

Fixed Points of Belief Propagation – An Analysis via Polynomial Homotopy Continuation

Christian Knoll, *Student Member, IEEE*, Dhagash Mehta, Tianran Chen, and Franz Pernkopf, *Senior Member, IEEE*

Abstract—Belief propagation (BP) is an iterative method to perform approximate inference on arbitrary graphical models. Whether BP converges and if the solution is a unique fixed point depends on both the structure and the parametrization of the model. To understand this dependence it is interesting to find *all* fixed points. In this work, we formulate a set of polynomial equations, the solutions of which correspond to BP fixed points.

To solve such a nonlinear system we present the numerical polynomial-homotopy-continuation (NPHC) method. Experiments on binary Ising models and on error-correcting codes show how our method is capable of obtaining all BP fixed points. On Ising models with fixed parameters we show how the structure influences both the number of fixed points and the convergence properties. We further assess the accuracy of the marginals and weighted combinations thereof. Weighting marginals with their respective partition function increases the accuracy in all experiments. Contrary to the conjecture that uniqueness of BP fixed points implies convergence, we find graphs for which BP fails to converge, even though a unique fixed point exists. Moreover, we show that this fixed point gives a good approximation, and the NPHC method is able to obtain this fixed point.

Index Terms—Graphical models, belief propagation, probabilistic inference, sum-product algorithm, Bethe free energy, phase transitions, inference algorithms, dynamical equations.



1 INTRODUCTION

JOINT distributions over many random variables (RVs) are often modeled as probabilistic graphical models. Belief propagation (BP) is a prominent tool to determine marginal distributions of such models. On tree-structured models the marginals are exact, but BP provides only an approximation on graphs with loops. Despite the lack of guarantee for convergence, BP has been successfully used for models with many loops, including applications in computer vision, medical diagnosis systems, and speech processing [1], [2], [3]. However, instances of graphs do exist where BP fails to converge. A deeper understanding of the reasons for convergence of BP, and whether and how its non-convergence relates to the number of fixed points may therefore be crucial in understanding BP.

A precise relation among the uniqueness of fixed points, convergence rate, and accuracy is yet to be theoretically understood [4]. Sufficient conditions for uniqueness of fixed points were refined by accounting for both the potentials as well as the structure of the model [4], [5]. On graphs with a single loop [6] and on small grid graphs [7] accuracy and convergence rate are related; this does, however, not necessarily hold for all graphs [8]. As a consequence using

provably convergent variants of BP can still lead to accurate approximations [9], [10]. Changing the BP update schedule can also help to achieve convergence [11], [12], [13]. One can further increase the accuracy of BP by considering multiple fixed points [2]. Survey propagation [14] and its efficient approximation scheme [15] represent distributions over BP messages and marginalize over all obtained fixed points.

In this work we are interested in the relation among convergence properties, the number of fixed points, and the accuracy of these fixed points. To get deeper insights into the behavior of BP we aim to find *all* fixed points – including unstable ones. If BP converges, however, it does only provide a *single* fixed point. In order to find all fixed points, we reformulate the update rules of BP as a system of polynomial equations.

There are indeed several methods to solve such a system of polynomial equations. Numerical solvers (e.g., Newton’s method) are well established, but their ability in obtaining the solutions strongly depends on the initial point. Moreover, such methods only find a single solution at a time, and do not guarantee to find all solutions even with different initial guesses. Symbolic methods, on the other hand, are guaranteed to find *all* solutions. The Gröbner basis method [16], [17] is widely used, but it is limited to systems with rational coefficients and it suffers from fast growing run time and memory complexity. Moreover the method has a limited scalability in parallel computation. In this paper, we describe the numerical polynomial homotopy continuation (NPHC) method [18], [19] that overcomes all the above mentioned problems of both iterative and symbolic methods, yet guarantees to find *all* solutions of the system. Over the last few decades, this class of methods has been proven to be robust, efficient, and highly parallelizable. We

- Christian Knoll (christian.knoll.c@ieee.org) and Franz Pernkopf (pernkopf@tugraz.at) are with the Signal Processing and Speech Communication Laboratory, Graz University of Technology.
- Dhagash Mehta (Dhagash.B.Mehta.11@nd.edu) is with the Department of Applied and Computational Mathematics and Statistics, University of Notre Dame.
- Tianran Chen (ti@nranchen.org) is with the Department of Mathematics and Computer Science, Auburn University at Montgomery.

exploit the sparsity of our polynomial system and compute a tight upper bound on the number of complex solutions. This is an essential step to reduce the computational complexity by orders of magnitude so that the systems tackled in this work can be solved in practice.

We apply the NPHC method to different realizations of the Ising model on complete and grid graphs¹, although it can be used for any other graph structure. These models are appealing because they are well studied in the physics literature [20], [21]. On these graphs we obtain all BP fixed points and show how the number of fixed points changes at critical regions (i.e., phase transitions) in the parameter space and how the fixed points behave under varying potentials.

Recently we also performed stability analysis of all fixed points on Ising graphs with uniform potentials in [22]. We observed that the existence of non-vanishing local fields helps to achieve convergence and increases the accuracy of the best fixed point. We further showed that damping does not improve convergence properties on bipartite graphs.

We further use the obtained fixed points to estimate the approximation of the marginal distribution which we compare to the exact marginal distribution. Our main (empirical) observations are: (i) on loopy graphs BP does not necessarily converge to the best possible fixed point, (ii) on some graphs where BP does not converge we are able to show that there is a unique fixed point but if we enforce convergence to this unique fixed point the obtained marginals still give a good approximation, and (iii) convex combinations of multiple fixed points drastically improve the accuracy. We further apply the NPHC method to analyze error-correcting code and show that BP performs better for densely connected variables.

The paper is structured as follows: Section 2 provides a brief background on probabilistic graphical models and free energy approximations. In Section 3, we reformulate the message passing equations, and introduce the NPHC method that guarantees to find all BP fixed point solutions. Our experimental results are presented and discussed in Section 4. Finally, we conclude the paper in Section 5.

2 BACKGROUND

In this section, we briefly introduce probabilistic graphical models and the BP algorithm, and fix our notations. For an in-depth treatment of these topics we refer the reader to [3], [23], [24].

2.1 Probabilistic Graphical Models

We consider a finite set of N discrete random variables $\mathbf{X} = \{X_1, \dots, X_N\}$ taking values from the binary alphabet $x_i \in \mathbb{S} = \{-1, +1\}$. Let us consider the joint distribution $P(\mathbf{X} = \mathbf{x})$ and the corresponding undirected graph $G = (\mathbf{X}, \mathbf{E})$ where $\mathbf{X} = \{X_1, \dots, X_N\}$ is the set of nodes and \mathbf{E} is the set of edges. Between the RVs and the nodes a one-to-one correspondence holds. An interaction between two nodes X_i and X_j , $i \neq j$ is represented by an undirected edge

1. A complete graph is an undirected graph where each pair of nodes is connected by an edge. A grid graph, or lattice graph, has all edges aligned along the 2D square lattice. Examples are depicted in Fig. 2.

$e_{i,j} \in \mathbf{E}$. We sometimes consider the graph $G' = (\mathbf{X}, \mathbf{E}')$ where $\mathbf{E}' \subset \mathbf{E}$ such that two nodes are connected by one edge at most, i.e., $e_{i,j} \in \mathbf{E}' \implies e_{j,i} \notin \mathbf{E}'$. Let the set of neighbors of X_i be defined by $\partial(X_i) = \{X_j \in \mathbf{X} \setminus X_i : e_{i,j} \in \mathbf{E}\}$. The joint probability of an undirected graphical model factorizes to

$$P(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \prod_{l=1}^L \Phi_{C_l}(C_l), \quad (1)$$

where the potentials Φ_{C_l} are specified over the maximal cliques C_l of the nodes [1, p.105]. If we restrict all potentials to consist of at most two variables: then, the joint distribution is factorized according to

$$P(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \prod_{(i,j) : e_{i,j} \in \mathbf{E}'} \Phi_{X_i, X_j}(x_i, x_j) \prod_{i=1}^N \Phi_{X_i}(x_i), \quad (2)$$

where $Z \in \mathbb{R}_+^*$ denotes the strictly positive normalization constant to guarantee a valid probability distribution. The first product runs over all edges and the second product runs over all nodes; pairwise potentials and local evidence are denoted as Φ_{X_i, X_j} and Φ_{X_i} respectively. We use a shorthand notation for the marginal probabilities $P(x_i) = P(X_i = +1)$ and $P(\bar{x}_i) = P(X_i = -1)$ where no ambiguities occur.

2.2 Belief Propagation

Belief propagation (BP) is an algorithm that approximates marginal probabilities (or beliefs) $\tilde{P}(X_i = x_i)$. The marginals are approximated by recursively updating messages between random variables. This update rule is guaranteed to converge and return the exact marginals on graphs without loops [1]. Note that this procedure has been discovered in different fields independently: belief propagation for probabilistic reasoning [1], the sum-product-algorithm in information theory [25], [26], and the Bethe-Peierls approximation in statistical mechanics [21].

The messages from X_i to X_j of state x_j at iteration $n+1$ are given by the following update rule:

$$\mu_{i,j}^{n+1}(x_j) = \alpha_{i,j}^n \sum_{x_i \in \mathbb{S}} \Phi_{X_i, X_j}(x_i, x_j) \Phi_{X_i}(x_i) \prod_{X_k \in \Gamma_{i \setminus j}} \mu_{k,i}^n(x_i), \quad (3)$$

where $\Gamma_{i \setminus j} = \partial(X_i) \setminus \{X_j\}$. Loosely speaking BP collects all messages sent to X_i , except for X_j , and multiplies this product with the local potential $\Phi_{X_i}(x_i)$ and the pairwise potential $\Phi_{X_i, X_j}(x_i, x_j)$. The sum over both states of X_i is sent to X_j . In practice the messages are often normalized by $\alpha_{i,j}^n \in \mathbb{R}_+^*$ so as to sum to one [27].

Lemma 1. *Messages being sent from node X_i to X_j , $i \neq j$, represent probabilities – provided all messages are initialized to be positive.*

Proof. Positive potentials in (3) guarantee that all messages remain positive at every iteration. Consequently, a normalization term $\alpha_{i,j}^n$ exists so that $\sum_{x_j \in \mathbb{S}} \mu_{i,j}^{n+1}(x_j) = 1$. \square

The set of all messages at iteration n is given by $\underline{\mu}^n = \{\mu_{i,j}^n(x_j) : e_{i,j} \in \mathbf{E}\}$. In a similar manner we collect all normalization terms in $\underline{\alpha}^n$. Let the mapping of all messages

induced by (3) be denoted as $\underline{\mu}^{n+1} = \mathcal{BP}\{\underline{\mu}^n\}$. If all successive messages show the same value (up to some predefined precision), that is $\underline{\mu}^n \cong \underline{\mu}^{n+1}$, then BP is converged. We refer to converged messages and the associated normalization terms as *fixed points* ($\underline{\mu}^*$, $\underline{\alpha}^*$).

