Inference algorithm for finite-dimensional spin glasses: Belief propagation on the dual lattice

Alejandro Lage-Castellanos and Roberto Mulet
Department of Theoretical Physics and Henri-Poincaré Group of Complex Systems, Physics Faculty, University of Havana, La Habana, Codigo Postal 10400, Cuba

Federico Ricci-Tersenghi
Dipartimento di Fisica, INFN–Sezione di Roma 1, and CNR–IPCF, UOS di Roma, Università La Sapienza, Piazzale Aldo Moro 5, I-00185 Roma, Italy

Tommaso Rizzo
Dipartimento di Fisica and CNR–IPCF, UOS di Roma, Università La Sapienza, Piazzale Aldo Moro 5, I-00185 Roma, Italy
(Received 18 February 2011; revised manuscript received 15 September 2011; published 24 October 2011)

Starting from a cluster variational method, and inspired by the correctness of the paramagnetic ansatz [at high temperatures in general, and at any temperature in the two-dimensional (2D) Edwards-Anderson (EA) model] we propose a message-passing algorithm—the dual algorithm—to estimate the marginal probabilities of spin glasses on finite-dimensional lattices. We use the EA models in 2D and 3D as benchmarks. The dual algorithm improves the Bethe approximation, and we show that in a wide range of temperatures (compared to the Bethe critical temperature) our algorithm compares very well with Monte Carlo simulations, with the double-loop algorithm, and with exact calculation of the ground state of 2D systems with bimodal and Gaussian interactions. Moreover, it is usually 100 times faster than other provably convergent methods, as the double-loop algorithm. In 2D and 3D the quality of the inference deteriorates only where the correlation length becomes very large, i.e., at low temperatures in 2D and close to the critical temperature in 3D.

DOI: 10.1103/PhysRevE.84.046706 PACS number(s): 05.10.−a, 75.50.Lk

I. INTRODUCTION

Given a joint probability distribution $P(\sigma)$ over a large number N of variables, $\sigma \equiv \{s_1, \ldots, s_N\}$, a common inference problem is that of computing marginal probabilities, like the following [1]:

$$p_i(s_i) \equiv \sum_{\sigma \setminus s_i} P(\sigma).$$

In the general (and the most interesting) case, this problem cannot be solved exactly in a time growing subexponentially with the size N. We have to use some kind of approximation in order to compute marginals in a time growing linearly with N. The approximation schemes used so far are mainly adopted from the field of statistical mechanics, where mean-field-like approximations are standard and well-controlled tools [2].

We are interested in models with disorder, and we will focus on spin glasses in the present paper. When working with an ensemble of models or problems, the results obtained by the statistical mechanics tools refer to average quantities, i.e., those of the typical samples. In other words one is not concerned with the behavior of a specific sample, but rather one looks at the whole ensemble. On the contrary when doing inference one is interested in a single sample and thus the above approximation schemes need to be converted into an algorithm that can be run on such a specific sample. Computing marginals on a given sample clearly gives more information than computing averages over the ensemble.

To our knowledge, effective (i.e., linear time) algorithms for computing marginals can be essentially divided into two broad classes: stochastic local search algorithms, e.g. Monte Carlo Markov chains [3], which roughly sample the configuration space according to $P(\sigma)$, and algorithms based on some kind of mean-field approximation, e.g., belief propagation [4]. The former are exact in the long run, but the latter can be much more useful if an approximated answer is required in a short time. Unfortunately, the latter also have some additional drawbacks due to the mean-field nature of the underlying approximation, e.g., spurious phase transitions, that may prevent the proper convergence of the algorithm.

One more reason why this latter class of algorithms has a limited scope of application is that the convergence of the algorithm may strongly depend on the presence of short loops in the network of interactions between variables. In this sense the successful application of one of these algorithms to models defined on regular lattices (which have many short loops) would be a major achievement.

In this paper we introduce a fast algorithm for computing marginals in two- and three-dimensional spin glass models defined by the Hamiltonian (further details are given below)

$$\mathcal{H}(\sigma) = -\sum_{\langle i,j \rangle} J_{ij} s_i s_j. \tag{1}$$

The first nontrivial mean-field approximation for the above model corresponds to the Bethe-Peierls approximation scheme and the belief propagation (BP) algorithm. Unfortunately when the BP algorithm is run on a given spin glass sample defined on a D-dimensional lattice, it seems to provide exactly the same output as if it were run on a random regular graph with fixed degree $2D$: that is, for $T \gtrsim T_{\text{Bethe}}$ it converges to a solution with all null local marginals ($\langle s_i \rangle = 0$), while for $T \lesssim T_{\text{Bethe}}$ it does not converge to a fixed point.

The next step in the series of mean-field approximations (also known as Kikuchi approximations or the cluster variational method) is to consider joint probability distributions

of the four spins belonging to the same plaquette. Under this approximation an algorithm has been derived in [5] which is called generalized belief propagation (GBP). To our knowledge, this algorithm has been applied to two-dimensional (2D) spin glasses, but only in the presence of an external magnetic field, which is known to improve the convergence properties of GBP. Our experience says that, in running plain GBP on a generic sample of a 2D spin glass without external field, a fixed point is reached only for high enough temperatures. The lack of convergence of GBP [and other similar message-passing algorithms (MPAs)] is a well-known problem, whose solution is in general far from being understood. For this reason, a new class of algorithms has been recently introduced [6,7], which are proved to converge to a fixed point: these algorithms use a double-loop iterative procedure (to be compared to the single loop in GBP) and for this reason they are usually quite slow.

The algorithm we are going to introduce is as fast as BP and provides marginals as accurately as double-loop GBP. Actually we will show that in a wide range of parameters the algorithm converges and yields the same marginals as the double-loop algorithm. As we already said, if the external fields are too small GBP has serious problems of convergence. On the other hand if the external fields are strictly zero we expect that the variables are not biased. In our algorithm this property is enforced by hand and leads to some additional conditions on the messages of GBP that result in a considerable extension of the range of parameters over which it converges. As we will see, the algorithm can be viewed as BP on a dual lattice.

Most of the simulations reported in the present paper regard the 2D Edwards-Anderson model. Nevertheless, the algorithm we are proposing is of a general kind and can be used even on lattices in larger dimensions (an account of partial results in this direction is provided in the last part of the present paper). The fact that many exact algorithms exist for solving efficiently models on a 2D lattice is not relevant for our purposes. We have to compare our algorithm to those belonging to the same class: fast and approximate inference algorithms for general dimensions. It is not our intention to propose an algorithm competing with exact algorithms for planar models.

The rest of the work is organized as follows. In Sec. II we derive the GBP equations for the 2D Edwards-Anderson model. In Sec. II A we rewrite these equations in terms of fields, a notation that has a nicer physical interpretation and that we are going to use in the rest of the work. Section III presents our algorithm, inspired by the paramagnetic ansatz to the GBP equations. In Sec. IV we show the results of running this algorithm on the 2D Edwards-Anderson model. There we compare its performance with Monte Carlo simulations, with the double-loop algorithm, and with exact calculations of the ground state of systems with bimodal and Gaussian interactions. Then, in Sec. V we generalize our message-passing equation to general dimensions and present some results for the 3D Edwards-Anderson model. Finally, some conclusions are drawn in Sec. VI.

II. GENERALIZED BELIEF PROPAGATION ON THE 2D EDWARDS-ANDERSON MODEL

Here we present the GBP equations for the Edwards-Anderson (EA) model on a 2D square lattice; we refer the reader to [5,8] for a more general introduction. In our case (as well as in many other cases) GBP is equivalent to Kikuchi's approximation, known as the cluster variational method (CVM) [9]. We will try a presentation as physical as possible.

Consider the 2D EA model consisting of a set $\sigma = \{s_1, \ldots, s_N\}$ of N Ising spins $s_i = \pm 1$ located at the nodes of a 2D square lattice with periodic boundary conditions, interacting with a Hamiltonian

$$\mathcal{H}(\sigma) = -\sum_{\langle i,j \rangle} J_{ij} s_i s_j, \tag{2}$$

where the sum runs over all couples of neighboring spins (first neighbors on the lattice). The J_{ij} are the coupling constants between spins and are supposed to be fixed for any given instance of the model. If the interactions are not random variables, i.e., $J_{ij} = J$, then the 2D ferromagnet is recovered. We will focus on the two most common disorder distributions: bimodal interactions, $P(J) = \frac{1}{2}\delta(J-1) + \frac{1}{2}\delta(J+1)$ and Gaussian interactions $P(J) = \exp(-J^2/2)/\sqrt{2\pi}$.

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

$$P(\sigma) = \frac{e^{-\beta \mathcal{H}(\sigma)}}{Z}.$$

The direct computation of the partition function

$$Z = \sum_{\sigma} e^{-\beta \mathcal{H}(\sigma)}$$

or any marginal distribution $p(s_i, s_j) = \sum_{\sigma \setminus s_i, s_j} P(\sigma)$ is a time-consuming task, impossible in practice, since it involves the addition of 2^N terms, and therefore an approximation is required.[1]

The idea of the region graph approximation to the free energy [5] is to replace the real distribution $P(\sigma)$ by a reduced set of its marginals. The 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 to all neighboring site marginals $p_{ij}(s_i, s_j)$ (Bethe), then to all square plaquette marginals $p_{ijkl}(s_i, s_j, s_k, s_l)$, 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 $b_i(s_i)$, $b_{ij}(s_i, s_j)$, etc. obtained from the minimization of a region-based free energy. In the region graph approximation to the free energy, a set of regions, i.e., sets of variables and their interactions, is defined, and a free energy is written in terms of the beliefs at each region. The cluster variational method [9] does a similar job, but instead of starting from an arbitrary choice of regions, it starts by defining the set of largest regions, and smaller regions are defined recursively by the intersections of bigger regions. In this sense, the CVM is a specific choice of region graph approximation to the free energy.

[1] The 2D case is special: indeed, thanks to the small genus topology, the partition function Z can be computed efficiently. However we are interested in developing an algorithm for the general case, and we will not make use of this peculiarity.

