Coarse-grained belief propagation for simulation of interacting quantum systems at all temperatures

Ersen Bilgin[1],[*] and David Poulin[2],[†]

[1]*Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA*
[2]*Département de Physique, Université de Sherbrooke, Sherbrooke, Québec, Canada*
(Received 26 October 2009; revised manuscript received 6 January 2010; published 9 February 2010)

We continue our numerical study of quantum belief propagation initiated in [Phys. Rev. A **77**, 052318 (2008)]. We demonstrate how the method can be expressed in terms of an effective thermal potential that materializes when the system presents quantum correlations, but is insensitive to classical correlations. The thermal potential provides an efficient means to assess the precision of belief propagation on graphs with no loops. We illustrate these concepts using the one-dimensional quantum Ising model and compare our results with exact solutions. We also use the method to study the transverse field quantum Ising spin glass for which we obtain a phase diagram that is largely in agreement with the one obtained in [Phys. Rev. B **78**, 134424 (2008).] using a different approach. Finally, we introduce the coarse-grained belief propagation (CGBP) algorithm to improve belief propagation at low temperatures. This method combines the reliability of belief propagation at high temperatures with the ability of entanglement renormalization to efficiently describe low-energy subspaces of quantum systems with local interactions. With CGBP, thermodynamic properties of quantum systems can be calculated with a high degree of accuracy at all temperatures.

PACS number(s): 05.30.−d, 02.70.−c, 05.50.+q, 75.50.Lk

I. INTRODUCTION

The most interesting phenomena in condensed matter physics occur when a large number of quantum particles are put into interaction. However, even the most simplified models of these interactions rarely admit analytical solutions. For this reason, numerical methods must be called into play. As direct numerical simulations require resources that scale exponentially with the number of particles, approximation schemes are needed to understand these systems at their thermodynamic limit.

Starting with White's density matrix renormalization group (DMRG) method,[1] families of states have been introduced that accurately describe the low-energy sector of locally interacting quantum systems using few parameters.[2–4] In particular, matrix product states approximate ground states of one-dimensional gapped Hamiltonians[5,6] and projected entangled pair states are accurate in higher dimensions at finite temperature.[7,8]

Even though these states require few parameters, finding the right set of parameters for a given system remains a formidable task.[9] Only in special cases, such as simulating imaginary time evolution within a variational set of states,[10] or using iterative minimization procedures,[3] have successful heuristics been devised that solve this problem.

In this article, we present *coarse-grained belief propagation* (CGBP), which combines the strength of two such heuristics. The first method, devised by Vidal, is a refinement of real space renormalization[11] called *entanglement renormalization* (ER).[12,13] The strength of this method comes from the realization that entanglement is organized on different length-scales in the ground state of many systems of interest, including critical systems[13–15] and systems with topological order.[16,17] As a consequence of this organization, entanglement can be efficiently removed from the state by a sequence of local coarse-graining transformations. This leads to a very

efficient scheme for finding low-energy states of local Hamiltonians, making ER a very effective method for low temperatures.

The second method is belief propagation (BP), which is a well-studied and widely used method to solve inference problems involving a large number of correlated random variables (see, e.g., Refs. 18–20). This algorithm is exact on trees where it essentially reduces to a transfer matrix solution. On more general graphs, it can be described as performing a constrained minimization of the Bethe free energy of the system.[20] Therefore, it is often a very good approximation on graphs containing no small loops. On graphs with small loops, generalizations of BP can also provide reliable approximations.[20,21]

Belief propagation was recently generalized to quantum theory by us and others.[22–25] The gist of the method consists of solving the system exactly on a small cluster and using this solution to compute effective thermal Hamiltonians on the neighboring clusters. This procedure is repeated iteratively until it produces a correction to the bare Hamiltonian that accurately mimics a system of infinite size. The method is very accurate at high temperatures and on trees or graphs with only large loops.[25] However, BP becomes unreliable at low temperatures because the clusters must be larger or equal to the range of the effective Hamiltonian, which grows like the inverse of the temperature.

Coarse-grained belief propagation combines the features of BP and ER. Starting at high temperatures where plain BP is accurate, temperature is lowered until coarse graining the lattice by eliminating its shortest length-scale degrees of freedom becomes favorable. Coarse graining discards some high-energy states, which results in a systematic error in the thermal state. On the other hand, it increases the effective size of the clusters, making BP more accurate. The coarse-graining procedure is continued as temperature is lowered to zero where plain ER is accurate. An estimate of the error

caused by BP can be used to determine the temperatures at which each coarse-graining procedure should be performed.

In Sec. II, we review the BP method and describe how to estimate its accuracy. Based on this technique, Sec. III presents results obtained for the transverse field quantum Ising spin glass studied in Ref. 23. Finally, in Sec IV we present the CGBP algorithm and benchmark it using an infinite quantum Ising chain.

