INSTITUTE OF PHYSICS PUBLISHING JOURNAL OF PHYSICS A: MATHEMATICAL AND GENERAL

J. Phys. A: Math. Gen. **39** (2006) 1265–1283 doi:10.1088/0305-4470/39/6/004

Survey propagation at finite temperature: application to a Sourlas code as a toy model

B Wemmenhove and H J Kappen

SNN, Radboud University Nijmegen Geert Grooteplein 21, 6525 EZ, Nijmegen, The Netherlands

E-mail: b.wemmenhove@science.ru.nl

Received 24 August 2005, in final form 28 November 2005
Published 25 January 2006
Online at stacks.iop.org/JPhysA/39/1265

Abstract

In this paper we investigate a finite temperature generalization of survey propagation, by applying it to the problem of finite temperature decoding of a biased finite connectivity Sourlas code for temperatures lower than the Nishimori temperature. We observe that the result is a shift of the location of the dynamical critical channel noise to larger values than the corresponding dynamical transition for belief propagation, as suggested recently by Migliorini and Saad for LDPC codes. We show how the finite temperature 1RSB SP gives accurate results in the regime where competing approaches fail to converge or fail to recover the retrieval state.

PACS numbers: 89.70.+c, 89.90.+n, 05.50.+q

1. Introduction

In the recent past, there have been a number of developments in the field of information processing and computer science that were triggered by insights originating from statistical mechanics. A number of them were initiated after the discovery of the close relationship between the sum-product algorithm or loopy-belief propagation [1] in computer science on the one hand, and the Bethe approximation, the cavity method and the replica method in statistical mechanics on the other hand. [2, 3]. The belief propagation algorithm is a method to compute approximations of the microstate probability distribution or the corresponding marginals of variables that interact on a given graph. This objective arises in various statistical inference problems with which the machine learning community is concerned, such as stereo vision [4], optimization problems [5], image restoration [6], digital zoom and shape-from-shading [7], error correcting codes for telecommunication [8, 9], human movement modelling [10], protein folding [11] evaluating positions in the game of Go [12]. On the other hand, in the field of statistical mechanics, the main interest is the macroscopic behaviour of systems, and

thus the objective is the calculation of macroscopic observables describing ensemble averages over graphs distributed according to a given measure. In the limit of infinite system size, the macroscopic observables for a given instance of a graph coincides with the ensemble average in many models of interest.

The cavity method for sparse random graphs may be viewed as the macroscopic counterpart of belief propagation. If belief propagation converges, the cavity method describes its macroscopic statistics in the limit of infinite system size [3, 13]. A more careful interpretation is needed when belief propagation ceases to converge, and in this case more exotic statistical mechanics techniques can play a crucial role in generalizations of the belief propagation algorithm.

When belief propagation (BP) breaks down, one may distinguish between different possible causes for this failure. One of these possible causes is replica symmetry breaking, related to the breaking of ergodicity. This is reflected by a free energy landscape exhibiting a structure of a large number of free energy valleys that are disconnected, in the sense that the correlations between them are smaller than correlations within one valley, and the free energy barriers between valleys are extensive. A generalization of the cavity method on sparse graphs has been developed to take into account a first step of replica symmetry breaking [15, 16] over the last years. Although this scheme was designed for the computation of macroscopic observables, i.e. averages over ensembles, it could be adapted to yield an algorithm that gives valuable new information on specific instances of hard optimization problems: survey propagation [5, 17].

Survey propagation (SP) has proved to be a successful algorithm which can cope with hard constraint satisfaction problems such as random three-satisfiability and graph colouring [18]. These problems become hard when the solution space breaks apart in disconnected components, such that the relative sizes of the basins of attraction corresponding to these solutions for algorithms like ordinary belief propagation vanish and many suboptimal fixed points appear. Loosely speaking, the basic idea of survey propagation is to gather the statistics of the fixed points that correspond to solutions of the constraint satisfaction problem[1]. These 'surveys' (probability distributions associated with each link of the graph) may be used for decimation of the original constraint satisfaction problem, until the problem is reduced to a sufficiently easy one that can be solved with a conventional algorithm. One of the reasons why SP is so successful is the fact that it corresponds to a zero-temperature problem in statistical mechanics, having the consequence that the 'surveys,' which represent probability distributions, have a very efficient parametrization. This makes it possible to study very large graph instances. A more careful inspection of SP for K-SAT reveals that the algorithm can be mapped onto a BP algorithm on an extended state space [14], in which boolean variables, besides having a probability of being 0 or 1, can be 'indecisive' (the so-called joker state).

The 1-step replica symmetry broken (1RSB) version of the cavity method is applicable to finite temperature cases as well [15], suggesting a possible potential for the algorithmic counterpart on a specific instance. The benefit aimed at is that in cases where the free energy landscape is complex and belief propagation itself does not converge, the surveys provide sufficient information leading to improved marginal approximations.

However, various difficulties, some practical and some technical, must be taken into account, which we will shortly list here, to be discussed more thoroughly later in the text:

[1] Strictly speaking, there are a number of subtle differences, such as the assumption that the cavity fields take integer values. Note that nonconvergence of BP is not the only problem which is cured by SP: algorithms that are guaranteed to converge to local minima of the Bethe free energy in contrast to SP do not find solutions to three-sat problems in the hard-sat phase [19], i.e. the local minima one arrives at are not likely to be a groundstate.

- The surveys we will be dealing with are not efficiently parametrized, since they can take values on a continuous interval. We will be forced to sample from distributions, which severely limits the system sizes. For infinite values of the connectivity of the corresponding graphs, more efficient representations become possible, as suggested in [20, 21].
- Given that we have a limited system size, we must remember that the approach is likely to be sensitive to finite size effects. One of the issues is the optimization of the generalized free energy with respect to the replica symmetry breaking parameter m. This may be interpreted as zooming in at the states with the lowest free energy, where, moreover, the 'complexity' (entropy of states) vanishes in a truly 1RSB scenario. It is questionable what remains of this idea for a small system size.
- It is not straightforward to test the performance of the algorithm: if the output of the algorithm is a set of approximate single-spin marginal probabilities, how can we judge the quality of this approximation? Comparing to exact results is possible for very small system sizes only ($N = O(10^2)$).
- We should take into account the possibility of further steps of replica symmetry breaking, which might prevent the algorithm from converging, or cause malfunction in a different manner.

The above complications show that our expectations regarding an algorithm based on the 1RSB generalization of the cavity method should be modest. However, in the current paper we present encouraging results, indicating that there is at least a possible type of application, albeit for future generations of fast computers, in the field of error correcting codes. This point has very recently been recognized and suggested by Migliorini and Saad [22] for LDPC codes, where the authors argued that a 1RSB-type algorithm might increase the dynamical limit of the channel noise for decoding. In this paper, we for the first time propose such an algorithm for the biased finite connectivity Sourlas code with 3-spin interactions. We show how a finite temperature survey propagation algorithm gives accurate results in the regime where competing approaches fail to converge or fail to recover the encoded message.

The suggested algorithm may be used for finite temperature decoding of encrypted messages, where the single spin marginals are the basis for inference of original message bit values. A corresponding performance measure of the algorithm is the bit error rate, which is easily calculated. Thereby the required knowledge of exact values of marginals is circumvented and one is not limited to small system sizes. The model we will study is a biased finite connectivity Sourlas code with 3-spin interactions. This model is closely related to a general 3-spin model with external fields.

In the following sections we will introduce the general model class of Ising 3-spin systems with external fields, we will shortly review belief propagation, the cavity method and its 1RSB generalization, and finally describe the finite temperature generalization of survey propagation. Exact comparison of marginals is only possible for limited system sizes ($N = 45$ using a junction tree algorithm [1]), which indeed turned out to be too small to see any improvement in a few cases we studied. Therefore we consider, as a toy model, finite temperature decoding of the finite connectivity Sourlas error correcting code based on 3-spin interactions. As was also pointed out for the case of a CDMA model in [20, 21], when decoding happens without the knowledge of the channel noise, one may suffer from replica symmetry breaking effects if the adopted estimate of the channel noise is below the Nishimori temperature [23]. The appearance of metastable states disrupts the behaviour of simple message passing algorithms before the onset of replica symmetry breaking of the equilibrium state [24]. As mentioned above, in [22] it was proposed that a 1RSB-type algorithm might be able to increase the critical noise value for the decoding dynamics as compared to RS-type decoding algorithms.