For the 2D EA model, we will consider the expansion of the free energy in terms of the marginals at three levels of regions: single sites (or spins), links, and plaquettes. By a plaquette we mean the square basic cell of the 2D lattice. This choice of regions corresponds to the CVM having the square plaquettes as the biggest regions. The free energy of the system is therefore written as

$$F = \sum_R c_R F_R,$$

where R runs over all regions considered, and the free energy in a particular region depends on the marginals at that level, $b_R(\sigma_R)$:

$$\beta F_R = \sum_{\sigma_R} [b_R(\sigma_R)\beta E_R(\sigma_R) + b_R(\sigma_R)\ln b_R(\sigma_R)].$$

The symbol σ_R refers to the set of spins in region R, while E_R is the energy contribution in that region. The counting numbers c_R (also called Moebius coefficients) are needed to ensure that bigger regions do not overcount the contribution in free energy of smaller regions, and follow the prescription

$$c_R = 1 - \sum_{R' \supset R} c_{R'}, \tag{3}$$

where R' is any region containing completely region R, such as, e.g., a plaquette containing a link or a link containing a site. In the case of the 2D lattice, the biggest regions are the square plaquettes, and therefore $c_{\text{plaq}} = 1$, while the link regions have $c_{\text{link}} = 1 - 2c_{\text{plaq}} = -1$ (as each of them is contained in two plaquette regions), and finally the spin regions have $c_{\text{site}} = 1 - 4c_{\text{plaq}} - 4c_{\text{link}} = 1$ (as each spin belongs to four links and four plaquettes). So the actual approximation for the EA model on a 2D square lattice is

$$\begin{aligned}\beta F = &\sum_{\mathcal{P}} \sum_{\sigma_\mathcal{P}} b_\mathcal{P}(\sigma_\mathcal{P}) \ln \frac{b_\mathcal{P}(\sigma_\mathcal{P})}{\exp[-\beta E_\mathcal{P}(\sigma_\mathcal{P})]} \quad \text{(plaquettes)} \\ &- \sum_{L} \sum_{\sigma_L} b_L(\sigma_L) \ln \frac{b_L(\sigma_L)}{\exp[-\beta E_L(\sigma_L)]} \quad \text{(links)} \\ &+ \sum_{i} \sum_{s_i} b_i(s_i) \ln \frac{b_i(s_i)}{\exp[-\beta E_i(s_i)]} \quad \text{(sites)},\end{aligned} \tag{4}$$

where the sums run over all plaquettes, links, and sites, respectively. Please note that we are using the following notation for region indices: lower case for sites, upper case for links, and upper case calligraphic for plaquettes. The energy term $E_i(s_i)$ in the site contribution is relevant only when an external field acts over spins and can be taken as zero in our case, since no external field is considered. Notice that whenever the interactions are included in more than one region (in our case in link and plaquette regions), the counting numbers guarantee that the exact thermodynamical energy $U = \sum_\sigma P(\sigma)\mathcal{H}(\sigma)$ is obtained when the beliefs are the exact marginals of the Boltzmann distribution. On the other hand, the entropy contribution is intrinsically approximated, since the cutoff in the region sizes imposes a certain kind of factorization of $P(\sigma)$ in terms of its marginals (see [8] for an explanation of the region graph approximation in terms of cumulant expansions of the entropy).

The next step in the method is to compute the beliefs from the minimization condition of the free energy. However, an unrestricted minimization will generally produce inconsistent solutions, since the beliefs (marginals) are not independent, as they are related by the marginalization conditions

$$\begin{aligned} b_i(s_i) &= \sum_{\sigma_{L\setminus i}} b_L(\sigma_L) = \sum_{s_j} b_L(s_i,s_j)\,, \\ b_L(\sigma_L) &= b_L(s_i,s_j) = \sum_{\sigma_{\mathcal{P}\setminus L}} b_\mathcal{P}(\sigma_\mathcal{P}) = \sum_{s_k,s_l} b_\mathcal{P}(s_i,s_j,s_k,s_l), \end{aligned} \tag{5}$$

where $\sigma_L = \{s_i,s_j\}$ and $\sigma_\mathcal{P} = \{s_i,s_j,s_k,s_l\}$. In order to minimize under the constraints in Eq. (5) and under the normalization condition for each belief, a set of Lagrange multipliers should be added to the free energy in Eq. (4). There are different ways of choosing the Lagrange multipliers [5], and each of them will produce a different set of self-consistency equations. We choose the so called parent to child scheme (see Sec. IX A in [5]), in which constraints in Eq. (5) are imposed by two sets of Lagrange multipliers: $\mu_{L\to i}(s_i)$ relating the belief at link L to that at site i, and $\nu_{\mathcal{P}\to L}(\sigma_L)$ relating the one at plaquette $\mathcal{P}$ to the one at link L.

With constraints (5) enforced by Lagrange multipliers, the free energy stationary conditions for the beliefs are the following:

$$\begin{aligned} b_i(s_i) &= \frac{1}{Z_i} \exp\left(-\beta E_i(s_i) - \sum_{L\supset i}^{4} \mu_{L\to i}(s_i)\right), \\ b_L(\sigma_L) &= \frac{1}{Z_L} \exp\left(- \beta E_L(\sigma_L) - \sum_{\mathcal{P}\supset L}^{2} \nu_{\mathcal{P}\to L}(\sigma_L) \right. \\ &\qquad \left. - \sum_{i\subset L}^{2} \sum_{\substack{L'\supset i \\ L'\neq L}}^{3} \mu_{L'\to i}(s_i)\right), \\ b_\mathcal{P}(\sigma_\mathcal{P}) &= \frac{1}{Z_\mathcal{P}} \exp\left(- \beta E_\mathcal{P}(\sigma_\mathcal{P}) - \sum_{L\subset\mathcal{P}}^{4} \sum_{\substack{\mathcal{P}'\supset L \\ \mathcal{P}'\neq\mathcal{P}}}^{1} \nu_{\mathcal{P}'\to L}(\sigma_L) \right. \\ &\qquad \left. - \sum_{i\subset\mathcal{P}}^{4} \sum_{\substack{L\supset i \\ L\not\subset\mathcal{P}}}^{2} \mu_{L\to i}(s_i)\right), \end{aligned} \tag{6}$$

where the notation $L \supset i$ refers to all links containing site i and $\mathcal{P} \supset L$ to all plaquettes containing link L. The upper indices in the sums are written just to help in understanding how many terms are in each sum for the 2D case. The precise meaning of the indices in each summation can be understood from the graphical representation in Fig. 1. Lagrange multipliers are shown as arrows going from parent regions to child regions: simple arrows correspond to $\mu_{L\to i}$ and triple arrows to $\nu_{\mathcal{P}\to L}$. Let us consider, for instance, the belief in a link region $b_L(\sigma_L)$, depicted in the central picture of Fig. 1: the sum of the two Lagrange multipliers $\nu_{\mathcal{P}\to L}(\sigma_L)$ corresponds to the triple arrows from plaquettes on the left and right of the link L, while the double sum over the three $\mu_{L\to i}(s_i)$ and the three $\mu_{L\to j}(s_j)$ corresponds to the six arrows acting over the two spins.

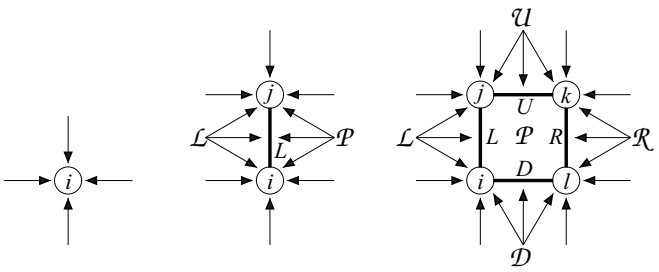

FIG. 1. Schematic representation of belief equations (6). Lagrange multipliers are depicted as arrows, going from parent regions to child regions.

In Eq. (6), the Z_R are normalization constants, and the terms $E_{\mathcal{P}}(\sigma_{\mathcal{P}}) = E_{\mathcal{P}}(s_i,s_j,s_k,s_l) = -(J_{ij}s_is_j + J_{jk}s_js_k + J_{kl}s_ks_k + J_{li}s_ls_i)$ and $E_L(\sigma_L) = E_L(s_i,s_j) = -J_{ij}s_is_j$ are the corresponding energies in plaquettes and links; they are represented in Fig. 1 by bold lines (interactions) between circles (spins). In our case $E_i(s_i)$ is zero since no external field is acting on the spins.

The Lagrange multipliers are fixed by the constraints they were supposed to enforce, Eq. (5), and they must satisfy the following set of self-consistency equations:

$$\exp[-\mu_{L\to i}(s_i)] = \sum_{s_j}\exp\left[-\beta E_{L\setminus i}(s_i,s_j) - \sum_{\mathcal{P}\supset L}^{2}\nu_{\mathcal{P}\to L}(s_i,s_j) - \sum_{\substack{L'\supset j\\ L'\neq L}}^{3}\mu_{L'\to j}(s_j)\right],$$

$$\exp[-\nu_{\mathcal{P}\to L}(s_i,s_j) - \mu_{D\to i}(s_i) - \mu_{U\to j}(s_j)] = \sum_{s_k,s_l}\exp\left[-\beta E_{\mathcal{P}\setminus L}(s_i,s_j,s_k,s_l) - \sum_{\substack{L'\in\mathcal{P}\\ L'\neq L}}^{3}\sum_{\substack{\mathcal{P}'\supset L'\\ \mathcal{P}'\neq\mathcal{P}}}^{1}\nu_{P'\to L'}(\sigma_{L'}) - \sum_{\substack{L'\supset k\\ L'\not\subset\mathcal{P}}}^{2}\mu_{L'\to k}(s_k) - \sum_{\substack{L'\supset l\\ L'\not\subset\mathcal{P}}}^{2}\mu_{L'\to l}(s_l)\right]. \quad (7)$$

Again, to help in understanding these equations, we provide in Fig. 2 their graphical representation. Note that there is one of these equations for every pair link-site and every pair plaquette-link in the graph. With $E_{\mathcal{P}\setminus L}$ we refer to interactions in plaquette $\mathcal{P}$ that are not in link L.

For each link L in the 2D lattice, there are two link-to-site multipliers $\mu_{L\to i}(s_i)$ and $\mu_{L\to j}(s_j)$. For each plaquette there are four plaquette-to-link multipliers $\nu_{\mathcal{P}\to L}(s_i,s_j)$, corresponding to the four links contained inside the plaquette. Let N be the number of spins in the lattice; there are $2N$ links and N plaquettes. So the originally intractable problem of computing marginals, has been replaced by the problem of solving a set of $4N + 4N$ coupled equations for Lagrange multipliers like those in Eq. (7). Once these equations are solved, the approximation for the marginals is obtained from Eq. (6) for

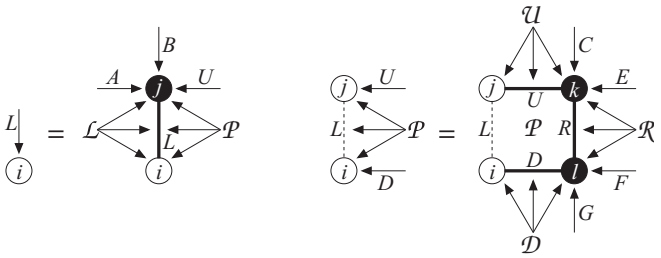

FIG. 2. Message-passing equations (7), shown schematically. Messages are depicted as arrows, going from parent regions to child regions. On any link J_{ij}, represented as bold lines between spins (circles), a Boltzmann factor $e^{\beta J_{ij}s_is_j}$ exists. Dark circles represent spins to be traced over. Messages from plaquettes to links, $\nu_{\mathcal{P}\to L}(s_i,s_j)$, are represented by a triple arrow, because they can be written in terms of three parameters U, u_i, and u_j, defining the correlation $\langle s_is_j\rangle$ and magnetizations $\langle s_i\rangle$ and $\langle s_j\rangle$, respectively.

the beliefs, and all thermodynamic quantities are derived from them as in Eq. (4).