II. BELIEF PROPAGATION

A. Classical setting

Consider a system composed of N classical spins on a chain with Hamiltonian $H(x_1, \ldots, x_N) = \Sigma_{\langle i,j \rangle} h(x_j, x_j)$, where $\langle i, j \rangle$ denotes neighboring sites. The partition function for this system is defined as

$$\mathcal{Z}(\beta) = \sum_{x_1, \ldots, x_N} e^{-\beta H(x_1, \ldots, x_N)}, \tag{1}$$

$$= \sum_{x_1, \ldots, x_N} e^{-\beta h_{N,N-1}} \cdots e^{-\beta h_{3,2}} e^{-\beta h_{2,1}}, \tag{2}$$

where we use the shorthand $h_{i,j} = h(x_i, x_j)$. A brute force calculation of this quantity requires summing over an exponential (in N) number of terms. However, taking advantage of the local structure of the Hamiltonian, the sum can be rearranged as follows:

$$\sum_{x_N} e^{-\beta h_{N-1,N}} \left[\cdots \sum_{x_2} \left(e^{-\beta h_{2,3}} \sum_{x_1} e^{-\beta h_{1,2}} \right) \cdots \right].$$

Now, the sums can be performed sequentially with each sum only involving a small number of terms,[26] allowing the computation of $\mathcal{Z}$ in a time proportional to N. This leads to an iterative rule where messages $m_{i \to j}$ are exchanged along the edge of the chain connecting two sites i, j with the update rule $m_{i \to i+1}(x_{i+1}) = \Sigma_{x_i} e^{-\beta h(x_i, x_{i+1})} m_{i-1 \to i}(x_i)$. With the initialization $m_{0 \to 1}(x_1) = 1$, we see that $\mathcal{Z}(\beta) = \Sigma_{x_N} m_{N-1 \to N}(x_N)$.

This procedure can be generalized to arbitrary graphs by defining the update rule for the message $m_{i \to j}(x_j)$, passed from site i to a neighboring site j, to be

$$m_{i \to j}(x_j) = \sum_{x_i} e^{-\beta h(x_i, x_j)} \prod_{k \in \mathcal{N}(i) \backslash j} m_{k \to i}(x_i), \tag{3}$$

where $\mathcal{N}(i) \backslash j$ denotes the set of neighbors of site i other than j. On a tree, these messages will converge to their final values after a number of iterations equal to tree's diameter, where one iteration is defined to be to be the updating of all messages simultaneously. The one- and two-body beliefs $b_j(x_j) = \frac{1}{\mathcal{Z}(\beta)} \Pi_{k \in \mathcal{N}(j)} m_{k \to j}(x_j)$ and $b_{i,j}(x_i, x_j) = \frac{1}{\mathcal{Z}(\beta)} \Pi_{k \in \mathcal{N}(i)} m_{k \to i}(x_i) \Pi_{k' \in \mathcal{N}(j)} m_{k' \to j}(x_j) e^{-\beta h(x_i, x_j)}$ are equal to the reduced one- and two-body distribution, respectively, and the partition function can be evaluated from any one of their normalization. When the underlying graph contains loops, BP is no longer exact but often provides accurate approximation to the true marginal states and partition function.

B. Quantum setting

Consider now the quantum case $H = \Sigma_{\langle i,j \rangle} h_{i,j}$ where $h_{i,j}$ are Hermitian operators acting on site i and j of a chain. Defin-

ing the $\odot$ product between positive operators $A \odot B = e^{\log(A) + \log(B)}$, we can write the partition function in a form very similar to Eq. (2)

$$\mathcal{Z}(\beta) = \text{Tr}(e^{-\beta H}), \tag{4}$$

$$= \text{Tr}(e^{-\beta h_{N,N-1}} \odot \cdots e^{-\beta h_{3,2}} \odot e^{-\beta h_{2,1}}), \tag{5}$$

with the sums replaced by traces and products by $\odot$ products. This expression cannot be reorganized like its classical counterpart because—unlike ordinary products—the $\odot$ product does not obey a distributive law in the sense that $\text{Tr}_a(e^{-\beta h_{c,b}} \odot e^{-\beta h_{b,a}}) \neq e^{-\beta h_{c,b}} \odot \text{Tr}_a(e^{-\beta h_{b,a}})$.

However, the distributive law holds when the chain forms a quantum Markov network,[24] i.e., $I(a:c|b) = 0$ where the quantity $I(a:c|b) = S(a,c) + S(b,c) - S(b) - S(a,b,c)$ is known as the conditional mutual information and S is the von Neumann entropy. Intuitively, this condition means that all correlations between a and c are mediated through b. This does not prevent a and c from being correlated, but implies that all information learned about a by measuring c can equivalently be obtained by measuring b instead. While this condition does not hold for generic quantum interactions, it becomes approximately true when the "Markov shield" b is sufficiently thick.

To understand this behavior physically, consider again a one-dimensional chain with nearest-neighbor interactions $H = \Sigma_i h_{i,i+1}$. The thermal Gibbs state at inverse temperature β is given by $\rho = e^{-\beta H} / \mathcal{Z}(\beta)$. The reduced state of spins $2, 3, \ldots, N$ is obtained by tracing out the first spin of the chain, i.e., $\rho_{2,\ldots,N} = \text{Tr}_1(\rho)$. We can formally define the *effective thermal Hamiltonian* H_{eff} acting on sites $2, \ldots, N$ up to normalization by

$$\rho_{2,\ldots,N} = e^{-\beta H_{\text{eff}}}. \tag{6}$$

In other words, H_{eff} is the traceless part of $-\frac{1}{\beta} \log(\rho_{2,\ldots N})$ where log is the principal matrix logarithm. Similar effective Hamiltonians have also been studied in Ref. 27 in the context of reduced density matrices of fermionic and bosonic ground states. We also define the *effective thermal potential* $V_{\text{eff}} = H_{\text{eff}} - \Sigma_{i=2}^{N-1} h_{i,i+1}$ as the term added to the bare Hamiltonian on sites 2 to N due to the presence of site 1.