2. The 3-spin model

Although the formalism may be set up in a more general fashion, we will be focusing on systems of N binary variables $\sigma_i \in \{-1, 1\}$, $i \in \{1, 2, \ldots, N\}$ that interact on a random graph, where the order of the interaction is 3. The equilibrium microstate probability distribution will be described by the Boltzmann measure

$$p(\boldsymbol{\sigma}) = \frac{\exp[-\beta H(\boldsymbol{\sigma})]}{\sum_{\boldsymbol{\tau}} \exp[-\beta H(\boldsymbol{\tau})]}, \tag{1}$$

where the Hamiltonian is given by

$$H(\boldsymbol{\sigma}) = \sum_{\nu} J_\nu \sigma_{\nu_1} \sigma_{\nu_2} \sigma_{\nu_3} + \sum_i \theta_i \sigma_i, \tag{2}$$

where $J_\nu \in \{-1, 1\}$. Greek letters μ, ν will label interaction vertices (factors in a factor graph representation), ν_i label spins in $V(\nu)$, the set of site indices contained in interaction ν. We will be discussing random graphs with fixed connectivity, i.e. the number of interactions for spin i, $|V(i)|$, will be fixed and equal for all i, where $V(i)$ is the set of all interactions containing i, i.e. we will choose $|V(i)| = k\,\forall i$ from now on. Consequently, the sum over ν in the Hamiltonian (2) runs from 1 to $kN/3$, the total number of interactions.

In a statistical inference problem, we assume the values of $\{J_\nu\}$, $\{\theta_i\}$ and β are known, and the objective is the calculation of the partition sum, or of a marginalized probability, e.g.

$$p(\sigma_i) = \sum_{\boldsymbol{\sigma}\setminus i} p(\boldsymbol{\sigma}). \tag{3}$$

In general the calculation of this quantity requires $O(2^N)$ operations, and thus for large N one resorts to approximations.

In the regime of finite connectivity, the average loop length is of order $\log(N)$, such that for large system sizes, correlations that propagate through loops may be neglected if the couplings are of order N^0. In this regime the Bethe approximation, which is exact on tree graphs, performs well. Consequently, the cavity method gives rather accurate results, as long as the system is in a single state, and belief propagation is a good candidate for inference on a (large) single instance [13]. For large values of β (low temperature values), replica symmetry breaking occurs, and a single-state cavity method description is no longer accurate. Belief propagation ceases to converge, or gets trapped in a sub-optimal free energy minimum [24], and does not give any information. In the absence of external fields, it has been shown that the 3-spin model with interactions $J_\mu \in \{-1, 1\}$ is self-consistently described within the 1RSB cavity framework [25], for a fixed number of interactions per spin.

3. The methods

3.1. The cavity method

The cavity method [26, 15] in its simple form is based on two assumptions:

- The values of the spins $\{\sigma_j\}$ that are neighbours on the graph of site i are only nontrivially correlated through the interaction $\nu \in V(j) \cap V(i)$ with site i.
- In equilibrium, the system is in a single state (there is just one dominant free energy valley), i.e. there is a one to one correspondence between the states of neighbouring spins.

As a consequence of these assumptions, in the absence of their interactions ν with spin σ_0, we may write the probabilities for $2(k-1)$ neighbours $\{\sigma_j\}$ of σ_0 as independent, parametrized in the following way:

$$p_j^{\nu}(\sigma_j) \sim \exp\left(\beta h_j^{\nu}\sigma_j\right). \tag{4}$$

Then linking the $2(k-1)$ spins to spin 0, we may write

$$p_0^{\mu}(\sigma_0) \sim \sum_{\{\boldsymbol{\sigma}_{\setminus\sigma_0}\}} \exp\left\{\beta\sigma_0\left(\sum_{\nu\in V(0)_{\setminus\mu}}\left[J_\nu\sigma_{\nu_2}\sigma_{\nu_3}+\theta_0\right]\right)+\beta\left[\sum_{\nu\in V(0)_{\setminus\mu}}\left[h_{\nu_1}^{\nu}\sigma_{\nu_1}+h_{\nu_2}^{\nu}\sigma_{\nu_2}\right]\right]\right\}. \tag{5}$$

It then follows that in this state the corresponding field h_0 associated with spin 0 (which again misses an interaction with one other pair of spins, interaction μ) is given by

$$h_0^{\mu} = \sum_{\nu\in V(0)_{\setminus\mu}} \frac{1}{\beta}\tanh^{-1}\left[\tanh(\beta J_\nu)\tanh\left(\beta h_{\nu_1}^{\nu}\right)\tanh\left(\beta h_{\nu_2}^{\nu}\right)\right]+\theta_0. \tag{6}$$

On a sparse random graph with random interaction values, the macroscopic statistics will be governed by the distribution of cavity fields $W(h)$ over the graph, which, due to the statistical independence of cavity spins, follows from iterating equation (6) in distribution:

$$W(h) = \int\prod_{i=1}^{k-1}[\mathrm{d}h_i\mathrm{d}g_i\,W(h_i)W(g_i)] \times\left\langle\delta\left\{h-\frac{1}{\beta}\sum_{i=1}^{k-1}\tanh^{-1}[\tanh(\beta J)\tanh(\beta h_i)\tanh(\beta g_i)]+\theta\right\}\right\rangle_{J,\theta}, \tag{7}$$

where the brackets denote averages over the interactions J and random external fields θ. This equation may be solved numerically with a 'population dynamics' algorithm [15]. From the stationary distribution under this iteration algorithm, all macroscopic quantities can be derived.

3.2. Belief propagation

Belief propagation in essence is the above procedure applied to a specific instance of a (sparse) graph. Now the procedure is focused on the computation of local quantities, all the cavity fields associated with the graph, i.e. one for each link. By iterating equation (6) for a given graph, one avoids averaging over the interaction values J and the external fields θ. If the algorithm converges to a fixed point, from the resulting cavity fields one may compute the single node 'beliefs'

$$b_i(\sigma_i) = \tfrac{1}{2}(1+m_i\sigma_i) \tag{8}$$

which are approximations for the marginal probabilities

$$b_i(\sigma_i) \simeq p(\sigma_i). \tag{9}$$

Their expressions in terms of the cavity fields are

$$m_i = \tanh(\beta H_i), \tag{10}$$

where

$$H_i = \sum_{\mu\in V(i)} \frac{1}{\beta}\tanh^{-1}\left[\tanh(\beta J_\mu)\tanh\left(\beta h_{\mu_1}^{\mu}\right)\tanh\left(\beta h_{\mu_2}^{\mu}\right)\right]+\theta_i. \tag{11}$$

The 3-spin beliefs, approximations for the 3-spin marginals are computed from the cavity fields in a similar way.

Belief propagation has been rediscovered several times, and has a close relation with the Bethe approximation. Fixed points of belief propagation are minima of the Bethe free energy, which is exact on tree-like graphs, but is often a good approximation when a graph contains loops [27].

For lower temperatures and for graphs with interactions inducing frustration effects, the cavity method and belief propagation break down. Belief propagation often ceases to converge in these cases. One would expect a similar breakdown of the cavity method in this case, but since the order parameter of the cavity method is a macroscopic object, the behaviour is quite different generally. The distribution of cavity fields may still be stationary under the population dynamics iteration in parameter regimes where belief propagation does not converge. However, this stationary distribution obviously does not represent the distribution of locally consistent beliefs. Thus the interpretation of this result is a more complicated issue in terms of local quantities[2]. It is suggested [13] that stability analysis through the iteration of perturbations in the distribution might reveal the location of the AT line [28], at which the replica symmetric approximation becomes unstable. For the SK model [29] on a dense graph this has been shown [13] to be the case indeed. Several heuristic [30] or rigorous statements [31, 32] have been made concerning convergence of belief propagation in relation to the location in the phase diagram of Ising spin-type systems.

