Journal of Statistical Mechanics: Theory and Experiment
An IOP and SISSA journal

Characterizing and improving generalized belief propagation algorithms on the 2D Edwards–Anderson model

Eduardo Domínguez[1], Alejandro Lage-Castellanos[1], Roberto Mulet[1], Federico Ricci-Tersenghi[2] and Tommaso Rizzo[3]

[1] Department of Theoretical Physics and 'Henri-Poincaré-Group' of Complex Systems, Physics Faculty, University of Havana, La Habana, CP 10400, Cuba
[2] Dipartimento di Fisica, INFN—Sezione di Roma 1 and CNR—IPCF, UOS di Roma, Università La Sapienza, Piazzale Aldo Moro 5, 00185 Roma, Italy
[3] Dipartimento di Fisica and CNR—IPCF, UOS di Roma, Università La Sapienza, Piazzale Aldo Moro 5, 00185 Roma, Italy
E-mail: eduardo@fisica.uh.cu, lage@fisica.uh.cu, mulet@fisica.uh.cu, federico.ricci@roma1.infn.it and tommaso.rizzo@inwind.it

Received 7 October 2011
Accepted 17 November 2011
Published 12 December 2011

Online at stacks.iop.org/JSTAT/2011/P12007
doi:10.1088/1742-5468/2011/12/P12007

Abstract. We study the performance of different message passing algorithms in the two-dimensional Edwards–Anderson model. We show that the standard belief propagation (BP) algorithm converges only at high temperature to a paramagnetic solution. Then, we test a generalized belief propagation (GBP) algorithm, derived from a cluster variational method (CVM) at the plaquette level. We compare its performance with BP and with other algorithms derived under the same approximation: double loop (DL) and a two-way message passing algorithm (HAK). The plaquette-CVM approximation improves BP in at least three ways: the quality of the paramagnetic solution at high temperatures, a better estimate (lower) for the critical temperature, and the fact that the GBP message passing algorithm converges also to nonparamagnetic solutions. The lack of convergence of the standard GBP message passing algorithm at low temperatures seems to be related to the implementation details and not to the

1742-5468/11/P12007+23$33.00

appearance of long range order. In fact, we prove that a gauge invariance of the constrained CVM free energy can be exploited to derive a new message passing algorithm which converges at even lower temperatures. In all its region of convergence this new algorithm is faster than HAK and DL by some orders of magnitude.

Keywords: cavity and replica method, disordered systems (theory), message-passing algorithms, statistical inference

ArXiv ePrint: 1110.1259

Contents

1. Introduction

The 2D Edwards–Anderson (EA) model in statistical mechanics is defined by a set $\sigma = \{s_1 \cdots s_N\}$ of N Ising spins $s_i = \pm 1$ placed on the nodes of a 2D square lattice, and random interactions $J_{i,j}$ at the edges, with a Hamiltonian

$$\mathcal{H}(\sigma) = -\sum_{\langle i,j \rangle} J_{i,j} s_i s_j,$$

where $\langle i, j \rangle$ runs over all couples of neighboring spins (first neighbors on the lattice). The $J_{i,j}$ are the magnetic interchange constants between spins and are supposed fixed for any given instance of the system, and the spins s_i are the dynamic variables. We will focus on one of the most common disorder types, the bimodal interactions $J = \pm 1$ with equal probabilities.

The statistical mechanics of the EA model, at a temperature $T = 1/\beta$, is given by the Gibbs–Boltzmann distribution

$$P(\sigma) = \frac{\mathrm{e}^{-\beta \mathcal{H}(\sigma)}}{Z} \qquad \text{where } Z = \sum_{\sigma} \mathrm{e}^{-\beta \mathcal{H}(\sigma)}.$$

The direct computation of the partition function Z, or any marginal probability distribution like $p(s_i, s_j) = \sum_{\sigma \setminus s_i, s_j} P(\sigma)$ is a time consuming task, unattainable in

J. Stat. Mech. (2011) P12007

general, and therefore an approximation is required. We are interested in fast algorithms for inferring such marginal distributions. Actually for the 2D EA model, thanks to the graph planarity, algorithms computing Z in a time polynomial in N exist. However we are interested in very fast (i.e. linear in N) algorithms that can be used also for more general models, e.g. the EA model in a field or defined on a 3D cubic lattice. For these more general cases a polynomial algorithm is very unlikely to exist and some approximations are required.

A simple and effective mean field approximation is the one due to Bethe [1], in which the marginals over the dynamic variables, like $p(s_i)$, are obtained from the minimization of a variational free energy in a self-consistent way. The Bethe approximation is exact for a model without loops in the interaction network, which unfortunately is far from being the usual case in physics. In the context of finite dimensional lattices, Kikuchi [2] derived an extension of this approximation to larger groups of variables, which accounts for short loops exactly, and is usually referred to as the cluster variational method (CVM).

The interest in spin glasses, with quenched random disorder, brought a new testing ground for both approximations. In particular, the Bethe approximation (exact on trees) has been the starting point of many useful theoretical and applied developments. It is at the basis of the cavity method, which allows a restatement of replica theory in probabilistic terms for finite connectivity systems [3]. The Bethe approximation is connected to well known algorithms in computer science, namely belief propagation [4] and the sum–product algorithm [5]. A major achievement of this confluence between computer science and statistical mechanics has been the conception of the survey propagation algorithm [6, 7], inspired by the cavity method and the replica symmetry breaking [3, 8, 9], that shows great performance on hard optimization problems [6, 7, 10, 11]. Statistical mechanics clarified the relation between phase transitions and easy–hard transitions in optimization problems, and allowed the statistical characterization of the onset of the hard phase [12]–[14], as well as the analytical description of search algorithms based on BP [15, 16].

The correctness of the Bethe approximation and the related algorithms is, however, linked to the lack of topological correlations in the interactions (random graphs are locally tree like), since the approximation is exact only on tree topologies. This is a strong limitation for physical purposes, since tree topologies or random graphs are not the common situation. The Bethe approximation performs poorly in finite dimensional lattices, and the associated algorithms are usually nonconvergent at low temperatures.

Recently the cluster variational method (CVM) has been reformulated in a broader probabilistic framework called *region-based* approximations to free energy [17] and connected to a generalized belief propagation (GBP) algorithm to find the stationary points of the free energy. It extends the Bethe approximation by considering correlations in larger regions, allowing, in principle, short loops to be taken into account accurately. In [17] it was shown that stable fixed points of the GBP message passing algorithm correspond to stationary points of the approximated CVM free energy, while the converse is not necessarily true. Furthermore, the GBP message passing is not guaranteed to converge at all. Prompted by this lack of convergence, a new kind of provably convergent algorithms for minimizing the CVM approximated free energy, known as double loop (DL) algorithms [18, 19], has been developed, at the cost of a drastic drop off in speed.

GBP has been applied in the last decade to inference problems [20]–[22], consistently outperforming BP. In particular, the image reconstruction problems [20, 23] are based on

J. Stat. Mech. (2011) P12007

a 2D lattice structure, but, at variance with 2D EA model, the interactions among nearby spins (pixels) are ferromagnetic, and the damaged image is used as an external field. Both factors help convergence of GBP algorithms. An analysis of the CVM approximation using GBP algorithms on single instances of finite dimensional disordered models of physical interest, like the EA model, has not been made so far.

The Edwards–Anderson model in 2D has been largely studied by other methods (see [24, 25] and references therein), suggesting that it remains paramagnetic all the way down to zero temperature, lacking any thermodynamic transition at any finite T, although at low T there are metastable states of very long lifetime, leading to very slow dynamics. Based on this fact, a paramagnetic version of the GBP on the 2D EA model was studied recently in [26]. The modern approach to the CVM was introduced in [27] and the following discussion may be better understood in connection with that paper; we also mention a related scheme developed in [28]. On the other hand a combination of the replica method and the CVM has been recently presented in [29], opening the way to the introduction of a generalized survey propagation algorithm. However, the implementation of the latter algorithm on finite dimensional lattices is computationally very demanding, and should be preceded by the study of the original CVM approximation and GBP algorithm. The first application of the CVM to random systems and spin glasses was presented in [30]–[32]; see [29] for a discussion of those early results and more recent studies.

In this paper we study the convergence properties of the GBP message passing algorithm and the performance of the CVM approximation on the 2D EA model. After the introduction of the region-based free energy in section 2 and the message passing algorithm in terms of cavity fields, we compute the critical (inverse) temperature $T_{\rm CVM} \simeq 0.82$ ($\beta_{\rm CVM} \simeq 1.22$) of the plaquette-CVM approximation in section 3, improving the Bethe estimate $T_{\rm Bethe} = 1.51$ ($\beta_{\rm Bethe} \simeq 0.66$) by roughly a factor of 2. The CVM average case temperature, however, does not clearly correspond to the single instance behavior of the GBP message passing algorithm, as is shown in section 4. At variance with belief propagation, GBP converges to spin glass solutions (below $T_{\rm SG} \simeq 1.27$, above $\beta_{\rm SG} \simeq 0.79$), and stops converging near $T \simeq 1.0$, before the average case prediction $T_{\rm CVM}$. In section 5 we show that this convergence problem depends on the implementation details of the message passing algorithm, and can be improved by a simultaneous update of message. In order to do so the gauge invariance of the message passing equations has to be fixed. In section 6 we compare the solutions and the performance of GBP with three other algorithms for the minimization of the CVM free energy: double loop [19], two-way message passing [19], and the dual algorithm [26]. In terms of the CVM free energy, the paramagnetic solution is in general the one to be chosen, except for a small interval in temperatures where the spin glass solution has a lower free energy. Our results are summarized in section 7.