Minimization of a region graph approximation to the free energy, like that in Eq. (4) with constraints Eq. (5), or equivalently solution of the set of self-consistent equations in Eq. (7), is still a nontrivial task. Let us consider two ways of doing it. The first method is the "direct" minimization of the constrained free energy, using a double-loop algorithm [7]. This method is quite solid, since it guarantees convergence to an extremal point of the constrained free energy, but it may be very slow to converge. The second method, which is generally faster but is not guaranteed to converge, is the family of the so called message-passing algorithms, in which the Lagrange multipliers are interpreted as *messages* $\nu_{\mathcal{P}\to L}(\sigma_L)$ going from plaquettes to links, and messages $\mu_{L\to i}(s_i)$ from links to sites. The self-consistency equations (7) can be viewed as the update rules for the messages in the left hand side, in terms of those in the right hand side. A random order updating of the messages in the graph by Eq. (7) (message passing) can reach a fixed point solution, and therefore an extremal point of the constrained free energy [5]. Next, we show explicitly what the message-passing equations look like in terms of fields.

A. From multipliers to fields

A particularly useful way of representing the multipliers (messages), with a nice physical interpretation, is the one used in [10], which we adopt here. In full generality [10,11], these multipliers can be written in terms of effective fields:

$$-\mu_{L\to i}(s_i) = \beta\, u_{L\to i}\, s_i, \quad (8)$$

$$-\nu_{\mathcal{P}\to L}(s_i,s_j) = \beta\,(U_{\mathcal{P}\to L}\, s_i\, s_j + u_{\mathcal{P}\to i}\, s_i + u_{\mathcal{P}\to j}\, s_j). \quad (9)$$

In particular, the field u corresponds to the cavity field in the Bethe approximation [5]. Use of Lagrange multipliers, messages, or fields is essentially equivalent. We will often refer to fields as u messages to emphasize their role in a message-passing algorithm, and we will refer to the self-consistency equations (7) as the message-passing equations.

This parametrization of the multipliers has proved useful to other endeavors, like the extension of the replica theory to general region graph approximations [11]. Here, all the relevant information in the Lagrange multipliers is translated

to "effective fields" u and (U,u_a,u_b). Notice that in this representation every single field u corresponds to an arrow in the schematic message-passing equations in Fig. 2. In particular, the messages going from plaquettes to links are characterized by three fields (U,u_a,u_b), and the field U acts as an effective interaction term that adds directly to the energy terms in the Boltzmann factor. For instance, the first message-passing Eq. (7) is

$$\begin{aligned}\exp[\beta u_{L\to i}s_i] = \sum_{s_j}\exp\{\beta[&(u_{\mathcal{P}\to i}+u_{\mathcal{L}\to i})\,s_i\\ &+(U_{\mathcal{P}\to L}+U_{\mathcal{L}\to L}+J_{ij})\,s_is_j\\ &+(u_{\mathcal{P}\to j}+u_{\mathcal{L}\to j}+u_{A\to j}\\ &+u_{B\to j}+u_{U\to j})\,s_j]\},\end{aligned}\tag{10}$$

where the indices refer to the notation used in Fig. 2 and J_{ij} is the interaction coupling constant between spins s_i and s_j. This equation naturally defines the updating rule for the message $u_{L\to i}$:

$$\begin{aligned}u_{L\to i} = \hat{u}(&u_{\mathcal{P}\to i}+u_{\mathcal{L}\to i},U_{\mathcal{P}\to L}+U_{\mathcal{L}\to L}+J_{ij},u_{\mathcal{P}\to j}+u_{\mathcal{L}\to j}\\ &+u_{A\to j}+u_{B\to j}+u_{U\to j}),\end{aligned}\tag{11}$$

where

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

Note that the usual cavity equation for fields in the Bethe approximation [12–14] is recovered if all contributions from plaquettes $\mathcal{P}$ and $\mathcal{L}$ are set to zero.

Working in a similar way for the second equation in (7) we end up with the updating rule for the message $(U_{\mathcal{P}\to L},u_{\mathcal{P}\to i},u_{\mathcal{P}\to j})$ sent from any given plaquette region $\mathcal{P}$ to one of its child links L (see right picture in Fig. 2):

$$\begin{aligned}U_{\mathcal{P}\to L} &= \frac{1}{4\beta}\ln\frac{K(1,1)K(-1,-1)}{K(1,-1)K(-1,1)},\\ u_{\mathcal{P}\to i} &= u_{\mathcal{D}\to i}-u_{D\to i}+\frac{1}{4\beta}\ln\frac{K(1,1)K(1,-1)}{K(-1,1)K(-1,-1)},\\ u_{\mathcal{P}\to j} &= u_{\mathcal{U}\to j}-u_{U\to j}+\frac{1}{4\beta}\ln\frac{K(1,1)K(-1,1)}{K(1,-1)K(-1,-1)},\end{aligned}\tag{12}$$

where

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

Equations (11) and (12) are equivalent to the equations in (7), once multipliers (messages) are parametrized in terms of fields. For instance, note that the μ multipliers in the left hand side of the second equation in (7) appear now subtracted in the right hand side of Eq. (12).

The field notation is more comprehensible and has a clear physical meaning. Each plaquette $\mathcal{P}$ is telling its child links L that they should add an effective interaction term $U_{\mathcal{P}\to L}$ to the real interaction J_{ij}, due to the fact that spins s_i and s_j are also interacting through the other three links in the plaquette $\mathcal{P}$. The fields u act like magnetic fields upon spins, and the complete $\nu_{\mathcal{P}\to L}(s_i,s_j)$ message is characterized by the triplet $(U_{\mathcal{P}\to L},u_{\mathcal{P}\to i},u_{\mathcal{P}\to j})$, and will be referred to from now on as a Uuu message. Furthermore, it is clear that some fields enter directly in the message-passing equations, like $u_{\mathcal{P}\to i}$ and $u_{\mathcal{L}\to i}$ in Eq. (11) and $u_{\mathcal{D}\to i}$ and $u_{\mathcal{U}\to j}$ in Eq. (12). Also note that, since our model has no external field, the fields u break the symmetry of the original Hamiltonian whenever they are nonzero. For instance, in the ferromagnetic model, when all $J_{ij}=J$, these fields are zero at high temperature and become nonzero at Kikuchi's critical temperature $T=1.42$ [9], implying a spontaneous magnetization in the ferromagnet.

III. THE DUAL APPROXIMATION FOR THE PARAMAGNETIC PHASE

Unfortunately, the iterative message-passing algorithm for solving the GBP equations (11) and (12) often does not converge on finite-dimensional lattices. While this is expected if long-range correlations are present, it is rather disappointing that it happens also in the paramagnetic phase, where one would like to find easily the solution to the model. Here we are going to focus only on the paramagnetic phase and propose an improved solution algorithm based on physical assumptions.

In the paramagnetic phase of any spin model defined by the Hamiltonian in Eq. (2), that is with no external field, variables have no bias or magnetization: this in turn implies that in the solution all u fields must be zero, and only U fields should be fixed self-consistently to nonzero values.

This paramagnetic solution has some interesting properties. First, it is always a solution of the GBP equations, since Eqs. (11) and (12) are self-consistent with all $u=0$. This means that, starting from unbiased messages (all $u=0$) the iterative GBP algorithm keeps this property. Second, the paramagnetic ansatz is correct, from the GBP perspective, at least at high enough temperatures, meaning that even if we start with biased messages ($u\neq 0$), the iterative algorithm converges to all $u=0$ at high temperatures. And last, but not least, the well studied physical behavior for the 2D EA model with zero-mean random interactions J_{ij} is expected to remain always paramagnetic, i.e., to have no transition to a spin glass phase at any finite temperature [15]. Therefore, the ansatz $u=0$ is both physically plausible and algorithmically desirable.

Under the paramagnetic ansatz, which we shall also call the dual approximation for a reason to be explained soon, the message-passing equation (11) is irrelevant, as it is always satisfied given that $\hat{u}(0,U,0)=0$, while Eq. (12) now turns into (see Fig. 3)

$$\begin{aligned}U_{\mathcal{P}\to L} &= \widehat{U}(U_{\mathcal{U}\to U},U_{\mathcal{R}\to R},U_{\mathcal{D}\to D})\\ &= \frac{1}{\beta}\operatorname{arctanh}[\tanh\beta(U_{\mathcal{U}\to U}+J_{jk})\tanh\beta(U_{\mathcal{R}\to R}+J_{kl})\\ &\quad\times\tanh\beta(U_{\mathcal{D}\to D}+J_{li})].\end{aligned}\tag{13}$$

The only relevant messages now are those associated with the multipliers $\nu_{\mathcal{P}\to L}(s_i,s_j)=\beta U_{\mathcal{P}\to L}s_is_j$, and they will be referred to as U messages. Equation (13) can be interpreted as a correlation message-passing equation, giving the new interaction field $\widehat{U}$ that a certain link will experience as

a consequence of the correlations transmitted around the plaquette. The belief equations (6) also simplify. Obviously $b(s_i) = 0.5$ for every spin in the graph, and the link and plaquette beliefs are

$$b_L(s_i,s_j) = \frac{1}{Z_L} e^{\beta(U_{\mathcal{L}\to L}+U_{\mathcal{P}\to L}+J_{ij})s_i s_j},$$
$$b_{\mathcal{P}}(s_i,s_j,s_k,s_l) = \frac{1}{Z_{\mathcal{P}}} e^{\beta(U_{\mathcal{L}\to L}+J_{ij})s_i s_j+\beta(U_{\mathcal{U}\to U}+J_{jk})s_j s_k+\beta(U_{\mathcal{R}\to R}+J_{kl})s_k s_l+\beta(U_{\mathcal{D}\to D}+J_{li})s_l s_i}. \tag{14}$$

As already mentioned in Sec. II, the entropy in Eq. (4) is always approximated. The quality of the estimates obtained for the marginals will depend on the quality of this approximation. The bigger the correlation length in the system, the more inaccurate the entropy will be. Therefore, at low temperatures (even in a paramagnetic phase) the minimization of the approximated free energy will give worse estimates for the marginals.