The cavity method has been extended to deal with cases in which replica symmetry breaking occurs, and thus it is worthwhile to modify this method and apply it to the case of single instance statistical inference. In the same manner, survey propagation and its generalization for nonzero energy density [33] were developed for optimization problems. The main difference in the following discussion will be the application to problems at nonzero noise level, i.e. at finite temperature.

We will note here explicitly that this approach will not cure problems arising from short-range correlations, e.g. in regular lattices. A possible approach for this type of problems, which we shall not discuss here, is to take into account correlations between larger numbers of spins in the free energy. This idea was first introduced by Kikuchi [34] and leads to the cluster variation method [35].

3.3. 1-step replica symmetry breaking

The cavity method at the level of 1-step replica symmetry breaking is usually applied at zero temperature [16, 17, 36, 37]. The reason for this is that at $T = 0$ on the one hand the equations simplify, and on the other hand, there are more known applications in optimization problems for which a single instance generalization of this technique is fruitful (e.g., 3-sat, graph colouring).

Here we will be applying the finite temperature version of the cavity method, of which we will just discuss the basic assumptions, and refer the reader to [15] and [16] for a more thorough description. The main assumptions, to be compared with the assumptions of the cavity method in its simple (replica symmetric) form, are summarized in the following.

In contrast with the 'single-state' cavity method, one now assumes that an exponential number of states contribute to the Boltzmann measure. At each link of the graph, instead of regarding just one cavity field, one takes into account a distribution of cavity fields. The

[2] From a global point of view the equations are equivalent to the replica symmetric approximation of the macroscopic observables.

density of states at a given free energy is of an exponential form

$$\rho(F) = \exp[m\beta(F - F_{\mathrm{R}})], \tag{12}$$

where F_{R} is the reference energy and m is the so-called replica symmetry breaking parameter. It can be shown that this assumption implies that the free energy and local fields of the relevant (lowest free energy) states are uncorrelated, in the sense that the joint distribution of free energy values and local field values factorizes. Since, on the contrary, local free energy differences are direct functions of the local fields, as a consequence at each iteration the distribution is reweighted by a factor $\exp[-m\beta\Delta F]$, where ΔF is the difference between the free energies of the state before and after the iteration. Observables will first be summed over all states α weighted by their corresponding Boltzmann factors $\exp(-\beta F^{\alpha}) \big/ \sum_{\gamma} \exp(-\beta F^{\gamma})$, the result of which should be averaged over the distribution of free energies (12).

In practice, the calculation of observables is numerically implemented by an advanced population dynamics algorithm, in which N populations of M fields, each representing a distribution at a given link, are iterated. This procedure can be executed for all values of $m \in [0, 1]$, yielding a generalized free energy $F(m)$. Selecting the equilibrium state is achieved by demanding $\partial F(m)/\partial m = 0$ [26], which is equivalent to zooming in at the generalized free energy of states for which the entropy of states (or 'complexity') becomes non-extensive [38].

The numerical cost of the procedure described above is considerable, since the cavity field distributions are represented in terms of (large) sample populations. Moreover, the extremization of the generalized free energy with respect to m may be difficult, when the variation of $F(m)$ as a function of m is small. At zero temperature, for a large model class, the situation simplifies a lot, since the representation of cavity field distributions is much more efficient. In the case of discrete values for the local Hamiltonian (e.g., any spin glass with $\pm J$ interactions, k-sat energy terms) the cavity fields can be assumed to take integer values, such that the only degrees of freedom are coefficients of a small number of delta peaks. For random k-satisfiability, a single number parametrizes the field distribution for each link on the graph. In the SAT phase, where the energy is zero, the reweighting of distributions by the free energy shift as described above effectively either leaves the distribution invariant, or rules out the state completely (contradictory states), leading to the very efficient algorithm of survey propagation on single instances of random graphs representing the satisfiability problem.

4. Finite temperature survey propagation

The goal of this paper in general is an investigation of the benefits of a survey propagation algorithm for finite temperature. The type of problem we address is the inference of marginal probabilities for a spin-glass-type model displaying replica symmetry breaking effects. Along the lines of the 1-step RSB cavity method for finite temperature, we iterate distributions of cavity fields, where instead of sampling external fields and interaction values at each iteration, we let the distributions evolve on a given instance of a sparse graph, thus fixing the value of θ and J locally. For a given value of the replica symmetry breaking parameter, the iteration equations are the following:

$$W_0^{\mu}\left(h_0^{\mu}\right) = \int \prod_{\nu \in V(0)\setminus\mu} \left[\mathrm{d}h_{\nu_1}^{\nu} W_{\nu_1}^{\nu}\left(h_{\nu_1}^{\nu}\right) \mathrm{d}g_{\nu_2}^{\nu} W_{\nu_2}^{\nu}\left(g_{\nu_2}^{\nu}\right)\right] \exp\left(-m\beta\Delta F\left(\{J_{\nu}\}, \left\{h_{\nu_1}^{\nu}\right\}, \left\{g_{\nu_2}^{\nu}\right\}\right)\right)$$
$$\times\, \delta\left\{h_0^{\mu} - \sum_{\nu \in V(0)\setminus\mu} \frac{1}{\beta} \tanh^{-1}\left[\tanh(\beta J_{\nu}) \tanh\left(\beta h_{\nu_1}^{\nu}\right) \tanh\left(\beta g_{\nu_2}^{\nu}\right)\right] - \theta_0\right\}, \tag{13}$$

where

$$\Delta F\left(\{J_\nu\},\{h^\nu_{\nu_1}\},\{g^\nu_{\nu_2}\}\right)=\log 2+\sum_{\nu\in V(0)_{\setminus\mu}}\log\left\{\frac{\cosh(\beta J_\nu)}{\cosh\left(\beta u\left(J_\nu,h^\nu_{\nu_1},g^\nu_{\nu_2}\right)\right)}\right\}$$
$$+\log\cosh\left\{\beta\left[\sum_{\nu\in V(0)_{\setminus\mu}}u\left(J_\nu,h^\nu_{\nu_1},g^\nu_{\nu_2}\right)+\theta_0\right]\right\} \tag{14}$$

and

$$u\left(J_\nu,h^\nu_{\nu_1},g^\nu_{\nu_2}\right)=\frac{1}{\beta}\tanh^{-1}\left[\tanh(\beta J_\nu)\tanh\left(\beta h^\nu_{\nu_1}\right)\tanh\left(\beta g^\nu_{\nu_2}\right)\right]. \tag{15}$$

From the distributions of cavity fields, one may construct distributions of 'belief' fields

$$\hat{W}_0(H_0)=\int\prod_{\mu\in V(0)}\left[\mathrm{d}h^\mu_{\mu_1}W^\mu_{\mu_1}\left(h^\mu_{\mu_1}\right)\mathrm{d}g^\mu_{\mu_2}W_{\mu_2}\left(g^\mu_{\mu_2}\right)\right]\exp\left(-m\beta\Delta\hat{F}\left(\{J_\mu\},\{h^\mu_{\mu_1}\},\{g^\mu_{\mu_2}\}\right)\right)$$
$$\times\delta\left\{H_0-\sum_{\mu\in V(0)}\frac{1}{\beta}\tanh^{-1}\left[\tanh(\beta J_\mu)\tanh\left(\beta h^\mu_{\mu_1}\right)\tanh\left(\beta g^\mu_{\mu_2}\right)\right]-\theta_0\right\} \tag{16}$$

where

$$\Delta\hat{F}\left(\{J_\mu\},\{h^\mu_{\mu_1}\},\{g^\mu_{\mu_2}\}\right)=\log 2+\sum_{\mu\in V(0)}\log\left\{\frac{\cosh(\beta J_\mu)}{\cosh\left(\beta u\left(J_\mu,h^\mu_{\mu_1},g^\mu_{\mu_2}\right)\right)}\right\}$$
$$+\log\cosh\left\{\beta\left[\sum_{\mu\in V(0)}u\left(J_\mu,h^\mu_{\mu_1},g^\mu_{\mu_2}\right)+\theta_0\right]\right\}. \tag{17}$$

The marginal probability $p(\sigma_0)$ for the value of the spin at site 0 is now given by

$$p(\sigma_0)=\frac{1}{2}\left[1+\sigma_0\int\mathrm{d}H_0\hat{W}_0(H_0)\tanh(\beta H_0)\right]. \tag{18}$$

Following [15], we can deduce expressions for the generalized free energy of the system $F(m)$ and the derivative of the generalized free energy $\partial F(m)/\partial m$. These are given in (appendix A).