2. Generalized belief propagation on EA 2D

Given that a detailed derivation of the plaquette-GBP message passing equations for the 2D Edwards–Anderson model was presented in [26], here we only summarize this derivation, skipping unnecessary details.

J. Stat. Mech. (2011) P12007

The idea of the *region-based* free energy approximation [17, 33] is to mimic the exact (Boltzmann–Gibbs) distribution $P(\sigma)$ by a reduced set of its marginals. A hierarchy of approximations is given by the size of such marginals, starting with the set of all single spin marginals $p_i(s_i)$ (mean field), then following with all neighboring site marginals $p(s_i, s_j)$ (Bethe approximation), then to all square plaquette marginals $p(s_i, s_j, s_k, s_l)$, and so on. Since the only way of knowing such marginals exactly is the unattainable computation of Z, the method pretends to approximate them by a set of beliefs $b_i(s_i)$, $b_L(s_i, s_j)$, $b_{\mathcal{P}}(s_i, s_j, s_k, s_l)$, etc obtained from a minimization of a region-based free energy.

Following the derivation in [26], the plaquette level approximated free energy for the 2D EA model is given as the contribution of all plaquettes, links and spins in the 2D lattice:

$$\begin{aligned} -\beta F = &\sum_{\mathcal{P}} \sum_{\sigma_{\mathcal{P}}} b_{\mathcal{P}}(\sigma_{\mathcal{P}}) \log \frac{b_{\mathcal{P}}(\sigma_{\mathcal{P}})}{\exp(-\beta E_{\mathcal{P}}(\sigma_{\mathcal{P}}))} \qquad \text{plaquettes} \\ &- \sum_{L} \sum_{\sigma_L} b_L(\sigma_L) \log \frac{b_L(\sigma_L)}{\exp(-\beta E_L(\sigma_L))} \qquad \text{links} \\ &+ \sum_{i} \sum_{s_i} b_i(s_i) \log \frac{b_i(s_i)}{\exp(-\beta E_i(s_i))} \qquad \text{spins,} \end{aligned} \tag{1}$$

where the symbol $\sigma_R = (s_1, \ldots, s_k)$ stands for the set of spins in region R, while $E_R(\sigma_R) = -\sum_{\langle i,j \rangle \in R} J_{i,j} s_i s_j$ stands for the energy contribution in that region. The energy term $E_i(s_i)$ in the spin contribution is only relevant when an external field acts over the spins, and will be neglected from now on.

An unrestricted minimization of the free energy (1) in terms of its beliefs produces incongruent results. Beliefs are only meaningful as an approximation to the correct marginals if they obey the marginalization constraints $b_i(s_i) = \sum_{s_j} b_L(s_i, s_j)$ and $b_L(s_i, s_j) = \sum_{s_k, s_l} b_P(s_i, s_j, s_k, s_l)$. This marginalization is enforced by the introduction of Lagrange multipliers (see [17] for a general introduction, and [26] for this particular case) in the free energy expression. There is one Lagrange multiplier $\mu_{L \to i}(s_i)$ for every link L and spin $i \in L$, and a Lagrange multiplier $\nu_{\mathcal{P} \to L}(s_i, s_j)$ for each plaquette $\mathcal{P}$ and link $L \in \mathcal{P}$. In terms of these Lagrange multipliers, the stationary condition of the approximated free energy is achieved with

$$\begin{aligned} b_i(s_i) &= \frac{1}{Z_i} \exp\left(-\beta E_i(s_i) - \sum_{L \supset i}^{4} \mu_{L \to i}(s_i) \right), \\ b_L(\sigma_L) &= \frac{1}{Z_L} \exp\left(-\beta E_L(\sigma_L) - \sum_{\mathcal{P} \supset L}^{2} \nu_{\mathcal{P} \to L}(\sigma_L) - \sum_{i \subset L}^{2} \sum_{\substack{L' \supset i \\ L' \neq L}}^{3} \mu_{L' \to i}(s_i) \right), \\ b_{\mathcal{P}}(\sigma_{\mathcal{P}}) &= \frac{1}{Z_{\mathcal{P}}} \exp\left(-\beta E_{\mathcal{P}}(\sigma_{\mathcal{P}}) - \sum_{L \subset \mathcal{P}}^{4} \sum_{\substack{\mathcal{P}' \supset L \\ \mathcal{P}' \neq \mathcal{P}}}^{1} \nu_{\mathcal{P}' \to L}(\sigma_L) - \sum_{i \subset \mathcal{P}}^{4} \sum_{\substack{L \supset i \\ L \not\subset \mathcal{P}}}^{2} \mu_{L \to i}(s_i) \right). \end{aligned} \tag{2}$$

J. Stat. Mech. (2011) P12007

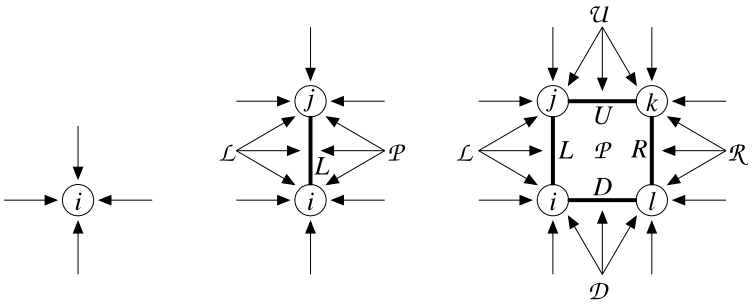

Figure 1. Schematic representation of the belief equations (2). Lagrange multipliers are depicted as arrows, going from parent regions to child regions. Italic capital letters are used to denote plaquettes, simple capital letters denote links, and lower case letters denote spins.

A graphical representation of these equations is given in figure 1. Lagrange multipliers are shown as arrows going from parent regions to children. Take, for example, the middle equation for the belief in link regions $b_L(\sigma_L) = b_L(s_i, s_j)$. The sum of the two Lagrange multipliers $\nu_{\mathcal{P}\to L}(s_i, s_j)$ corresponds to the triple arrows on both sides of the link in the middle of figure 1, while the two sums over three messages $\mu_{L'\to i}(s_i)$ correspond to the three arrows acting over the top (j) and bottom (i) spins, respectively. In equations (2), the Z_R are normalization constants. The terms $E_{\mathcal{P}}(\sigma_{\mathcal{P}}) = E_{\mathcal{P}}(s_i, s_j, s_k, s_l) = -(J_{i,j}s_is_j + J_{j,k}s_js_k + J_{k,l}s_ks_l + J_{l,i}s_ls_i)$ and $E_L(s_i, s_j) = -J_{i,j}s_is_j$ are the corresponding energies in plaquettes and links respectively, and are represented in the diagram by the lines (interactions) between circles (spins); zero since no field is acting upon the spins.

The Lagrange multipliers can be parametrized in terms of the cavity fields u and (U, u_a, u_b) as

$$-\mu_{L\to i}(s_i) = \beta u_{L\to i}\, s_i, \tag{3}$$

$$-\nu_{\mathcal{P}\to L}(s_i, s_j) = \beta(U_{\mathcal{P}\to L}\, s_is_j + u_{\mathcal{P}\to i}\, s_i + u_{\mathcal{P}\to j}\, s_j). \tag{4}$$

In particular, the field $u_{L\to i}$ corresponds to the cavity field in the Bethe approximation [17]. The choice of these parametrizations is the reason for the use of single and triple arrows in figures 1 and 2. In particular, the messages going from plaquettes to links are characterized by three fields $(U_{\mathcal{P}\to L}, u_{\mathcal{P}\to i}, u_{\mathcal{P}\to j})$, and the capital $U_{\mathcal{P}\to L}$ acts as an effective interaction term.

The Lagrange multipliers are related among them by the constraints they are supposed to impose (see [26]). In terms of the cavity fields and using the notation in figure 2, link-to-spin cavity fields will be related by

$$u_{L\to i} = \hat{u}(u_{\mathcal{P}\to i} + u_{\mathcal{L}\to i}, U_{\mathcal{P}\to L} + U_{\mathcal{L}\to L} + J_{ij}, u_{\mathcal{P}\to j} + u_{\mathcal{L}\to j} + u_{A\to j} + u_{B\to j} + u_{U\to j}), \tag{5}$$

where

$$\hat{u}(u, U, h) \equiv u + \frac{1}{2\beta}\log\frac{\cosh\beta(U+h)}{\cosh\beta(U-h)}.$$

Note that the usual cavity equation for fields in the Bethe approximation [3] is recovered if all contributions from plaquettes $\mathcal{P}$ and $\mathcal{L}$ are set to zero.

J. Stat. Mech. (2011) P12007

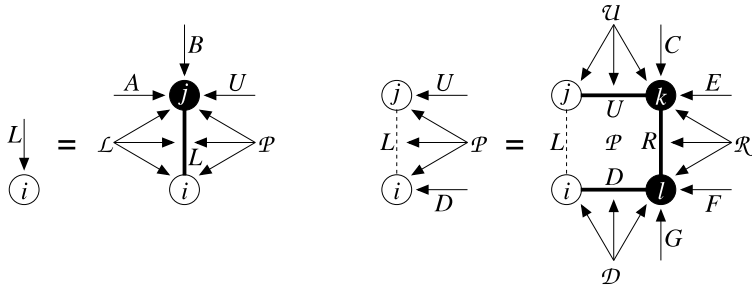

Figure 2. The message passing equations (5) and (6), shown schematically. Messages are depicted as arrows, going from parent regions to child regions. On any link $J_{i,j}$, represented as a bold line between spins (circles), a Boltzmann factor $e^{\beta J_{i,j} s_i s_j}$ exists. Dark circles represent spins to be traced over. Messages from plaquettes to links $\nu_{P\to L}(s_i, s_j)$ are represented by triple arrows, because they can be written in terms of three parameters U, u_i and u_j, defining the correlation $\langle s_i s_j\rangle$ and magnetizations $\langle s_i\rangle$ and $\langle s_j\rangle$, respectively.