At high temperatures, the effective thermal potential is short ranged. Consider for instance the one-dimensional Ising model with transverse field on an infinite chain $H = \Sigma_{i=-\infty}^{\infty} \sigma_i^z \sigma_{i+1}^z + B \sigma_i^x$. At zero temperature, this model exhibits a phase transition at the critical transverse field value $B = 1$. Figure 1 shows the value of the effective potential $V_{\text{eff}} = -\frac{1}{\beta} \log \text{Tr}_{-\infty,\ldots,0}(e^{-\beta H}) - H_{\text{bare}}$ obtained from cutting the critical Ising chain in half, i.e., tracing out spins $-\infty$ to 0 from the thermal state of an infinite chain. What is plotted is (an upper bound to) the operator norm of the cumulants of V_{eff}: $V_{\text{eff}}^j = \lim_{N \to \infty} \text{Tr}_{j+1,\ldots,N}(V_{\text{eff}} - \Sigma_{k=1}^{j-1} V_{\text{eff}}^k) / 2^{N-j}$ with $V_{\text{eff}}^1 = \lim_{N \to \infty} \text{Tr}_{2,\ldots,N}(V_{\text{eff}}) / 2^{N-1}$. We see that V_{eff} has a very short range, in fact much shorter than the correlation length in the system.

The effective hamiltonian is expected to be shorter ranged than the correlations in the system because it only responds to quantum correlations. For classical systems (e.g., when
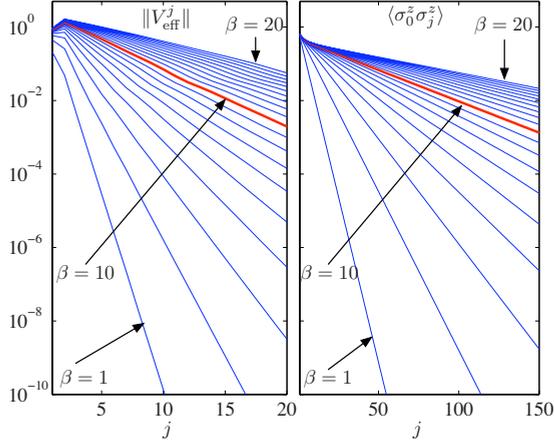
FIG. 1. (Color online) Upper bound on $\|V_{\mathrm{eff}}^j\|$ (left) and correlations (right) of the Ising chain with critical transverse field. These exact values are obtained from a Jordan-Wigner transform. Note that $\|V_\infty^j\|$ decays much faster than the correlations (notice the different length scales). The values for $\beta=10$ are in a different color as an aid to the eye.

FIG. 2. Calculating $m_{4\rightarrow5}$ from $m_{3\rightarrow4}$ in an iteration of BP algorithm with $l=4$. In the first step $e^{h_{4,5}}$ is added to $m_{3\rightarrow4}$ using the $\odot$ product. Then, the first spin is traced out yielding $m_{4\rightarrow5}$.

$B=0$ in transverse Ising), V_{eff} acts nontrivially only on one qubit at the boundary, even though longer-ranged classical correlations are present. Turning on a transverse field B creates genuine quantum correlations (i.e., entanglement) and increases the range of V_{eff}. Similar effect can be obtained by cooling the system since the system approaches a classical state at high temperature.

The distributive law becomes approximately true when the distance between the traced-out site and the end of the cluster is sufficiently large to support the effective thermal potential. Thus, we can reorganize the sum of Eq. (5) as in the classical case, but keeping an l site Markov blanket between the traced-out site and the end of the cluster:

$$\mathcal{Z}(\beta)=\mathrm{Tr}_{1,\dots,N}(e^{-\beta h_{N-1,N}}\dots\odot e^{-\beta h_{1,2}}), \tag{7}$$

$$\approx\mathrm{Tr}_{2,\dots,N}[e^{-\beta h_{N-1,N}}\dots\odot e^{-\beta h_{l+1,l+2}}\odot$$
$$\times\mathrm{Tr}_1(e^{-\beta h_{l,l+1}}\odot\ \dots\ \odot e^{-\beta h_{1,2}})], \tag{8}$$

$$\approx\mathrm{Tr}_{3,\dots,N}[e^{-\beta h_{N-1,N}}\dots\odot e^{-\beta h_{l+2,l+3}}\odot$$
$$\times\mathrm{Tr}_2(e^{-\beta h_{l+1,l+2}}\odot m_{l\rightarrow l+1})], \tag{9}$$

$$\approx\dots, \tag{10}$$

$$\approx\mathrm{Tr}_{N-l,\dots,N}(e^{-\beta h_{N-1,N}}\odot m_{N-1\rightarrow N}), \tag{11}$$

where we have defined $m_{i\rightarrow i+1}=\mathrm{Tr}_{i-l+1}(e^{-\beta h_{i,i+1}}\odot m_{i-1\rightarrow i})$, which is an operator acting on sites $i-l+2$ to $i+1$ (see Fig. 2). These equations require manipulating operators on up to $l+1$ spins, so their complexity increases exponentially with the window size l.