Note that at least one fixed point always exists if all potentials are positive [28]. Existence of fixed points, however, is not sufficient to guarantee convergence; in fact BP may be trapped in limit cycles or show chaotic behavior [29]. If the messages oscillate, one can try to achieve convergence by damping [9], i.e., by replacing the messages with a weighted average of the last messages so that $\underline{\mu}^{n+1} = (1 - \epsilon)\mathcal{BP}\{\underline{\mu}^n\} + \epsilon\underline{\mu}^n$: it follows, that any fixed point of BP with damping is a fixed point of BP without damping as well.

If BP converges to a fixed point, then the marginals are approximated by the normalized product

$$\tilde{P}(X_i = x_i) = \frac{1}{Z_i} \Phi_{X_i}(x_i) \prod_{X_k \in \partial(X_i)} \mu_{k,i}^*(x_i), \quad (4)$$

where $Z_i \in \mathbb{R}_+^*$ is required so that $\sum_{\mathbb{S}} P(X_i = x_i) = 1$. Similar the pairwise marginal of two nodes (X_i, X_j) , connected by an edge, is defined by the product of all incoming messages times and all factors involved:

$$\tilde{P}(X_i = x_i, X_j = x_j) = \frac{1}{Z_{i,j}} \Psi_{X_i, X_j}(x_i, x_j) \prod_{X_k \in \Gamma_{i \setminus j}} \mu_{k,i}^*(x_i) \prod_{X_l \in \Gamma_{j \setminus i}} \mu_{l,j}^*(x_j), \quad (5)$$

where $\Psi_{X_i, X_j}(x_i, x_j) = \Phi_{X_i}(x_i)\Phi_{X_j}(x_j)\Phi_{X_i, X_j}(x_i, x_j)$ and $Z_{i,j} \in \mathbb{R}_+^*$ is the normalization term.

2.3 Free Energy Approximations

Over the years a fruitful connection between computer science and statistical mechanics was established (cf. [8], [21], [30], [31]). In particular, the relationship between stationary points of the *Bethe free energy* \mathbb{F}_B and fixed points of BP led to a deeper understanding of BP. We briefly discuss important insights and present differences in notations to circumvent any confusions.

In the Ising model, a popular statistical mechanics model, often used for evaluation of BP, each node X_i has an associated spin taking values in $\mathbb{S} = \{-1, +1\}$. We define the corresponding energy function [21, p.44] by assigning a coupling $J_{i,j} \in \mathbb{R}$ to each edge $e_{i,j} \in \mathbb{E}$ and some local field $\theta_i \in \mathbb{R}$ acting on each node $X_i \in \mathbb{X}$. Note that we drop the subscripts of $J_{i,j}$ and θ_i whenever they are the same for all edges and nodes. Let the local and pairwise Ising potentials of state x_i be $\Phi_{X_i}(x_i) = \exp(\theta_i x_i)$ and $\Phi_{X_i, X_j}(x_i, x_j) = \exp(J_{i,j} x_i x_j)$. Then by plugging these potentials into (2) the joint distribution is equal to the Boltzmann distribution

$$P(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \exp \left(\beta \cdot \sum_{(i,j) : e_{i,j} \in \mathbb{E}'} J_{i,j} x_i x_j + \sum_{i=1}^N \theta_i x_i \right). \quad (6)$$

We omit the term of the inverse temperature β by choosing $\beta = 1$ for the rest of this work. The Bethe free energy is a function of the marginals and the pairwise marginals

$$\begin{aligned} \mathbb{F}_B(P(X_i), P(X_i, X_j)) = & \sum_{(i,j) : e_{i,j} \in \mathbb{E}'} \sum_{x_i, x_j} P(X_i = x_i, X_j = x_j) \cdot \\ & \ln \frac{P(X_i = x_i, X_j = x_j)}{\Phi_{X_i, X_j}(x_i, x_j)} - \sum_{X_i} \sum_{x_i} P(X_i = x_i) \ln \Phi_{X_i}(x_i) \\ & - \sum_{X_i} (|\partial(X_i)| - 1) \sum_{x_i} P(X_i = x_i) \ln P(X_i = x_i), \end{aligned} \quad (7)$$

and relates to the *Bethe partition function* according to

$$Z_B(P(X_i), P(X_i, X_j)) = \exp(-\mathbb{F}_B(P(X_i), P(X_i, X_j))). \quad (8)$$

An excellent overview of free energy approximations from a variational perspective and how this relates to BP can be found in [28], [32].

If the Ising model is not on a path graph, critical regions in the parameter space can exist where phase transitions occur [20, Ch.12]. If all nodes $X_i \in \mathbb{X}$ have equal degree $|\partial(X_i)| = d + 1$, these phase transitions can be determined by replacing the graph with a Cayley tree² of degree d [33]. Therefore let

$$p(J, d) = \begin{cases} d \operatorname{arctanh} \sqrt{\frac{d \cdot w - 1}{d/w - 1}} - \operatorname{arctanh} \sqrt{\frac{d-1/w}{d-w}} & \text{if } J > \operatorname{arctanh}(d) \\ d \operatorname{arctanh} \sqrt{\frac{d \cdot w - 1}{d/w - 1}} + \operatorname{arctanh} \sqrt{\frac{d-1/w}{d-w}} & \text{if } J < \operatorname{arctanh}(d) \end{cases} \quad (9)$$

where $w = \tanh |J|$. Also, in accordance with [33] and [20, p.247-255], let the phase transitions partition the parameter space (J, θ) into the following three distinct regions (I), (II), and (III) as :

$$(J, \theta) \in (I) \text{ if } J > 0, J > p(J, d) \text{ and } |\theta| \leq p(J, d), \quad (10)$$

$$(J, \theta) \in (II) \text{ if } J < 0, J < -p(J, d) \text{ and } |\theta| < p(J, d), \quad (11)$$

$$(J, \theta) \in (III) \text{ if } (J, \theta) \notin (I) \text{ and } (J, \theta) \notin (II). \quad (12)$$

In the literature one distinguishes three different interactions on Ising grids. First, if all couplings $J > 0$ the model is *ferromagnetic*: for ferromagnetic models BP converges to a unique stationary point inside (III). This stationary point becomes unstable and two additional stationary points emerge inside (I) [21], [28]. Second, if all couplings $J < 0$ the model is *antiferromagnetic*: here, BP only converges inside (III) and not inside (II) [34]. Finally, *spin glasses* are models containing both positive and negative couplings.

Depending on the graph structure spin glasses and antiferromagnetic models allow for frustrations; i.e., the marginals minimizing the Bethe free energy may correspond to a degenerate joint distribution [29], [21, pp.45].

2. A Cayley tree is an infinite tree without loops that captures interactions of a cyclic finite size graph.

2.4 Stationary Points of Bethe Free Energy

As stationary points of \mathbb{F}_B and of Z_B (cf. (8)) correspond to BP fixed points, one can try to obtain stationary points of \mathbb{F}_B directly, instead of performing BP. For tree-structured graphs and one-loop graphs \mathbb{F}_B is a convex function [5]. For general graphs, however, convexity breaks down and \mathbb{F}_B has multiple local minima. Note that stable fixed points (cf. Sec. 4.5) of BP do always correspond to local minima of \mathbb{F}_B ; the converse however need not be the case, i.e., unstable fixed points can be either local maxima or local minima of \mathbb{F}_B [35]. Sufficient conditions for convexity of \mathbb{F}_B are often used to make statements regarding uniqueness of BP fixed point solutions. This seems to be a rather limiting point of view – it is possible to add loops to a former tree-structured model without changing the distribution [5, pp.2391]. Hence, the number of BP fixed points does not only depend on the structure of the graph but also depends on the potentials.

Energy functions with many local minima can be decomposed into a convex and a concave problem (this decomposition is in general not unique) [10]. Alternatively one can construct convex surrogates and minimize these convex functions instead [36]. There is a variety of methods that obtain the marginal probabilities by minimizing \mathbb{F}_B . Belief optimization [37] follows the negative gradient of \mathbb{F}_B and guarantees the marginalization constraints to be satisfied; i.e., the minimization takes place along the edges of the Bethe polytope. An efficient way to obtain an approximate fixed point is presented in [38]. A projection scheme in the minimization task allows for a fully polynomial-time approximation on sparse graphs with $\max(|\partial(X)|) = \mathcal{O}(\log N)$. This algorithm is further improved in [8] to return an approximation of the best (stable) fixed point of BP. An alternative method to approximate and combine all stable solutions of BP (i.e., local minima of \mathbb{F}_B) is presented in [15].

However all of these methods rely on the idea of finding minima of \mathbb{F}_B ; and consequently fail to obtain (unstable) fixed points that correspond to local maxima of \mathbb{F}_B .

3 SOLVING BP FIXED POINT EQUATIONS

We reformulate the message update rules (3) as a system of polynomial equations whose solutions are the fixed points of BP. Whether such systems can be solved in practice depends largely on the chosen method. We list a variety of approaches and describe NPHC that can be applied in practice to obtain *all* BP fixed points.