Similarly, by imposing the marginalization of the beliefs at plaquettes onto their child links, we find the self-consistent expression for the plaquette-to-link cavity fields:

$$\begin{aligned} U_{\mathcal{P}\to L} &= \hat{U}(\#) = \frac{1}{4\beta}\log\frac{K(1,1)K(-1,-1)}{K(1,-1)K(-1,1)},\\ u_{\mathcal{P}\to i} &= -u_{D\to i} + \hat{u}_i(\#) = u_{\mathcal{D}\to i} - u_{D\to i} + \frac{1}{4\beta}\log\frac{K(1,1)K(1,-1)}{K(-1,1)K(-1,-1)},\\ u_{\mathcal{P}\to j} &= -u_{U\to j} + \hat{u}_j(\#) = u_{\mathcal{U}\to j} - u_{U\to j} + \frac{1}{4\beta}\log\frac{K(1,1)K(-1,1)}{K(1,-1)K(-1,-1)}, \end{aligned} \tag{6}$$

where

$$\begin{aligned} K(s_i,s_j) = \sum_{s_k,s_l} \exp[\beta((&U_{\mathcal{U}\to U} + J_{jk})s_j s_k + (U_{\mathcal{R}\to R} + J_{kl})s_k s_l + (U_{\mathcal{D}\to D} + J_{li})s_l s_i\\ &+ (u_{\mathcal{U}\to k} + u_{C\to k} + u_{E\to k} + u_{\mathcal{R}\to k})s_k + (u_{\mathcal{R}\to l} + u_{F\to l} + u_{G\to l} + u_{\mathcal{D}\to l})s_l)] \end{aligned}$$

and the symbol # stands for all incoming fields on the right-hand side of the equations. The functions $\hat{u}(u, U, h)$ and $[\hat{U}(\#), \hat{u}_i(\#), \hat{u}_j(\#)]$ will be used in section 3 for the average case calculation.

For a given system of size N (number of spins) there are $2N$ links and N square plaquettes, and therefore there are $4N$ plaquette-to-link fields $[U_{\mathcal{P}\to L}, u_{\mathcal{P}\to i}, u_{\mathcal{P}\to j}]$, and $4N$ link-to-spin fields $u_{L\to i}$. At the stationary points of the free energy their values are related by the set of $4N + 4N$ equations (5) and (6).

The set of $4N + 4N$ self-consistent equations is also called the message passing equations when they are used as update rules for fields in the message passing algorithm, or cavity iteration equations in the context of cavity calculations. The field notation is more comprehensible than the original Lagrange multiplier notation, and has a clear physical meaning: each plaquette is telling its child links that they should add an effective interaction term $U_{P\to L}$ to the direct interaction $J_{i,j}$, due to the fact that the spins s_i and s_j are also interacting through the other three links in the plaquette. Terms u_i act like

J. Stat. Mech. (2011) P12007

magnetic fields upon spins, and the complete $\nu(s_i, s_j)$-message is characterized by the triplet $(U_{i,j}, u_i, u_j)$.

3. Critical temperature of the plaquette-CVM approximation

In this section we revisit the method used in [29] to compute the critical temperature at which the CVM approximation develops a spin glass phase. By a spin glass phase we mean a phase characterized by nonzero local magnetizations $m_i = \tanh(\beta \sum_L^4 u_{L\to i})$ and nearly zero total magnetization $m = (1/N)\sum_i m_i \simeq 0$ (remember that we work with no external field). The 2D EA model is paramagnetic down to zero temperature, but spin glass like solutions can appear in the CVM approximation due to its mean field character. We correct one of the conclusions reached in [29], where we fail to observe the appearance of the spin glass phase in the CVM approximation to the 2D Edwards–Anderson model. We follow an average case approach, which is similar in spirit but different from the single instance stability analysis made in [34] for the Bethe approximation (belief propagation).

The average case calculation is a mathematical technique developed in [3] to study the typical solutions of cavity equations in disordered systems, with a deep and fundamental connection to the replica trick [9]. When applied to the plaquette-CVM approximation [29], we end up with two equations, in which fields (messages) are now replaced by functions of fields $q(u)$ and $Q(U, u_1, u_2)$, and the interactions are averaged out. As a consequence of the homogeneity of the 2D lattice and the averaging over local disorder $J_{i,j}$, all plaquettes, links, and spins in the graph are now equivalent, and we only need to study one of them to characterize the whole system.

More precisely, the average case self-consistent equation for the distribution $q(u)$ is given by

$$\begin{aligned} q(u_i) = \mathbb{E}_J \int & \mathrm{d}q\,(u_{A\to j})\,\mathrm{d}q\,(u_{B\to j})\,\mathrm{d}q\,(u_{U\to j}) \\ & \times \mathrm{d}Q\,(U_{\mathcal{P}\to L}, u_{\mathcal{P}\to i}, u_{\mathcal{P}\to j})\,\mathrm{d}Q\,(U_{\mathcal{L}\to L}, u_{\mathcal{L}\to i}, u_{\mathcal{L}\to j}) \\ & \times \delta(u_i - \hat{u}(\#)), \end{aligned} \tag{7}$$

with $\hat{u}(\#)$ as defined in the right-hand side of equation (5), and $\mathrm{d}f(x) \equiv f(x)\,\mathrm{d}x$.

The corresponding self-consistent equation for $Q(U, u_1, u_2)$ is

$$\begin{aligned} \int\int Q(U, u_a, u_b) q(u_i - u_a) q(u_j - u_b)\,\mathrm{d}u_a\,\mathrm{d}u_b & \\ = \mathbb{E}_J \int & \mathrm{d}q\,(u_{C\to k})\,\mathrm{d}q\,(u_{E\to k})\,\mathrm{d}q\,(u_{F\to l})\,\mathrm{d}q\,(u_{G\to l})\,\mathrm{d}Q\,(U_{\mathcal{U}\to U}, u_{\mathcal{U}\to j}, u_{\mathcal{U}\to k}) \\ & \times \mathrm{d}Q\,(U_{\mathcal{R}\to R}, u_{\mathcal{R}\to k}, u_{\mathcal{R}\to l})\,\mathrm{d}Q\,(U_{\mathcal{D}\to D}, u_{\mathcal{D}\to l}, u_{\mathcal{D}\to i}) \\ & \times \delta(U - \hat{U}(\#))\,\delta(u_i - \hat{u}_i(\#))\,\delta(u_j - \hat{u}_j(\#)), \end{aligned} \tag{8}$$

where the notation corresponds to equation (6). In both equations (7) and (8) the expression $\mathbb{E}_J = \int \mathrm{d}J\,P(J)\cdots$ stands for the average over the quenched randomness.

At high temperatures we expect the fixed point equations (5) and (6) to yield a paramagnetic solution. Such a solution is characterized by link-to-site messages $u = 0$, and plaquette-to-link messages $(U, u_1, u_2) = (U, 0, 0)$. If we impose this ansatz on the fields, we recover the paramagnetic or dual algorithm of [26] for the single instance message

J. Stat. Mech. (2011) P12007

passing, and the paramagnetic average case study of [29] for the average case. Let us remember that the 2D EA model is expected to have no thermodynamic transition at any finite temperature, and hence it remains paramagnetic all the way down to $T = 0$. Following [29], in the average case the paramagnetic solution has the form

$$q(u) = \delta(u), \qquad Q(U, u_1, u_2) = Q(U)\delta(u_1)\delta(u_2).$$

Equation (7) is always satisfied when $q(u) = \delta(u)$, whatever $Q(U)$. Equation (8) can be solved self-consistently for $Q(U)$:

$$Q(U) = \mathbb{E}_J \int \mathrm{d}Q\,(U_{\mathcal{U}})\,\mathrm{d}Q\,(U_{\mathcal{R}})\,\mathrm{d}Q\,(U_{\mathcal{D}}) \times \delta\left(U - \frac{1}{\beta}\mathrm{arctanh}[\tanh\beta(J_{\mathcal{U}} + U_{\mathcal{U}})\tanh\beta(J_{\mathcal{R}} + U_{\mathcal{R}})\tanh\beta(J_{\mathcal{D}} + U_{\mathcal{D}})]\right) \tag{9}$$

and the average free energy and all other relevant functions can be derived in terms of it (see [29]).

On the other hand, a general (not paramagnetic) solution of the average case equations (7) and (8) is very difficult, since it involves the deconvolution of distributions $q(u)$ on the left-hand side of equation (8) in order to update $Q(U, u_1, u_2)$ by an iterative method. A critical temperature can be found, however, by an expansion in small u around the paramagnetic solution. We can focus on the second moments of the distributions

$$a = \int q(u)u^2\,\mathrm{d}u,$$

$$a_{i\,j}(U) = \int\int Q(U, u_1, u_2)\,u_i\,u_j\,\mathrm{d}u_1\,\mathrm{d}u_2, \qquad \text{where } i, j \in \{1, 2\},$$

and check whether the paramagnetic solution ($a = 0$ and $a_{ij}(U) = 0$) is locally stable. To do this we expand equations (7) and (8) to second order, and we obtain the following linearized equations:

$$a = K_{a,a}a + \int \mathrm{d}U'\,K_{a,a_{11}}(U')a_{11}(U') + \int \mathrm{d}U'\,K_{a,a_{12}}(U')a_{12}(U'),$$

$$a\,Q(U) + a_{11}(U) = K_{a_{11},a}(U)a + \int \mathrm{d}U'\,K_{a_{11},a_{11}}(U, U')a_{11}(U') + \int \mathrm{d}U'\,K_{a_{11},a_{12}}(U, U')a_{12}(U'),$$

$$a_{12}(U) = K_{a_{12},a}(U)a + \int \mathrm{d}U'\,K_{a_{12},a_{11}}(U, U')a_{11}(U') + \int \mathrm{d}U'\,K_{a_{12},a_{12}}(U, U')a_{12}(U').$$

The actual values of the K_{a_x,a_y} come from the expansion in small u of the original equations (see equation (90) in [29] for an example).