Note that the above equations are just slightly more restrictive than for the case in which one is interested in averages over an ensemble of graphs, the original 1RSB cavity method. Numerically, the only difference for the ensemble is that at each iteration, the set $V(0)_{\setminus\mu}$ is a random subset of a population of cavity field distributions, and each value of the interaction J^ν, $\nu\in V(0)_{\setminus\mu}$ is randomly sampled from the distribution of interactions.

5. The 3-spin model as a biased finite connectivity Sourlas-code for a binary symmetric channel

The 3-spin model with finite connectivity may be regarded as an error correcting code of the Sourlas type [39]. The Sourlas code for p-spin interactions was originally studied for a fully connected model. In the limit $p\to\infty$, it is equivalent to the random energy model [40] and the Shannon limit is approached. The benefit of looking at finite connectivity codes and finite p is the fact that the transmission rate, which vanishes for the dense graph and $p\to\infty$, is now finite [41]. However, the drawback is the fact that the bit error rate does not vanish, such

that for practical purposes its potential is limited. Codes of the Gallagher type (low-density parity check codes) [42] are more promising from that perspective, since they have vanishing bit error rate and approach the Shannon limit.

The optimal decoding temperature for error correcting codes can be shown to coincide with the temperature corresponding to the 'channel noise' [23]. Since this temperature is equal to the so-called 'Nishimori temperature,' it can be shown (under certain conditions, including the absence of external fields or 'biases') that this temperature always corresponds to a replica symmetric phase in the phase diagram [43], suggesting that belief propagation may be an appropriate decoding algorithm.

In real life, however, the channel noise might be an unknown quantity. If the estimate of the channel noise in this case is too small, the decoding is likely to happen in a parameter region where RSB effects play a role, in which BP either does not converge, or gives bad results. One would expect that in this case, the 1RSB generalization might be a robust alternative [21, 20].

A more detailed inspection of the generic behaviour of error correcting codes tells us that decoding problems are caused by the so-called dynamical phase transition [24]: in a low temperature regime above some critical channel noise value p_d, the fragmentation of the free energy landscape prevents the algorithm from finding the ferromagnetic state. Above an even larger noise value $p_c > p_d$, the free energy of a replica symmetry broken spin glass state is lower than the ferromagnetic free energy, such that decoding is not possible at all. Thus, the only region in which a 1RSB algorithm might improve the decoding performance, is in the region $p_d < p < p_c$, thereby shifting the dynamical phase transition to a higher value.

In the presence of external fields, the model may be interpreted as a *biased* Sourlas code. The location in the phase diagram where improvement upon conventional decoding is to be expected, is close to the critical channel noise, for low decoding temperature. Since in the case of a random bias we might not have a sharp phase transition, we will have to be careful in the selection of a suitable parameter regime.

In order to select a parameter regime in which our algorithm might perform better than BP, we will perform a short macroscopic analysis of the model. We will here shortly discuss the cavity analysis of the model for 3-spin interactions and random external fields. In [41], the model was studied for uniformly biased messages (a uniform external field) using the replica method. In principle, their result applies to the present case of random biases, due to gauge symmetry, for the appropriate choice of a message bias, but for clarity we will here discuss the derivation of this result to make the connection explicit.

5.1. RS analysis of biased code for 3-spin interactions

It is well known that an unbiased Sourlas code has such a small basin of attraction that BP does not converge to the 'ferromagnetic' or 'retrieval' state for any order of the interaction larger than 2, unless prior knowledge is taken into account in the initialization [41]. Therefore, we have added random external fields, or random biases to a fraction of the sites, which have the effect of breaking the local symmetry.

First of all, a locally biased message $\boldsymbol{\xi} \in \{-1, 1\}^N$ is generated, according to

$$p(\boldsymbol{\xi}|\boldsymbol{\theta}) = \prod_i \frac{\exp[\beta_p \theta_i \xi_i]}{2\cosh[\beta_p \theta_i]}. \tag{19}$$

The external fields or biases will be generated from the distribution

$$p(\boldsymbol{\theta}) = \prod_i p(\theta, i), \tag{20}$$

where we choose

$$p(\theta_i) = c(+)\delta(\theta_i - \theta) + c(-)\delta(\theta_i + \theta) + c(0)\delta(\theta_i) \tag{21}$$

with θ being constant and of course $c(+) + c(-) + c(0) = 1$.

Given the message, the received code $\mathbf{J} \in \{-1, 1\}^{Nk/p}$ is generated according to

$$p(\mathbf{J}|\boldsymbol{\xi}) = \prod_{\mu} \frac{\exp\left[\beta_p J^{\mu} \xi_{\mu_1} \xi_{\mu_2} \xi_{\mu_3}\right]}{2\cosh(\beta_p)}. \tag{22}$$

Note that the above choices correspond to a binary symmetric channel (BSC), for which each transmitted bit is likely to be flipped to its inverse with equal probability p_{ch} given by

$$\beta_p = \frac{1}{2} \log \frac{1 - p_{\mathrm{ch}}}{p_{\mathrm{ch}}}. \tag{23}$$

The original message is inferred according to

$$\hat{\xi}_i = \mathrm{sgn}(\langle \sigma_i \rangle_{\beta}), \tag{24}$$

where the average is taken with respect to the Boltzmann measure

$$p(\boldsymbol{\sigma}) = \frac{\exp[-\beta H(\boldsymbol{\sigma})]}{\sum_{\tau} \exp[-\beta H(\boldsymbol{\tau})]}. \tag{25}$$

Thus, for the finite temperature decoding at inverse temperature β, the knowledge of the marginal $p(\sigma_i)$ is required. Since we want to describe the macroscopic behaviour of the code, we use the cavity method to average over the disorder[3] (the messages $\boldsymbol{\xi}$ and the interaction values $\mathbf{J}$).

Here we summarize the equations to be iterated using a population dynamics algorithm (see appendix B for details)

$$\begin{aligned}
W(h) &= \sum_{\alpha} q(\theta_\alpha) \int \prod_{i=1}^{k-1} [\mathrm{d}u_i \, Q(u_i)] \delta\left[h - \sum_{i=1}^{k-1} u_i - \theta_\alpha\right] \\
Q(u) &= \sum_{J} \frac{\mathrm{e}^{\beta_p J}}{2\cosh(\beta_p)} \int \mathrm{d}h \, \mathrm{d}g \, W(h) W(g) \delta\left[u - \frac{1}{\beta} \tanh^{-1}[\tanh(\beta h)\tanh(\beta g)\tanh(\beta J)]\right] \\
\hat{W}(H) &= \sum_{\alpha} q(\theta_\alpha) \int \prod_{i=1}^{k} [\mathrm{d}u_i \, Q(u_i)] \delta\left[H - \sum_{i=1}^{k} u_i - \theta_\alpha\right] \\
\mu &= \int \mathrm{d}H \, \hat{W}(H) \mathrm{sgn}(H).
\end{aligned} \tag{26}$$

where

$$q(\theta_\alpha) = c(\alpha) \frac{e^{\beta_p \theta_\alpha}}{\cosh(\beta_p \theta_\alpha)}. \tag{27}$$

From these equations we compute the bit error rate $\rho = (1 - \mu)/2$ as a function of the channel noise, for connectivity $k = 4$ at inverse temperature $\beta = 5$, where we initialize the population dynamics algorithm with random positive fields uniformly on the interval $0 \leqslant h \leqslant 3$, or randomly on a symmetric interval $-3 \leqslant h \leqslant 3$.

We chose the fraction of external fields according to the values $c(0) = 0.8$, $c(+) = c(-) = 0.1$ and $\theta = 1.5$, such that the magnitude at low temperatures might compete with the local interaction terms at sites where $\theta \neq 0$. For a temperature sufficiently below the Nishimori temperature ($\beta_{\mathrm{Nishimori}} = \beta_p$), i.e. $\beta = 5$, the results are reported in figure 1. Clearly there

[3] Note that using the replica method in the replica symmetric framework one generally arrives at identical equations.