Like in the classical setting Eq. (3), these message passing rules can be generalized to arbitrary graphs, enabling the computation of one- and two-body beliefs from which vari-
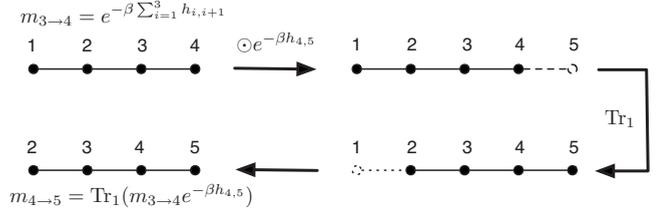
ous quantities of interest such as energy can be computed. The method can also be adapted to estimate all correlation functions.

Because of computational limitations, the effective thermal potential V_{eff} cannot be computed exactly. Instead, we can only estimate its value on a cluster of finite size. Thus, the main source of error in our method is due to the truncation of V_{eff}. We can assess the error by evaluating the portion of V_{eff} we discard. On a chain for instance, the error caused on the estimate of the beliefs is

$$\frac{1}{\mathcal{Z}(\beta)}\left\{e^{-\beta H_{\mathrm{eff}}}-\exp\left[-\beta\left(H_{\mathrm{eff}}-\sum_{j>l}V_{\mathrm{eff}}^j\right)\right]\right\}\approx\beta\sum_{j>l}\langle V_{\mathrm{eff}}^j\rangle, \tag{12}$$

for $\beta V_{\mathrm{eff}}^j\ll1$. Making the assumption (see Fig. 1) that $\|V_{\mathrm{eff}}^j\|$ decreases exponentially with j, we estimate this quantity by

$$\beta\langle V_{\mathrm{eff}}^{l+1}\rangle\approx\beta\langle V_{\mathrm{eff}}^l\rangle\left(\frac{\langle V_{\mathrm{eff}}^l\rangle}{\langle V_{\mathrm{eff}}^{l-1}\rangle}\right), \tag{13}$$

which is our final error estimate. With the same reasoning, similar estimates can be derived for the error associated with different observables.

Figure 3 compares this error estimate to the true error produced by BP for the evaluation of the energy density of the critical one-dimensional Ising chain. Clearly, the error estimate accurately bounds the true error. This figure also illustrates the power of BP by comparing the accuracy with
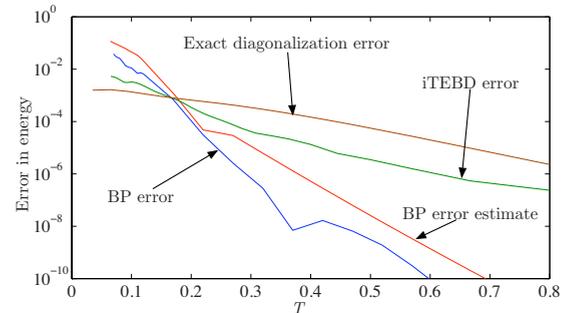

FIG. 3. (Color online) Error on the energy density as a function of temperature for the critical Ising chain. For QBP with window size $l=10$, we show the error estimate of Eq. (13) and the true error obtained by comparison with analytical solution. Also shown is the finite-size error for exact diagonalization of a 11-site chain and error caused by iTEDB (Ref. 10) with parameters $\delta T=0.001$ and $\chi=150$, both of which require equivalent computational resources.
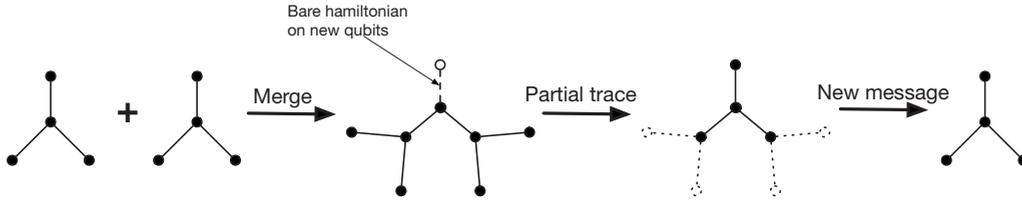
FIG. 4. Schematic illustration of the procedure to calculate the message from a node to its parent. First, two messages from its children are merged using the $\odot$ product. The bare hamiltonian term relating the current node to its parent is also added. Tracing out the leaves of this eight-spin message gives the message to the node's parent.

which BP estimates the energy density of an infinite chain to what is achieved by brute force diagonalization with equivalent computational resources. BP largely outperforms diagonalization down to temperatures of order 0.1. This behavior is expected because, as seen on Fig. 1 (red line), the range of the effective thermal potential becomes larger than the window size ($l=10$) at this temperature.

III. DISORDERED SYSTEM REVISITED

The use of belief propagation in physics originates in the study of disordered systems—spin glasses—where it is more often referred to as the "cavity method."[28] Along with the sliding window algorithm outlined in the previous section, we presented in[24,25] a second distinct way to generalize BP to the quantum setting. This method, which we named "replica BP," maps the quantum system to a classical system with one additional spatial dimension of length equal to the inverse temperature β. The edges of a graph become ribbons on which classical BP can be employed for sufficiently low β. This technique was independently introduced by Laumann *et al.*[23] for the study of the transverse field quantum Ising spin glass. The continuous imaginary time limit of this procedure was later studied in Ref. 29.

The numerical results obtained in Ref. 25 suggest that, for a given amount of computational power, sliding window BP is much more accurate than replica BP. In this section, we revisit the spin-glass model of Ref. 23 using sliding window BP, and apply the method outlined in the previous section to estimate the accuracy of our results.