We cannot solve these equations analytically because we do not have an analytical expression of $Q(U)$ for the paramagnetic solution at all temperatures. By discretizing the values of U uniformly in $(-U_{\max}, U_{\max})$, i.e. $U = i\Delta U$ with $i \in [-I_{\max}, I_{\max}]$, we can transform the continuous set of equations to a system of the form

$$\vec{a} = \mathbf{K}(\beta) \cdot \vec{a}, \tag{10}$$

J. Stat. Mech. (2011) P12007

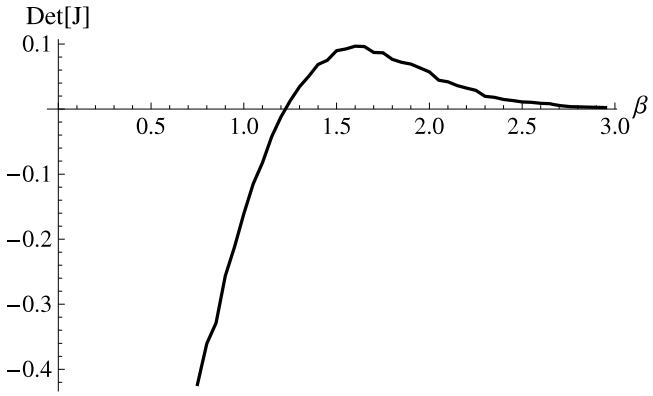

Figure 3. Determinant of the Jacobian $\mathbf{J} = \mathbf{I} - \mathbf{K}(\beta)$ as a function of the inverse temperature β. The critical inverse temperature is $\beta_{\rm CVM} \simeq 1.22$.

where the vector of the second moments $\vec{a} = (a, a_{11}(U), a_{12}(U))$ has the form

$$\vec{a} = (a, a_{11}(-U_{\max}), a_{11}(-U_{\max} + \Delta U), \ldots, a_{11}(U_{\max} - \Delta U), a_{11}(U_{\max}), \\ a_{12}(-U_{\max}), a_{12}(-U_{\max} + \Delta U), \ldots, a_{12}(U_{\max} - \Delta U), a_{12}(U_{\max})).$$

$\mathbf{K}(\beta)$ is a $(2I_{\max} + 1) \times (2I_{\max} + 1)$ matrix, that stands for the discrete representation of the integrals on the right-hand side of the linearized equations, and depends on the inverse temperature via the solution $Q(U)$ of equation (9).

The paramagnetic solution $\vec{a} = 0$ always satisfies the homogeneous equation (10). The stability criterion for the paramagnetic solution is the singularity of the Jacobian $\det(\mathbf{I} - \mathbf{K}(\beta)) = 0$. When such a condition is satisfied, a nonparamagnetic solution continuously arises from the paramagnetic one, since a flat direction appears in the free energy.

Numerically, we worked with a discretization of $2I_{\max} + 1 = 41$ points between $(-U_{\max} = -3.5, U_{\max} = 3.5)$. The paramagnetic solution $Q(U)$ is found solving equation (9) by an iterative method at every temperature, and then used to compute the elements of the $\mathbf{K}(\beta)$ matrix. In figure 3 we show the determinant of the Jacobian matrix $\mathbf{J} = \mathbf{I} - \mathbf{K}(\beta)$. The critical inverse temperature derived from this analysis is $\beta_{\rm CVM} \simeq 1.22$ for the appearance of a flat direction in the free energy.

In [29] $\beta_{\rm CVM}$ was thought to be infinite (zero temperature) because an incomplete range of the values of β was examined. The critical temperature found here is below the Bethe critical temperature $\beta_{\rm Bethe} \simeq 0.66$, and therefore improves the Bethe approximation by roughly a factor of 2, since the 2D EA model is likely to remain paramagnetic at all finite temperatures. At variance with the Bethe approximation, the single instance behavior of the message passing is not so clearly related to the average case critical temperature, as we show in section 4.

4. The performance of GBP on the 2D EA model

Before studying GBP message passing for the plaquette-CVM approximation, let us check what happens to the simpler Bethe approximation and the corresponding message passing algorithm known as belief propagation (BP) in the 2D EA model. When running BP at

J. Stat. Mech. (2011) P12007

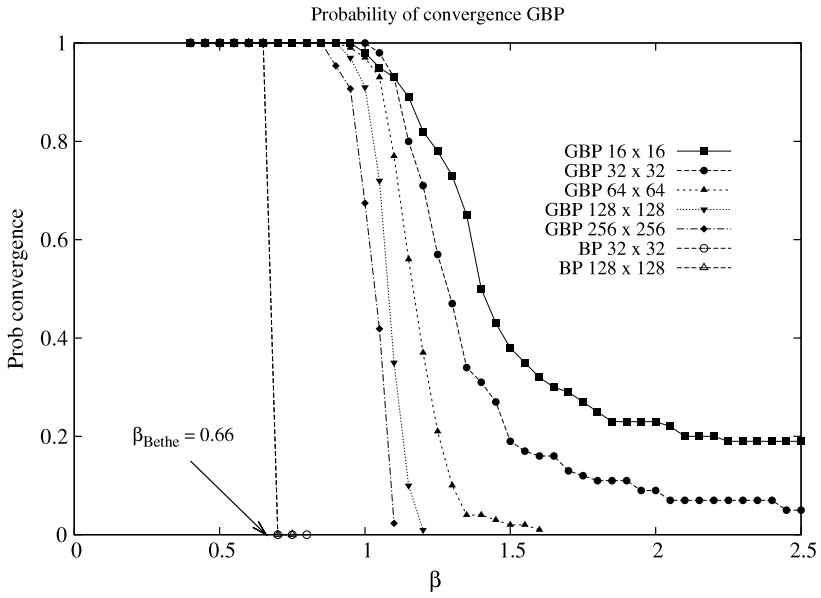

Figure 4. Probability of convergence of BP and GBP on a 2D EA model, with random bimodal interactions, as a function of the inverse temperature $\beta = 1/T$. The Bethe spin glass transition is expected to occur at $\beta_{\text{Bethe}} \simeq 0.66$ on a random graph with the same connectivity. The BP message passing algorithm on the 2D EA model stops converging very close to that point. Above that temperature, BP equations converge to the paramagnetic solution, i.e. all messages are trivial, $u = 0$. Below the Bethe temperature (nearly) the Bethe instability takes messages away from the paramagnetic solution, and the presence of short loops is thought to be responsible for the lack of convergence. On the other hand, the GBP equations converge at lower temperatures, but eventually stop converging as well.

high temperatures (above $T_{\text{Bethe}} = 1/\beta_{\text{Bethe}} \simeq 1.51$) in a typical instance of the model with bimodal interactions, we find the paramagnetic solution (given by all fields $u = 0$), and, therefore, the system is equivalent to a set of independent interacting pairs of spins, which is only correct at infinite temperature. The Bethe temperature T_{Bethe} (computed in the average case and exact on acyclic graphs[4]), seems to mark precisely the point where BP stops converging (see figure 4). Indeed messages flow away from zero below T_{Bethe}, and convergence of the BP message passing algorithm is not achieved anymore. So, the Bethe approximation is disappointing when applied to single instances of the Edwards–Anderson model: either it converges to a paramagnetic solution at high temperatures, or it does not converge at all below T_{Bethe}.

The natural question arises as to what extent the GBP message passing algorithm for the plaquette-CVM approximation is also nonconvergent below its critical temperature, and whether this temperature coincides with the average case one. To check this we used GBP message passing equations (5) and (6), with a damping factor of 0.5 in the

[4] The Bethe temperature T_{Bethe} is the one at which a nontrivial spin glass solution appears for a random regular Bethe lattice with connectivity $K = 4$. The Bethe lattice looks locally like a tree.