3.1 Reformulation of Belief Propagation

For all messages $\mu_{i,j}(x_j) : e_{i,j} \in \mathbf{E}$ the residual, i.e., the difference between two successive message values $\underline{\mu}^{n+1} - \underline{\mu}^n$ and the message normalization constraints $\underline{\alpha}^n$ are given by the following system of polynomial equations:

$$\mathbf{F}(\underline{\mu}, \underline{\alpha}) = \begin{cases} -\mu_{i,j}^n(x_j) + \alpha_{i,j}^n \sum_{x_i \in \mathbb{S}} \Phi_{X_i, X_j}(x_i, x_j) \Phi_{X_i}(x_i) \prod_{X_k \in \Gamma_i \setminus j} \mu_{k,i}^n(x_i) \\ -\mu_{i,j}^n(\bar{x}_j) + \alpha_{i,j}^n \sum_{x_i \in \mathbb{S}} \Phi_{X_i, X_j}(x_i, \bar{x}_j) \Phi_{X_i}(x_i) \prod_{X_k \in \Gamma_i \setminus j} \mu_{k,i}^n(x_i) \\ \mu_{i,j}^n(x_j) + \mu_{i,j}^n(\bar{x}_j) - 1. \end{cases} \quad (13)$$

This system of polynomial equations consist of

$$M = 2|\mathbf{E}| \cdot (|\mathbb{S}| + 1) \quad (14)$$

equations $(f_1(\underline{\mu}, \underline{\alpha}), \dots, f_M(\underline{\mu}, \underline{\alpha}))$. To solve such a polynomial system, it is advantageous to consider it defined over complex variables rather than real variables, in order to apply the NPHC method. Let the set of solutions over complex variables, without accounting for multiplicity, be

$$V(\mathbf{F}) = \{(\underline{\mu}, \underline{\alpha}) \in \mathbb{C} : f_m(\underline{\mu}, \underline{\alpha}) = 0 \text{ for all } f_m \in \mathbf{F}(\underline{\mu}, \underline{\alpha})\}. \quad (15)$$

We further define the set of solutions $V_{\mathbb{R}_+}^*(\mathbf{F}) \subset V(\mathbf{F})$ over strictly positive real numbers.

Theorem 1 (Fixed Points of BP). *Let $(\underline{\mu}, \underline{\alpha})$ be some set of messages and normalization terms. Then, $(\underline{\mu}, \underline{\alpha})$ is a fixed point solution of BP, if and only if $(\underline{\mu}, \underline{\alpha}) \in V_{\mathbb{R}_+}^*(\mathbf{F})$.*

Proof. First we show that $(\underline{\mu}, \underline{\alpha}) \in V_{\mathbb{R}_+}^*(\mathbf{F})$ is sufficient to characterize fixed point solutions. All messages are positive and represent probabilities (Lemma 1). Further, it follows from (13) that $BP\{\underline{\mu}^n\} - \underline{\mu}^n = 0$, which constitutes a fixed point solution (cf. Section 2).

Conversely, consider some fixed point messages and the corresponding normalization coefficients $(\underline{\mu}^*, \underline{\alpha}^*)$, it then follows by definition that $BP\{\underline{\mu}^*\} = \underline{\mu}^*$ and consequently $\mathbf{F}(\underline{\mu}, \underline{\alpha}) = 0$. \square

Corollary 1.1. *Consider a graph with strictly positive potentials. Then, the solution set $V_{\mathbb{R}_+}^*(\mathbf{F})$ is nonempty. Moreover, if Φ_{X_i, X_j} and Φ_{X_i} are Ising potentials, then $V_{\mathbb{R}_+}^*(\mathbf{F})$ is always nonempty.*

Proof. For non-negative potentials the average energy is bounded from below and \mathbb{F}_B has at least one minimum [28, Th. 4]. Minima of the constrained \mathbb{F}_B correspond to BP fixed point solutions, the existence of which implies non-emptiness of $V_{\mathbb{R}_+}^*(\mathbf{F})$ by Theorem 1. \square

3.2 Solving Systems of Polynomial Equations

Solving systems of (nonlinear) polynomial equations is a classic problem in computational mathematics, and a great variety of approaches have been developed such as iterative solvers, symbolic methods, and homotopy methods.

One basic method for solving systems of nonlinear equations is Newton's method which is an iterative solver that progressively refines an initial guess to reach a solution. Such iterative solvers find a *single* solution in the vicinity of an already known initial guess. However, when the initial guess is not sufficiently close to a solution, they may diverge or even exhibit chaotic behavior. Moreover, it is difficult, to obtain the full set of solutions with these methods. We are interested in the entire set of (positive) real solution $V_{\mathbb{R}_+}^*(\mathbf{F})$,

i.e., in obtaining *all* fixed points; therefore iterative solvers are not useful in the current setting.

From a completely different point of view, symbolic methods [16], [17] (e.g. Gröbner basis method, Wu’s method, and methods of sparse resultant) rely on symbolic manipulation of the equations and successive elimination of variables to obtain a simpler but equivalent form of the equations. In a sense, these methods can be considered as generalizations of the Gaussian elimination method for linear systems into nonlinear settings. In the past several decades, symbolic methods, especially the Gröbner basis method, have seen substantial development. But the method is limited to systems with rational coefficients and has a worst case complexity that is double exponential in the variables [39], [40]. This and the limited scalability in parallel computation limits the application of symbolic methods to smaller systems. In fact, the Gröbner did not converge for any of our experiments.

3.3 Numerical Polynomial Homotopy Continuation

Another important approach for solving a system of nonlinear equations is the *numerical polynomial homotopy continuation* (NPHC) method [18], [19]. The “target” system $\mathbf{F}(\underline{\mu}, \underline{\alpha})$ in (13), which we intend to solve, is continuously deformed into a closely related “starting system” $\mathbf{Q}(\underline{\mu}, \underline{\alpha}) = (q_1(\underline{\mu}, \underline{\alpha}), \dots, q_M(\underline{\mu}, \underline{\alpha}))$ that is trivial to solve. With an appropriate construction, the corresponding solutions also vary continuously under this deformation forming “solution paths” that connect the solutions of the trivial system to the desired solutions of the target system.

For instance, a basic form of linear homotopy for the target system can be given by

$$H(\underline{\mu}, \underline{\alpha}, t) = (1 - t)\mathbf{Q}(\underline{\mu}, \underline{\alpha}) + \gamma t\mathbf{F}(\underline{\mu}, \underline{\alpha}) = 0. \quad (16)$$

Clearly, at $t = 0$ the homotopy reduces to the starting system $\mathbf{Q}(\underline{\mu}, \underline{\alpha})$ and at $t = 1$ it reduces to $\mathbf{F}(\underline{\mu}, \underline{\alpha})$. As t varies continuously from $t = 0$ to $t = 1$, the homotopy represents a deformation from the starting system to the target system. For a generic complex γ , the above procedure is guaranteed to find *all* isolated complex solutions of $\mathbf{F}(\underline{\mu}, \underline{\alpha})$ [19, pp.91].

Even though only positive real solutions $V_{\mathbb{R}_+}^*(\mathbf{F})$ are of interest in the current study, the homotopy method benefits from extending the domain to the field of complex numbers. Only then smooth solution paths emerge towards the desired solutions. Each of these solution paths can be tracked independently making this approach *pleasantly parallelizable*; this is essential in dealing with large polynomial systems.

3.4 Polyhedral Homotopy

More advanced “nonlinear” homotopies can be constructed where the parameter t appears in nonlinear form. Among a great variety of polynomial homotopy constructions, the *polyhedral homotopy method*, developed by B. Huber and B. Sturmfels [41], is particularly suited as it is capable of finding *all* isolated nonzero complex solutions³, which must

3. Here, “nonzero complex solutions” refer to complex solutions of a system of polynomial equations where each variable is nonzero. A solution is considered to be isolated if it has no degree of freedom, i.e., there is an open set containing it but no other solutions.

include *all* BP fixed points $V_{\mathbb{R}_+}^*(\mathbf{F})$. This advantage and the level of parallel scalability are the main motivation for applying the polyhedral homotopy method

In applying the NPHC method to solve (13), the choice of $\mathbf{Q}(\underline{\mu}, \underline{\alpha})$ (the trivial system of equations that the target system is deformed into) plays an important role in the overall efficiency of the approach since different choices of $\mathbf{Q}(\underline{\mu}, \underline{\alpha})$ may induce a vastly different number of solution paths one has to track. The crucial part is to come up with a good upper bound on the number of solutions and to create an appropriate start system. Once this is solved, every solution can be tracked completely independently in parallel in order to reach all desired solutions.

In each equation of (13) only few of the monomials are present, i.e., the update equations of BP imply a sparse system of equations [41]. This holds even if the graph is non-sparse. In our experiments, we observed that despite the rather high *total degree*⁴ d_t [19, pp.118], each equation in (13) contains only few of the monomials the number of solution paths to be tracked to solve (13) is relatively small. The number of solution paths one has to track when using the polyhedral homotopy method for solving a system of polynomial equations is given by the so-called *Bernstein-Kushnirenko-Khovanskii (BKK) bound*: fixing the list of monomials that appear in the polynomial system, it is an important yet surprising fact in algebraic geometry that for almost all choices of the coefficients (in the probabilistic sense), the number of isolated nonzero complex solutions is a fixed number which only depends on the list of monomials. This number is known as the BKK bound [42], [43], [44]. Intermediate steps in the determination of the BKK bound are reused to create an appropriate start system. Using the fully parallel implementation Hom4PS-3 [45] of the polyhedral homotopy method, we compute the BKK bound, that is tight in all our experiments, and obtain *all* isolated positive solutions.

The polyhedral homotopy method exploits the structure of (13), but also requires some subtle steps. Rather than presenting all technical details we present an illustrative example to explain the underlying principles. For more details we refer the reader to the excellent overview papers [18], [46], [47] or to [19, Sec.8.5.4] and the references therein.

3.4.1 Illustrative Example

The essential steps in solving polynomial system with the polyhedral homotopy method are: first, compute a root count based on mixed volume computations [18, Section 3]; secondly, come up with an easy to solve start system [18, Section 4]; and finally, solve the start system and track the solution paths to the target system [18, Section 1].

(i) *Root Count*: Consider the example system in two

4. The total degree of a system of polynomial equations is the product of the degrees of each equation. It is a basic fact in algebraic geometry that the total number of isolated complex solutions a polynomial system has is bounded by its total degree (i.e., Bezout bound). Therefore the total degree serves as a crude measure of the complexity of the polynomial system.




Fig. 1: (a) Polytope S_1 , (b) polytope S_2 , and (c) Minkowski sum $S_1 + S_2$

unknowns $\underline{x} = \{x_1, x_2\}$ and with 4 solutions taken from [19, p.142]:

$$\mathbf{F}(\underline{x}) = \begin{cases} 1 + ax_1 + bx_1^2x_2^2 \\ 1 + cx_1 + dx_2 + ex_1x_2^2. \end{cases} \quad (17)$$

The total degree of this system of equations is $d_t = 4 \cdot 3 = 12$, which serves as an upper bound on the actual number of solutions. If the system of equations is sparse, the BKK bound serves as a much tighter bound.

Every equation $f_m \in \mathbf{F}(\underline{x})$ has an associated polytope S_m which is the convex hull of the exponent vectors for all monomials of f_m . For f_1 the polytope is

$$S_1 = \{(0,0)(1,0)(2,2)\}, \quad (18)$$

which has a graphical representation in Fig. 1a. Similar for f_2 the polytope $S_2 = \{(0,0)(1,0)(0,1)(1,2)\}$ is shown in Fig. 1b.