The dual algorithm we are proposing to study the paramagnetic phase of the EA model is a standard message-passing algorithm for the U messages, which works as follows.

(1) Start with all U messages null
(2) **repeat**
(3) Choose randomly one plaquette $\mathcal{P}$ and one of its child links L
(4) Update the field $U_{\mathcal{P}\to L}$ according to Eq. (13) as in Fig. 3
(5) **until** The last change for any U message is less than ϵ (we use typically $\epsilon = 10^{-10}$)
(6) **return** The beliefs $b_L(s_i,s_j)$ defined in Eq. (14) for every pair of neighboring spins

Some damping factor $\gamma \in [0,1)$ can be added in the update step $U_{\mathcal{P}\to L} = \gamma U_{\mathcal{P}\to L} + (1-\gamma)\widehat{U}$ in order to help convergence.

A. Mapping to the dual model

It is worth noticing that Eq. (13) is nothing but the BP equation for the corresponding dual model (hence the name of the algorithm). The dual model has a binary variable $x_{ij} \equiv s_i s_j$ on every link of the original model, and the original coupling constants play now the role of an external polarizing (eventually random) field

$$\mathcal{H}_{\text{dual}}(\vec{x}) = -\sum_{\langle i,j\rangle} J_{ij} x_{ij}.$$

This Hamiltonian looks like the sum of independent variables, but this is not the case. The dual variables $x_{ij} = \pm 1$ must satisfy a constraint for each cycle (or closed path) in the original graph, enforcing that their product along the cycle must be equal to 1. On a regular lattice any closed path can be expressed in terms of elementary cycles of four links (the plaquettes) and so it is enough to enforce the constraint on every plaquette: $x_{ij}x_{jk}x_{kl}x_{li} = 1$. The Gibbs-Boltzmann probability distribution for the dual model is then given by

$$P(\vec{x}) = \frac{1}{Z} e^{-\beta\mathcal{H}_{\text{dual}}(\vec{x})} \prod_{\langle i,j,k,l\rangle} \delta_{x_{ij}x_{jk}x_{kl}x_{li},\,1}, \tag{15}$$

where the product runs over all elementary plaquettes.

The model described by the probability measure in Eq. (15) can be viewed as a constraint satisfaction problem with a nonuniform prior (given by $e^{-\beta\mathcal{H}_{\text{dual}}(\vec{x})}$). It is straightforward to derive the BP equations for such a problem. Indeed by defining the marginal for the variable x_{ij} on link L in the presence of the only neighboring plaquette $\mathcal{P}$ as $(1 + x_{ij}\tanh\beta U_{\mathcal{P}\to L})/2 \propto \exp(\beta U_{\mathcal{P}\to L} x_{ij})$, the BP equations read

$$\begin{aligned}
&\frac{1}{2}(1 + x_{ij}\tanh\beta U_{\mathcal{P}\to L}) \\
&\quad\propto \sum_{x_{jk},x_{kl},x_{li}} e^{\beta U_{\mathcal{U}\to U} x_{jk}} e^{\beta J_{jk} x_{jk}} e^{\beta U_{\mathcal{R}\to R} x_{kl}} e^{\beta J_{kl} x_{kl}} e^{\beta U_{\mathcal{D}\to D} x_{li}} e^{\beta J_{li} x_{li}} \\
&\qquad\times \delta_{x_{jk}x_{kl}x_{li},\,x_{ij}} \\
&\quad\propto \sum_{\substack{x_{jk},x_{kl},x_{li}:\\ x_{jk}x_{kl}x_{li}=x_{ij}}} [1 + x_{jk}\tanh\beta(U_{\mathcal{U}\to U} + J_{jk})][1 + x_{kl}\tanh\beta \\
&\qquad\times (U_{\mathcal{R}\to R} + J_{kl})][1 + x_{li}\tanh\beta(U_{\mathcal{D}\to D} + J_{li})] \\
&\quad= 1 + x_{ij}\tanh\beta(U_{\mathcal{U}\to U} + J_{jk})\tanh\beta(U_{\mathcal{R}\to R} + J_{kl}) \\
&\qquad\times \tanh\beta(U_{\mathcal{D}\to D} + J_{li}).
\end{aligned} \tag{16}$$

In the second summation the terms containing one or two x variables sum to zero, while the other two terms are those written in the last expression. Equating the first and the last expressions, this equation is manifestly equal to Eq. (13).

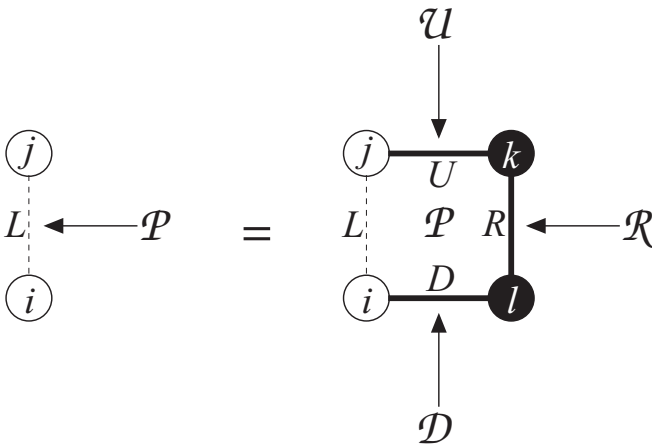

FIG. 3. Message passing of correlation messages in the dual approximation. In the right hand side the trace is taken over the black spins.

B. Average case solution

GBPs in general, and the dual approximation in particular, are methods for the study of the thermodynamic properties of a given problem. However, in the limit of large systems ($N \to \infty$, the thermodynamic limit), we expect a typical behavior to arise in some observables. This is the so called *self-averaging* property of disordered systems. By "typical" we mean that almost every realization of the interactions J_{ij} will result in a system whose thermodynamic properties (free energy, energy, entropy) are very close to the average values.

Normally, in disordered systems, we cope with the $N \to \infty$ limit and with the average over the random J_{ij} by the replica method. The application of the replica trick to region graph approximations is a challenging task [11]. However, we can still grasp the average case behavior with a cavity average case solution of the dual message-passing equations, at the price of neglecting the local structure of the graph (beyond plaquettes).

The idea is to represent the set of U messages flowing in any given graph, by a population of messages $Q(U)$. Then the message-passing Eq. (13) is used to obtain such population in a self-consistent way. More precisely, in every iteration three messages U_1, U_2, U_3 are randomly drawn from the population $Q(U)$ and a new message $U_0 = \widehat{U}(U_1, U_2, U_3)$ is computed by Eq. (13) using three couplings randomly selected from $P(J)$. The obtained message U_0 is put back into the population, and the iteration is repeated many times, until the population stabilizes.

Once we have the self-consistent population of messages, we can compute the average energy

$$E_{\rm av} = \langle -J_{ij} \tanh \beta(J_{ij} + U_1 + U_2) \rangle_{Q(U_1), Q(U_2), P(J_{ij})} \qquad (17)$$

by a random sampling of the population and of the interactions. The average case solution is supposed to be very good whenever the network of interactions has no or few short loops. This is not the case in any finite-dimensional lattice, since there the short loops (plaquettes) are abundant. Nonetheless, the average case solution gives a reasonably good approximation to the single-instance results in 2D and 3D, as shown in the next section.

IV. RESULTS ON THE 2D EDWARDS-ANDERSON MODEL

Message-passing algorithms work fine in the high-temperature regime ($T > T_c$) of models defined on random topologies: this is the reason why these methods have been successfully applied in random constraint satisfaction problems, like random satisfiability or random coloring [16–19]. However, when used on regular finite-dimensional lattices, they can experience difficulties even in the paramagnetic phase, because the presence of short loops spoils message-passing convergence.

It is well known [20,21] that on a random graph of fixed degree (connectivity) $c = 4$ the cavity approximation gives a paramagnetic result above $T_{\rm Bethe} \simeq 1.52$ (i.e., $\beta_{\rm Bethe} \simeq 0.66$) with all cavity fields $u_i = 0$. Below the Bethe critical temperature, this solution becomes unstable to perturbations, and we expect many solutions to appear with nontrivial messages $u_i \neq 0$. The presence of many solutions in the message-passing equations is connected to the existence of many thermodynamic states in the Gibbs-Boltzmann measure, or, equivalently, to the presence of replica symmetry breaking. The appearance of such a spin glass phase is also responsible for the lack of convergence of message-passing equations, since the intrinsic locality of the message-passing equations fails to coordinate distant regions of the graph (which are now long-range correlated). As a consequence, the application of BP to the 2D EA model (which also has fixed degree $c = 4$) still finds the paramagnetic phase at high temperatures, but below $T_{\rm Bethe}$, the Bethe instability takes the message-passing iteration away from the $u = 0$ solution and does not allow the messages to converge to a fixed point (i.e., the algorithm wanders forever). In Fig. 4 we show the convergence probability for the BP message-passing equations in the 2D EA model.

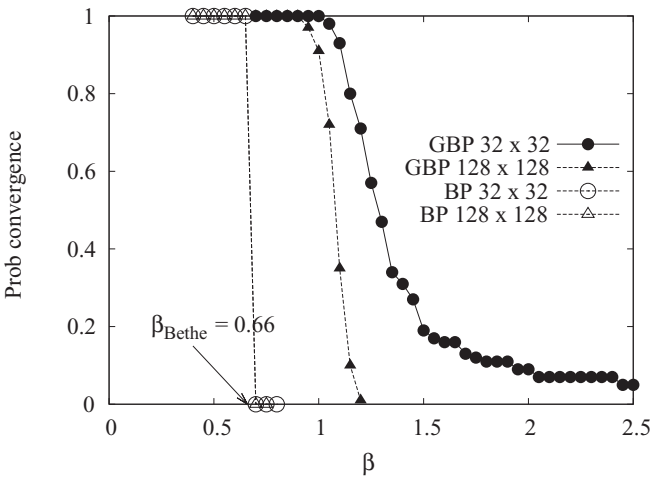

FIG. 4. Convergence probability of BP (Bethe approximation) and GBP on a 2D square lattice, as a function of inverse temperature. Data points are averages over 100 systems with random bimodal interactions. System sizes are $N = L^2$ with $L = 32,128$ and a damping factor $\gamma = 0.5$ has been used in the iteration of the message-passing equations. The Bethe spin glass transition is expected to occur at $\beta_{\rm Bethe} \simeq 0.66$ ($T_{\rm Bethe} \simeq 1.52$) for a random graph with the same connectivity as the 2D square lattice. Notably, that temperature also marks the convergence threshold for BP equations in the 2D square lattice. GBP, on the contrary, reaches lower temperatures, but eventually stops converging.