J. Stat. Mech. (2011) P12007

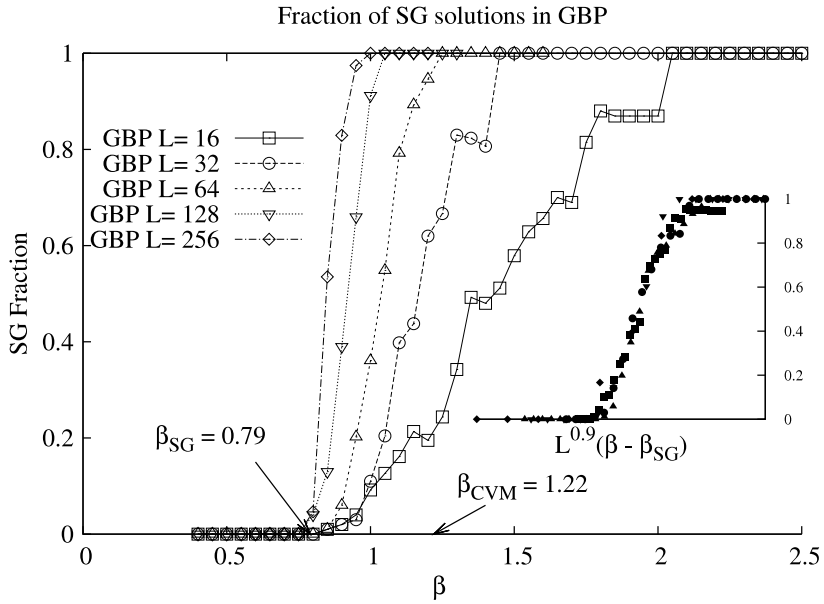

Figure 5. The data points correspond to the fraction of SG solutions in a population of 100 systems of sizes 16^2, 32^2, 64^2, 128^2, 256^2 respectively. At high temperatures (low β) GBP message passing converges always to the paramagnetic solution. The average case critical inverse temperature $\beta_{\rm CVM} \simeq 1.22$ does not corresponds to the single instance behavior, as the spin glass solutions in GBP appear around $\beta_{\rm SG} \simeq 0.79$. The inset shows that all the data collapsed if plotted as a function of the scaling variable $L^{0.9}(\beta - 0.79)$, where the exponent 0.9 and the critical inverse temperature $\beta_{\rm SG} \simeq 0.79$ are obtained from best data collapse.

link-to-site fields u:

$$u^{\rm new}_{L\to i} = 0.5\, u^{\rm old}_{L\to i} + 0.5\, \hat{u}(\#).$$

We will make the distinction between two types of solution for the GBP algorithm. The high temperature or paramagnetic solution is characterized by zero local magnetization of spins, $m_i = \sum_{s_i} s_i b_i(s_i) = \tanh(\beta \sum_L^4 u_{L\to i}) = 0$. At low temperatures, following the average case analysis, a nonparamagnetic or spin glass solution should appear, characterized by nonzero local magnetizations, but roughly null global magnetization. The temperature at which nonzero local magnetizations appear will be called $T_{\rm SG} = 1/\beta_{\rm SG}$.

Figure 4 shows that GBP is able to converge below the Bethe critical temperature, but stops converging before the CVM average case critical temperature $\beta_{\rm CVM} \simeq 1.22$. Furthermore, figure 5 shows that even before it stops converging, GBP finds a spin glass solution in most instances.

The inset to figure 5 shows a collapse of the data points for different system sizes using the scaling variable $L^{0.9}(\beta - 0.79)$, which gives an estimate $\beta_{\rm SG} \simeq 0.79$ (the exponent of 0.9 is obtained from the best data collapse). Since $\beta_{\rm SG} \simeq 0.79$ is well below the average case inverse critical temperature $\beta_{\rm CVM} \simeq 1.22$, the relevance of the latter to the behavior of GBP on single samples is questionable. By a similar data collapse procedure, we estimate the nonconvergence temperature for the GBP message passing algorithm to be $\beta_{\rm conv} \simeq 0.96$ (see figure 9), which is again far away from the average case prediction $\beta_{\rm SG}$.

J. Stat. Mech. (2011) P12007

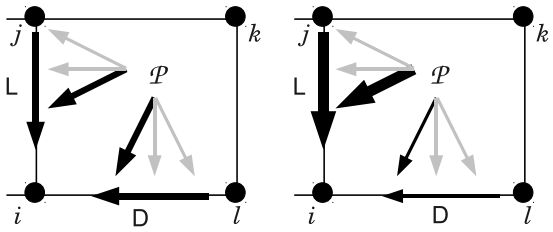

Figure 6. Null modes of the plaquette-CVM free energy in terms of fields. The small-u fields that act over a given spin i inside a plaquette can be shifted by an arbitrary amount δ as in equation (11) without changing the self-consistent (message passing) equations.

So, beyond the simple Bethe approximation, we found three different temperatures in the CVM approximation: $\beta_{\mathrm{SG}} \simeq 0.79 < \beta_{\mathrm{conv}} \simeq 0.96 < \beta_{\mathrm{CVM}} \simeq 1.22$, corresponding respectively to the appearance of spin glass solutions, to the lack of convergence on single instances, and to the average case prediction for the critical temperature.

We can summarize three main differences between the properties of BP and GBP. At high temperatures (below $\beta_{\mathrm{SG}} \simeq 0.79$) GBP gives a quite good approximation of the marginals [26], namely the paramagnetic solution with nontrivial correlation fields $U \neq 0$, while BP treats the system as a set of independent pairs of linked spins. Furthermore, this naive approach is all that BP can do for us, since above $\beta_{\mathrm{Bethe}} \simeq 0.66$ it no longer converges. GBP, on the other hand, is not only able to converge beyond β_{Bethe}, but it is also able to find spin glass solutions above β_{SG}. The third difference between the two algorithms is that the nonconvergence of BP seems to occur exactly at the same temperature where a spin glass phase should appear (and arguably because of it), while the GBP convergence problems appear deep into the spin glass phase. The lack of convergence of GBP, however, seems to depend strongly on the implementation details, as we show next.

5. Gauge invariance of the GBP equations

The convergence properties of the GBP message passing are sensitive to the implementation details, e.g. the damping value in the update equations, and this is not an inherent property of the CVM (or region-graph) approximation. We might try, for instance, to update simultaneously all *small-u* fields pointing towards a given spin, hoping to gain some more stability in the message passing algorithm. When trying to do this we find out that there is a freedom in the choice of these fields that has no effect over the fixed point solutions. This freedom (similar to the one noticed in [35]) is the result of having introduced unnecessary Lagrange multipliers to enforce marginalization constraints that were already indirectly enforced.

Consider, for instance, the messages shown in figure 6. If the belief on a plaquette $b_P(s_i, s_j, s_k, s_l)$ correctly marginalizes to the beliefs of two of its child links $b_L(s_i, s_j)$ and $b_D(s_l, s_i)$, and one of those beliefs marginalizes to the common spin $b_i(s_i) = \sum_{s_j} b_L(s_i, s_j)$, it is inevitable that the second link D also marginalizes to the same belief on s_i, since $b_i(s_i) = \sum_{s_j} b_L(s_i, s_j) = \sum_{s_j, s_l, s_k} b_{\mathcal{P}}(s_i, s_j, s_k, s_l) = \sum_{s_l} b_D(s_l, s_i)$. Therefore the Lagrange multiplier that was introduced to force this last marginalization is not needed. This

J. Stat. Mech. (2011) P12007

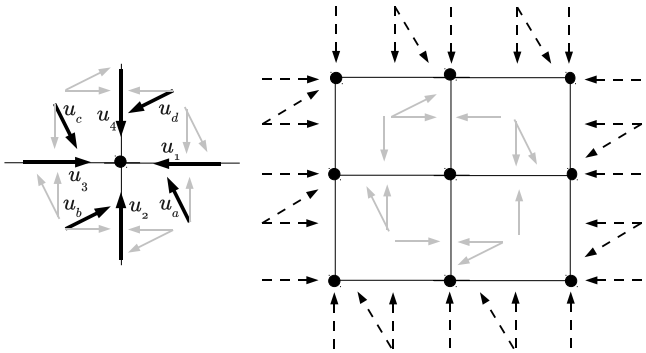

Figure 7. In the left diagram, all eight *small-u* messages pointing to the central spin are highlighted with bold face. They are four link-to-site u-messages, and four plaquette-to-link u_{left}-messages. They have linear dependence among them. The right diagram shows four plaquettes around a spin, and the messages that contribute in a nonlinear way to the aforementioned eight messages. The idea of GBP + GF is to compute the nonlinear contributions to the message passing equations, and then assign the values of the u-messages in order to satisfy their linear relations.

redundancy is a general feature of GBP equations when there are more than two levels of regions (plaquettes, links, and spins in our case).

The consequence of having introduced unnecessary multipliers is a gauge invariance on the field (message) values. Such an invariance can be better understood by looking at the GBP equations at infinite temperature: for $\beta = 0$ the nonlinear parts of the message passing equations (5) and (6) disappear, but there is still a set of linear equations to be satisfied for the small-u messages with infinitely many nontrivial solutions. These solutions correspond, however, to the same physical paramagnetic solution, since the total field $h_i = \sum_L^4 u_{L\to i}$ and the magnetizations $m_i = \tanh(\beta h_i)$ are always zero. It is easy to check that once we have a solution of the message passing equations (5) and (6) at any temperature, we can change by an arbitrary amount δ any group of four u-messages inside a plaquette (figure 6) pointing to the same spin as

$$\begin{aligned} u_{L\to i} &\to u_{L\to i} + \delta, \qquad & u_{\mathcal{P}_L\to i} &\to u_{\mathcal{P}_L\to i} + \delta, \\ u_{D\to i} &\to u_{D\to i} - \delta, \qquad & u_{\mathcal{P}_D\to i} &\to u_{\mathcal{P}_D\to i} - \delta, \end{aligned} \tag{11}$$

and still all self-consistent equations are satisfied.

This local null mode of the standard GBP equations can be avoided by arbitrarily setting to zero one of the four small-u fields entering equation (11). We choose to fix the gauge by removing the right small-u field in every plaquette-to-link field $(U, u_{\text{left}}, u_{\text{right}})$, as shown in figure 7. Once the gauge is fixed, the fields are uniquely determined, and we can try to implement the simultaneous updating of all *small-u* fields around a given spin, hopefully improving the convergence.