Some important operation on polytopes are: the Minkowski sum $S_1 + S_2 = \{s_1 + s_2 : s_1 \in S_1, s_2 \in S_2\}$, and the computation of volumes, denoted by $V(S_m)$. Note that this approach generalizes to higher dimensions, therefore we refer to $V(S_m)$ as volume even if the polytopes lie in the two-dimensional space. The computation of the mixed volume $M(S_1, S_2)$ is a combinatorial problem that is especially comprehensible in the case of two equations where

$$M(S_1, S_2) = V(S_1 + S_2) - V(S_1) - V(S_2). \quad (19)$$

For a generalization to higher dimensions see [19, p.140]. The polytope of $S_1 + S_2$ is illustrated in Fig. 1c: the mixed volume is then obtained by subtracting $V(S_1)$ and $V(S_2)$ from $V(S_1 + S_2)$, which equals the sum of all gray areas (known as mixed cells) in Fig. 1c.

It is straightforward to see that each parallelogram has volume equal to 2; consequently the BKK bound equals $M(S_1, S_2) = 4$ and provides a tight bound on the number of solutions.

(ii) *Start System*: The BKK-bound does not immediately lead to the solutions of an appropriate start system $\mathbf{Q}(\underline{x})$ with $q_m(x) = \sum_{a \in S_m} c_{m,a} \underline{x}^a$, where $\underline{x}^a = x_1^{a_1} \cdot x_2^{a_2}$ and $c_{m,a}$ are random coefficients.

However, the mixed volume computation can also be accomplished by introducing a lifting $\omega_m = \{\omega_m(a) : a \in S_m\}$ for each f_m . Thereby we increase the dimension of the polytope S_m to \hat{S}_m by adding one component to each exponent-vector

a . This component is obtained by the lifting function $\omega_m(a)$. In our example we choose the lifting values $\omega_1 = \{0, 0, 0\}$ and $\omega_2 = \{0, 1, 1, 3\}$; these values are obtained by the inner products $\omega_1(a) = (0, 0) \circ (a_1, a_2)$ and $\omega_2(a) = (1, 1) \circ (a_1, a_2)$. The polytopes are lifted accordingly so that $\hat{S}_1 = \{(0, 0, 0)(0, 1, 0)(2, 2, 0)\}$, $\hat{S}_2 = \{(0, 0, 0)(1, 0, 1)(0, 1, 1)(1, 2, 3)\}$ and $\hat{S}_1 + \hat{S}_2 = \{(0, 0, 0)(0, 1, 0)(2, 2, 0)(0, 2, 1)(0, 1, 1)(3, 2, 1)(3, 4, 3)\}$. Then the faces in the lower hull of $\hat{S}_1 + \hat{S}_2$ correspond to cells shown in Fig. 1c, which is known as a fine mixed subdivision.

These liftings and the random coefficients $c_{m,a}$ form the homotopy $\hat{\mathbf{Q}}(\underline{x}, t)$ with $\hat{q}_m = \sum_{a \in S_m} c_{m,a} \underline{x}^a t^{\omega_m(a)}$ such that

$$\hat{\mathbf{Q}}(\underline{x}, t) = \begin{cases} 1 + c_{1,1}x_1 + c_{1,2}2x_1^2x_2^2 \\ 1 + c_{2,1}x_1t + c_{2,2}x_2t + c_{2,3}x_1x_2^2t^3. \end{cases} \quad (20)$$

By closer inspection, however, it is not possible to identify the starting points because $\hat{q}_2(\underline{x}, t = 0) = 1$. This problem can be resolved according to [41, Lemma 3.1]: i.e., initial values are obtained by solving a binomial system for every cell that contributes to the mixed volume computation (gray cells in Fig. 1c). One can increase t and track the solution paths to $\hat{\mathbf{Q}}(\underline{x}, t = 1) = \mathbf{Q}(\underline{x})$.

(iii) *Target System*: Finally we have all 4 solutions to $\mathbf{Q}(\underline{x})$. Now what remains is to construct a *linear* homotopy according to (16) and increase t starting $t = 0$. At $t = 1$ the homotopy reduces to $\mathbf{F}(\underline{x})$ and provides the desired solutions.

4 EXPERIMENTS

In various experiments we apply the NPHC method to system (13) and obtain all fixed points by first finding all the isolated non-zero complex solutions. We consider ferromagnetic, antiferromagnetic, and spin glass models on fully connected graphs and grid graphs shown in Fig. 2. We evaluate and compare the accuracy of all fixed points obtained by NPHC; details for our evaluation criteria are presented in Sec. 4.1. Further we present the evolution of the fixed points over the parameter space in Sec. 4.5 and present the implications on the accuracy to better understand under which settings BP can be expected to provide good results. The capability of NPHC to obtain all fixed points allows for a thorough stability analysis. This inspired some theoretical investigations on Ising graphs with uniform parameters [22]. There we show why convergence properties degrade with growing graph size and why non-vanishing fields help to achieve convergence. The limited influence of damping on bipartite graphs is also explained.

The systems considered in this work are simply too large to be solved with symbolic methods or even with the NPHC method based on a linear homotopy. The BKK bound, however, takes into account the sparsity of the systems $\mathbf{F}(\underline{\mu}, \underline{\alpha})$, induced by the structure of the graph, and reduces the number of solution paths to be tracked so that the problem can be solved in practice. We present a detailed runtime analysis in Sec. 4.6. Note that neither the structure of (13) nor the number of complex solutions in $V(F)$ does change if the underlying structure of the graph is the same [48], [49];




Fig. 2: Structure of Ising graphs considered: (a) grid graph with 9 RVs; (b) fully connected graph.

but, depending on the potentials, the number of solutions in $V_{\mathbb{R}^+}^*(\mathbf{F})$ may change.

Finally, the performance of BP-decoding for low error-correcting codes is analyzed by obtaining all fixed points with NPHC for the associated factor graph and by comparing them with the exact solutions in Sec. 4.7.

4.1 Evaluation Criteria

We apply (4) to the positive real solutions to obtain the marginals and compare them with marginals obtained by an implementation of BP without damping [50]; we further compare these results with the exact marginals obtained by the junction tree algorithm [51]. To evaluate the correctness of approximated marginals we average the mean squared error (MSE) over all N nodes. For binary RVs we can apply symmetry properties of the probability mass function $P(x_i) = 1 - P(\bar{x}_i)$, so that

$$\text{MSE} = \frac{2}{N} \sum_{i=1}^N |P(x_i) - \tilde{P}(x_i)|^2. \quad (21)$$

The mean magnetization $\langle m \rangle = \frac{1}{N} \sum_{i=1}^N m_i$ describes the response of the system to the field θ [21], where we parametrize binary RVs by the magnetization $m_i = P(X_i = 1) - P(X_i = -1)$. Note that the difference between the mean magnetization and the approximate mean magnetization $\langle \tilde{m} \rangle$ is the sum over all marginal errors

$$\langle m \rangle - \langle \tilde{m} \rangle = \frac{2}{N} \sum_{i=1}^N P(x_i) - \tilde{P}(x_i). \quad (22)$$

It follows that the mean belief $\langle P(\mathbf{X}) \rangle$ (i.e., the expectation of binary RVs with $\mathbb{S} = \{0, 1\}$ averaged over all nodes) relates to the mean magnetization by $\langle P(\mathbf{X}) \rangle = \frac{1}{N} \sum_{i=1}^N P(x_i) = \frac{1}{2}(\langle m \rangle + 1)$.

Combining properly weighted marginals can lead to accuracy-improvements [14], [15]. As the NPHC method provides all fixed point solutions we combine multiple solutions to evaluate how much a weighted combination of – either all or just some – marginals increases the accuracy.

Therefore we consider the partition function and evaluate: the fixed point maximizing \tilde{Z}_B , a combination of all fixed points weighted by \tilde{Z}_B , and a weighted combination of fixed points that are local maxima of \tilde{Z}_B . Therefore we obtain the approximated marginals $\tilde{P}(\mathbf{X} = \mathbf{x})$ for every fixed

point according to (4) and (5); the associated approximate Bethe partition function $\tilde{Z}_B(\tilde{P}(X_i), \tilde{P}(X_i, X_j))$ is obtained by (7) and (8). Let $\tilde{\mathcal{P}}_{\text{ALL}} = \{(\tilde{P}(\mathbf{X} = \mathbf{x}), \tilde{Z}_B) : (\underline{\mu}, \underline{\alpha}) \in V_{\mathbb{R}^+}^*(\mathbf{F})\}$ be the set of all marginals and the associated Bethe partition function (obtained at fixed points). Accordingly, we define $\tilde{\mathcal{P}}_{\text{STABLE}}$ for locally stable solutions only. Then, if multiple fixed points exist, the weighted marginals are combined so that

$$\tilde{P}_{\text{ALL}}(\mathbf{X} = \mathbf{x}) = \frac{1}{\sum_{\tilde{Z}_B \in \tilde{\mathcal{P}}_{\text{ALL}}} \tilde{Z}_B} \sum_{(\tilde{P}(\mathbf{X}=\mathbf{x}), \tilde{Z}_B) \in \tilde{\mathcal{P}}_{\text{ALL}}} \tilde{P}(\mathbf{X} = \mathbf{x}) \cdot \tilde{Z}_B. \quad (23)$$

Further we consider the combination of stable fixed points $\tilde{\mathcal{P}}_{\text{STABLE}}(\mathbf{X} = \mathbf{x})$ and the marginals that maximize the Bethe partition function:

$$\tilde{P}_{\text{MAX}}(\mathbf{X} = \mathbf{x}) = \underset{\tilde{P}}{\text{argmax}} \tilde{Z}_B(\tilde{P}(X_i), \tilde{P}(X_i, X_j)). \quad (24)$$

4.2 Grid Graph with Random Factors (Spin Glass)

Consider a grid graph of size $N = 3 \times 3$ (Fig. 2a) with randomly distributed parameters. All pairwise and local potentials are sampled uniformly; i.e., $(J_{i,j}, \theta_i) \sim \mathcal{U}(-K, K)$. The larger the support of the uniform distribution is, the more difficult the task of inference becomes. For $K = 3$ inference is sufficiently difficult (cf. [13]).