On the other hand, a straightforward GBP parent-to-child implementation does not fully overcome this problem. At high temperatures, the parent-to-child equations converge to a paramagnetic solution with all $u = 0$ and nontrivial $U \neq 0$, which turns out to be the same solution found by our dual algorithm. When the temperature decreases, the convergence properties of the algorithm worsen and are sensitive to tricks like damping and bounding in the fields. A thorough discussion of these properties is left for future work, but let us summarize: typically the algorithm stops converging at low temperatures, somewhere below $T_{\rm Bethe}$, as shown in Fig. 4.

So, in general, BP and GBP equations are not simple to use in finite-dimensional systems at low enough temperatures: this warning was already reported in Refs. [7,8,22]. Indeed a different method for extremizing the constrained free energy named the double-loop algorithm [6,7] was developed to overcome such difficulties. As mentioned earlier, the double loop guarantees convergence of the beliefs, on any topology,

with or without short loops. Given the convergence problems in GBP, researchers typically resort to double-loop algorithms to extremize region graph approximations to the free energy, below the Bethe critical temperature.

In order to make a fair comparison with our dual algorithm, we have used an optimized code for GBP and double-loop algorithms: the open source LIBDAI library written in C++ [23].

The first interesting result of our work is that our dual algorithm converges at all temperatures, just as the double-loop algorithm does. The reason why it converges is that there are no u messages, so the Bethe instability does not affect our message-passing iteration.

The second relevant result of our dual algorithm is the fact that it finds the same solution found by the double-loop algorithm at all temperatures. In other words, the direct extremization of the region graph approximation to the free energy Eq. (4) via a double-loop algorithm finds a paramagnetic solution characterized by the beliefs $b_i(s_i) = 0.5$ and $b_L(s_i,s_j) = \frac{1}{z} e^{-\beta \tilde{J}_{ij} s_i s_j}$; and the effective interactions $\tilde{J}_{ij}$ found by the double-loop algorithm are exactly equal to those found with our dual algorithm, $\tilde{J}_{ij} = J_{ij} + U_{\mathcal{P}\to L} + U_{\mathcal{L}\to L}$. This means that beliefs and correlations found by the two algorithms are identical: $\langle s_i s_j \rangle_{\text{double loop}} = \langle s_i s_j \rangle_{\text{dual}}$.

The third result is that the running times of our dual algorithm are nearly four orders of magnitude smaller than those required by the double-loop implementation in LIBDAI, at least in a wide range of temperatures (see Fig. 5). More precisely, the convergence time of the dual algorithm grows exponentially with $\beta = 1/T$, but still, in the relevant range of temperatures where the region graph approximation is a good approximation (not too low temperatures), the running time is always roughly a factor of 10^4 smaller than that for the double loop.

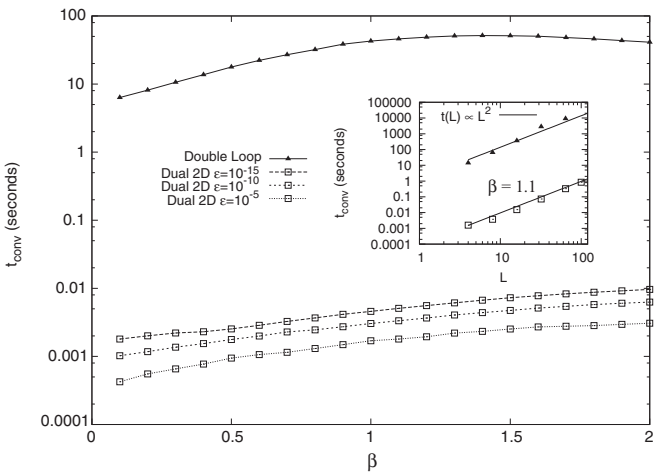

FIG. 5. Running times of the double-loop algorithm [7,23] (LIBDAI) and the dual algorithm averaged over ten realizations of a 2D 8×8 EA model with Gaussian interactions. Generally the double-loop algorithm requires a time four orders of magnitude larger than that used by the dual algorithm. Three different precision goals were used for the dual algorithm, $10^{-5}, 10^{-10}, 10^{-15}$, while the precision of the double-loop algorithm is 10^{-9}. The inset shows the behavior of the running times for both algorithms versus the system size $L = \sqrt{N}$. The growth is linear in N, as expected.

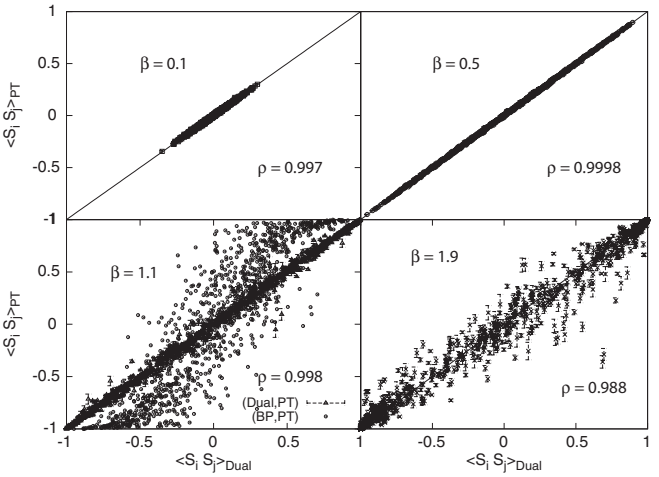

FIG. 6. Comparison between the correlations $\langle s_i s_j \rangle_{\text{dual}}$ obtained by the dual algorithm and the nearly exact correlations obtained by a parallel tempering simulation. We used a 64×64 EA model with Gaussian interactions. In the left lower plot the trivial inference $\langle s_i s_j \rangle_{\text{BP}} = \tanh(\beta J_{i,j})$ is also plotted for comparison purposes. Notice that this is the correlation resulting from BP, when it converges to the paramagnetic $u = 0$ solution. At each temperature the data correlation coefficient ρ is reported.

A. Dual approximation vs Monte Carlo simulations

The fact that our dual algorithm provides the same results (and much faster) than the double-loop algorithm is good news. Essentially it is telling us that we are not losing anything by restricting the space of possible messages, as far as the region graph approximation is concerned. However, the ultimate comparison for the approximation has to be done with the exact marginals and correlations. In Fig. 6 we show a comparison between the exact correlations $C_{ij,\text{PT}} = \langle s_i s_j \rangle_{\text{PT}}$ of neighboring spins obtained with a parallel tempering (PT) Monte Carlo simulation, and the dual approximation estimate for the same two-spin correlations $C_{ij,\text{dual}} = \langle s_i s_j \rangle_{\text{dual}}$. The agreement between $C_{ij,\text{PT}}$ and $C_{ij,\text{dual}}$ is essentially perfect at high temperatures, and it becomes weaker as the temperature is decreased. The reason for the discrepancies is obviously the fact that we are using an approximation in which the collective behavior of spins is accounted for exactly only until the plaquette level; more distant correlations are approximated, and these correlations become more important at low temperatures. In particular, the correlation length of the 2D EA at $\beta = 2.0$ is already above 10 [15], and therefore the local inference method performs poorly.

However, the average mean error between the correlations inferred from the dual algorithm and those obtained by Monte Carlo (PT) simulation decreases with increasing system size at any fixed temperature. In Fig. 7 the two-point and four-point correlation errors, defined as

$$\Delta_2 = \sqrt{\frac{\sum_{\langle i,j \rangle} (C_{ij,\text{PT}} - C_{ij,\text{dual}})^2}{\sum_{\langle i,j \rangle} C_{ij,\text{PT}}^2}},$$
$$\Delta_4 = \sqrt{\frac{\sum_{\langle i,j,k,l \rangle} (C_{ijkl,\text{PT}} - C_{ijkl,\text{dual}})^2}{\sum_{\langle i,j,k,l \rangle} C_{ijkl,\text{PT}}^2}}, \tag{18}$$

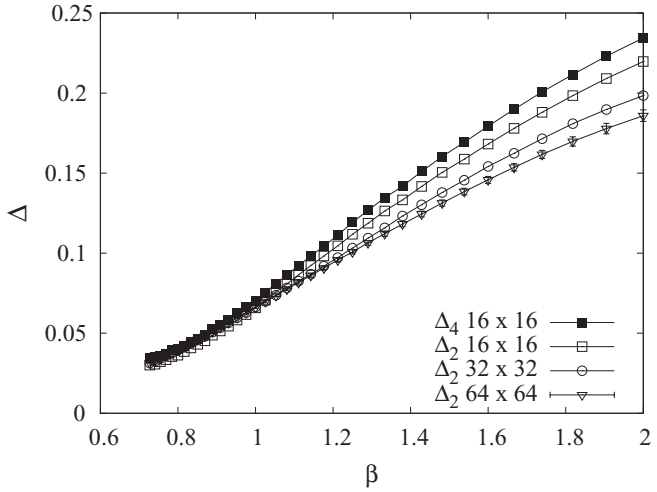

FIG. 7. Average errors Δ_2 and Δ_4 between the nearly exact two-point and four-point correlations (obtained by MC calculation) and their dual estimates. The average has been taken over ten Edwards-Anderson samples of three different sizes and bimodal interactions. The quality of the inference becomes worst as the temperature goes down (higher β), but it improves for larger systems.

are shown as functions of the inverse temperature. The sum goes over the pairs of first neighbors in the case of Δ_2 and over the groups of four spins in a square plaquette for Δ_4. Four-point correlations appear to be slightly worse than two-point correlations. For clarity in the plot, only the data for the four-point correlation in a 16^2 system are plotted, but the behavior is similar to that of the two-point correlations. The inferred correlations worsen at lower temperatures but the errors diminish with increasing system size.

