

# Approximate inference on planar graphs using Loop Calculus and Belief Propagation

**Vicenç Gómez**  
**Hilbert J. Kappen**  
*Department of Biophysics*  
*Radboud University Nijmegen*  
*6525 EZ Nijmegen, The Netherlands*

V.GOMEZ@SCIENCE.RU.NL  
 B.KAPPEN@SCIENCE.RU.NL

**Michael Chertkov**  
*Theoretical Division and Center for Nonlinear Studies*  
*Los Alamos National Laboratory*  
*Los Alamos, NM 87545*

CHERTKOV@LANL.GOV

**Editor:**

## Abstract

We introduce novel results for approximate inference on planar graphical models using the loop calculus framework. The loop calculus (Chertkov and Chernyak, 2006a) allows to express the exact partition function of a graphical model as a finite sum of terms that can be evaluated once the belief propagation (BP) solution is known. In general, full summation over all correction terms is intractable. We develop an algorithm for the approach presented in Chertkov et al. (2008) which represents an efficient truncation scheme on planar graphs and a new representation of the series in terms of Pfaffians of matrices. We analyze the performance of the algorithm for the partition function approximation for models with binary variables and pairwise interactions on grids and other planar graphs. We study in detail both the loop series and the equivalent Pfaffian series and show that the first term of the Pfaffian series for the general, intractable planar model, can provide very accurate approximations. The algorithm outperforms previous truncation schemes of the loop series and is competitive with other state-of-the-art methods for approximate inference.

**Keywords:** belief propagation, loop calculus, approximate inference, partition function, planar graphs.

## 1. Introduction

Graphical models are popular tools widely used in many areas which require modeling of uncertainty. They provide an effective approach through a compact representation of the joint probability distribution. The two most common types of graphical models are Bayesian Networks (BN) and Markov Random Fields (MRFs).

The partition function of a graphical model, which plays the role of normalization constant in a MRF or probability of evidence (likelihood) in a BN is a fundamental quantity which arises in many contexts such as hypothesis testing or parameter estimation. Exact computation of this quantity is only feasible when the graph is not too complex, or equiv-

alently, when its tree-width is small. Currently many methods are devoted to approximate this quantity.

The belief propagation (BP) algorithm (Pearl, 1988) is at the core of many of these approximate inference methods. Initially thought as an exact algorithm for tree graphs, it is widely used as an approximation method for loopy graphs (Murphy et al., 1999; Frey and MacKay, 1998). The exact partition function is explicitly related to the BP approximation through the loop calculus framework introduced by Chertkov and Chernyak (2006a). Loop calculus allows to express the exact partition function as a finite sum of terms (loop series) that can be evaluated once the BP solution is known. Each term maps uniquely to a subgraph, also denoted as a generalized loop, where the connectivity of any node within the subgraph is *at least* degree two. Summation of the entire loop series is a hard combinatorial task since the number of generalized loops is typically exponential in the size of the graph. However, different approximations can be obtained by considering different subsets of generalized loops in the graph.

It has been shown empirically (Gómez et al., 2007; Chertkov and Chernyak, 2006b) that truncating this series may provide efficient corrections to the initial BP approximation. More precisely, whenever BP performs satisfactorily which occurs in the case of sufficiently weak interactions between variables or short-range influence of loops, accounting for only a small number of terms is sufficient to recover the exact result (Gómez et al., 2007). On the other hand, for those cases where BP requires many iterations to converge, many terms of the series are required to improve substantially the approximation. A formal characterization of the classes of tractable problems via loop calculus still remains as an open question.

A step toward this goal has been done in Chertkov et al. (2008) where it was shown that for any graphical model, summation of a certain subset of terms can be mapped to a summation of weighted perfect matchings on an extended graph. For planar graphs (graphs that can be embedded into a plane without crossing edges), summation of the subset can be performed in polynomial time evaluating the Pfaffian of a skew-symmetric matrix associated with the extended graph. Furthermore, the full loop series can be expressed as a sum over certain Pfaffian terms, where each Pfaffian term accounts for a large number of loops and is solvable in polynomial time as well.