The spins are located at the vertices of a degree-3 Cayley tree. The Hamiltonian has Ising coupling between neighboring spins and a transverse field B:

$$H = \sum_{\langle ij \rangle} \sigma_i^z \sigma_j^z + B \sum_i \sigma_i^x + \sum_{i \in \text{Boundary}} r_i \sigma_i^z. \quad (14)$$

The last term is a random parallel boundary field introduced to create frustration in the system. The strength of the boundary fields r_i are chosen at random uniformly in $[-1,1]$. The quantity of interest in this setting is the Edwards-Anderson order parameter $q_{EA} = \langle \sum_j \langle \sigma_j^z \rangle^2 \rangle_Q$ on the lattice at the thermodynamic limit, where $\langle \cdot \rangle_Q$ refers to the quench average over the random boundary field configurations. Note that this order parameter is defined along the axis perpendicular to the applied external field B, but parallel to the random boundary field. Hence, q_{EA} is zero in the paramagnetic phase. It becomes nonzero on the onset of the glassy phase where the

system settles into one of many metastable randomly polarized state.

Our numerical simulations are performed on a tree of depth 12. To further reduce the finite-size effects, the Edwards-Anderson (EA) order parameter is only evaluated on the central spin of the lattice, away from the boundary. The order parameter can be computed from the single-body belief b_j obtained from BP. On a chain, we would start from one end and propagate messages to the other end of the chain. On the tree, we start from the leaves and propagate messages toward the middle. The basic message passing step is demonstrated in Fig. 4. First, two messages from the children of a given node are combined using the $\odot$ product. The bare Hamiltonian term relating the node to its parent is then added. Tracing out the leaves of this eight-spin message gives the message to the node's parent. This procedure is repeated until the central site is reached where three messages are joined to produce the belief. The quench average is obtained by repeating this procedure 100 times with different boundary field configurations.

Because the graph contains no loops, the error estimate presented in the previous section is reliable. However, note that the statistical fluctuations of the quench average are not included in this error estimate. The statistical fluctuations of the average of q_{EA} over the many instances of boundary fields range from 2% at low temperatures to 14% at high temperatures. Therefore, the main source of error in parts of Fig. 5 is the statistical fluctuations which are not shown in the plot, and can be systematically reduced by increasing the sample size.

Figure 5 shows the EA order parameter q_{EA} in the transverse field-temperature diagram, along with the estimated BP error. The glassy/paramagnetic phase transition line agrees with results of[23] up to the statistical fluctuations above $T \gtrsim 0.3$. At low temperatures ($T \lesssim 0.3$), our results indicate a phase transition line with a decreasing value of B as T is lowered. We could not think of any physical mechanism that could explain this behavior. Moreover, this happens in a region of the phase diagram where the BP error is high. Hence, we suspect that the true phase transition line has a monotonous behavior in temperature and that the glassy phase persists all the way to zero temperature for a transverse field $B \lesssim 1.65$. This conclusion and, more generally, our entire phase diagram is in very good agreement with that of.[23]

IV. COARSE-GRAINED BELIEF PROPAGATION

We have seen that BP provides reliable estimates of thermal expectation values as long as the effective thermal po-
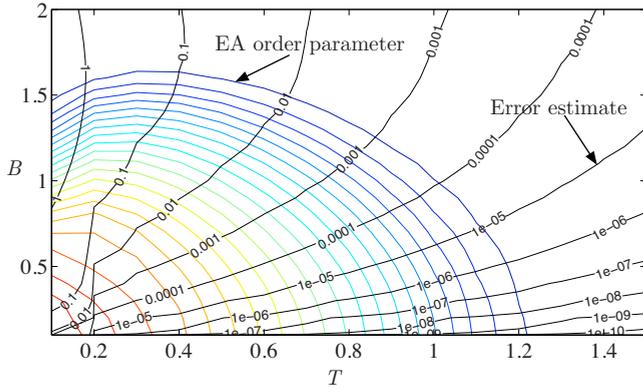
FIG. 5. (Color online) The Edwards-Anderson order parameter for transverse field quantum Ising spin glass on a degree-3 Cayley tree with random parallel boundary fields. The total depth of the tree is 12, and the plots show the average of 100 instances of random boundary fields. Values of q_{EA} range from 0 to 1 and contours are equally spaced. The error on the order parameter is estimated using the procedure outlined at the end of Sec. II B.

tential V_{eff} is sufficiently short ranged to be tracked numerically. As seen e.g., in Fig. 1, the range of V_{eff} grows linearly with β, so BP becomes unreliable at low temperatures. To probe lower temperatures, one needs to increase the window size l, which is not feasible because resources scale as $O(2^l)$. On the other hand, as we lower the temperature, high-energy excitations become increasingly irrelevant. This fact leads to efficient algorithms for zero temperature simulations such as entanglement renormalization and DMRG. These algorithms become rapidly inaccurate at finite temperature because they are only able to keep track of a small number of eigenstates. In this section, we will return to one-dimensional systems and describe a method that interpolates between BP at high T to ER at $T=0$. Before we do so, we briefly review ER, see[12,13] for a detailed description.