Given the good correspondence between the correlations under the dual approximation and the true correlations, we expect a good estimate for the energy too. In Fig. 8 we show

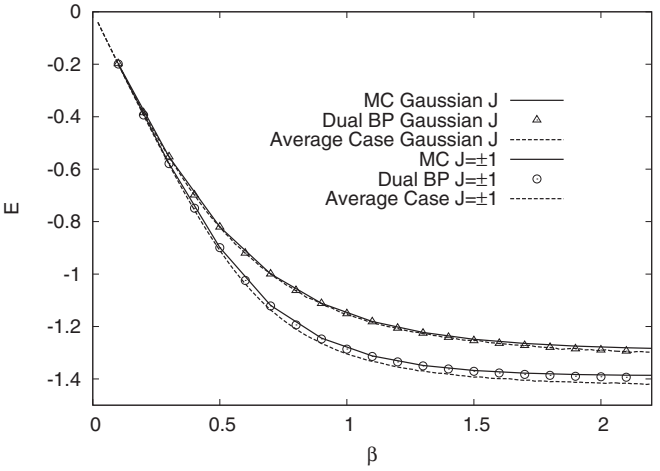

FIG. 8. Energy as a function of the inverse temperature β for a 64×64 2D EA model, with both types of interaction, Gaussian and bimodal. Full lines represent the exact thermodynamic energy as obtained by a Monte Carlo simulation, points are the energies obtained under the dual approximation, and dashed lines are the average case energies.

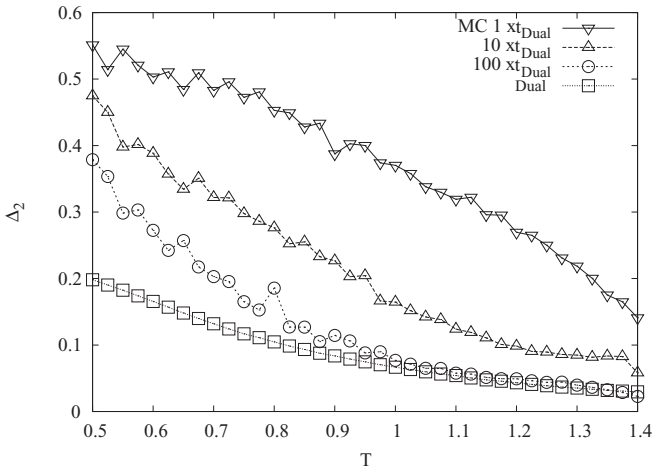

FIG. 9. Error Δ_2 made by a Monte Carlo simulation (with parallel tempering) for the estimation of first-neighbor correlations in a 2D EA model of size 64^2. The Monte Carlo simulation is run for a time that is 1, 10, and 100 times the convergence time of the dual algorithm. The error made by the dual approximation is also reported and is lower in the whole range of temperatures analyzed, suggesting that the dual approximation is a better choice when only an approximated result is required in a short time for large systems.

with points the energy under the dual approximation and with full lines the Monte Carlo exact energy: the data are indeed very close. The dashed lines show the average case energy for the dual approximation, Eq. (17). In spite of the fact that the average case does not take into account the local structure of the lattice, the average case energy is quite close to the single-instance one.

Concluding this section on the comparison between the dual algorithm and the Monte Carlo method (which is the standard general purpose inference method), we emphasize that the dual algorithm is not able to provide the exact answer, because of the underlying approximation, even if run for a very long time. However it is able to provide a very good approximate result in a very short time. To quantify this statement, we show in Fig. 9 the error Δ_2 achieved by the Monte Carlo method (using the parallel tempering algorithm) when run for the same time required by the dual algorithm to converge: the error achieved by the Monte Carlo simulation is at least twice that obtained by the dual algorithm for any temperature in the range considered, $0.5 \leqslant T \leqslant 1.4$. Moreover, improving the error Δ_2 by running a longer Monte Carlo simulation is not easy, since an increase of running times by a factor of 100 is required to obtain equivalent performances at least for high temperatures ($T \gtrsim 1$). So in all those cases where an approximate inference is required in a very short time, the present dual algorithm greatly outperforms the standard Monte Carlo methods.

B. Ground state configuration in 2D

The good agreement between the correlations found by the dual algorithm and those found in a Monte Carlo simulation, for the 2D EA model, suggests that we should check whether the inferred correlations can be used down to $T = 0$. More

precisely, using the correlations obtained by our dual algorithm at low temperatures ($\beta = 2.0$), we try to compute a ground state configuration by a decimation procedure. For the 2D EA model there are fast and exact algorithms for finding the ground state (GS) [24]. Our intention is not to propose a new (approximated) algorithm for the search for the GS, but rather to perform a severe test of the quality of the correlations implied by our dual algorithm. Here we present and test the decimation only in the 2D EA, leaving for the future more challenging cases (like the 3D EA model, where no exact polynomial algorithm for computing the GS exists).

The idea is to freeze iteratively the relative position $s_i s_j$ of those interacting spins that are more strongly correlated (this is done by setting $J_{ij} \to \pm\infty$), and rerunning the dual algorithm until convergence every time one pair of spins is frozen. Note that freezing the relative position of spins is equivalent to freezing the dual variable $x_{ij} = s_i s_j$. The freezing procedure is very simple, but for the fact that one has to check that frozen links must be consistent with a spin configuration. More precisely, frozen x_{ij} variables must satisfy the requirement that on any closed loops the product is 1,

$$\prod_{ij}^{\text{closed loop}} x_{ij} = 1. \tag{19}$$

For very short loops the satisfaction of this condition is automatically induced by the dual algorithm: for example, if three links on a plaquette freeze, the fourth link is immediately frozen to a value satisfying the condition in Eq. (19). However, for longer loops (like the one shown in Fig. 10), the propagation of these constraints by the dual algorithm is not perfect, since the information degrades with distance beyond the plaquette level. Then we need to enforce the constraints of Eq. (19) by a proper algorithm. At each stage of the freezing process, we define the clusters of frozen links as follows: if two frozen links share a spin, then they belong to the same cluster. In Fig. 10 a cluster of frozen links is represented by bold lines. Notice that, once a spin is fixed in a cluster, all other spins are fixed as well by the frozen correlations. On the other hand, different clusters of spins can have arbitrary relative orientations.

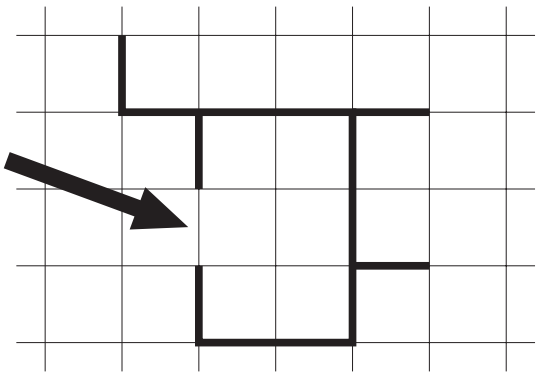

FIG. 10. Even if the link marked by the arrow is not the most polarized link according to the marginals provided by the dual algorithm, the spins it connects are fully correlated by the fact that they belong to a cluster of frozen links (bold lines). Therefore, the marked link must be immediately fixed accordingly.

Consider now the situation depicted in Fig. 10 and focus on the value of the correlation between the two spins connected by the link marked by the arrow. From the fact that these two spins belong to the same cluster of frozen links (shown as bold links in Fig. 10) we know they are perfectly correlated, however, by running the dual algorithm we could get a weak value for this correlation and then proceed by freezing a different link. A set of suboptimal choices of this kind may finally produce a configuration of frozen links where the constraints in Eq. (19) are not all satisfied. In order to avoid these constraint violations we force any link whose spins are already part of the same cluster to be polarized accordingly. The freezing algorithm, therefore, works as follows.

(1) **repeat**
(2) Run the Dual algorithm until convergence (at a low enough temperature)
(3) Find the link L with largest finite $\tilde{J}_L = J_L + U_{\mathcal{P}\to L} + U_{\mathcal{L}\to L}$
(4) Freeze that link by setting $J_L \leftarrow \text{sgn}\tilde{J}_L)\,\infty$
(5) **if** Link L is connected to clusters **C** and **C**′ of frozen links **then**
(6) merge clusters **C**, **C**′ and link L in a unique cluster
(7) **else**
(8) **if** Link L is connected to a single cluster **C** of frozen links **then**
(9) add link L to cluster **C**
(10) **else**
(11) Create a new cluster with link L
(12) **end if**
(13) **end if**
(14) **for all** Nonfrozen links L' at the boundaries of a cluster of frozen links **do**
(15) **if** Link L' shares both spins with the same cluster **then**
(16) Freeze link L' accordingly {to avoid violations of constraint}
(17) **end if**
(18) **end for**
(19) **until** All links are frozen
(20) **return** The spin configuration obtained by setting one spin and fixing the rest according to frozen links

The results obtained with this freezing procedure are good, considering the simplicity of the method. In Fig. 11 we compare the resulting ground state energies per spin with the exact solutions obtained using a web service running an exact solving algorithm [24]. We used an ensemble of 100 EA models on the 2D square lattice with Gaussian interactions (so the ground state is not degenerate) and with bimodal interactions of sizes 16×16, 32×32, and 64×64 (32×32 not shown in Fig. 11).

While most of the time the algorithm does not find the exact ground state (especially for large system sizes), it finds a nearby solution, with an average error of $\overline{E_{\text{dual}} - E_{\text{GS}}} \sim 0.003$ for the bimodal system and ~ 0.002 for the Gaussian one. Furthermore, the fraction of pairs of linked spins that share their relative position between the exact solution and the solution found by the dual + freezing algorithm remains constant for growing system sizes ($\sim 86\%$ for the bimodal and $\sim 94\%$ for the Gaussian). Keep in mind that the EA