According to (14) the system of equations consists of $M = 72$ equations and 72 unknowns. More specifically, (13) consists of 24 linear (i.e., normalization constraints), 40 quadratic, and 8 cubic equations; the total degree bounds the number of solutions by $d_t = 1^{24} \cdot 2^{40} \cdot 3^8 = 7.2 \cdot 10^{15}$. Tracking of such an amount of solution paths is not feasible in practice, even with a parallel implementation of the NPHC method. The system of equations in (13), however, is sparse. We can exploit this sparsity that is induced by the graph structure if we consider the BKK bound and reduce the computational complexity. The number of complex solutions for this graph is bounded by BKK = 608. After creating a suitable start system the problem is straightforward to solve with the NPHC method. It actually turns out that the BKK bound is tight for all graphs considered.

In particular we evaluated 100 grid graphs with random factors: on 99 graphs BP converged after at most 10^4 iterations. Although the grid graph has multiple loops and the constrained \mathbb{F}_B is not necessarily convex [5, Corr.2], we observe that for all 100 graphs NPHC obtains a unique positive real solution corresponding to a unique BP fixed point.

4.3 Grid Graph with Uniform Factors

We further analyze the convergence properties of BP on grid graphs of size $N = 3 \times 3$ (Fig. 2a) with constant potentials among all nodes and edges; i.e., for all edges $J_{i,j} = J$ and for all nodes $\theta_i = \theta$. We apply BP and NPHC for 1681 graphs in the parameter region $(J, \theta) \in \{-2, -1.9, \dots, 1.9, 2\}$.

The number of solutions in $V_{\mathbb{R}^+}^*(\mathbf{F})$ is presented in Fig. 3a. In the well-behaved region (III) of the parameter space a unique fixed point exists, whereas 3 fixed points exist in (I) and (II) – this is in accordance with statistical

TABLE 1: MSE OF MARGINALS AND COMBINED MARGINALS OBTAINED BY BP AND NPHC FOR THE GRID GRAPH WITH UNIFORM FACTORS.

| Parameters | | Fixed Points | | Combined | | |
|-----------------|----------------------|--------------|-------|----------|---------------------|---------------------|
| Couplings | Local Field | BP | NPHC | MAX | ALL | STABLE |
| $J \in [-2, 2]$ | $\theta \in [-2, 2]$ | 0.197 | 0.016 | 0.037 | 0.005 | 0.004 |
| $J \in (I)$ | $\theta \in (I)$ | 0.010 | 0.010 | 0.076 | $9.2 \cdot 10^{-4}$ | $1.0 \cdot 10^{-9}$ |
| $J \in (II)$ | $\theta \in (II)$ | 0.836 | 0.050 | 0.126 | 0.003 | $1.5 \cdot 10^{-6}$ |
| $J \in (III)$ | $\theta \in (III)$ | 0.006 | 0.006 | 0.006 | 0.006 | 0.006 |

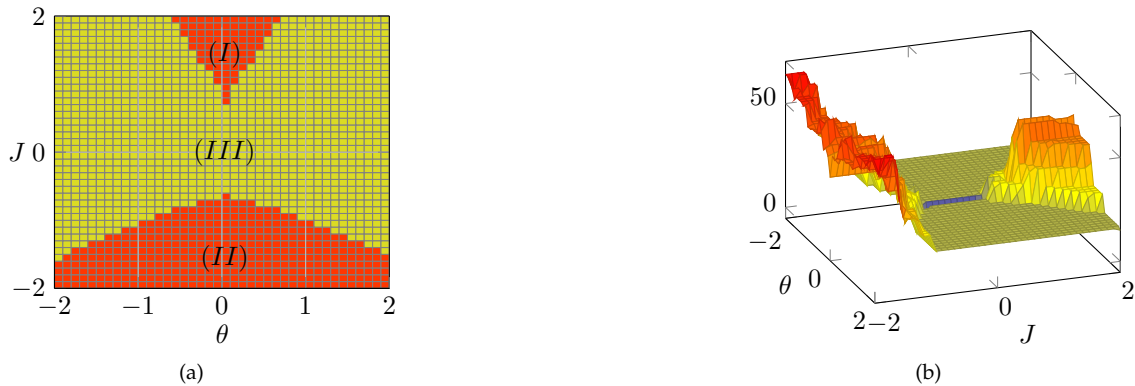


Fig. 3: Number of fixed points on the grid graph of size $N = 3 \times 3$: (a) number of positive real fixed points (yellow: unique fixed point, red: three fixed points); (b) number of real fixed points. The increase in number of both real solutions and positive real solutions, indicates a phase transition.

mechanics⁵ [21, p.43]. Interestingly, we observe a close relation between the onset of phase transitions⁶ and the increase in the number of real solutions in Fig. 3b. Most of these solutions, however, correspond to negative message values which violate Lemma 1 and are not feasible.

BP converges to some fixed point on all 1681 graphs within at most 10^4 iterations. This raises the question: what is the approximation error of BP if it converges to the best possible fixed point? Or, speaking in terms of free energies, how large is the gap between the global minimum of the constrained \mathbb{F}_B and the minimum of the Gibbs free energy? To answer this question we evaluate the correctness of the approximated marginals by computing the MSE between the exact and the approximated marginals according to (21). The results are presented in Table 1.

Averaged over all graphs we can see that BP does not necessarily converge to the best solution. For NPHC we present the MSE for the fixed point with the lowest MSE; this highlights the existence of fixed points, which give *more* accurate approximation than BP. Looking at all parameter regions separately we can see that BP does converge to the global optimum in (III), as well as in (I). In the antiferromagnetic region (II) BP converges to a fixed point that does not necessarily give the best possible approximation.

If we consider regions with multiple fixed point solutions (i.e., (I) and (II)) it becomes obvious that the fixed point maximizing \tilde{Z}_B is not necessarily the best one; this is

5. Note that the graph under consideration is of finite size and thus $|\partial(X_i)|$ varies among the nodes. As a consequence the partitioning according to (10) - (12), are only approximations.

6. We use the term phase transition for the finite-size manifestation of the phase transition in the thermodynamical sense.

especially surprising as BP obtains the best possible fixed point solution inside (I). It turns out that for $\theta = 0$ identical initialized messages $\mu_{i,j}^0(x_j) = \mu^0$ correspond to a fixed point that, although potentially being unstable, equals the exact solution where $P(\bar{x}_i) = P(x_i)$ for all X_i (cf. convergence results in Fig. 4a).

Inspired by these observations, one should not only consider the fixed point maximizing \tilde{Z}_B , i.e., $\tilde{P}_{\text{MAX}}(\mathbf{X} = \mathbf{x})$, but rather obtain multiple fixed points by NPHC and combine them. Indeed, especially inside region (I) a combination of *all* fixed point solutions according to (23), i.e., $\tilde{P}_{\text{ALL}}(\mathbf{X} = \mathbf{x})$ increases the accuracy of the approximation. If we combine the stable solution only, i.e., $\tilde{P}_{\text{STABLE}}(\mathbf{X} = \mathbf{x})$, the accuracy increases even more and gives the *most* accurate approximation over the entire parameter space.

4.4 Fully Connected Graph with Uniform Factors

We consider a fully connected Ising model with $|\mathbf{X}| = 2 \times 2$ binary RVs (Fig. 2b). The system of equations (13) consists of 36 equations in 36 unknowns and has its number of solutions bounded by the total degree $d_t = 1^{12} \cdot 2^{24} = 16.8 \cdot 10^6$. Similar to Sec. 4.2 a much tighter bound of 120 is provided by the BKK bound. Among all four nodes we apply uniform factors $J_{i,j} = J$ and $\theta_i = \theta$. This type of graph is particularly interesting because one can derive exact conditions where phase transitions occur (cf. Section 2.3).

For $(J, \theta) \in (III)$ BP has a unique fixed point, which is a stable attractor in the whole message space [34]. In (I) three fixed points satisfy (13), one of which is unstable and a local minimum of \tilde{Z}_B . Both other fixed points are local maxima of \tilde{Z}_B and BP converges to one of them. In

TABLE 2: MSE OF MARGINALS AND COMBINED MARGINALS OBTAINED BY BP AND NPHC FOR THE FULLY CONNECTED GRAPH WITH UNIFORM FACTORS.

| Parameters | | Fixed Points | | Combined | | |
|-----------------|----------------------|--------------|-------|----------|-------|---------------------|
| Couplings | Local Field | BP | NPHC | MAX | ALL | STABLE |
| $J \in [-2, 2]$ | $\theta \in [-2, 2]$ | 0.069 | 0.007 | 0.011 | 0.003 | 0.0027 |
| $J \in (I)$ | $\theta \in (I)$ | 0.034 | 0.033 | 0.070 | 0.002 | $2.0 \cdot 10^{-8}$ |
| $J \in (II)$ | $\theta \in (II)$ | 0.304 | 0.004 | 0.004 | 0.004 | 0.004 |
| $J \in (III)$ | $\theta \in (III)$ | 0.003 | 0.003 | 0.003 | 0.003 | 0.003 |




Fig. 4: Fully connected graph with $N = 2 \times 2$. (a) Convergence of BP: for the blue region BP did convergence – for the red region it did not converge after $4 \cdot 10^5$ iterations. (b) Number of fixed points (yellow: unique fixed point, red: three fixed points). (c) Number of real solutions – note the sudden increase at the onset of phase transitions.

the antiferromagnetic case BP does only converge inside (III) and not inside (II) as shown in Fig. 4a. We can see two interesting effects: first, in Fig. 4c the number of real solutions increases at the onset of phase transitions; secondly, even though the convergence of BP breaks down at the phase transition (Fig. 4) a unique fixed point exists inside (II) (Fig. 4b) that gives an accurate approximation (cf. Table 2).

Similar as in Sec. 4.3 we assess the MSE of the marginals obtained by NPHC and BP; averaged over all graphs, and for each distinct region in Table 2. Indeed, inside region (II) the marginals obtained by NPHC give a much better approximation than BP does.

Furthermore, note that the fixed point maximizing \tilde{Z}_B does not always give the best approximation. For $\theta = 0$ BP is initialized at the unstable fixed point and obtains the exact marginals (cf. Sec. 4.3). If multiple fixed points exist, a weighted combination of all marginals according to (23) increases the accuracy – only considering stable solutions, i.e., $\tilde{P}_{\text{STABLE}}(\mathbf{X} = \mathbf{x})$, gives the *most* accurate approximations.