The approach of Chertkov et al. (2008) builds on classical results from 1960s by Kasteleyn (1963); Fisher (1966) and other physicists who addressed the question of counting the number of perfect matchings on a planar grid, also known as the dimer problem in the statistical physics literature (a dimer correspond to a colored edge of the graph, and a valid dimer configuration consists of exactly one dimer per any edge of the graph). The key result of Kasteleyn (1963); Fisher (1966) can be summarized as follows: the partition function of a *planar graphical model defined in terms of binary variables* can be mapped to a weighted perfect matching problem and calculated in polynomial time under the restriction that interactions only depend on agreement or disagreement between the signs of their variables. Such a model is known in statistical physics as the Ising model *without external field*. Notice that exact inference for a general binary graphical model on a planar graph (that is Ising model with external field) is intractable (Barahona, 1982).

Recently, some methods for inference over graphical models, based on the works of Kasteleyn and Fisher, have been introduced. Globerson and Jaakkola (2007) obtained upper bounds on the partition function for non-planar graphs with binary variables by

decomposition of the partition function into a weighted sum over partition functions of spanning tractable (zero field) planar models. The resulting problem is a convex optimization problem and, since exact inference can be done in each planar *sub*-model, the bound can be calculated in polynomial time.

Another example is the work of Schraudolph and Kamenetsky (2008) which provides a framework for exact inference on a restricted class of planar graphs using the approach of Kasteleyn and Fisher. More precisely, they showed that any joint probability function defined on binary variables can be expressed in a functional form without external fields by adding a new auxiliary node linked to all the existing nodes. Under this transformation, single-variable external fields can be allowed for a subset  $\mathcal{B}$  of variables. If the graphical model is  $\mathcal{B}$ -outerplanar, which means that there exists a planar embedding in which the subset  $\mathcal{B}$  of the nodes lie on the same face, the techniques of Kasteleyn and Fisher can still be applied.

Contrary to the two aforementioned approaches which rely on exact inference on a tractable planar model, the loop calculus directly leads to a framework for approximate inference on general planar graphs. Truncating the loop series according to Chertkov et al. (2008) already gives the exact result in the zero external field case. In the general planar case, however, this truncation may result into an accurate approximation that can be incrementally corrected by considering subsequent terms in the series.

In the next Section we review the main theoretical results of the loop calculus approach for planar graphs and introduce the proposed algorithm. In Section 3 we provide experimental results on approximation of the partition function for regular grids and other types of planar graphs. We focus on a planar-intractable binary model with symmetric pairwise interactions but nonzero single variable potentials. The source code used to derive these results is freely available at <http://www.mbfys.ru.nl/staff/v.gomez/>. We end this manuscript with conclusions and future work in Section 4.

## 2. Belief Propagation and loop Series for Planar Graphs

We consider the Forney graph representation, also called general vertex model (Forney, 2001; Loeliger, 2004), of a probability distribution  $p(\boldsymbol{\sigma})$  defined over a vector  $\boldsymbol{\sigma}$  of binary variables (vectors are denoted using bold symbols). Forney graphs are associated with general graphical models which subsume other factor graphs, e.g. those correspondent to BNs and MRFs. In Appendix A we show how to convert a factor graph model to its equivalent Forney graph representation.

A binary Forney graph  $\mathcal{G} := (\mathcal{V}, \mathcal{E})$  consists of a set of nodes  $\mathcal{V}$  where each node  $a \in \mathcal{V}$  represents an interaction and each edge  $(a, b) \in \mathcal{E}$  represents a binary variable  $ab$  which take values  $\sigma_{ab} := \{\pm 1\}$ . We denote  $\bar{a}$  the set of neighbors of node  $a$ . Interactions  $f_a(\boldsymbol{\sigma}_a)$  are arbitrary functions defined over typically small subsets of variables where  $\boldsymbol{\sigma}_a$  is the vector of variables associated with node  $a$ , i.e.  $\boldsymbol{\sigma}_a := (\sigma_{ab_1}, \sigma_{ab_2}, \dots)$  where  $b_i \in \bar{a}$ .

The joint probability distribution of such a model factorizes as:

$$p(\boldsymbol{\sigma}) = Z^{-1} \prod_{a \in \mathcal{V}} f_a(\boldsymbol{\sigma}_a), \quad Z = \sum_{\boldsymbol{\sigma}} \prod_{a \in \mathcal{V}} f_a(\boldsymbol{\sigma}_a), \quad (1)$$

where  $Z$  is the normalization factor, also called the partition function.