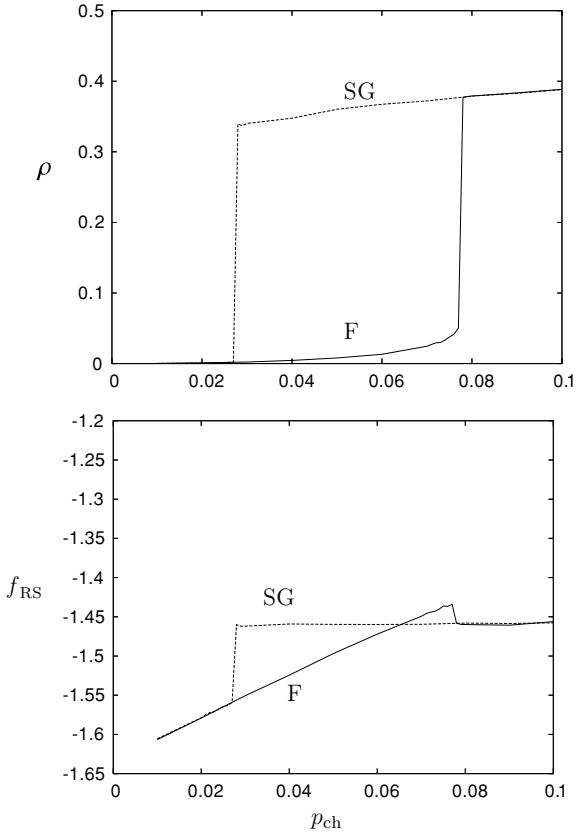

Figure 1. Top: bit error rate ρ as a function of channel noise for a code with $k = 4$, $\beta = 5$ and bias $\theta = 1.5$ with ferromagnetic initial conditions versus spin-glass initial conditions. Bottom: the RS free energies corresponding to the top picture. In both figures the label F denotes ferromagnetic initial conditions and SG denotes the branch with spin-glass-type initial conditions (see text).

is a jump in the bit error rate as a function of the channel noise around $p_{\mathrm{ch}} = 0.028$ for 'spin-glass' initial conditions, and around $p_{\mathrm{ch}} = 0.078$ for ferromagnetic initial conditions. We may interpret the former line as a dynamical phase transition for BP. Comparing the free energy values in both cases, we find that the replica symmetric approximation of the spin-glass free energy is nearly constant over a large interval, whereas the ferromagnetic free energy (which is expected to be a replica symmetric indeed) is an increasing function up to the point where the population dynamics collapse to the spin-glass state. The RS critical channel noise corresponds to the point where the two lines cross, suggesting that retrieval is not at all possible above this point, around $p_{\mathrm{ch}} = 0.065$. Within the RS framework, therefore, the jump at $p_{\mathrm{ch}} = 0.078$ is not the physical critical channel noise.

Figure 2 displays the location of the various transition points for a range of temperatures below the Nishimori temperature within the here described RS framework. Within a 1RSB framework the results are likely to change, since in general the spin-glass free energy will be higher than the RS-approximated result. From the above RS analysis, we expect BP to fail for p_{ch} above the dynamical transition line, and we expect improvement using the 1RSB treatment above the RS dynamical transition.

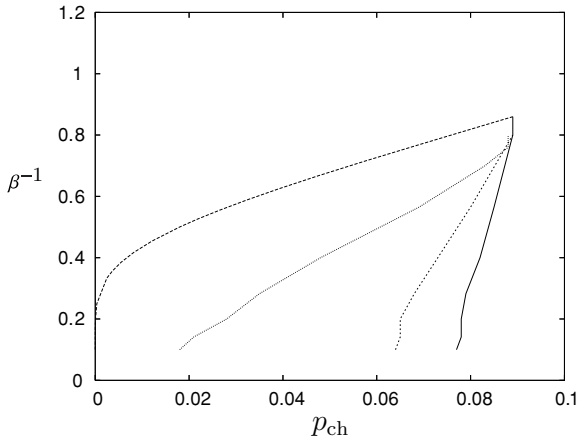

Figure 2. Replica symmetric phase diagram of the finite connectivity Sourlas code for $k = 4, p = 3, c(0) = 0.8, c(+) = c(-) = 0.1, \theta = 1.5$. From top to bottom the lines indicate the Nishimori temperature, the dynamic transition, the RS critical transition and the line along which the population dynamics algorithm ceases to find the ferromagnetic state.

6. Numerical results for single instance decoding

We have tested a number of algorithms on single instances of a Sourlas code corresponding to the same parameter values as above, i.e. with $\beta = 5$, way below the Nishimori line in the critical region. Ordinary BP was able to find the ferromagnetic state roughly up to the value of $p = 0.028$, the point where the randomly initiated population dynamics algorithm ceases to converge. Beyond that point, BP ceased to converge too, even when we took into account the prior knowledge corresponding to the external field in the initialization. However, it turned out that a 'damped' version of BP (i.e., at each timestep $(t + 1)$ replacing the field h_t by $(1 - \epsilon)h_t + \epsilon h_{t+1}$) improved the performance dramatically, again with random initial conditions. In case this algorithm did not converge either, we applied a time-averageing procedure, along the lines of [44, 45]. The resulting algorithm was able to find an optimal ferromagnetic-type state (i.e., with bit error rate comparable to the lower branch in figure 1) for nearly all instances with noise $p_{\rm ch} < 0.07$.

For $p_{\rm ch} = 0.07$, we ran our 1RSB algorithm on 20 instances of a Sourlas code of $N = 999$ bits (note that we need kN to be divisible by 3) and compared the results to those of damped time-averaged BP (TABP). Additionally, we compared the results to those of a so-called 'double-loop' algorithm, which is guaranteed to converge to a local minimum of the Bethe free energy [46][4]. For the sample size of the cavity field distributions we chose $M = 1024$, and we iterated for 50 updates per field distribution to get rid of transients before iterating another 50 updates for the calculation of observables. These numbers imply a considerable numerical cost, limiting the number of experiments we could perform in a reasonable amount of time. We ran the algorithm for the values $m = 0, 0.05, 0.1, 0.15$ and 0.2 of the replica symmetry breaking parameter. TABP was damped with $\epsilon = 0.05$, the resulting marginals were averages over the last 2500 (sequential) updates of 10 000 in total. For each instance we repeated the latter procedure with random initial conditions ten times.

[4] Note that this algorithm is much more successful for regular lattices, where it provides an implementation of the cluster variation method. Here it is only used as a comparison since it will converge to a Bethe free energy minimum even when BP or damped BP does not.

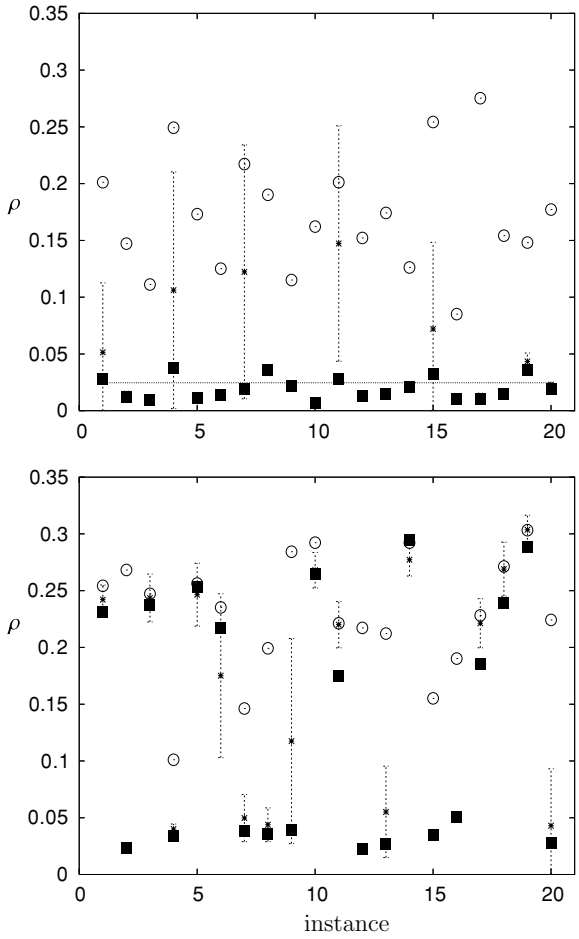

Figure 3. $k = 4, \beta = 5$ single-instance bit error rate results for 20 systems of size $N = 999$, with channel noise $p_{\mathrm{ch}} = 0.07$ (top) $p_{\mathrm{ch}} = 0.08$ (bottom). Open circles: double-loop algorithm, full squares: our 1RSB algorithm, stars with error bars: ten runs of damped TABP, full square fall on top when not visible. Top figure horizontal line: population dynamics (ferromagnetic initial conditions) RS result.