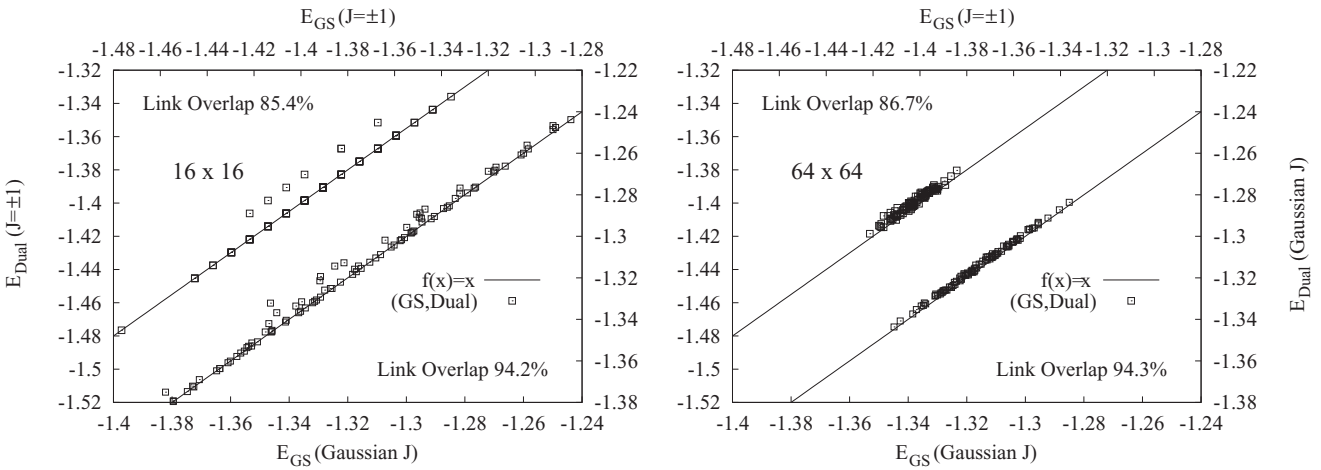

FIG. 11. Correlation between the ground state energy per spin obtained by the dual + freezing algorithm and by an exact method for $N = 16 \times 16$ (left) and $N = 64 \times 64$ (right) systems. In each plot the top left points correspond to 100 bimodal systems $J_{ij} = \pm 1$, while the right bottom points correspond to 100 systems with Gaussian interactions. In both cases the dual + freezing algorithm finds a state close in energy to the ground state. For bimodal interactions, the degeneracy of the ground state reduces the expected link overlap with the exact ground state solution $\text{Prob}(s_i^{\text{dual}} s_j^{\text{dual}} = s_i^{\text{exact}} s_j^{\text{exact}}) = 86\%$. For Gaussian interactions, the ground state is not degenerate and therefore the average link overlap is very high (94%). The line $f(x) = x$ is shown to guide the eye. Kindly note that two set of axes are being used.

systems with bimodal interactions have degenerate ground states; therefore solutions nearby in energy need not be too close in the configuration space.

Even if the dual algorithm converges quite fast, the decimation procedure used in this section requires running the algorithm after every freezing of the dual variables, making the dual + freezing algorithm quite slow compared to the exact algorithms for the ground state. The performance of this algorithm in more interesting cases, like the 3D EA model, is left for future work.

V. GENERALIZATION TO OTHER DIMENSIONS

Let us now consider the region-graph-based approximation to the free energy for a generic D-dimensional (hyper)cubic lattice, using the same hierarchy of regions: square plaquettes, links, and spins. After computation of the counting numbers for a general D-dimensional lattice, see Eq. (3), the free energy approximation becomes

$$\begin{aligned}\beta F = &\sum_{\mathcal{P}} \sum_{\sigma_{\mathcal{P}}} b_{\mathcal{P}}(\sigma_{\mathcal{P}}) \ln \frac{b_{\mathcal{P}}(\sigma_{\mathcal{P}})}{\exp[-\beta E_{\mathcal{P}}(\sigma_{\mathcal{P}})]} \quad \text{(plaquettes)} \\ &- (2D-3) \sum_{L} \sum_{\sigma_L} b_L(\sigma_L) \ln \frac{b_L(\sigma_L)}{\exp[-\beta E_L(\sigma_L)]} \quad \text{(links)} \\ &+ (2D^2 - 4D + 1) \sum_i \sum_{s_i} b_i(s_i) \\ &\times \ln \frac{b_i(s_i)}{\exp[-\beta E_i(s_i)]} \quad \text{(spins)}. \end{aligned} \tag{20}$$

Plaquettes are still the biggest regions considered and so have counting number 1, but now each link is contained in $2(D-1)$ plaquettes, and each spin is in $2D$ links and $2D(D-1)$ plaquettes. The message-passing equations for the dual algorithm in D dimensions are then

$$\begin{aligned} U_{\mathcal{P}\to L} = \frac{1}{\beta} \operatorname{arctanh} \Bigg[&\tanh\beta \left(\sum_i^{2(D-1)-1} U_{\mathcal{U}_i \to U} + J_U \right) \\ &\times \tanh\beta \left(\sum_i^{2(D-1)-1} U_{\mathcal{R}_i \to R} + J_R \right) \\ &\times \tanh\beta \left(\sum_i^{2(D-1)-1} U_{\mathcal{D}_i \to D} + J_D \right) \Bigg], \end{aligned} \tag{21}$$

where $\mathcal{U}_i$ ($\mathcal{R}_i$ and $\mathcal{D}_i$) are the $2(D-1)-1$ plaquettes containing the link U (R and D) excluding plaquette $\mathcal{P}$.

In the high-temperature phase, this dual approximation with all $u = 0$ should still be a valid approach for any dimensionality D. At low temperatures, however, the EA model in more than two dimensions has a spin glass phase transition and, therefore, we expect the dual approximation to become poorer, as it cannot account for a very long correlation length and a nontrivial order parameter.

By running the dual algorithm for the 3D EA model we have found a divergence of U fields around $\beta \simeq 0.39$ for bimodal couplings and around $\beta \simeq 0.41$ for Gaussian couplings. This divergence is due to the fact the U fields get too strongly self-reinforced under iteration. This divergence does not come as a surprise, given that it happens also when one studies the simpler pure ferromagnetic Ising model. However, in the ferromagnetic model the temperature at which U fields diverge is always below the critical temperature and so the dual algorithm still provides a very good description of the entire paramagnetic phase.

Unfortunately, in the 3D EA model the divergence of U fields takes place well above the critical temperature (which is $T_c \simeq 1.12$ for bimodal coupling and $T_c \simeq 0.95$ for Gaussian

couplings; see Ref. [25] for a summary of critical temperatures in 3D spin glasses), and this would make the dual algorithm of very little use. We have studied the origin of this divergence and have found a general principle for reducing the divergence of U fields due to self-reinforcement, thus improving the convergence properties of the dual algorithm. The idea is the following. When the dual approximation is written as a constraint satisfaction problem with a nonuniform prior [see Eq. (15)], the constraints may be redundant. This is the case for the 3D cubic lattice: indeed both the number of links (i.e., variables in the dual problem) and the number of plaquettes (i.e., constraints in the dual problem) are $3N$. So, if constraints were independent, the entropy would be null at $\beta = 0$ and negative for $\beta > 0$ (and this is clearly absurd). The solution to the apparent paradox is that constraints are not independent: actually only $2/3$ of these are independent, and the remaining third is uniquely fixed by the value of the former. In this way the correct entropy is recovered at $\beta = 0$, given that a problem with $3N$ unbiased binary variables subject to $2N$ independent parity-check constraints has entropy $N \ln(2)$. The dependence among constraints can be easily appreciated by looking at the six plaquettes around a cube: if five of the six constraints are satisfied, then the sixth one is automatically satisfied and redundant.

The general rule for improving the convergence of the dual algorithm is to remove redundant constraints (this principle is similar to the maxent-normal property of region-based free energy approximations [5]). Redundant constraints have no role in determining the fixed point values for the beliefs (since they are redundant), but during the iterations they provide larger fluctuations to messages and may be responsible for the lack of convergence. In practice, on a 3D cubic lattice, we may remove redundant constraints in many different ways: the basic rule states that one constraint (i.e., a plaquette) should be removed for each elementary cube, otherwise if a cube remains with its six plaquettes at least one redundant constraint will exist. We are going to present data obtained by removing all constraints corresponding to plaquettes in the xy plane.

The dual algorithm for the 3D EA model on the cubic lattice with no redundant constraints converges for any temperature above $T \simeq 0.8$, and so we can use it to study the entire paramagnetic phase. The lack of convergence deep in the spin glass phase is to be expected. Just as in the 2D case, the dual algorithm (when it converges) still finds the same solution obtained by a double-loop algorithm, and again it finds the solution nearly 100 times faster (see Fig. 12). The double-loop algorithm has the apparent advantage of converging at any temperature, even at very low ones. However, deep in the spin glass phase, where the underlying paramagnetic approximation is clearly inaccurate, we believe that an algorithm (like the dual one) that stops converging is providing an important warning that something wrong is probably happening. Such a warning would be lacking in using a double-loop algorithm.

In Fig. 13 the correlations predicted by the dual approximation and those obtained by a parallel tempering Monte Carlo simulation are compared. At high temperatures the correspondence is quite good, but not as good as in 2D. However, it is important to stress that the 3D EA model is much more difficult to simulate than the 2D case: there is no

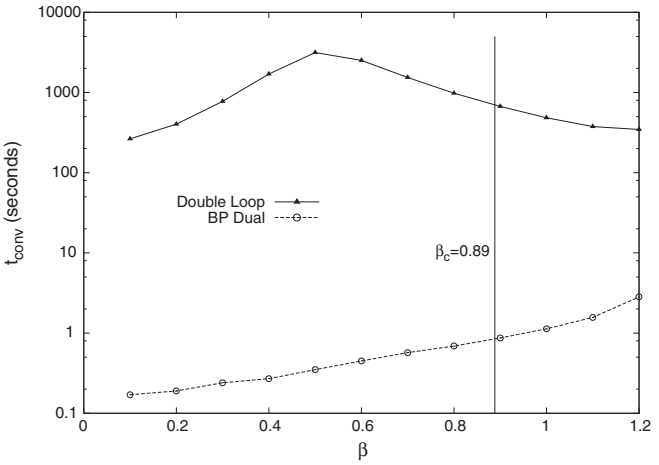

FIG. 12. Running times of the double-loop algorithm [7,23] (LIBDAI) and of the dual algorithm on an $8 \times 8 \times 8$ EA model with bimodal interactions ($J_{ij} = \pm 1$). The dual algorithm is generally several orders of magnitude faster and returns the same solution as does the double-loop algorithm.

fast algorithm for sampling the configuration space according to the Gibbs distribution as in the 2D case [26], and Monte Carlo methods require huge thermalization times, especially in the vicinity of the spin glass transition and in the spin glass phase.