In the left diagram of figure 7 all messages involving the central spin are represented, and in bold face those that act precisely upon that spin. These messages enter linearly in the message passing equations of each other (see equations (5) and (6)). Therefore, the self-consistent equations they should satisfy at the fixed points can be written as (using

J. Stat. Mech. (2011) P12007

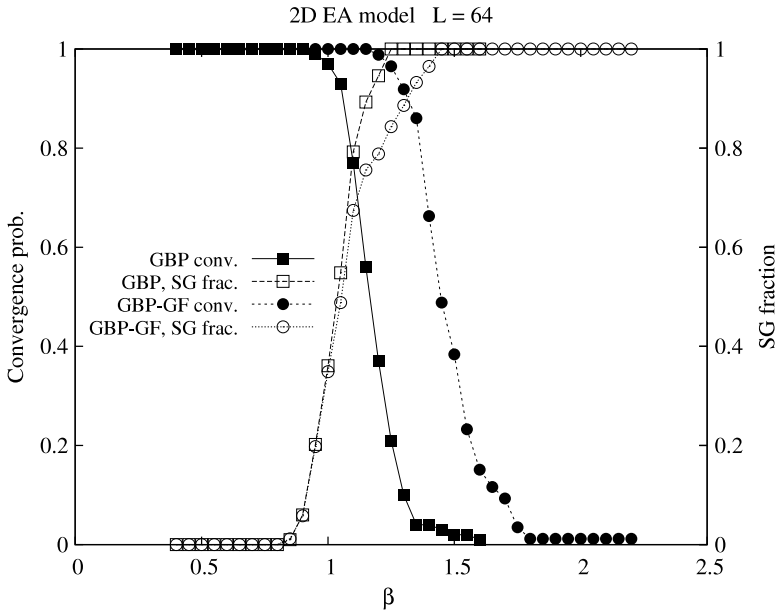

Figure 8. Convergence probability of GBP and GBP + GF as a function of β. The solution found by either iteration method is always the same (when both converge), but GBP + GF reaches lower temperatures while converging. The fraction of spin glass solutions found by either algorithm shows that GBP + GF sees the same spin glass transition temperature. The fraction of spin glass solutions is always given with respect to the amount of convergent solutions.

the notation of figure 7)

$$\begin{aligned} u_1 &= u_a + \mathrm{NL}_1, & u_a &= u_b - u_2 + \mathrm{NL}_a, \\ u_2 &= u_b + \mathrm{NL}_2, & u_b &= u_c - u_3 + \mathrm{NL}_b, \\ u_3 &= u_c + \mathrm{NL}_3, & u_c &= u_d - u_4 + \mathrm{NL}_c, \\ u_4 &= u_d + \mathrm{NL}_4, & u_d &= u_a - u_1 + \mathrm{NL}_d, \end{aligned} \tag{12}$$

where the NL stand for the nonlinear contributions to the corresponding equation. As a consequence, the values of the eight u-messages pointing to the central spin can be assigned precisely by a linear transformation for any given values of the nonlinear contributions. This gauge fixed updating method, that we will call GBP + GF, updates all u-messages around a spin simultaneously and in such a way that they are consistent with each other via the message passing equations.

The right diagram in figure 7 shows the messages entering the nonlinear parts. Taking the eight u-messages as zero, the nonlinear contributions are the right-hand sides of the message passing equations involved. With the nonlinear parts computed, the system of equations (12) is solved for the u-variables multiplying the nonlinearity vector by the corresponding matrix. The eight u-messages are then updated, usually with a damping factor. The update of the U correlation fields is carried out as in the original GBP method, via equation (6), since it does not depend on the u-messages that are being updated.

Figure 8 shows the probability of convergence versus the inverse temperature for GBP and GBP + GF, and also the fraction of the solutions found that correspond to

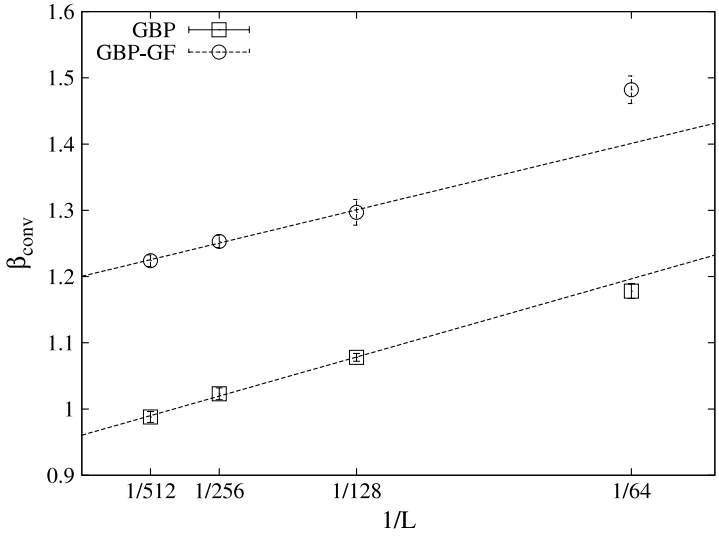

Figure 9. Estimate of the nonconvergence temperature for different system sizes using standard GBP (squares) and gauge fixed GBP (circles). As shown, with the gauge fixed procedure the nonconvergence extrapolated temperature is quite close to the average case prediction, $\beta_{\mathrm{CVM}} \simeq 1.22$. Each data point corresponds to the average of the nonconvergence temperature over many realizations of the disorder: 10 realizations for the 512×512 systems, 20 for the 256×256 systems, and 100 for the others.

a spin glass phase. Let us emphasize here that GBP and GBP + GF are not different approximations, but different methods to find the same fixed point solution by message passing. They are expected to find the same solutions, and in fact they do. At high temperatures both methods converge to the paramagnetic solution, with all null local magnetizations $m_i = \tanh(\beta \sum_L^4 u_{L\to i}) = 0$. The standard message passing update of the GBP equations hardly converges above $\beta_{\mathrm{conv}} \simeq 0.96$, while the GBP + GF method reaches lower temperatures, $\beta_{\mathrm{conv-GF}} \simeq 1.2$, as can be seen in figure 9. Furthermore, the GBP + GF allows us to work in a range of temperatures where most solutions are spin glass like. This proves that the nonconverging temperature found for GBP, $\beta_{\mathrm{conv}} \simeq 0.96$, is not a feature of the CVM approximation, but a characteristic of the message passing method used, and can be outperformed by other message passing schemes, like GBP + GF. Note in figure 9 that the nonconvergence inverse temperature of GBP + GF, $\beta_{\mathrm{conv-GF}} \simeq 1.2$, is quite close to the average case prediction for the critical temperature, $\beta_{\mathrm{CVM}} \simeq 1.22$. Whether this is accidental or not is still unclear. Since the average case instability should describe the breakdown of the paramagnetic phase, and the lack of convergence in single instances occurs while already in a nonparamagnetic phase, it seems far fetched assuming that both critical behaviors are related.

5.1. Gauge fixed average case stability

The disagreement between the average case critical temperature, β_{CVM}, and the one observed in the single instance, β_{SG}, can be due to a number of reasons. First, the average case calculation assumes that the cavity fields are uncorrelated. However, in our

J. Stat. Mech. (2011) P12007

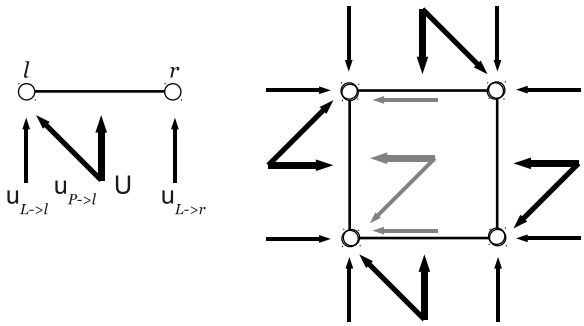

Figure 10. Left: the set of four messages that we compute jointly by population dynamics. Right: the population dynamic step consists in taking four quadruplets at random from the population (those in black), and computing a new quadruplet (the one in gray inside the plaquette) using randomly selected interactions J_{ij} on the plaquette.

case, messages participating in the cavity iteration are very close to each other in the lattice, and thus correlated. Furthermore, GBP does not have the equivalent of a Bethe lattice for BP, i.e. a model in which the correlation between cavity messages is close to zero by construction. The second reason for a failure of the average case prediction is that the transition we observe in single instances might be due to the almost inevitable appearance of ferromagnetic domains in large systems (Griffith instability). The third, and most obvious, reason is that the gauge invariance was not accounted for in the average case calculation.

Reproducing the method of section 3 to obtain an average case prediction of the critical temperature for gauge fixed GBP is not straightforward. The reason is that the link-to-spin messages u should fulfil two different equations: their own original equation (5), and the implicit equation derived from the fact that the gauge is fixed and one of the fields in the plaquette-to-link message (U, u, u) is set to zero.

However, a different average case calculation is possible. We can represent the messages flowing in the lattice by a population of quadruplets $(u_{L_l\to l}, u_{\mathcal{P}\to l}, U_{\mathcal{P}\to lr}, u_{L_r\to r})$, where one of the original messages is absent because the gauge has been fixed (see the left panel in figure 10). Given any four of these quadruplets of messages around a plaquette, we can compute, using the message passing equations, the new messages inside the plaquette (see the right panel in figure 10). The new population dynamics consists in picking four of these quadruplets out of the population at random, then computing the new quadruplet (using also random interactions in the plaquette) and finally putting it back in the population. After several steps, the population stabilizes either to a paramagnetic solution (where all $u = 0$ and only $U \neq 0$), or to a nonparamagnetic one (where also $u \neq 0$).