6.1. Results

In figure 3, we report the results of the above experiments for $p_{\mathrm{ch}} = 0.07$ and $p_{\mathrm{ch}} = 0.08$. The results of the double-loop algorithm, not surprisingly, clearly have the largest bit-error rate, since the algorithm is not designed to find a global minimum of the Bethe free energy. Generically, for small values of the replica symmetry breaking parameter m, the 1RSB algorithm settles in a suboptimal spin glass state, with a bit-error rate comparable to the double-loop result. The resulting generalized free energy, which increases with m, should be maximized with respect to m. However, before reaching a maximum, the generalized free energy drops to the ferromagnetic free energy, the corresponding bit-error rate dropping simultaneously. The optimal value of m is selected by choosing the smallest m after the generalized free energy drops. This behaviour was found for LDPC codes macroscopically in [22], but was not yet applied to real instances. For $p_{\mathrm{ch}} = 0.07$, we observe that our

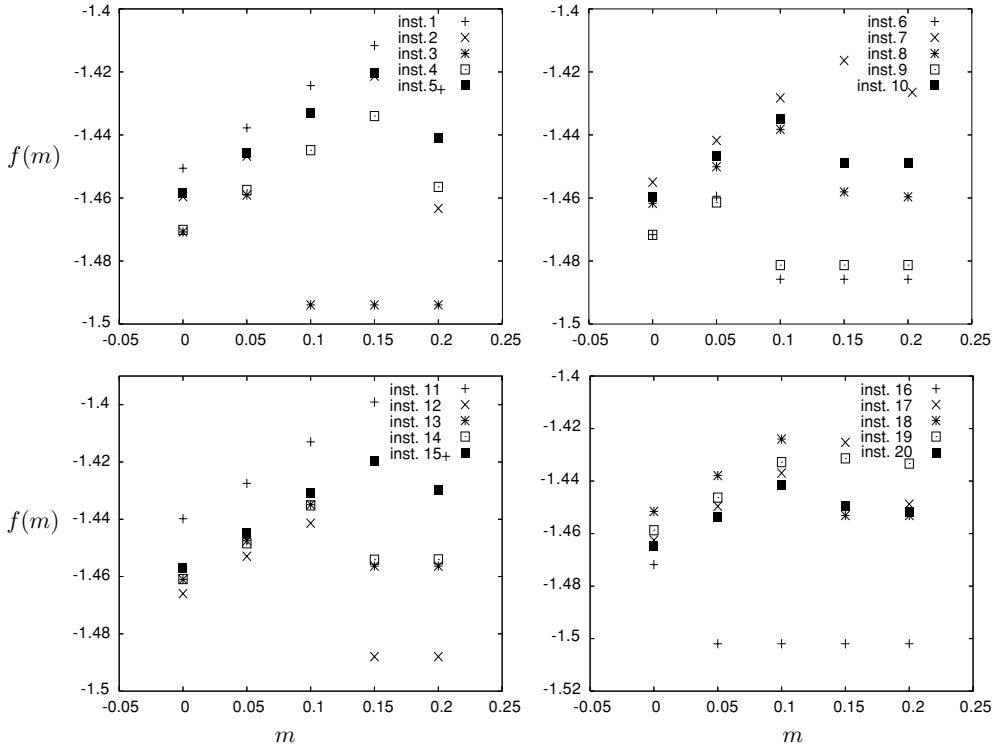

Figure 4. Generalized free energy densities as a function of the replica symmetry breaking parameter m for all instances at $p_{\rm ch} = 0.07$.

algorithm finds the ferromagnetic state for each instance at some value of m, which still fluctuates between different instances due to finite size effects. We conclude that the damped TABP algorithm performs surprisingly well on some instances. However, whether it finds a near-optimal state depends on each individual run in general. All the instances 1, 4, 7, 11, 15 and 19 suffer from initialization dependence and have small probability of convergence to the ferromagnetic state. The 1RSB algorithm, in contrast, is robust in the sense that it finds the optimal state for some m, regardless of initial conditions. Note that the error bars of the damped TABP algorithm are somewhat misleading, since they indicate the standard deviation, although the distribution of bit-error rates is not symmetric.

In figure 4, we have plotted all values of the generalized free energy as a function of the values of m for which we ran the algorithm, illustrating the behaviour described above. Note that the maximum of the spin glass free energy indeed seems to extrapolate to values higher than the approximate RS value (around -1.46, see figure 1), indicating that the critical channel noise is likely to be increased in the 1RSB framework.

For $p_{\rm ch} = 0.08$, the ferromagnetic state was not always found by the 1RSB algorithm, indicating either that this is above the 1RSB dynamical transition, or we should look in a larger range of m. However, when damped TABP finds a ferromagnetic state, in all cases the 1RSB algorithm finds it robustly. There was one exception to the latter point, in instance 6, where TABP found the ferromagnetic state in one of the ten runs. When we increased the value of m further, however, the algorithm stopped converging (i.e., the distributions of fields started fluctuating heavily). A further increase of the sampling size might be helpful here.

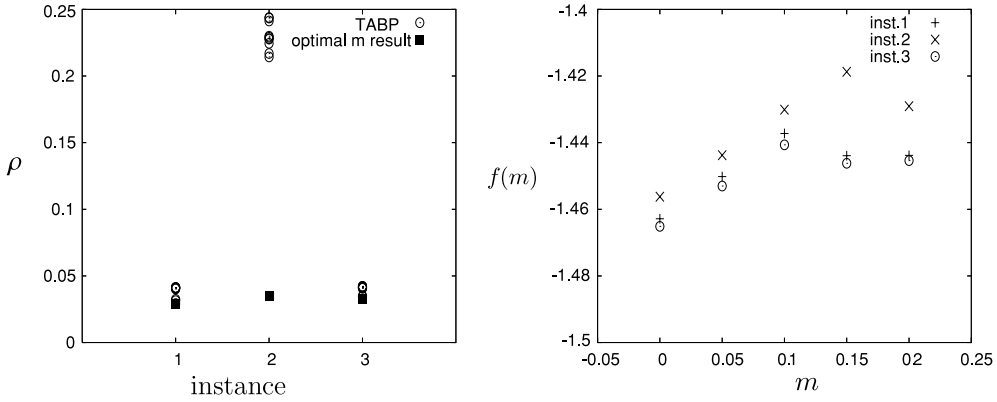

Figure 5. Left: error bit rate comparison between the 1RSB algorithm and ten runs of TABP at $p_{\mathrm{ch}} = 0.075$ and $N = 4998$. Right: corresponding generalized free energy density as a function of m.

Finally, to reduce finite size effects further, we ran the 1RSB algorithm for just three instances for a large code of $N = 4998$ at the noise value $p_{\mathrm{ch}} = 0.075$. Results are reported in figure 5. Still we observe different results for each instance, but again the 1RSB algorithm finds the ferromagnetic state in all three cases. This result is spectacular in comparison to the performance of the TABP algorithm in the second instance, where it did not find the ferromagnetic state in any of the ten runs.

7. Discussion

In this paper we have presented results of applying a finite temperature generalization of survey propagation on sparse graphs. This algorithm, though numerically expensive, seems particularly of interest for finite temperature decoding of error correcting codes at temperatures beneath the Nishimori line close to the critical channel noise. We have illustrated this by regarding a finite connectivity, partly biased Sourlas code with 3-spin interactions. As predicted in [22], the dynamical transition of belief propagation is pushed towards larger values of the channel noise for temperatures beneath the Nishimori temperature. This feature is also reminiscent of zero temperature results obtained for survey propagation in the UNSAT phase [33].