4.5 Fixed Point Evolution

To get further insights we investigate the dependence of the accuracy of the fixed points on the parameters (J, θ) and compare the fixed point solutions obtained by NPHC to the exact solution. Therefore, we fix the values of $\theta \in \{0, 0.1, 0.5\}$ and vary $J \in [-2, 2]$. We illustrate the mean magnetization (cf. Sec. 4.1) of the exact solution and the approximate mean magnetization of all fixed points obtained

by NPHC for both the grid graph and the fully connected graph in Fig. 5-6.

The exact solution (red) is obtained by the junction tree algorithm [51]; solutions to (13) are obtained by NPHC and are depicted by blue dots (stable) and by green dots (unstable). We further illustrate the fixed point that maximizes \tilde{Z}_B , i.e., $\tilde{P}_{\text{MAX}}(\mathbf{X} = \mathbf{x})$, in black.

A fixed point is locally stable if a neighborhood exists such that messages inside this neighborhood converge to the fixed point [52, pp.170] and locally unstable otherwise. To analyze local stability one has to obtain all fixed points $(\underline{\mu}, \underline{\alpha}) \in V_{\mathbb{R}_+}^*(\mathbf{F})$ with the NPHC method first. Then BP is linearized in every fixed point by taking the partial derivatives of every messages with respect to all other messages, i.e., by analyzing the Jacobian matrix. Finally, the eigenvalues of the Jacobian matrix determine the stability of the fixed point. In practice one often uses damping, i.e., to replace messages with a weighted average of older messages: this may change the stability of the fixed points, but does not change the solutions $V_{\mathbb{R}_+}^*$ of (13) (cf. Sec. 2.2). We discuss some implications of the local stability analysis here, but focus mainly on the accuracy of the fixed points. For a detailed stability analysis of BP on Ising models we refer the reader to [22], [34].

The worst case scenario in terms of accuracy is shown in Fig. 5a for the grid graph with $N = 9$ binary RVs. For $\langle \tilde{m} \rangle = 0$ a fixed point exists that equals the exact solution. This fixed point, however, is only stable for $(J, \theta) \in (III)$. As $|J|$ increases to the onset of a phase transition this fixed point becomes unstable and two additional solutions

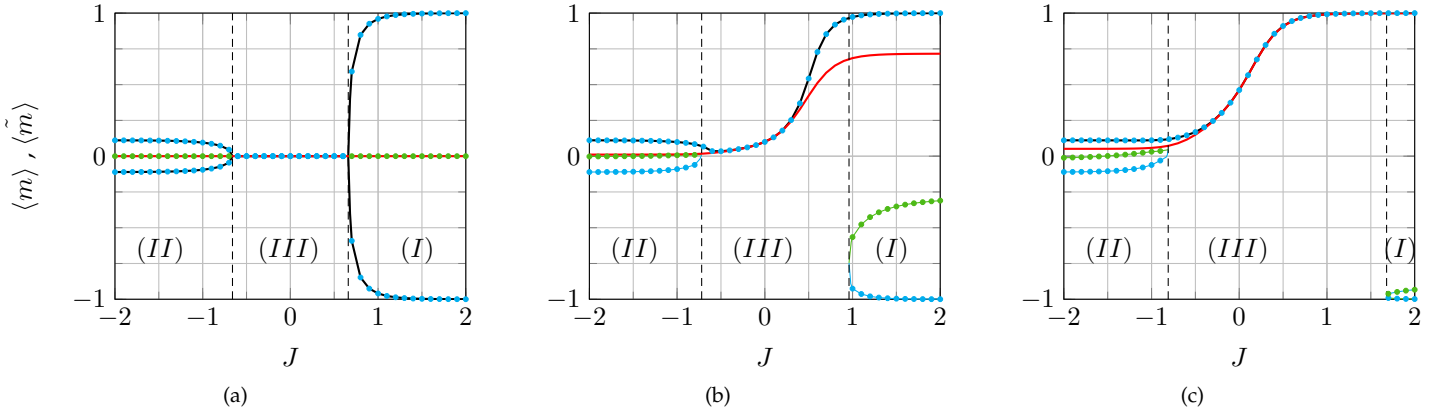


Fig. 5: Results for the grid graph of size $N = 3 \times 3$; mean magnetization $\langle m \rangle$ and $\langle \tilde{m} \rangle$ for $J \in [-2, 2]$ and for: (a) $\theta = 0$, (b) $\theta = 0.1$, and (c) $\theta = 0.5$. The exact solution is illustrated in red. All fixed points obtained by NPHC are depicted by blue dots (stable) and by green dots (unstable). The fixed point maximizing the partition function is illustrated in black.

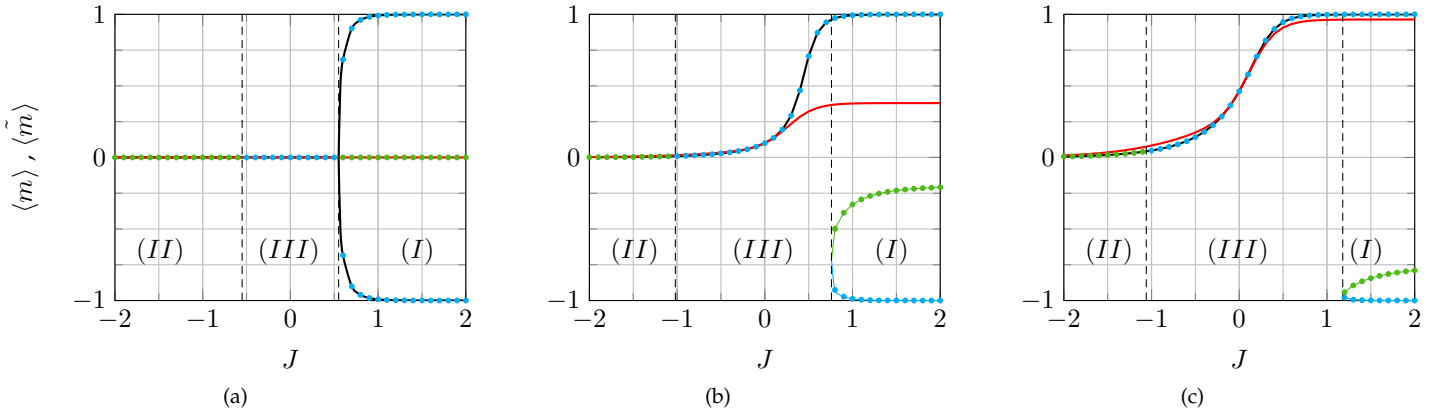


Fig. 6: Results for the fully connected graph of size $N = 2 \times 2$; mean magnetization $\langle m \rangle$ and $\langle \tilde{m} \rangle$ for $J \in [-2, 2]$ and for: (a) $\theta = 0$, (b) $\theta = 0.1$, and (c) $\theta = 0.5$. The exact solution is illustrated in red. All fixed points obtained by NPHC are depicted by blue dots (stable) and by green dots (unstable). The fixed point maximizing the partition function is illustrated in black.

emerge. These additional solutions are symmetric, stable, and guarantee the convergence of BP on this graph (Fig. 5a).

For $\theta \neq 0$ a unique stable fixed point exists inside (III). If we gradually increase J until $(J, \theta) \in (I)$ two additional fixed points emerge, one of which is unstable (Fig. 5b). Note that the fixed point maximizing \tilde{Z}_B (black) remains stable for all values of $J \in [-2, 2]$. An increase in θ (Fig. 5c) does enlarge the region where a unique fixed point exists (Fig. 3a) and further increases the accuracy of the fixed point maximizing \tilde{Z}_B . For $(J, \theta) \in (II)$ similar behavior is observed: i.e., for small values of θ the unstable fixed point has the highest accuracy, but as θ increases, the accuracy of the fixed point maximizing \tilde{Z}_B increases as well.

For the fully connected graph with $N = 4$ binary RVs (Fig. 2b) results are shown in Fig. 6. For $\theta = 0$ and large values of J the fixed point with $\langle m \rangle = 0$ is unstable and is accompanied by two symmetric, stable fixed points (Fig. 6a). In contrast to the grid graph a unique fixed point exists for $(J, \theta) \in (II)$; this fixed point, however, is unstable (Fig. 4b). Therefore, despite the existence of an accurate fixed point, BP does not converge (Fig. 4a).

For $\theta \neq 0$ the non-convergent region (II) is reduced, but the problem of a unique unstable fixed point persists. In contrast to the grid graph, the fully connected graph allows for frustrations with purely antiferromagnetic interactions. These points at a close connection between the existence of frustrations and the existence of a unique unstable solution. For $(J, \theta) \in (I)$ the fully connected graph behaves similar to the grid graph, i.e., the accuracy of the fixed point maximizing \tilde{Z}_B increases as θ increases (Fig. 6b-6c).

Our main findings are: First, increasing the field θ does lead to better convergence properties (cf. [22]) and increases the accuracy of the fixed point maximizing \tilde{Z}_B . Secondly, for $\theta \neq 0$ the fixed point maximizing \tilde{Z}_B is unique and varies continuously under a change of J . Finally, the stable fixed points are close to being symmetric, i.e., $\tilde{P}_{\text{stable},1}(\mathbf{X} = \mathbf{x}) \cong \tilde{P}_{\text{stable},2}(\mathbf{X} = \bar{\mathbf{x}})$. Consequently combining – either all or only stable – fixed points will not lead to good approximations unless a proper weighting by \tilde{Z}_B is applied. Applying a proper weighting, however, leads to accurate approximations (cf. Table 1-2).

TABLE 3: RUNTIME COMPARISON BETWEEN BP AND NPHC (IN SECONDS).

| | Grid Graph: Random (J, θ) | | Grid Graph: Uniform (J, θ) | | Fully Conn. Graph: Uniform (J, θ) | |
|-------------|-------------------------------------|----------|--------------------------------------|----------|---------------------------------------------|----------|
| | total | parallel | total | parallel | total | parallel |
| Mixed Vol. | 1364.5 | 11.2 | 1311.1 | 12.0 | 0.27 | – |
| Path Track. | 69.0 | 0.97 | 70.5 | 1.0 | 3.77 | – |
| Post Proc. | 29.9 | 1.40 | 43.6 | 4.1 | 2.3 | – |
| NPHC | 13.57 | | 17.1 | | 6.34 | |
| BP | $3.6 \cdot 10^{-3}$ | | $0.7 \cdot 10^{-3}$ | | 0.03 | |

4.6 Runtime Analysis

The time for solving (13) is presented in Table 3 for grid graphs with random factors (Sec. 4.2), grid graphs with uniform factors (Sec. 4.3), and fully connected graph with uniform factors (Sec. 4.4). Comparing the overall computation time of the NPHC method to BP it becomes obvious that NPHC is no alternative in terms of computational efficiency. It is however the only method which is guaranteed to obtain *all* fixed point solutions – we were not able to apply the Gröbner basis method beyond a single-cycle graph with $N = 4$. For the experiments we utilized a cluster-system with 160 CPUs: if we compare the overall computation time to the actual computation time utilizing the parallel implementation it becomes obvious that NPHC benefits tremendously from the high degree of parallelization. The runtime of BP depends mainly on the number of iterations and less on the size of the graph. Consequently, the stability of fixed points directly effects the performance of BP (cf. non-convergent region in Fig. 4a). The NPHC method is much less sensitive to the stability of fixed point solutions; the mixed volume computation, which has the largest influence on the overall runtime, rather depends on the number of variables in (13). Note the mixed volume computation does not depend on the parameters. If one is interested in the fixed points for different parameter-sets on the same graph it would suffice to perform the mixed volume computation and start system creation only once; we did not do this to allow for a fair comparison.