In figure 11 we show the Edwards–Anderson order parameter $q_{\mathrm{EA}} = \sum_i m_i^2/N$ obtained at different temperatures using this population dynamics average case method. We find that q_{EA} becomes larger than zero at $\beta_{\mathrm{CVM-GF}} \simeq 0.81$, which is quite close to the inverse temperature $\beta_{\mathrm{SG}} \simeq 0.79$ where single instances develop nonzero local magnetizations and a spin glass phase. The correspondence between this average case result and the single instance behavior is very enlightening; indeed the average case

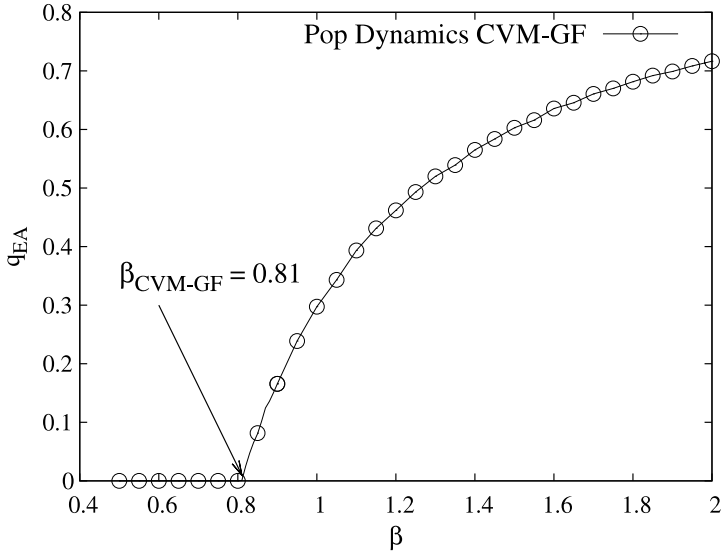

Figure 11. The Edwards–Anderson order parameter, see equation (13), obtained using a population of $N = 10^3$ messages, and running the population dynamics step $10^3 \times N$ times. In agreement with the single instance behavior, the transition between paramagnetic ($q_{\text{EA}} = 0$) and nonparamagnetic (spin glass) phases is found at $\beta \simeq 0.81$.

computation does not take into account correlations among quadruplets of messages and it is not sensitive to Griffith's singularities. So, the simplest explanation for the GBP + GF behavior on single samples of the 2D EA model is that quadruplets of messages arriving on any given plaquette are mostly uncorrelated and that at β_{SG} a true spin glass instability takes place (which is an artifact of the mean field like approximation). Please consider that under the Bethe approximation the SG instability occurs at $\beta_{\text{Bethe}} \simeq 0.66$, while the CVM approximation improves the estimate of the SG critical boundary to $\beta_{\text{SG}} \simeq 0.79$ (on single instances) and to $\beta_{\text{CVM-GF}} \simeq 0.81$ (in the average case).

6. Same approximation, four algorithms

It can be proved [17] that stable fixed points of the message passing equations correspond to stationary points of the region graph approximated free energy (or CVM free energy). The converse is not necessarily true, and some of the stationary points of the free energy might not be stable under the message passing heuristic. As we have seen, the message passing might not even converge at all. For a given free energy approximation (equation (1) in our case), there are other algorithms to search for stationary points, including other types of message passing and provably convergent algorithms. In this section we study two of these algorithms and show that they do find the same spin glass like transition at β_m, but have a different behavior at lower temperatures.

The one presented so far is the so called parent-to-child (PTC) message passing algorithm, in which Lagrange multipliers are introduced to force marginalization of larger (parent) regions onto their children. Other choices of Lagrange multipliers are possible [17], leading to the so called child-to-parent and two-way algorithms. Next we

J. Stat. Mech. (2011) P12007

test the following four algorithms for minimizing the plaquette-CVM free energy in typical instances of 2D EA.

- The double loop algorithm of Heskes *et al* [19] is a provably convergent algorithm that guarantees a step by step minimization of the free energy functional. It consists of two loops, the inner of which is a two-way message passing algorithm that we will call HAK. We use the implementation in the LibDai public library [36].
- The HAK message passing algorithm is a two-way message passing algorithm [19]. When it converges, it is usually faster than double loop.
- GBP parent-to-child is the message passing algorithm we have presented so far in this paper, and for which the simultaneous updating of cavity fields was introduced to help convergence. Nevertheless the following results were obtained using standard GBP PTC.
- The dual algorithm of [26] is the same GBP PTC setting all small fields $u = 0$, and carrying out message passing only in terms of correlation fields U (the first equation in (6)).

For the last three algorithms we use our own implementation in terms of cavity fields u and (U, u_a, u_b). The dual algorithm forces the solution of GBP to remain paramagnetic since all $u = 0$. This paramagnetic ansatz is especially suited for the 2D EA model since it is expected to be paramagnetic at any finite temperature (in the thermodynamical limit).

As shown in section 5.1, the GBP PTC message passing equation finds a paramagnetic solution in the 2D EA model at high temperatures, while below $T_{\rm SG} = 1/\beta_{\rm SG} \simeq 1.27$ it finds a spin glass like solution. By spin glass like we mean that the total field $h_i = \sum_L^4 u_{L\to i}$ and the magnetization $m_i = \tanh(\beta h_i)$ are nonzero and change from spin to spin. The order parameter

$$q_{\rm EA} = \frac{1}{N}\sum_i m_i^2 \tag{13}$$

is used to locate this phase. The critical temperature $T_{\rm SG}$, where $q_{\rm EA}$ becomes larger than zero, seems to be independent of message passing details, like damping or the use of gauge fixing for simultaneous updates of fields.

In figure 12 we show the free energy and the $q_{\rm EA}$ parameter of the solutions found by double loop, HAK and GBP PTC for two typical realizations of an $N = 16\times16$ EA system with bimodal interactions. The free energy of the dual approximation is subtracted to highlight the differences with respect to the paramagnetic solution. The figure shows that HAK and double loop do find the same spin glass solution as GBP PTC finds when going down in temperature. This solution is actually lower in free energy when it appears, but at even lower temperatures becomes subdominant compared to the paramagnetic one. The GBP PTC keeps finding the spin glass solutions while double loop and HAK switch back to the paramagnetic one. This is an interesting feature of double loop and in particular of HAK which is a fast message passing algorithm. By returning to the dual (paramagnetic) solution, HAK is also ensuring its convergence at low temperature [26], while GBP PTC gets lost in the irrelevant (and physically wrong) spin glass solution, and eventually stops converging.

However note that DL and HAK may stop finding the SG solution when this solution is still the one with lower free energy. Moreover, the lack of convergence of GBP can

J. Stat. Mech. (2011) P12007

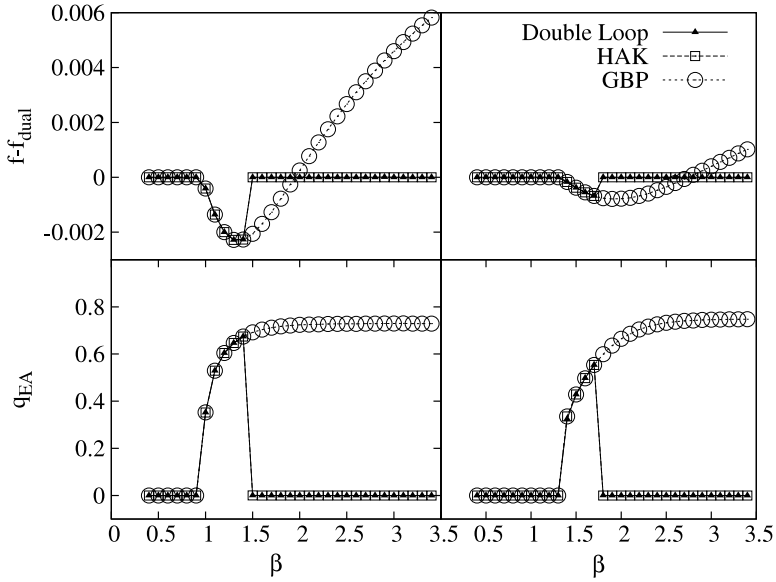

Figure 12. The free energy of the solutions found by the double loop algorithm, the HAK and the GBP PTC algorithm relative to the free energy of the paramagnetic solution (dual approximation), in a typical system in which GBP PTC finds a spin glass solution. At high temperatures all three algorithms find the same paramagnetic solution. Interestingly, there is a small range of temperatures where the spin glass solution found by GBP is actually the one that minimizes the free energy. However, at even lower temperatures the paramagnetic solution becomes again the correct one. While double loop and HAK switch back to the paramagnetic solution (even if at the wrong T), the GBP PTC gets stuck in the spin glass solution (and for this reason, it eventually stops converging).

be used as a warning that something wrong is happening with the CVM approximation, something that it is impossible to understand by looking at the behavior of provably convergent algorithms.

In figure 13 we compare the running times of double loop (LibDai [36]), HAK and GBP PTC (our implementation) for the two systems of figure 12. As expected, double loop is much slower than the message passing heuristics of HAK and GBP (please notice the log scale in the time axis). The peaks in the running times correspond to the transition points from the paramagnetic to the spin glass solution. Double loop and HAK have two peaks, the second corresponding to the transition back to the paramagnetic solution, while the GBP PTC has only the first peak.

7. Summary and conclusions

We studied the properties of the generalized belief propagation algorithm derived from a cluster variational method approximation to the free energy of the Edwards–Anderson model in 2D at the level of plaquettes. We compared the results obtained by parent-to-child GBP with the ones obtained by the dual (paramagnetic) algorithm [26] and by the HAK two-way algorithm [19] and the double loop provably convergent algorithm [19].