For the model we studied, it turned out that a damped version of time-averaged belief propagation performed particularly well too. However, the finite temperature survey propagation algorithm turned out to be more robust, and was able to find the retrieval state in cases where all other tested algorithms failed.

We feel that this result is an encouragement to further investigate the properties, in spite of its numerical cost, of the algorithm, which should certainly be tested on the more realistic LDPC codes.

Furthermore, the dependence of the performance on various parameters, like the temperature, the bias, and the system size should be investigated further. In principle, the simple phase-space structure of unbiased error correcting codes at the Nishimori line suggests that in practice it should be easy to infer the value of the channel noise by varying the temperature in a BP algorithm alone, thereby circumventing the need for time-consuming algorithms like the one proposed in this paper. The present analysis is therefore not intended

as the basis for a practical algorithm, but rather as a proof of principle to show (indirectly) that a finite temperature 1RSB method can accurately estimate marginal probabilities on problem instances where BP fails to converge.

An open relevant question is whether the retrieval states that are found in the hard regime are purely ferromagnetic, or correspond to an RSB mixed phase. In the former case, the cause of the superior dynamical properties of the survey propagation algorithm is less clear than in the latter case, where there is really a physical explanation. The data seem to suggest (we observed a slight difference between spin glass parameters q_0 and q_1, along with a slightly unstable cavity field distribution on a subset of links) that the instances in which the TABP algorithm does not find a retrieval state, the phase is actually a mixed one, but more analysis is needed to confirm this hypothesis.

Accordingly, the correspondence to results of a macroscopic analysis of the biased Sourlas code within the 1RSB framework may be further specified, thereby more carefully specifying the regime in which finite temperature survey propagation would be an appropriate algorithm. The latter analysis was omitted in the present study because of its numerically extensive nature.

Acknowledgments

The authors would like to thank Joris Mooij and Kees Albers for stimulating discussions. This research was financially supported by the Dutch Technology Foundation (STW).

Appendix A. Free energy expressions in the 1RSB algorithm

The expressions for the generalized free energy $F(m)$ and its derivative $\partial F(m)/\partial m$ are generalized from [15] in the same way as the iteration equations. Instead of a collection of N random samples of fields, constructed by sampling interaction values J and external fields θ at each iteration, we now have a distribution of cavity fields explicitly associated with each link, which is iterated using the actual fixed value of J and θ corresponding to the instance of the graph under consideration.

Distributions of free energy contributions for k interactions in which a site is involved $\left(\Delta F_i^{(1)}\right)$ and for each site linked to k interactions $\left(\Delta F_i^{(2)}\right)$ are constructed through the iteration equations

$$\begin{aligned}
-\beta\Delta F_i^{(1)}\left(\left\{J_\mu, h_{\mu_1}, h_{\mu_2}, h_{\mu_3}\right\}\right) &= \sum_{\mu\in V(i)} \log[\cosh(\beta J_\mu)] \\
&\quad + \sum_{\mu\in V(i)} \log\left[1+\tanh(\beta J_\mu)\tanh\left(\beta h_{\mu_1}\right)\tanh\left(\beta h_{\mu_2}\right)\tanh\left(\beta h_{\mu_3}\right)\right] \\
-\beta\Delta F_i^{(2)}\left(\left\{J_\mu, h_{\mu_1}, h_{\mu_2}\right\}, \theta_i\right) &= \log 2 + \sum_{\mu\in V(i)} \log\left\{\frac{\cosh(\beta J_\mu)}{\cosh\left(\beta u\left(J_\mu, h_{\mu_1}, h_{\mu_2}\right)\right)}\right\} \\
&\quad + \log\cosh\left\{\beta\left[\sum_{\mu\in V(i)} u\left(J_\mu, h_{\mu_1}, h_{\mu_2}\right) + \theta_i\right]\right\}.
\end{aligned} \tag{A.1}$$

The contribution of each site to the total free energy is (taking into account overcounting of link contributions)

$$\Delta F_i = -\tfrac{2}{3}\Delta F_i^{(1)} + \Delta F_i^{(2)} \tag{A.2}$$

in a replica symmetric setting. In the 1RSB framework, these quantities should be averaged over all states, weighted with their Boltzmann factors as in [15] and this result is then averaged over the distribution of free energies. The resulting expressions that are evaluated numerically are

$$f_i = -\tfrac{2}{3} F_i^{(1)} + F_i^{(2)} \tag{A.3}$$

with

$$F_i^{(1)} = -\frac{1}{m\beta} \log \frac{1}{M} \sum_{\alpha} \mathrm{e}^{-m\beta \Delta F_i^{(1)\alpha}} \qquad F_i^{(2)} = -\frac{1}{m\beta} \log \frac{1}{M} \sum_{\alpha} \mathrm{e}^{-m\beta \Delta F_i^{(2)\alpha}}. \tag{A.4}$$

In practice, the above quantities are computed from stable distributions of cavity fields parametrized by sample populations: for each site a sample population of both $\Delta F_i^{(1)}$ and $\Delta F_i^{(2)}$ are constructed according to the above equations. These samples may then be interpreted as the relevant states, and accordingly we label them by α and calculate the averaged free energy according to (A.3) and (A.4).

Appendix B. Cavity analysis of Sourlas code

In the following it will be useful to introduce sublattices: $I_\alpha = \{i|\theta_i = \theta_\alpha\}$ to take into account the fact that the sites are not equivalent due to their biases. The order parameter will thus be the collection of three sublattice distributions of cavity fields, given the message: $W_+(h|\xi)$, $W_-(h|\xi)$ and $W_0(h|\xi)$. The order parameter equations are given by

$$\begin{aligned} W_\alpha(h|\xi) &= \int \prod_{i=1}^{k-1} [\mathrm{d}u_i \, Q_\alpha(u_i|\xi)] \delta \left[h - \sum_{i=1}^{k-1} u_i - \theta_\alpha \right] \\ Q_\alpha(u|\xi) &= \sum_{\alpha',\alpha''} c(\alpha') c(\alpha'') \sum_J \sum_{\xi',\xi''} p(J|\xi,\xi',\xi'') p(\xi',\xi''|\alpha',\alpha'') \\ &\quad \times \int \mathrm{d}h \, \mathrm{d}g \, W_{\alpha'}(h|\xi') W_{\alpha''}(g|\xi'') \delta \left[u - \frac{1}{\beta} \tanh^{-1} [\tanh(\beta h) \tanh(\beta g) \tanh(\beta J)] \right] \end{aligned} \tag{B.1}$$

or

$$\begin{aligned} Q_\alpha(u|\xi) &= \sum_{\alpha',\alpha''} c(\alpha') c(\alpha'') \sum_J \sum_{\xi',\xi''} \frac{\exp(\beta_p [J\xi\xi'\xi'' + \xi'\theta_{\alpha'} + \xi''\theta_{\alpha''}])}{8 \cosh(\beta_p) \cosh(\beta_p \theta_{\alpha'}) \cosh(\beta_p \theta_{\alpha''})} \\ &\quad \times \int \mathrm{d}h \, \mathrm{d}g \, W_{\alpha'}(h|\xi') W_{\alpha''}(g|\xi'') \delta \left[u - \frac{1}{\beta} \tanh^{-1} [\tanh(\beta h) \tanh(\beta g) \tanh(\beta J)] \right]. \end{aligned} \tag{B.2}$$

These equations can be simplified by considering symmetries. First of all, we recognize that $Q_\alpha(u|\xi)$ does not actually depend on α, and obeys $Q(u|-\xi) = Q(-u|\xi)$. From this symmetry it follows that $W_\alpha(h|-\xi) = W_{-\alpha}(-h|\xi)$. As a result we may restrict to considering either $W_\alpha(h|+)$ or $W_\alpha(h|-)$. We will assume that $c(+) = c(-)$ from now on. This gives us

$$\begin{aligned} Q(u|+) &= \sum_J \frac{\mathrm{e}^{\beta_p J}}{2 \cosh(\beta_p)} \\ &\quad \times \int \mathrm{d}h \, \mathrm{d}g \left[c(0) W_0(h|+) + c(+) \frac{\mathrm{e}^{\beta_p \theta}}{\cosh(\beta_p \theta)} W_+(h|+) + c(-) \frac{\mathrm{e}^{-\beta_p \theta}}{\cosh(\beta_p \theta)} W_-(h|+) \right] \end{aligned}$$

$$\times \left[c(0) W_0(g|+) + c(+) \frac{e^{\beta_p \theta}}{\cosh(\beta_p \theta)} W_+(g|+) + c(-) \frac{e^{-\beta_p \theta}}{\cosh(\beta_p \theta)} W_-(g|+) \right]$$
$$\times \delta \left[u - \frac{1}{\beta} \tanh^{-1}[\tanh(\beta h) \tanh(\beta g) \tanh(\beta J)] \right]. \tag{B.3}$$