From a variational perspective, a fixed point of the BP algorithm represents a stationary point of the Bethe "free energy" approximation under proper constraints (Yedidia et al., 2000). In the Forney style notation:

$$\begin{aligned} Z^{BP} &= \exp(-F^{BP}), \\ F^{BP} &= \sum_a \sum_{\sigma_a} b_a(\sigma_a) \ln \left( \frac{b_a(\sigma_a)}{f_a(\sigma_a)} \right) - \sum_{b \in \bar{a}} \sum_{\sigma_{ab}} b_{ab}(\sigma_{ab}) \ln b_{ab}(\sigma_{ab}), \end{aligned} \quad (2)$$

where  $b_a(\sigma_a)$  and  $b_{ab}(\sigma_{ab})$  are the beliefs (pseudo-marginals) associated to each node  $a \in \mathcal{V}$  and variable  $ab$ . For graphs without loops, Equation (2) coincides with the Gibbs "free energy" and therefore  $Z^{BP}$  coincides with the exact partition function  $Z$ . If the graph contains loops,  $Z^{BP}$  is just an approximation critically dependent on how strong the influence of the loops is.

We introduce now some convenient definitions related to the loop calculus framework.

**Definition 1** A *generalized loop* in a graph  $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$  is any subgraph  $C = \langle V', E' \rangle$ ,  $V' \subseteq \mathcal{V}$ ,  $E' \subseteq (V' \times V') \cap \mathcal{E}$  such that each node in  $V'$  has degree two or larger.

For simplicity, we will use the term "loop", instead of "generalized loop", in the rest of this manuscript. Loop calculus allows to represent  $Z$  explicitly in terms of the BP approximation via the loop series expansion:

$$Z = Z^{BP} \cdot z, \quad z = \left( 1 + \sum_{C \in \mathcal{C}} r_C \right), \quad r_C = \prod_{a \in C} \mu_{a; \bar{a}_C}, \quad (3)$$

where  $\mathcal{C}$  is the set of all the loops within the graph. Each loop term  $r_C$  is a product of terms  $\mu_{a; \bar{a}_C}$  associated with every node  $a$  of the loop.  $\bar{a}_C$  denotes the set of neighbors of  $a$  within the loop  $C$ :

$$\mu_{a; \bar{a}_C} = \frac{\sum_{\sigma_a} b_a(\sigma_a) \prod_{b \in \bar{a}_C} (\sigma_{ab} - m_{ab})}{\prod_{b \in \bar{a}_C} \sqrt{1 - m_{ab}^2}}, \quad m_{ab} = \sum_{\sigma_{ab}} \sigma_{ab} b_{ab}(\sigma_{ab}). \quad (4)$$

In this work we consider planar graphs where all nodes are of degree not larger than three, that is  $|\bar{a}_C| \leq 3$ . We denote by *triplet* a node with degree three in the graph. In Appendix A we show that a graphical model can be converted to this representation at the cost of introducing auxiliary nodes.

**Definition 2** A *2-regular loop* is a loop in which all nodes have degree exactly two.

**Definition 3** The *2-regular partition function*  $Z_\emptyset$  is the truncated form of (3) which sums all 2-regular loops only: <sup>1</sup>

$$Z_\emptyset = Z^{BP} \cdot z_\emptyset, \quad z_\emptyset = 1 + \sum_{C \in \mathcal{C} \text{ s.t. } |\bar{a}_C|=2, \forall a \in C} r_C. \quad (5)$$

