The theory is mathematically well defined there (see for example book of Georgii).

However there is no physical principle that limits a priori the number of states that a system can have. It is believed that "glassy behavior" (wether the word refers to real glass or to models — that do not necessarily model real glass — such as the SK model) is associated to the occurrence of exponentially many states (in N). For such systems / models it is not really known how to define precisely "pure states". The following heuristic picture is often useful.

Recall the CW model. The free energy function at $h=0$, $K>1$.

$$\phi(m) = -\frac{K}{2}m^2 - H_2\left(\frac{1+m}{2}, \frac{1-m}{2}\right)$$

displays two minima that define two "pure states"

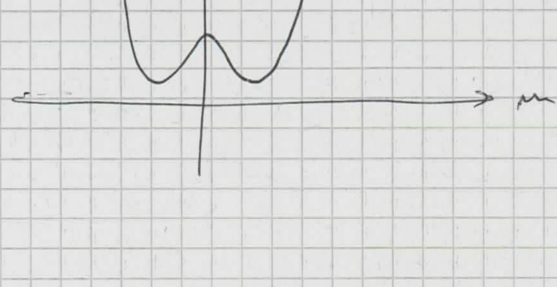

Consider now the Bethe free energy of some spin system (on a sparse graph say).

$$F\left[\{\tanh h_{i\to a},\ \tanh \hat{h}_{a\to i}\}\right] \text{ or } F\left[\nu_{i\to a}(x_i);\ \hat{\nu}_{a\to i}(x_i)\right].$$

The (near global) minima of F in the space $\{\tanh h_{i\to a},\ \tanh \hat{h}_{a\to i}\}$ qualify for what we will call "pure states". Glassy behavior is associated to the fact that the landscape of F can be very complex, displaying an exponentially large number of minima (as a function of N).

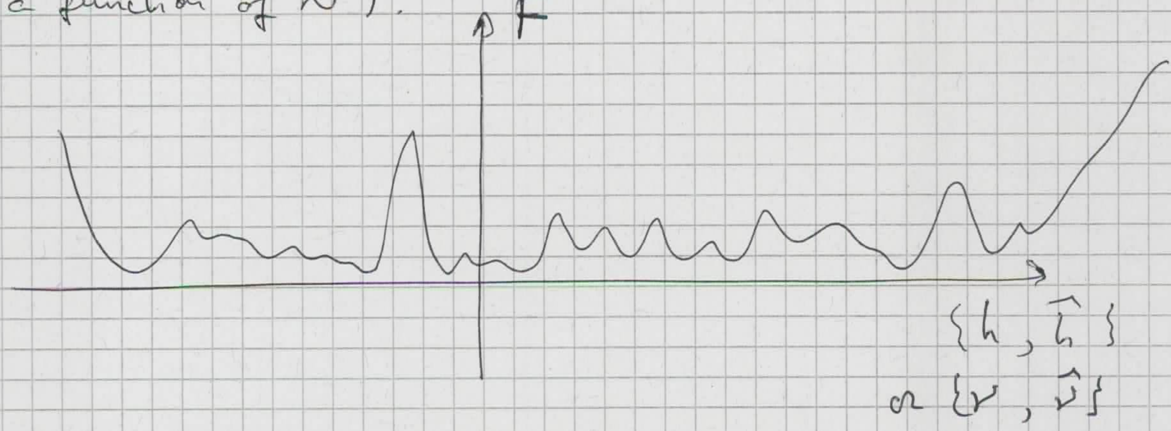

Let us briefly justify that the <u>free energy of a pure state should be given by the Bethe formula</u>. Take a sparse graph and select a node o