We are interested in the overlap with the message at the site in question. Therefore, we regard the joint distribution of field and message

$$W_\alpha(H, \xi) = p(\xi|\theta_\alpha) W_\alpha(H|\xi) = p(\xi|\theta_\alpha) \int \prod_{i=1}^{k} [\mathrm{d}u_i Q_\alpha(u_i|\xi)] \delta \left[H - \sum_{i=1}^{k} u_i - \theta_\alpha \right]. \tag{B.4}$$

If the sign of $\xi_i \tanh(\beta H_i)$ is positive, the message at site i will be decoded correctly. Thus, the overall bit-error rate is $(1 - \mu)/2$, where

$$\mu = \frac{1}{2} \sum_{\xi} \sum_{\alpha} c(\alpha) \int \mathrm{d}H W_\alpha(H, \xi) \mathrm{sgn}(\xi H). \tag{B.5}$$

Since $p(\xi|-\theta_\alpha) = p(-\xi|\theta_\alpha)$ and $W_\alpha(H|-\xi) = W_{-\alpha}(-H|\xi)$, we may deduce $W_\alpha(H, -\xi) = W_{-\alpha}(-H, \xi)$ and consequently

$$\mu = \int \mathrm{d}H \left[c(0) W_0(H|+) \right.$$
$$\left. + c(+) \frac{e^{\beta_p \theta}}{\cosh(\beta_p \theta)} W_+(H|+) + c(-) \frac{e^{-\beta_p \theta}}{\cosh(\beta_p \theta)} W_-(H|+) \right] \mathrm{sgn}(H) \tag{B.6}$$

All equations clearly close in terms of $Q(u)$, $W(h)$ and $W(H)$, where

$$\begin{aligned} Q(u) &= Q(u|+) \\ W(h) &= \sum_{\alpha} q(\theta_\alpha) W_\alpha(h|+) \\ W(H) &= \sum_{\alpha} q(\theta_\alpha) W_\alpha(H|+) \end{aligned} \tag{B.7}$$

with the effective measure

$$q(\theta_\alpha) = c(\alpha) \frac{e^{\beta_p \theta_\alpha}}{\cosh(\beta_p \theta_\alpha)} \tag{B.8}$$

leading to equations (26).

References

[1] Pearl J 1988 *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* 2nd edn (San Francisco: Morgan Kaufmann)

[2] Yedidia J S, Freeman W T and Weiss Y 2000 *Advances in Neural Information Processing Systems* vol 13 ed T K Leen, T G Dietterich and V Treps (Cambridge, MA: MIT Press) p 689

[3] Kabashima Y and Saad D 1998 *Europhys. Lett.* **44** 668

[4] Sun J, Li Y, Kang S B and Shum H-Y 2005 *IEEE Computer Society International Conference on Computer Vision and Pattern Recognition*

[5] Mezard M, Parisi G and Zecchina R 2002 *Science* **297** 812

[6] Tanaka K and Morita T 1995 *Phys. Lett.* A **203** 122

[7] Freeman W T, Pasztor E C and Carmichael O T 2000 *Int. J. Comput. Vis.* **40** 25

[8] Gallager R G 1963 *Low-Density Parity Check Codes* (Cambridge, MA: MIT Press)

[9] McElice R, MacKay D and Cheng J 1998 *J. Sel. Commun.* **16** 140

[10] Sigal L, Isard M I, Sigelman B H and Black M J 2004 *Advances in Neural Information Processing Systems* vol 16
[11] Yanover C and Weiss Y 2004 *Advances in Neural Information Processing Systems* vol 16
[12] Stern D H, Graepel T and MacKay D J C 2005 *Advances in Neural Information Processing Systems* vol 17 1353
[13] Kabashima Y 2003 *J. Phys. Soc. Japan* **72** 1645
[14] Braunstein A and Zecchina R 2004 *J. Stat. Mech. (JSTAT)* P06007
[15] Mézard M and Parisi G 2001 *Eur. Phys. J.* B **20** 217
[16] Mézard M and Parisi G 2003 *J. Stat. Phys.* **111** 1
[17] Mezard M and Zecchina R 2002 *Phys. Rev.* E **66** 056126
[18] Braunstein A, Mezard M and Zecchina R 2005 *Random Struct. Algorithms* **27** 201–26
Braunstein A, Mulet R, Pagnani A, Weigt M and Zecchina R 2003 *Phys. Rev.* E **68** 036702
[19] Pretti M 2005 *J. Stat. Mech. (JSTAT)* P11008
[20] Kabashima Y 2005 *J. Phys. Soc. Japan* **74** 2133
Kabashima Y 2005 *Preprint* cs.IT/0506062
[21] Neirotti J P and Saad D 2005 *Europhys. Lett.* **71** 866
[22] Migliorini G and Saad D 2005 *Preprint* cond-mat/0507458
[23] Nishimori H 2001 *Statistical Physics of Spin Glasses and Information Processing* (Oxford: Oxford University Press)
[24] Franz S, Leone M, Montanari A and Ricci-Tersenghi F 2002 *Phys. Rev.* E **66** 046120
[25] Franz S, Leone M, Ricci-Tersenghi F and Zecchina R 2001 *Phys. Rev. Lett.* **87** 127209
[26] Mézard M, Parisi G and Virasoro M A 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
[27] Yedidia J S, Freeman W T and Weiss Y, Constructing free energy approximations and generalized belief propagation algorithms *IEEE Trans. Inform. Theor.* (Also available as www.merl.com/papers/TR2002-35/) at press
[28] De Almeida J R L and Thouless D J 1978 *J. Phys. A: Math. Gen.* **11** 983
[29] Sherrington D and Kirkpatrick S 1975 *Phys. Rev. Lett.* **35** 1792
[30] Mooij J and Kappen H J K 2004 *Advances in Neural Information Processing Systems* vol 16
[31] Ihler A T, Fisher J W and Willsky A S 2005 *J. Mach. Learn. Res.* **6** 905
[32] Mooij J M and Kappen H J 2005 *Proc. 21st Annual Conf. on Uncertainty in Artificial Intelligence (UAI-05)*
[33] Battaglia D, Kolár M and Zecchina R 2004 *Phys. Rev.* E **70** 036107
[34] Kikuchi R 1951 *Phys. Rev.* **81** 988
[35] Pelizzola A 2005 *J. Phys. A: Math. Gen.* **38** R309
[36] Castellani T, Krzakala F and Ricci-Tersenghi F 2005 *Eur. Phys. J.* B **47** 99
[37] Montanari A, Parisi G and Ricci-Tersenghi F 2004 *J. Phys. A: Math. Gen.* **37** 2073
[38] Monasson R 1995 *Phys. Rev. Lett.* **75** 2847
[39] Sourlas N 1989 *Nature* **339** 693
[40] Derrida B 1981 *Phys. Rev.* B **24** 2613
[41] Vicente R, Saad D and Kabashima Y 1999 *Phys. Rev.* E **60** 5352
[42] Kabashima Y and Saad D 2004 *J. Phys.* A **37** R1–43
[43] Nishimori H and Sherrington D 2001 *Disordered and Complex Systems* ed P Sollich, A C C Coolen, L P Hughston and R F Streater (Melville, New York: American Institute of Physics)
[44] Mourik van J 2006 in preparation
[45] Bounkong S, Mourik van J and Saad D 2005 *Preprint* cond-mat/0507579
[46] Heskes T, Albers K and Kappen H J 2003 *Proc. Uncertainty in AI 2003* (San Francisco: Morgan Kaufmann) p 313