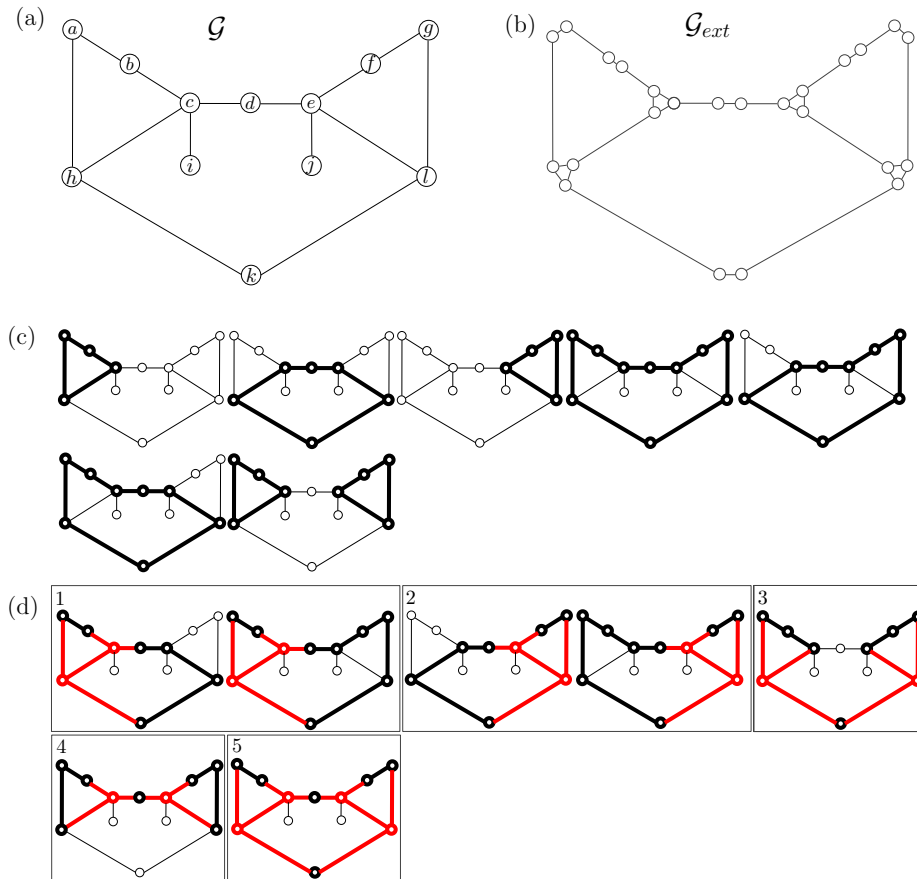


Figure 1: Example. **(a)** A Forney graph. **(b)** Corresponding extended graph. **(c)** Loops (in bold) included in the 2-regular partition function. **(d)** Loops (in bold and red) not included in the 2-regular partition function. Marked in red, the triplets associated with each loop. Grouped in gray squares, the loops considered in different subsets  $\Psi$  of triplets: (d.1)  $\Psi = \{c, h\}$ , (d.2)  $\Psi = \{e, l\}$ , (d.3)  $\Psi = \{h, l\}$ , (d.4)  $\Psi = \{c, e\}$  and (d.5)  $\Psi = \{c, e, h, l\}$  (see Section 2.2).

As an example, Figure 1a shows a small Forney graph and Figure 1c shows seven loops found in the corresponding 2-regular partition function.

### 2.1 Computing the 2-regular Partition Function Using Perfect Matching

In Chertkov et al. (2008) it has been shown that computation of  $Z_\theta$  can be mapped to a dimer/matching problem, or equivalently, to the computation of the sum of all weighted perfect matchings within another graph. A perfect matching is a subset of edges such that

1. Notice that this part of the series was called *single-connected partition function* in Chertkov et al. (2008). Here we prefer to call it 2-regular partition function because loops with more than one connected component are also included in this part of the series.