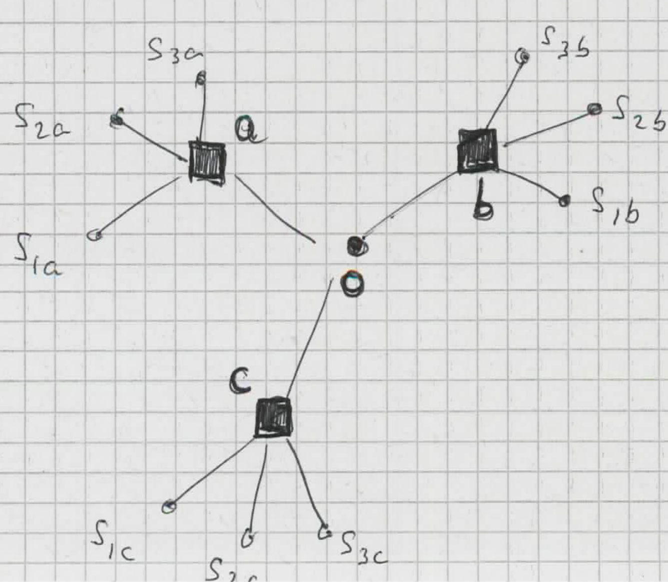

The marginal $\mu_0(s_0)$ satisfies by definition

$$\mu_0(s_0) \propto \sum f_a(s_{\partial a})\, f_b(s_{\partial b})\, f_c(s_{\partial c}) \underbrace{\sum_{\substack{\sim s_{\partial a \setminus 0} \\ s_{\partial b \setminus 0} \\ s_{\partial c \setminus 0}}} \prod_{d \neq a,b,c} f_d(s_{\partial d})}_{\mu_{cavity}(s_{\partial a \setminus 0},\, s_{\partial b \setminus 0},\, s_{\partial c \setminus 0})} .$$

Here μ_{cavity} is the marginal for the system with the "cavity" i.e

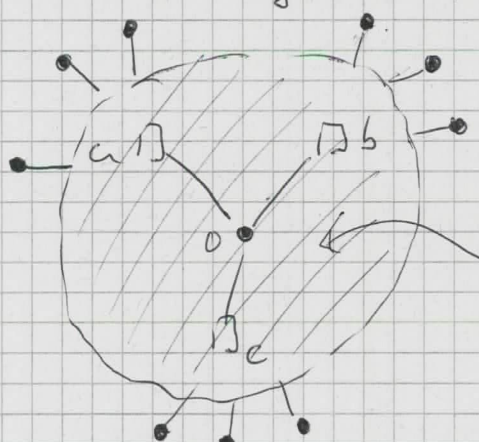

is removed. The Bethe approximation assumes that correlations decay in this system so that

$$\mu_{cavity}(s_{\partial a \setminus 0}, s_{\partial b \setminus 0}, s_{\partial c \setminus 0}) \approx \mu_{cav}(s_{\partial a \setminus 0})\, \mu_{cav}(s_{\partial b \setminus 0})\, \mu_{cav}(s_{\partial c \setminus 0})$$

We see that the validity of BP message passing and Bethe free energy hinges on the decay of correlations. Thus the Bethe free energy describes pure states.

4. "Level 1" Statistical mechanics.

The main idea is to consider the ensemble of pure states as a new statistical mechanical system. The new set of degrees of freedom is the space $\{h_{i\to a}, \hat{h}_{a\to i}\}$ and the have an "energy function" $F[\{h_{i\to a}, \hat{h}_{a\to i}\}]$.

For this system we can define a "Boltzman entropy", called the "complexity" ; (see lecture 2)

$$\Sigma(f) = \frac{1}{N} \ln \sum_{\substack{\{h_{i\to a}, \hat{h}_{a\to i}\} \\ \text{minima}}} \mathbb{1}\left(Nf - F[h_{i\to a}, \hat{h}_{a\to i}]\right).$$

and a "free energy" now called "level 1 RSB free energy",

$$\Phi(x) = \frac{1}{x} \ln \Xi(x) \quad \text{with}$$

$$\Xi(x) = \sum_{\substack{\{v_{i\to a}, \hat{v}_{a\to i}\} \\ \text{minima}}} e^{-xF[\{v_{i\to a}, \hat{v}_{a\to i}\}]}.$$

Here x is a kind of temperature called also "Parisi parameter". Recalling the link between temperature and Boltzmann entropy ($\beta = \frac{\partial S}{\partial E}$) it is not surprising that

$$x = \frac{\partial \Sigma}{\partial f}.$$

The complexity gives the number (log of number) of minima in the landscape at level Nf. The Parisi parameter gives the rate of growth of the number of minima as f varies.

Recall that $\Sigma(f)$ and $\Phi(x)$ are related by a Legendre transform (see lecture 2 or 3). Indeed

$$\Xi(x) = \sum_f e^{N\Sigma(f)} \cdot e^{-xNf}$$

$$\approx e^{N \max_f (\Sigma(f) - xf)}.$$

$$\Rightarrow \quad \boxed{x\,\Phi(x) = \max_f \left(\Sigma(f) - xf\right).}$$

If one wants to compute $\Sigma(f)$, one first computes $x\,\Phi(x)$ by the methods of statistical mechanics and then inverts the above Legendre transform.