A. Entanglement renormalization

Entanglement renormalization[12,13] builds on the multiscale renormalization ansatz (MERA) which asserts that certain degrees of freedom can be decoupled from the ground state of local Hamiltonians by unitary transformations acting on small spatial regions. A concrete example of this scheme is illustrated in Fig. 6 in the case of a one-dimensional lattice. The lattice is first partitioned into clusters each containing three consecutive sites. A disentangling transformation u [a unitary transformation on $(\mathbb{C}^\chi)^{\otimes 2}$] is applied on the boundary of each cluster in order to minimize the correlations be-
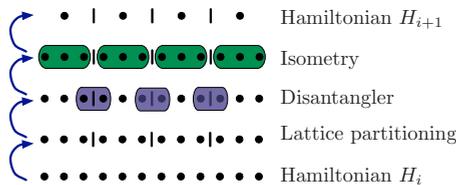

FIG. 6. (Color online) Schematics of entanglement renormalization for a ternary MERA of a one-dimensional lattice.

tween neighboring clusters. Finally, local degrees of freedom are discarded from each cluster by means of an isometry[30] v mapping $(\mathbb{C}^\chi)^{\otimes 3}$ to $\mathbb{C}^\chi$.

This procedure is applied repeatedly. At each iteration, the disentanglers and isometries transform the Hamiltonian H_i to a Hamiltonian H_{i+1} acting on a smaller lattice and retaining only the lowest eigenstates of the previous Hamiltonian. ER is halted when only a few sites remain in the lattice so it can be handled exactly numerically. The disentanglers and isometries are chosen in such a way that the final state (or more generally subspace) minimizes the energy of the initial Hamiltonian H_0. This minimization problem is in general hard, but good heuristics have been devised for it.[12,13]

B. Interpolated scheme: CGBP

Both ER and BP revolve around the idea that some correlations are short range in the state of interest. BP becomes exact when the conditional mutual information $I([-\infty, j] : [j+l, \infty] | [j+1, j+l-1])$ vanishes for a sufficiently large window size l. In other words, BP can work in the presence of arbitrary long range classical correlations but the purely quantum correlations must be short ranged. These quantum correlations tend to increase like the inverse temperature β (cf. Fig. 1), which limits BP to relatively high temperatures, unless the window size l can somehow be increased while keeping computational cost low. In contrast, ER makes use of disentanglers to eliminate short-range quantum correlations in the system and coarse grain the lattice. Because it only keeps a few low-energy states, it is limited to very low temperatures.

Coarse-grained belief propagation interpolates between these two methods and provides accurate thermal expectation values over a very large range of temperatures. At high temperature, CGBP reduces to ordinary BP. As the temperature is lowered, the error attributed to BP increases. At some temperature T_1, it becomes favorable to coarse grain the lattice using one step of ER. This procedure discards some high-energy states, leading to a systematic error, but it effectively increases the BP window length l by a constant factor (three in the ternary ER scheme illustrated in Fig. 6). This increase in the window length improves the accuracy of BP and compensates for the loss of high-energy states. As the temperature is lowered, the lattice is further coarse grained at T_2, T_3,... until the CGBP reduces to ordinary ER. Figure 7 illustrates this behavior for the one-dimensional critical quantum Ising model. Each coarse-graining level provides a reliable estimate only for a small temperature range, but the union of these ranges cover the entire temperature domain. Thus, CGBP provides accurate estimates of thermodynamical observables in temperature ranges that are accessible to neither BP nor ER.

To calculate the optimal coarse-graining temperatures T_i, we would need independent error assessments for ER and BP. Then, we could switch to a coarser lattice whenever the increase in window size compensates for the coarse graining of the Hamiltonian. However, we are not aware of a reliable method to estimate the error caused by ER. Instead of using error estimates, we simultaneously perform BP on two dif-
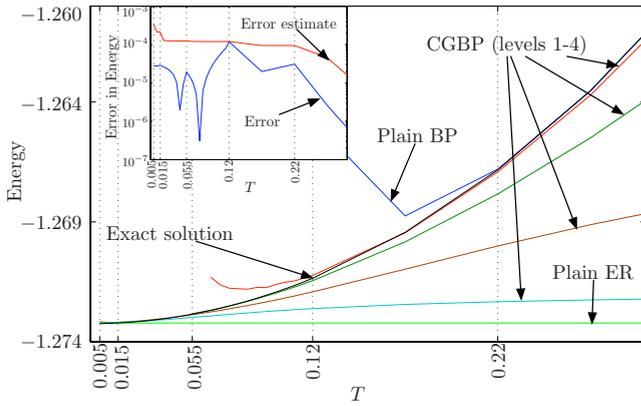
FIG. 7. (Color online) Energy vs temperature for the one-dimensional critical quantum Ising model. Presented are the plain BP results for $l=10$ (dark blue), CGBP results for various levels of coarse graining with $\chi=4$ and $l=5$ (various colors), the plain ER result $\chi=4$ (light green) along with the exact result obtained from Jordan-Wigner transformations (black). The labels on the temperature axis correspond to the switching temperatures between various levels of coarse graining in the CGBP algorithm. For example, plain BP is very accurate down to $T_1 \approx 0.22$, where BP combined with one level of coarse graining (red) becomes more accurate. The inset shows the absolute error in the combined CGBP result with respect to the exact solution (blue), and the error estimate calculated using the procedure given in Sec. IV B. While plain BP and plain ER are both very inaccurate for temperature range $0.005 \lesssim T \lesssim 0.22$, combining the two using CGBP yields very accurate results.

ferent coarse grained levels and determine the switching temperatures by comparing the results.