4.7 Error-Correcting Codes

One of the most prominent application areas where BP has a rich history of successful applications on cyclic graphs is iterative decoding. We keep this section as self-contained as the scope of this paper allows. For a thorough introduction we refer the interested reader to the textbooks [53], [54]; the connection between BP and decoding is further explained in great detail in [26], [55], [56].

We consider a binary symmetric channel (BSC) with a binary input $X_i \in \{0, 1\}$ and a binary output $Y_i \in \{0, 1\}$. The channel is specified by the error-probability ϵ , where transmitted bits are flipped with probability ϵ . That is $P(X_i = x_i | Y_i = y_i) = P(X_i = \bar{x}_i | Y_i = \bar{y}_i) = 1 - \epsilon$ and $P(X_i = x_i | Y_i = \bar{y}_i) = P(X_i = \bar{x}_i | Y_i = y_i) = \epsilon$ (cf. Fig. 7a). Additional, redundant bits help to detect and correct transmission errors. The aim of error-correcting codes is to reach the desired error-correcting performance while introducing as little redundancy as necessary, i.e., to operate as close as possible to the theoretical limit. Suppose we transmit a

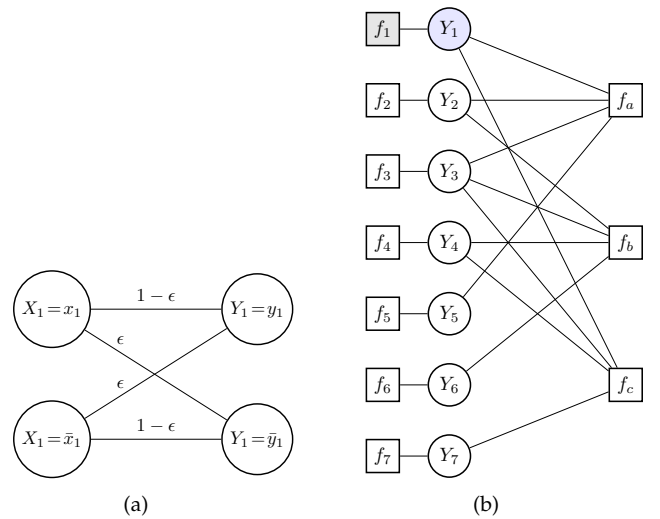


Fig. 7: (a) Binary symmetric channel with error probability ϵ ; (b) Factor graph for the $(7,4)$ Hamming code that corresponds to (31) where Y_1 is flipped.

codeword with block length $N = 7$ consisting of 4 source bits X_1, \dots, X_4 and three parity-check bits X_5, X_6, X_7 that satisfy

$$X_1 \oplus X_2 \oplus X_3 \oplus X_5 = 0, \quad (25)$$

$$X_2 \oplus X_3 \oplus X_4 \oplus X_6 = 0, \quad (26)$$

$$X_1 \oplus X_3 \oplus X_4 \oplus X_7 = 0, \quad (27)$$

where \oplus is an XOR, i.e., the sum in modulo-2 arithmetic. This linear irregular code is the $(7,4)$ Hamming code [53, Ch.1].

In this example we assume that the sent message is $\mathbf{X} = (0, 0, 0, 0, 0, 0, 0)^T$ and that exactly one bit suffers from a bit flip. For irregular codes the degree of the variables $\partial(Y_i)$ varies; therefore, we consider two scenarios: either $\mathbf{Y} = (1, 0, 0, 0, 0, 0, 0)$ or $\mathbf{Y} = (0, 0, 0, 0, 0, 1, 0)^8$.

It is often convenient to express a code in factorized form and represent it explicitly with a factor graph. A factor graph consists of variable nodes Y_i and factor nodes f_A where each factor f_A acts as a function on all variables connected $\mathbf{Y}_i = \{Y_i \in \partial(f_A)\}^9$. On a factor graph BP operates similar as introduced in Sec. 2.2. Now two types of messages are sent along every edge: factor-to-variable messages r_{f_A, Y_i} and variable-to-factor messages q_{Y_i, f_A} . All messages are iteratively updated according to

$$r_{f_A, Y_i}^{n+1}(y_i) = \sum_{\mathbf{y}_k: Y_k \in \partial(f_A) \setminus \{Y_i\}} f_A(\mathbf{Y}_k = \mathbf{y}_k, Y_i = y_i) \prod_{Y_k \in \partial(f_A) \setminus \{Y_i\}} q_{Y_k, f_A}^n(y_k), \quad (28)$$

$$q_{Y_i, f_A}^{n+1}(y_i) = \alpha_{Y_i, f_A}^n \prod_{f_B \in \{\partial(Y_i) \setminus f_A\}} r_{f_B, Y_i}^n(y_i), \quad (29)$$

7. Note that the properties of the BSC are independent of \mathbf{X} .

8. Normally the performance of a code is studied over an ensemble of sent codewords where each bit flips with probability ϵ .

9. Similar as in Sec. 2.2 we use $\partial(\cdot)$ to specify the neighbors of nodes and variables.




Fig. 8: Results for the (7,4) Hamming code. We compare the exact solution $P(X_i = 0|\mathbf{Y} = \mathbf{y})$ (red) to the approximate solution $\tilde{P}(X_i = 0|\mathbf{Y} = \mathbf{y})$ of NPHC (blue) as ϵ increases for: (a) Y_1 is flipped and (b) Y_6 is flipped.

where α_{Y_i, f_A} is chosen such that $q_{Y_i, f_A}^{n+1}(y_i) + q_{Y_i, f_A}^{n+1}(\bar{y}_i) = 1$. After all messages converged to a fixed point, the marginals of the variable nodes are approximated by the product of all incoming messages

$$P(Y_i = y_i) = \frac{1}{Z} \prod_{f_B \in \partial(Y_i)} r_{f_B, Y_i}^n(y_i). \quad (30)$$

Details of the factor graph representation can be found in [21, Ch.9] and [57].

Let us consider two types of factors: $f_i(Y_i) = P(X_i = x_i|Y_i = y_i)$ to model the BSC, and $f_A(\mathbf{Y}_i)$ to verify if all parity-checks are satisfied. Then $f_A(\mathbf{Y}_i) = 1$ if the sum of all arguments $\sum_{\mathbf{Y}_i} y_i$ is even and $f_A(\mathbf{Y}_i) = 0$ if the sum is odd. The conditional probability for $\mathbf{X} = \mathbf{x}$ to be the codeword, given the received codeword $\mathbf{Y} = \mathbf{y}$ then is

$$P(\mathbf{X} = \mathbf{x}|\mathbf{Y} = \mathbf{y}) = \frac{1}{Z} \prod_{i=1}^7 f_i(Y_i) \cdot f_a(Y_1, Y_2, Y_3, Y_5) \cdot f_b(Y_2, Y_3, Y_4, Y_6) \cdot f_c(Y_1, Y_3, Y_4, Y_7). \quad (31)$$

The corresponding factor graph representation is shown in Fig 7b.

Now we create a system of equations similar as in Sec. 3 and obtain all fixed points with NPHC. A unique fixed point exists for all settings – and this fixed point is stable, which justifies the application of BP on error-correcting codes. We further estimate the accuracy of the approximation; this relates to the question: If we communicate over a BSC, how vulnerable is BP decoding to an increased error probability? To answer this question we obtain the fixed points for $\epsilon \in [0, 0.5]$ and compare the exact solution obtained by the junction tree algorithm $P(X_i = 0|\mathbf{Y} = \mathbf{y})$ (red), to the approximate solution $\tilde{P}(X_i = 0|\mathbf{Y} = \mathbf{y})$ obtained by NPHC (blue) in Fig. 8¹⁰.

An error can be corrected by BP decoding if $\tilde{P}(Y_i = 0|\mathbf{X} = \mathbf{x}) > 0.5$. With exact decoding a single bit-flip can be corrected for $\epsilon < 0.21$. According to the fixed points obtained by NPHC, however, The fixed points obtained by NPHC however reveal that BP does not utilize the full potential of the code (Fig. 8). If Y_1 was corrupted the error

can be corrected for $\epsilon < 0.13$ (Fig. 8a). BP fails to correct the error if Y_6 was flipped for all values of ϵ (Fig. 8b). The reason therefore is a systematic error: the check-bit only has a single connection to the parity-check function f_c . According to (29) $q_{Y_6, f_c}^{n+1}(y_i) = \alpha_{Y_7, f_7} \cdot r_{f_7, Y_7}(y_i)$; therefore Y_6 does not incorporate any information from the remaining graph. To conclude, the higher the connectivity of a node, the more information of other bits is taken into account and the better the error-correction capability of BP and NPHC.

5 CONCLUSION

The NPHC method is presented as a tool to obtain all BP fixed point solutions. This work is an attempt to get a deeper understanding of BP, with potential implications for finding stronger conditions for uniqueness and convergence guarantees of BP fixed points.

One key feature of our framework is to come up with a favorable upper bound on the number of solutions. In particular the BKK bound utilizes the sparsity of the polynomial system induced by the graph structure and is tight in all our experiments. Finally, to obtain all fixed points, we create an appropriate start system and track all solution paths in parallel.

On fully connected graphs and grid graphs with binary Ising factors we show how the number of fixed points evolves over a large parameter region. While, in practice, fixed points have to be positive, we empirically showed that there is a close relation between the occurrence of phase transitions and an increase in the number of real solutions.