5. Level 1 Cavity Method or 1 RSB solution.

In this paragraph we show that $\Xi(x)$ is

the partition function of a new (auxiliairy) sparse graph model.

The constraint that we sum over cavity messages (or fields)

$$\{\nu_{i\to a}\ ,\ \hat{\nu}_{a\to i}\}$$

that are minima of F_{Bethe} is enforced through the indicator

functions enforcing BP fixed point equations :

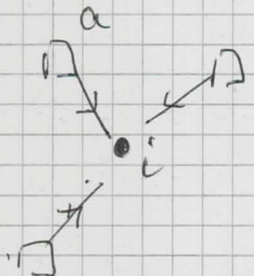

$$\prod_{a\in\partial i} \mathbb{1}\left(\hat{\nu}_{a\to i} \propto \prod_{j\in\partial a\setminus i} \nu_{j\to a}\right)$$

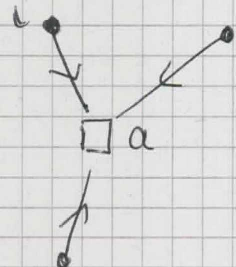

$$\prod_{i\in\partial a} \mathbb{1}\left(\nu_{i\to a} \overbrace{\propto \prod_{b\in\partial i\setminus a}}^{\sum_{\sim x_i} f_a(x_{\partial a})} \hat{\nu}_{b\to i}\right).$$

Recall also that

$$F_{Bethe} = \sum_a F_a[\{\nu_{i\to a}\}] + \sum_i F_i[\{\hat{\nu}_{a\to i}\}] - \sum_{ia} F_{ai}[\nu_{i\to a}, \hat{\nu}_{a\to i}]$$

Thus we have

$$\Xi(x) = \sum_{\vec{h},\,\vec{\hat{h}}} \left(\prod_a \psi_a\right)\left(\prod_i \psi_i\right)\left(\prod_{ia} \psi_{ai}\right).$$

$$\psi_a = e^{-xF_a[\{\nu_{i\to a}\}]} \cdot \prod_{i\in\partial a} \mathbb{I}\left(\nu_{i\to a} \propto \prod_{b\in\partial i\setminus a} \hat{\nu}_{b\to i}\right)$$

$$\psi_i = e^{-xF_i[\{\hat{\nu}_{a\to i}\}]} \cdot \prod_{a\in\partial i} \mathbb{I}\left(\hat{\nu}_{a\to i} \propto \underbrace{\cdots\cdots}_{\text{(see page 13)}}\right).$$

$$\psi_{ai} = e^{-xF_{ai}[\nu_{i\to a},\,\hat{\nu}_{a\to i}]}.$$

The auxiliary factor graph can be pictured as follows.