More precisely, let $\bar{x}_i(T)$ be the expectation value of some observable X at temperature T obtained by BP on the i^{th} level of coarse graining. To determine the switching temperature T_{i+1} from the i^{th} level of coarse graining to the $(i+1)^{th}$, we calculate both $\bar{x}_i$ and $\bar{x}_{i+1}$ as we slowly lower the temperature. At high temperatures, the dominating error is attributed to the discarded high-energy states, so $\bar{x}_i$ is more accurate than $\bar{x}_{i+1}$. On the other hand, at low temperatures, as the range of the effective thermal potential gets larger than the BP window, $\bar{x}_{i+1}$ becomes more accurate than $\bar{x}_i$. The two sources of errors are balanced when $|\bar{x}_i - \bar{x}_{i+1}|$ reaches a minimum (see Fig. 8), so T_i should be chosen at the position of this minimum. There can be exceptions to this rule that result from accidental crossings of $\bar{x}_{i+1}$ and $\bar{x}_i$. In that case, the error estimate for BP presented in Sec. II B can be used to discriminate between the multiple minima. Indeed, the switching should occur when the value of $|\bar{x}_i - \bar{x}_{i+1}|$ is the closest to the error attributed to BP on the i^{th} level because both numbers are estimates of the BP error. The switching temperatures for Fig. 7 were chosen following this method.

We can use the same reasoning to estimate the total error $\delta x(t)$ on our final estimate $\bar{x}(T)$ obtained by joining the $\bar{x}_i(T)$ over their respective range. We define $\delta x_i^{BP}(T)$ to be the error attributed to BP on the i^{th} level of coarse graining. This quantity can be estimated as described in Sec. II B. For $T > T_1$, we have $\delta x(T) = \delta x_0^{BP}(T)$ since BP is the only source of error. Between T_1 and T_2, there are two contributions to the error: the error $\delta x_1^{BP}(T)$ attributed to BP on the first coarse grained
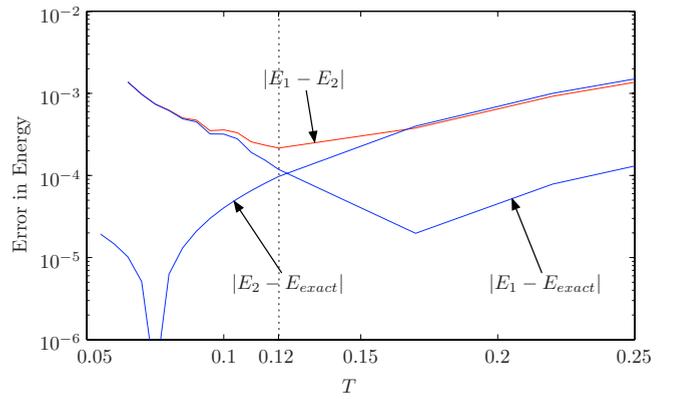

FIG. 8. (Color online) The exact error in the energies calculated at the first and second levels of coarse graining, $|E_1 - E_{exact}|$ and $|E_2 - E_{exact}|$ are plotted along with $|E_1 - E_2|$. Note that for both $|E_1 - E_{exact}|$ and $|E_2 - E_{exact}|$, MERA error dominates at high temperatures, and BP error dominates at low temperatures. The minimum of their difference $|E_1 - E_2|$ occurs when the BP error at level 1 of coarse graining is equal to the MERA error at level 2 of coarse graining. Furthermore, the high temperature portion of $|E_1 - E_2|$ is dominated by the MERA error at level 2 and the low-temperature portion of $|E_1 - E_2|$ is dominated by the BP error at level 1.

level and the error attributed to ER caused by discarding high-energy states. This second error decreases as temperature is lowered, and at $T = T_1$ it is equal to the BP error (this is how T_1 was defined). Thus, we obtain for $T_2 \leq T < T_1$ the bound $\delta x(T) \leq \delta x_0^{BP}(T_1) + \delta x_1^{BP}(T)$. More generally, for $T_{i+1} \leq T < T_i$ we get $\delta x(T) \leq \Sigma_{j \leq i} \delta x_{j-1}^{BP}(T_j) + \delta x_i^{BP}(T)$. See the inset of Fig. 7 for the error estimate of CGBP with $\chi=4$ and $l=5$.

In addition to CGBP with $\chi=4$ and $l=5$, we have also investigated various other combinations of χ and l. At equal computational costs—which are of $O(\chi^{3l})$—the results are qualitatively equivalent in the sense that they have roughly equivalent maximum error. However, varying χ and l at fixed computational cost can improve the results for a given temperature. For instance, at very low temperatures, the computations with a higher χ yield results with better accuracy. On the other hand, larger l and smaller χ perform better at high temperatures. Thus, the values of χ and l could also be varied dynamically in the simulation, but we leave out this possibility for the moment.

The results we obtain with CGBP compare favorably with results obtained by other methods using equivalent computational resources. This can be seen by comparing the inset of Fig. 7 to the various curves shown on Fig. 3. The CGBP result with $\chi=4$, $l=5$ has at least four digits of accuracy for the entire temperature range. The complexity of this simulation is equivalent to exact diagonalization of a 11-site chain, that produces a result accurate to roughly three digits. Many methods rely on a Trotter-Suzuki decomposition by discretizing imaginary time in intervals δT, which creates a systematic bias $O(\delta T^3)$ in the result. This is the case for instance of imaginary time dependent block decimation.[10] This algorithm has complexity $(\chi d^2)^3 T_{min} / \delta T$ where $d=2$ for the Ising model is the dimension of the spins. The parameters $\chi = 150$ and $\delta T = 0.001$ yield the same complexity as the $\chi=4$

$l=5$ CGBP algorithm, and produce an error of roughly 10^{-2} at sufficiently low temperatures.