In Fig. 14 we show the estimates for the energy obtained from the Monte Carlo simulation and the dual algorithm (both on a single sample and in the average case). The good agreement between the dual algorithm results on single samples and in the average case is telling us that U messages arriving at a given point on the lattice are uncorrelated to a very large extent. In other words, the effect of short loops in the lattice is not manifestly present in correlations between

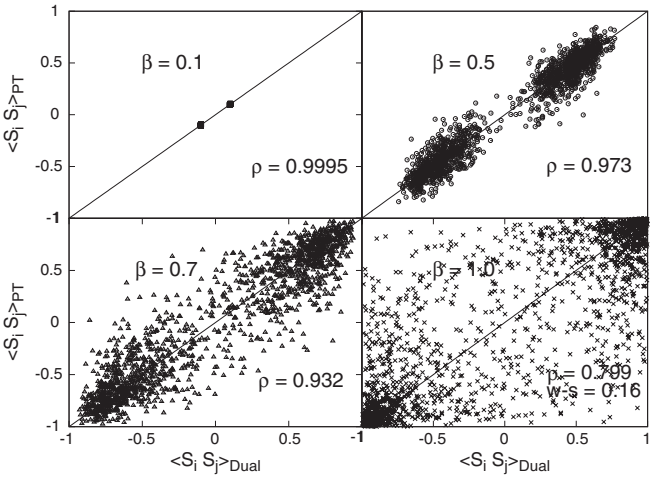

FIG. 13. Comparison between the correlations $\langle s_i s_j \rangle_{\text{dual}}$ obtained with the dual algorithm and the (nearly) exact correlations $\langle s_i s_j \rangle_{\text{PT}}$ obtained with a parallel tempering simulation in a 3D EA model of size $8 \times 8 \times 8$ with random bimodal interactions $J_{ij} = \pm 1$. At each temperature the correlation coefficient ρ is reported. For the lowest temperature shown, $\beta = 1.0$, we also report the fraction $w - s$ of pairs of spins such that $\langle s_i s_j \rangle_{\text{dual}} \langle s_i s_j \rangle_{\text{PT}} < 0$.

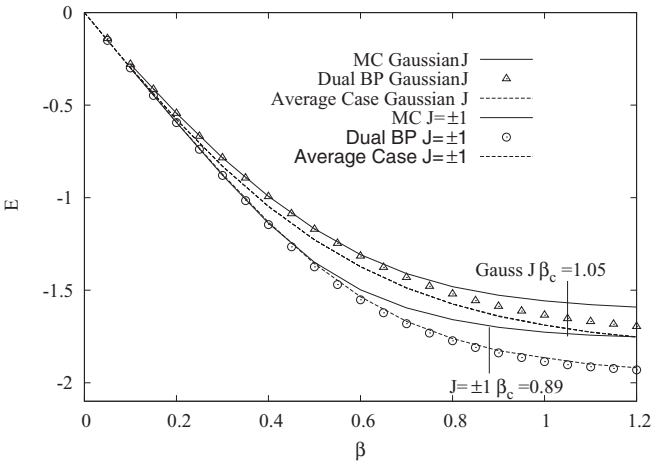

FIG. 14. The energy predicted by the dual approximation in the 3D EA model, compared to the average case energy, and the Monte Carlo simulation. We used an $8 \times 8 \times 8$ system with both types of random interaction, bimodal ($J_{ij} = \pm 1$) and Gaussian distributed.

messages. In contrast, the comparison between dual algorithm and Monte Carlo results is good only at high temperatures, and it degrades when the critical temperature is approached. This discrepancy can be understood as due to a growing correlation length in the EA model that diverges at the critical temperature: our dual approximation does not account for correlations beyond the plaquette level and so it becomes inevitably poorer when the correlation length diverges. However, in all those situations when the running time must be kept short, the dual algorithm is able to provide more precise marginals than the Monte Carlo method in the paramgnetic region. Indeed, in Fig. 15 we show the error Δ_2 obtained with the dual algorithm and with the Monte Carlo method (with parallel tempering

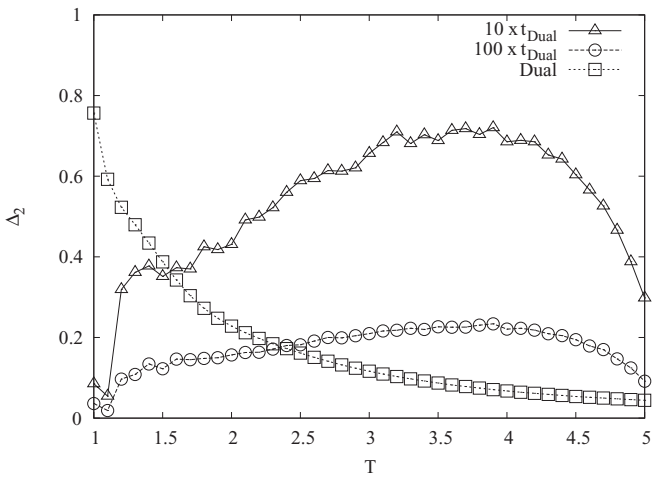

FIG. 15. Error Δ_2 made by a Monte Carlo simulation (with parallel tempering) for the estimation of first-neighbor correlations in a 3D EA model of size 8^3. The Monte Carlo simulation is run for a time that is 10 and 100 times the convergence time of the dual algorithm. The error made by the dual approximation is also reported and is lower than the Monte Carlo one in the high-temperature phase, not too close to the critical temperature.

algorithm) if the latter is run for a time that is 10 or 100 times the dual algorithm convergence time. Monte Carlo methods are to be preferred only if the running time can be long or if there is a spontaneous symmetry breaking (not accounted by the dual approximation). But in all those situation where there is no long-range order and the system is very large and/or the time at disposal is short (relative to the system size), the dual algorithm is to be preferred, since it outperforms Monte Carlo results.

VI. CONCLUSIONS

We have introduced a dual algorithm to compute marginal probabilities in the paramagnetic phase of frustrated spin models (e.g., spin glasses) on finite-dimensional lattices. Inspired by the fact that in a paramagnetic phase with no external field each variable is unbiased (i.e., local magnetizations are null), the dual algorithm is derived by adding such paramagnetic constraints in the GBP equations. While BP (i.e., Bethe approximation) and GBP algorithms have serious convergence problems at low temperatures even in the paramagnetic phase, the dual algorithm converges very fast in a much wider range thanks to these constraints. The dual algorithm can also be seen as BP on a dual lattice, where the interactions J_{ij} act as external fields on dual variables, thus improving convergence properties of the message-passing algorithm.

We have tested the dual algorithm for the Edwards-Anderson spin glass model with bimodal and Gaussian couplings on 2D (square) and 3D (cubic) lattices. The results are very encouraging, showing convergence in the whole paramagnetic phase (and even slightly in the frozen phase for the 3D EA model) and comparing very well with exact correlations measured in Monte Carlo simulations. A comparison with a double-loop algorithm (which is the state of the art among general purpose inference algorithms) shows that both algorithms find the same result, but our dual algorithm runs roughly 100 times faster. We also tried to push the dual approximation to the limit, and we used the correlations inferred from the dual algorithm to compute ground state configurations in the 2D EA model by a freezing procedure. Again, we showed that the ground states obtained in this way compare very well with those from exact computations.

The success of our proposal clearly shows that, as long as variables are not long-range correlated, the computation of correlations in a generic spin model can be done in a very fast way by means of message-passing algorithms, based on mean-field-like approximations. These kinds of inference algorithm do not provide in general an exact answer (unless they are used at very high temperatures or on locally treelike topologies), and so they can not be seen as substitutes for a Monte Carlo (MC) sampling. However, there are many situations where a fast and approximate answer is required more than a slow and exact answer. Let us just give a couple of examples of these situations. On the one side, if one needs to sample from very noisy data, an approximated inference algorithm whose level of approximation is smaller than the data uncertainty is as valid as a perfect MC sampler. On the other side, if one needs to use the inferred correlations as input for a second algorithm (as for the freezing algorithm in Sec. IV B) that will eventually

modify or correct these correlations, a fast and reasonably good inference is enough.

The promising results shown in the present work naturally ask for an improvement in several directions. For example, in the paramagnetic phase of a model defined on a 3D lattice, our inference algorithm could be improved by using the $2 \times 2 \times 2$ cube as the elementary region, instead of the plaquette. An even more important improvement would be to extend the applicability range of the algorithm to the low-temperature phase; but this requires a rather nontrivial modification, since in the low-temperature phase the assumption of zero local magnetization needs to be broken.

ACKNOWLEDGMENTS

F.R.-T. acknowledges financial support by the Italian Research Minister through the FIRB Project No. RBFR086NN1 on "Inference and Optimization in Complex Systems: From the Thermodynamics of Spin Glasses to Message Passing Algorithms."

[1] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge University Press, Cambridge, 2003).
[2] M. Mézard and A. Montanari, *Information, Physics, and Computation*, Oxford Graduate Texts (Oxford University Press, Oxford, 2009).
[3] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, 1999).
[4] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, San Francisco, 1988).
[5] J. Yedidia, W. T. Freeman, and Y. Weiss, IEEE Trans. Inf. Theory **51**, 2282 (2005).
[6] A. L. Yuille, Neural Comput. **14**, 1691 (2002).
[7] T. Heskes, C. Albers, and H. Kappen, in *Proceedings of the Conference on Uncertainty in Artificial Intelligence, 2003*, edited by Christopher Meek and Uffe Kjærulff (Morgan Kaufmann, San Francisco, 2003), p. 313.
[8] A. Pelizzola, J. Phys. A **38**, R309 (2005).
[9] R. Kikuchi, Phys. Rev. **81**, 988 (1951).
[10] Y. Kabashima, J. Phys. Soc. Jpn. **74**, 2133 (2005).
[11] T. Rizzo, A. Lage-Castellanos, R. Mulet, and F. Ricci-Tersenghi, J. Stat. Phys. **139**, 375 (2010).
[12] M. Mézard and G. Parisi, Eur. Phys. J. B **20**, 217 (2001).
[13] M. Mézard and G. Parisi, J. Stat. Phys. **111**, 1 (2003).
[14] T. Castellani, F. Krzakala, and F. Ricci-Tersenghi, Eur. Phys. J. B **47**, 99 (2005).
[15] T. Jörg, J. Lukic, E. Marinari, and O. C. Martin, Phys. Rev. Lett. **96**, 237205 (2006).
[16] M. Mézard and R. Zecchina, Phys. Rev. E **66**, 056126 (2002).
[17] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, J. Stat. Mech. (2008) P04004.
[18] F. Ricci-Tersenghi and G. Semerjian, J. Stat. Mech. (2009) P09001.
[19] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **89**, 268701 (2002).
[20] Y. Kabashima, J. Phys. Soc. Jpn. **72**, 1645 (2003).
[21] A. Pagnani, G. Parisi, and M. Ratiéville, Phys. Rev. E **68**, 046706 (2003).
[22] J. M. Mooij and H. J. Kappen, IEEE Trans. Inf. Theory **53**, 4422 (2007).
[23] J. M. Mooij *et al*., computer code LIBDAI, http://www.libdai.org/.
[24] M. Juergen *et al*., http://www.informatik.uni-koeln.de/spinglass/.
[25] H. G. Katzgraber, M. Körner, and A. P. Young, Phys. Rev. B **73**, 224432 (2006).
[26] C. K. Thomas and A. A. Middleton, Phys. Rev. E **80**, 046708 (2009).