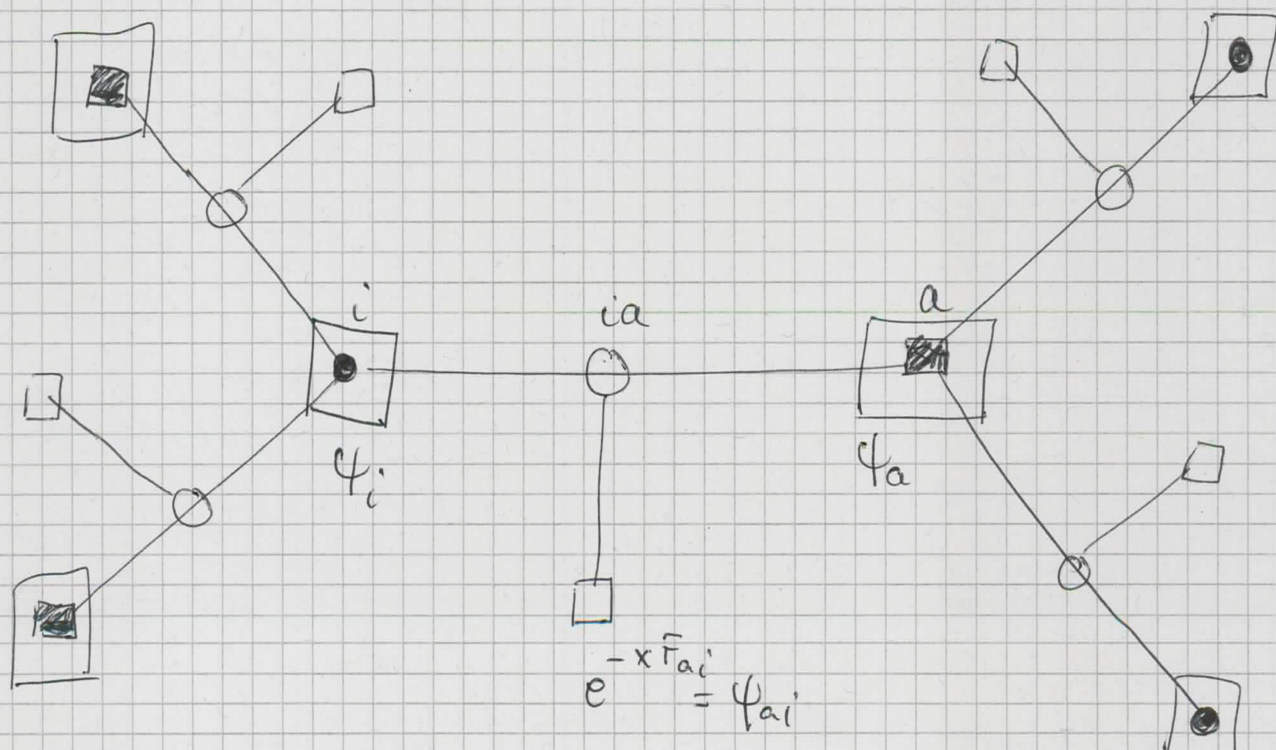

black ● & ■ : original factor graph.

blank ○ & □ : new factor graph.

The new factor graph is again sparse and the new partition function has again a factorized form. Thus we apply again the Bethe method, assuming that correlations of this new system decay. The 1RSB solution is nothing else than the new Bethe functional.

Note that variable nodes are now degree two nodes, thus they can be eliminated and one can work on the induced graph. But this is basically the original factor graph!.

(This means that everything can be expressed in terms of messages $Q_{ia\to a}$ and $Q_{ia\to i}$. We set $Q_{ia\to a} \equiv Q_{i\to a}$ and $Q_{ia\to i} = \hat{Q}_{a\to i}$).

The new "BP equations" called "survey propagation equations" are finally

$$Q_{i\to a}(\nu_{i\to a}) \propto \sum_{\{\hat{\nu}_{b\to i}\}_{b\neq a}} \mathbb{1}\Big(\nu_{i\to a} \propto \prod_{b\in\partial i\setminus a} \hat{\nu}_{b\to i}\Big)\, e^{-x(F_{i}-F_{ai})} \cdot \prod_{b\in\partial i\setminus a} \hat{Q}_{b\to i}(\hat{\nu}_{b\to i}).$$

$$\hat{Q}_{a\to i}(\hat{\nu}_{a\to i}) \propto \sum_{\{\nu_{j\to a}\}_{j\neq i}} \mathbb{1}\big(\hat{\nu}_{a\to i} \propto \dots\dots\big)\, e^{-x(F_{a}-F_{ai})} \cdot \prod_{j\in\partial a\setminus i} Q_{j\to a}(\nu_{j\to a}).$$

(see page 13).

One can also work out the Bethe expression for

$$x\,\phi_{\text{Bethe}}(x) = \sum_a F_a^{RSB} + \sum_i F_i^{RSB} - \sum_{ia} F_{ia}^{RSB}.$$

with

$$F_a^{RSB} = \log\Big\{ \sum_{\{\nu_{i\to a}\}_{i\in\partial a}} e^{x F_a(\{\nu_{i\to a}\})} \prod_{i\in\partial a} Q_{i\to a}(\nu_{i\to a}) \Big\}$$

$$F_i^{RSB} = \log\Big\{ \sum_{\{\hat\nu_{a\to i}\}_{a\in\partial i}} e^{x F_i(\{\hat\nu_{a\to i}\})} \prod_{a\in\partial i} \hat Q_{a\to i}(\hat\nu_{a\to i}) \Big\},$$

$$F_{ai}^{RSB} = \log\Big\{ \sum_{\nu_{i\to a},\hat\nu_{a\to i}} e^{x F_{ai}(\nu_{i\to a},\hat\nu_{a\to i})} Q_{ia}(\nu_{i\to a})\,\hat Q_{a\to i}(\hat\nu_{a\to i}) \Big\}.$$

Interprétation:

$Q_{i\to a}(\nu_{i\to a})$ = fraction of pure states that have cavity messages $\nu_{i\to a}$.

$\hat Q_{a\to i}(\hat\nu_{a\to i})$ = fraction of pure states that have cavity messages $\hat\nu_{a\to i}$.