J. Stat. Mech. (2011) P12007

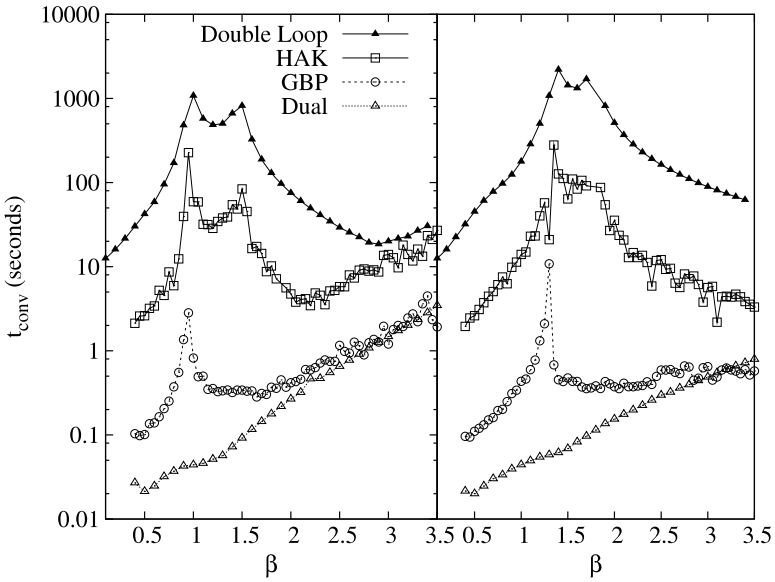

Figure 13. Convergence time in seconds for the double loop algorithm (full points) and standard message passing algorithms (empty points) for the plaquette-GBP approximation in two different realizations of a 16^2 Edwards–Anderson system. Message passing algorithms are typically faster, but not always convergent. The first cusp is related to the appearance of the spin glass solution, while the second cusp in the double loop algorithm is related to the switching back to the paramagnetic solution (see figure 12).

We found that the plaquette-CVM approximation (using parent-to-child GBP) is far richer than the Bethe (BP) approximation in the 2D EA model. BP converges only at high temperatures (above $T_{\rm Bethe} = 1/\beta_{\rm Bethe} = 1.51$), and in such a case it treats the system as a set of independent pairs of linked spins. GBP, on the other hand, makes a better prediction of the paramagnetic behavior of the model at high T, since it implements a message passing of correlation fields flowing from plaquettes to links in the graph. Furthermore with GBP the paramagnetic phase is extended to temperatures below $T_{\rm Bethe} = 1.51$ until $T_{\rm SG} = 1/\beta_{\rm SG} \simeq 1.27$, where spin glass solutions appear in the single instance implementation of the message passing algorithm. In contrast to the Bethe approximation, GBP is able to find spin glass solutions, and the standard message passing stops converging near $T_{\rm conv} \simeq 1$.

The average case calculation of the stability of the paramagnetic solution in the CVM approximation predicted that nonparamagnetic (spin glass) solutions should appear at lower temperatures $T_{\rm CVM} = 1/\beta_{\rm CVM} \simeq 0.82$. This average case result does not coincide with the single instance behavior of the standard GBP since it fails to mark both the point where GBP starts finding spin glass solutions, $T_{\rm SG}$, and the point where GBP stops converging, $T_{\rm conv}$.

However, the nonconvergence of GBP is not a feature of the CVM approximation, and is susceptible to changes from one implementation of the message passing to another. We showed that by fixing a hidden gauge invariance in the message passing equation, a simultaneous update of all cavity fields pointing to a single spin in the lattice improves the

convergence of the algorithm, without changing its speed drastically. Using gauge fixed GBP, the nonconvergence inverse temperature is moved to $T_{\mathrm{conv-GF}} \simeq 1.2$, quite close to the average case prediction T_{CVM} (whether this is only a coincidence is still not clear). Most importantly, the average case computation (population dynamics) with the gauge fixed identifies the same SG critical temperature $T_{\mathrm{CVM-GF}} \simeq 1.28$ measured on single samples (where $T_{\mathrm{SG}} \simeq 1.27$).

Finally we compared the fixed point solutions found by the GBP message passing with those found by the provably convergent double loop algorithm and the message passing heuristic of the two-way algorithm of [19]. All the algorithms find the same paramagnetic solutions at high T, while below T_{SG} they find a spin glass solution, in the sense that the local magnetizations are nonzero while the global magnetization is null. On decreasing the temperature double loop and HAK switch back from the spin glass to the paramagnetic solution, at the cost of factors of 10^2–10^3 and 10–10^2 respectively in running time, compared to GBP. Furthermore, the paramagnetic solution can always be found rapidly by the dual algorithm of [26], making these two algorithms (double loop and HAK) unnecessarily slow.

Although the thermodynamics of the 2D EA model is paramagnetic at low temperatures, the correlation length grows until eventually surpassing $L/2$ and therefore being effectively infinite for any finite size 2D system. In such a situation the nonparamagnetic solutions obtained by GBP can account for long range correlations, and presumably give better estimates for the correlations among spins than the paramagnetic solution obtained by HAK and double loop.

Establishment of the previous claim requires a detailed study of the quality of the CVM approximation at low temperatures (in the nonparamagnetic range) and its connections to the statics and dynamics of the 2D Edwards–Anderson model, which is already under study. Application of the CVM and GBP message passing to the Edwards–Anderson model in 3D is also appealing, since this model does have a spin glass behavior at low temperature.

References

[1] Bethe H A, 1935 *Proc. R. Soc.* A **150** 552
[2] Kikuchi R, 1951 *Phys. Rev.* **81** 988
[3] Mézard M and Parisi G, 2001 *Eur. Phys. J.* B **20** 217
[4] Pearl J, 1988 *Probabilistic Reasoning in Intelligent Systems* 2nd edn (San Francisco, CA: Morgan Kaufmann)
[5] Kschischang F R, Frey B J and Loeliger H, 1998 *IEEE Trans. Inf. Theory* **47** 498
[6] Mézard M and Zecchina R, 2002 *Phys. Rev.* E **66** 056126
[7] Mézard M, Parisi G and Zecchina R, 2002 *Science* **297** 812
[8] Mézard M and Parisi G, 2003 *J. Stat. Phys.* **111** 1
[9] Parisi G, Mézard M and Virasoro M A, 1987 *Spin Glass Theory and Beyond* (Singapore: World Scientific)
[10] Mulet R, Pagnani A, Weigt M and Zecchina R, 2002 *Phys. Rev. Lett.* **89** 268701
[11] Braunstein A, Mulet R, Pagnani A, Weigt M and Zecchina R, 2003 *Phys. Rev.* E **68** 036702
[12] Achlioptas D, Naor A and Peres Y, 2005 *Nature* **435** 759
[13] Krzakala F, Montanari A, Ricci-Tersenghi F, Semerjian G and Zdeborova L, 2007 *Proc. Nat. Acad. Sci.* **104** 10318
[14] Montanari A, Ricci-Tersenghi F and Semerjian G, 2008 *J. Stat. Mech.* P04004
[15] Montanari A, Ricci-Tersenghi F and Semerjian G, 2007 *Proc. 45th Annual Allerton Conf. on Communication, Control, and Computing* pp 352–9
[16] Ricci-Tersenghi F and Semerjian G, 2009 *J. Stat. Mech.* P09001
[17] Yedidia J, Freeman W T and Weiss Y, 2005 *IEEE Trans. Inf. Theory* **51** 2282

doi:10.1088/1742-5468/2011/12/P12007

J. Stat. Mech. (2011) P12007

[18] Yuille A L, 2002 *Neural Comput.* **14** 1691
[19] Heskes T, Albers C and Kappen H, 2003 *Proc. Conf. on Uncertainty in Artificial Intelligence, 2003* ed C Meek and U Kjærulff (San Francisco: Morgan Kaufmann) p 313
[20] Tanaka K, Inoue J and Titterington D M, 2003 *13th Proc. of the 2003 IEEE Signal Processing Society Workshop (17–19 Sept., 2003)* (Los Alamitos, CA: IEEE Computer Society Press) pp 329–38
[21] Albers C A, Heskes T and Kappen H J, 2007 *Genetics* **177** 1101
[22] Kappen H J, 2002 *Modeling Bio-Medical Signals* (Singapore: World Scientific) pp 3–16
[23] Tanaka K and Morita T, 1995 *Phys. Lett.* A **203** 122
[24] Jorg T, Lukic J, Marinari E and Martin O, 2006 *Phys. Rev. Lett.* **96** 237205
[25] Thomas C K and Middleton A A, 2009 *Phys. Rev.* E **80** 046708
[26] Lage-Castellanos A, Mulet R, Ricci-Tersenghi F and Rizzo T, 2011 *Phys. Rev.* E **84** 046706
[27] Morita T, 1984 *J. Stat. Phys.* **34** 319
[28] Morita T, 1989 *Physica* A **155** 73
[29] Rizzo T, Lage-Castellanos A, Mulet R and Ricci-Tersenghi F, 2010 *J. Stat. Phys.* **139** 375
[30] Katsura S and Fujiki S, 1980 *J. Phys. C: Solid State Phys.* **13** 4711
[31] Fujiki S and Katsura S, 1980 *J. Phys. C: Solid State Phys.* **13** 4723
[32] Katsura S and Nagahara I, 1980 *J. Phys. C: Solid State Phys.* **13** 4995
[33] Pelizzola A, 2005 *J. Phys. A: Math. Gen.* **38** R309
[34] Mooij J M and Kappen H J, 2005 *J. Stat. Mech.* P11012
[35] Weigt M, White R A, Szurmant H, Hoch J A and Hwa T, 2003 *Proc. Nat. Acad. Sci.* **106** 67
[36] Mooij J M, 2010 *J. Mach. Learn. Res.* **11** 2169

J. Stat. Mech. (2011) P12007