V. DISCUSSION

The concept of an effective thermal potential, obtained by tracing out sites from the thermal Gibbs state, gives a clear physical picture of the workings of belief propagation. Adding a site to a spin chain can modify the thermal state even far away from the added site, as far as the correlation length of the system. However, this effect can be mimicked by adding a short-range thermal potential to the Hamiltonian of the original spin chain. At finite temperature, we have seen that the range of this thermal potential can be much shorter than the correlation length of the system, making it more suitable for numerical simulations. This underlies the success of belief propagation and provides a means to assess its accuracy. We have illustrated these concepts and methods on the critical quantum Ising chain and the transverse field quantum Ising spin glass. For this last system, our findings are, within estimated error bars, in agreement with those of Ref. 23 obtained from a different belief propagation implementation.

At lower temperatures however, the range of the effective thermal potential becomes too large to handle numerically. For these temperature ranges, we have introduced the coarse-grained belief propagation algorithm by combining belief propagation with entanglement renormalization. Coarse graining discards some high-energy states, which leads to a systematic error in the thermal states. On the other hand, it increases the accuracy of BP by shortening the range of the effective thermal potential by a constant factor. CGBP seeks an optimal compromise between these two effects in order to accurately probe temperature regimes where neither ER nor BP are reliable. Thus, CGBP truly extends the domain of applicability of the two underlying approaches. Moreover, results obtained by CGBP compare favorably to other methods using equivalent computational resources.

The drawback of CGBP is that it inherits some intrinsic limitations of the underlying approaches. For instance, ER is applicable only when there exists a coarse-graining method which preserves the locality of the Hamiltonian. For graphs with exponential spreading structure, such as the Cayley tree studied in Sec. III, we are not aware of such coarse-graining procedures. This is the reason we have not implemented CGBP on that system. Another limitation comes from the shortcoming of BP on graphs containing many small loops such as two-dimensional lattices. Classically, this limitation can be alleviated using generalized BP.[20,21] We are currently working on combining such generalizations with ER to study lattices of higher dimension.

ACKNOWLEDGMENTS

*ersen@caltech.edu

†david.poulin@usherbrooke.ca

[1] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).

[2] I. Affleck, T. Kennedy, E. Lieb, and H. Tasaki, Commun. Math. Phys. **115**, 477 (1988).

[3] F. Verstraete and J. I. Cirac, arXiv:cond-mat/0407066 (unpublished).

[4] G. Vidal, Phys. Rev. Lett. **91**, 147902 (2003).

[5] M. B. Hastings, J. Stat. Mech. (2007) P08024.

[6] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).

[7] M. B. Hastings, Phys. Rev. B **73**, 085115 (2006).

[8] M. B. Hastings, Phys. Rev. B **76**, 035114 (2007).

[9] N. Schuch, I. Cirac, and F. Verstraete, Phys. Rev. Lett. **100**, 250501 (2008).

[10] G. Vidal, Phys. Rev. Lett. **98**, 070201 (2007).

[11] K. G. Wilson, Rev. Mod. Phys. **47**, 773 (1975).

[12] G. Evenbly and G. Vidal, Phys. Rev. B **79**, 144108 (2009).

[13] G. Vidal, Phys. Rev. Lett. **99**, 220405 (2007).

[14] L. Cincio, J. Dziarmaga, and M. M. Rams, Phys. Rev. Lett. **100**, 240603 (2008).

[15] G. Vidal, Phys. Rev. Lett. **101**, 110501 (2008).

[16] M. Aguado and G. Vidal, Phys. Rev. Lett. **100**, 070404 (2008).

[17] R. König, B. W. Reichardt, and G. Vidal, Phys. Rev. B **79**, 195123 (2009).

[18] S. Aji and R. McEliece, IEEE Trans. Inf. Theory **46**, 325 (2000).

[19] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, San Francisco, 1988).

[20] J. S. Yedidia, W. T. Freeman, and Y. Weiss, Mitsubishi Electric Research Laboratories Tech. Rep. No. TR-2001-22, 2002 (unpublished), pp. 239–236.

[21] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, IEEE Trans. Inf. Theory **49**, 1120 (2003).

[22] M. B. Hastings, Phys. Rev. B **76**, 201102(R) (2007).

[23] C. Laumann, A. Scardicchio, and S. L. Sondhi, Phys. Rev. B **78**, 134424 (2008).

[24] M. Leifer and D. Poulin, Ann. Phys. **323**, 1899 (2008).

[25] D. Poulin and E. Bilgin, Phys. Rev. A **77**, 052318 (2008).

[26] The procedure we are describing for a chain is usually referred to as the transfer matrix method in statistical physics.

[27] I. Peschel and V. Eisler, J. Phys. A: Math. Theor. **42**, 504003 (2009).

[28] M. Mézard and G. Parisi, Eur. Phys. J. B **20**, 217 (2001).

[29] F. Krzakala, A. Rosso, G. Semerjian, and F. Zamponi, Phys. Rev. B **78**, 134428 (2008).

[30] At the first iteration of ER, the disentangler would be a unitary transformation on $(\mathbb{C}^d)^{\otimes 2}$ and the isometry would map $(\mathbb{C}^d)^{\otimes 3}$ to $\mathbb{C}^d$, where d is the number of levels of the particles forming the lattice.