We empirically show an accuracy-gap between fixed points of BP and the best fixed points obtained by NPHC. In practice this justifies the exploration of multiple fixed points and selecting one that leads to the best approximation of the marginals. Moreover, we analyze the fixed point maximizing the partition function in detail: we show how it continuously deforms under varying parameters and how stability of this fixed point depends on the graph structure. The NPHC method further reveals for which parameters the fixed point maximizing the partition function does not correspond to the best fixed point obtained by NPHC. These observations motivate to consider a combination

10. Note for $\epsilon = 0.5$ the transmission is random

of weighted marginals. These weighted combinations provide strikingly accurate marginals whenever multiple fixed points exist.

When applied to graphs where BP does not converge, the NPHC method reveals that for some cases a unique fixed point does exist – consequently, uniqueness of BP fixed points is by no means sufficient to guarantee convergence of BP. Further we conjecture a close connection between the existence of frustrations and the existence of a unique unstable fixed point.

One requirement of NPHC to be efficient is a tight upper bound on the number of solutions. The determination of this bound limits our current investigations to relatively small graphs. In future we aim to exploit the graph structure in order to reduce the complexity of identifying an upper bound on the number of solutions.

ACKNOWLEDGMENTS

Research supported in part by NSF under Grant DMS 11-15587. This work was supported by the Austrian Science Fund (FWF) under the project number P28070-N33.

REFERENCES

- [1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, 1988.
- [2] D. Batra, P. Yadollahpour, A. Guzman-Rivera, and G. Shakhnarovich, “Diverse m-best solutions in markov random fields,” in *Computer Vision—ECCV 2012*. Springer, 2012, pp. 1–16.
- [3] F. Pernkopf, R. Peharz, and S. Tschitschek, *Introduction to Probabilistic Graphical Models*. Academic Press Library in Signal Processing, 2014.
- [4] J. Mooij and H. Kappen, “Sufficient conditions for convergence of the sum-product algorithm,” *IEEE Trans. on Information Theory*, vol. 53, no. 12, pp. 4422–4437, 2007.
- [5] T. Heskes, “On the uniqueness of loopy belief propagation fixed points,” *Neural Comp.*, vol. 16, no. 11, 2004.
- [6] Y. Weiss, “Correctness of local probability propagation in graphical models with loops,” *Neural Comp.*, vol. 12, no. 1, 2000.
- [7] A. Ihler, “Accuracy bounds for belief propagation,” in *Proceedings of UAI*, Jul. 2007, pp. 183–190.
- [8] A. Weller and T. Jebara, “Approximating the Bethe partition function,” in *Proceedings of UAI*, 2014.
- [9] K. Murphy, Y. Weiss, and M. Jordan, “Loopy belief propagation for approximate inference: An empirical study,” in *Proceedings of UAI*, 1999, pp. 467–475.
- [10] A. Yuille and A. Rangarajan, “The concave-convex procedure,” *Neural Comp.*, vol. 15, no. 4, pp. 915–936, 2003.
- [11] G. Elidan, I. McGraw, and D. Koller, “Residual belief propagation: Informed scheduling for asynchronous message passing,” in *Proceedings of UAI*, 2006.
- [12] C. Knoll, M. Rath, S. Tschitschek, and F. Pernkopf, “Message scheduling methods for belief propagation,” in *Machine Learning and Knowledge Discovery in Databases*. Springer, 2015, pp. 295–310.
- [13] C. Sutton and A. McCallum, “Improved dynamic schedules for belief propagation,” in *Proceedings of UAI*, 2007.
- [14] A. Braunstein, M. Mézard, and R. Zecchina, “Survey propagation: An algorithm for satisfiability,” *Random Structures & Algorithms*, vol. 27, no. 2, pp. 201–226, 2005.
- [15] C. Srinivasa, S. Ravanbakhsh, and B. Frey, “Survey propagation beyond constraint satisfaction problems,” in *Proceedings of AISTATS*, 2016, pp. 286–295.
- [16] D. Cox, J. Little, and D. O’Shea, *Ideals, Varieties, and Algorithms*. Springer, 1992, vol. 3.
- [17] —, *Using Algebraic Geometry*. Springer Science & Business Media, 2005, vol. 185.
- [18] T.-Y. Li, “Solving polynomial systems by the homotopy continuation method,” *Handbook of numerical analysis*, vol. 11, pp. 209–304, 2003.
- [19] A. Sommese and C. Wampler, *The Numerical Solution of Systems of Polynomials Arising in Engineering and Science*. World Scientific, 2005, vol. 99.
- [20] H. Georgii, *Gibbs Measures and Phase Transitions*. Walter de Gruyter, 2011, vol. 9.
- [21] M. Mezard and A. Montanari, *Information, Physics, and Computation*. Oxford Univ. Press, 2009.
- [22] C. Knoll and F. Pernkopf, “On loopy belief propagation – local stability analysis for non-vanishing fields,” *Submitted to UAI*, 2017.
- [23] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. MIT press, 2009.
- [24] M. Jordan, “Graphical models,” *Statistical Science*, pp. 140–155, 2004.
- [25] R. Gallager, *Information Theory and Reliable Communication*. Springer, 1968, vol. 2.
- [26] F. Kschischang, B. Frey, and H. Loeliger, “Factor graphs and the sum-product algorithm,” *IEEE Trans. on Information Theory*, vol. 47, no. 2, pp. 498–519, 2001.
- [27] A. Ihler, J. Fisher, and A. Willsky, “Loopy belief propagation: Convergence and effects of message errors,” in *Journal of Machine Learning Research*, 2005, pp. 905–936.
- [28] J. Yedidia, W. Freeman, and Y. Weiss, “Constructing free-energy approximations and generalized belief propagation algorithms,” *IEEE Trans. on Information Theory*, vol. 51, no. 7, pp. 2282–2312, 2005.
- [29] D. MacKay, “A conversation about the Bethe free energy and sum-product,” *Tech. Rep. of MERL*, 2001.
- [30] M. Welling and Y. Teh, “Approximate inference in Boltzmann machines,” *Artificial Intelligence*, vol. 143, no. 1, pp. 19–50, 2003.
- [31] S. C. Tatikonda and M. I. Jordan, “Loopy belief propagation and Gibbs measures,” in *Proceedings of UAI*, 2002, pp. 493–500.
- [32] M. J. Wainwright and M. I. Jordan, “Graphical models, exponential families, and variational inference,” *Foundations and Trends in Machine Learning*, vol. 1, no. 1–2, pp. 1–305, 2008.
- [33] N. Taga and S. Mase, “On the convergence of belief propagation algorithm for stochastic networks with loops,” 2004.
- [34] J. Mooij and H. Kappen, “On the properties of the Bethe approximation and loopy belief propagation on binary networks,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2005, no. 11, p. P11012, 2005.
- [35] T. Heskes et al., “Stable fixed points of loopy belief propagation are minima of the Bethe free energy,” in *NIPS*, vol. 15, 2003, pp. 359–366.
- [36] O. Meshi, A. Jaimovich, A. Globerson, and N. Friedman, “Convexifying the Bethe free energy,” in *Proceedings of UAI*, 2009, pp. 402–410.
- [37] M. Welling and Y. W. Teh, “Belief optimization for binary networks: A stable alternative to loopy belief propagation,” in *Proceedings of UAI*, 2001, pp. 554–561.
- [38] J. Shin, “Complexity of Bethe approximation.” in *Proceedings of AISTATS*, 2012, pp. 1037–1045.
- [39] E. Mayr and A. Meyer, “The complexity of the word problems for commutative semigroups and polynomial ideals,” *Advances in mathematics*, vol. 46, no. 3, pp. 305–329, 1982.
- [40] H. Möller and F. Mora, “Upper and lower bounds for the degree of Gröbner bases,” *EUROSAM 84*, pp. 172–183, 1984.
- [41] B. Huber and B. Sturmfels, “A polyhedral method for solving sparse polynomial systems,” *Mathematics of computation*, vol. 64, no. 212, pp. 1541–1555, 1995.
- [42] D. Bernstein, “The number of roots of a system of equations,” *Funkts. Anal. Pril.*, vol. 9, pp. 1–4, 1975.
- [43] A. Kushnirenko, “Newton polytopes and the bezout theorem,” *Funkts. Anal. Pril.*, vol. 10, no. 3, 1976.
- [44] A. Khovanski, “Newton polyhedra and the genus of complete intersections,” *Funkts. Anal. Pril.*, vol. 12, no. 1, pp. 51–61, 1978.
- [45] T. Chen, T.-L. Lee, and T.-Y. Li, “Hom4ps-3: A parallel numerical solver for systems of polynomial equations based on polyhedral homotopy continuation methods,” in *Mathematical Software—ICMS 2014*. Springer, 2014, pp. 183–190.
- [46] T.-Y. Li, “Numerical solution of multivariate polynomial systems by homotopy continuation methods,” *Acta numerica*, vol. 6, pp. 399–436, 1997.
- [47] T. Chen and T.-Y. Li, “Homotopy continuation method for solving systems of nonlinear and polynomial equations,” *Commun. Inf. Syst.*, vol. 15, no. 2, pp. 119–307, 2015.

- [48] T. Chen and D. Mehta, "On the network topology dependent solution count of the algebraic load flow equations," *arXiv preprint arXiv:1512.04987*, 2015.
- [49] T. Chen, D. Mehta, and M. Niemerg, "A network topology dependent upper bound on the number of equilibria of the Kuramoto model," *arXiv preprint arXiv:1603.05905*, 2016.
- [50] J. Mooij, "libdai: A free and open source c++ library for discrete approximate inference in graphical models," *The Journal of Machine Learning Research*, vol. 11, 2010.
- [51] S. Lauritzen and D. Spiegelhalter, "Local computations with probabilities on graphical structures and their application to expert systems," *Journal of the Royal Statistical Society. Series B*, pp. 157–224, 1988.
- [52] G. Teschl, *Ordinary Differential Equations and Dynamical Systems*, 2003.
- [53] D. J. MacKay, *Information Theory, Inference and Learning Algorithms*. Cambridge university press, 2003.
- [54] H. Wymeersch, *Iterative Receiver Design*. Cambridge University Press Cambridge, 2007, vol. 234.
- [55] F. R. Kschischang and B. J. Frey, "Iterative decoding of compound codes by probability propagation in graphical models," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 2, pp. 219–230, 1998.
- [56] S. M. Aji and R. J. McEliece, "The generalized distributive law," *IEEE transactions on Information Theory*, vol. 46, no. 2, pp. 325–343, 2000.
- [57] H.-A. Loeliger, "An introduction to factor graphs," *IEEE Signal Processing Magazine*, vol. 21, no. 1, pp. 28–41, 2004.