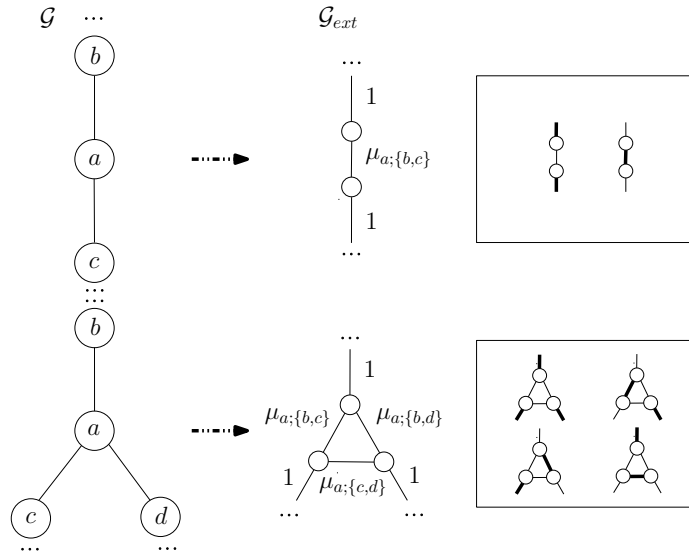


Figure 2: Fisher’s rules. **(Top)** A node  $a$  of degree two in  $\mathcal{G}$  is split in two nodes in  $\mathcal{G}_{ext}$ . **(Bottom)** A node  $a$  of degree three in  $\mathcal{G}$  is split in three nodes in  $\mathcal{G}_{ext}$ . The squares on the right indicate all possible matchings in  $\mathcal{G}_{ext}$  related with node  $a$ . Note that the rules preserve planarity.

each node neighbors exactly one edge from the subset. The weight of a matching is the product of weights of edges in the matching. The key idea of this mapping is to extend the original Forney graph  $\mathcal{G}$  into a new graph  $\mathcal{G}_{ext} := (\mathcal{V}_{\mathcal{G}_{ext}}, \mathcal{E}_{\mathcal{G}_{ext}})$  in such a way that each perfect matching in  $\mathcal{G}_{ext}$  corresponds to a 2-regular loop in  $\mathcal{G}$ . (See Figures 1b and c for an illustration). Under the condition of planarity, the sum of all weighted perfect matchings can be calculated in a polynomial time following Kasteleyn’s arguments. Here we reproduce these results with little variations and more emphasis on the algorithmic aspects.

Given a Forney graph  $\mathcal{G}$  and the BP approximation, we simplify  $\mathcal{G}$  and obtain the 2-core by removing nodes of degree one recursively. After this step,  $\mathcal{G}$  is either the null graph (and then BP is exact) or it is only composed of vertices of degree two or three.

To construct the extended graph  $\mathcal{G}_{ext}$  we split each node in  $\mathcal{G}$  according to the rules introduced by Fisher (1966) and illustrated in Figure 2. The procedure results in an extended graph of  $|\mathcal{V}_{\mathcal{G}_{ext}}| \leq 3|\mathcal{V}|$  nodes and  $|\mathcal{E}_{\mathcal{G}_{ext}}| \leq 3|\mathcal{E}|$  edges. It is easy to verify that each 2-regular loop in  $\mathcal{G}$  is associated with a perfect matching in  $\mathcal{G}_{ext}$  and, furthermore, this correspondence is *unique*. Consider, for instance, the vertex of degree three in the bottom of Figure 2. Given a 2-regular loop  $C$ , vertex  $a$  can appear in four different configurations: either node  $a$  does not appear in  $C$ , or  $C$  contains one of the following three paths:  $-b-a-c-$ ,  $-b-a-d-$  or  $-c-a-d-$ . These four cases correspond to node terms in a loop with values  $1$ ,  $\mu_{a;\{b,c\}}$ ,  $\mu_{a;\{b,d\}}$  and  $\mu_{a;\{c,d\}}$  respectively, and coincide with the matchings in  $\mathcal{G}_{ext}$  shown within the box on the bottom-right. An simpler argument applies to the vertex of degree two from the top portion of Figure 2.

Therefore, if we associate to each internal edge (new edge in  $\mathcal{G}_{ext}$  not in  $\mathcal{G}$ ) of each split node  $a$  the corresponding term  $\mu_{a;\bar{a}_C}$  of Equation (4) and to the external edges (existing edges already in  $\mathcal{G}$ ) weight 1, then the sum over all weighted perfect matchings defined on  $\mathcal{G}_{ext}$  is precisely  $z_\emptyset$ . The 2-regular partition function  $Z_\emptyset$  is obtained using Equation (5). Equivalently:

$$z_\emptyset = \sum \text{perfect matchings in } \mathcal{G}_{ext}.$$

Kasteleyn (1963) provided a method to compute this sum in polynomial time for planar graphs. We follow his approach. First, we create a *planar embedding* of  $\mathcal{G}_{ext}$ . A planar embedding of a graph divides the plane into disjoint regions that are bounded by sequences of edges in the graph. The regions are called *faces*. Second, we orient the edges of the planar embedding in such a way that for every face (except possibly the unbounded or external face) the number of clockwise oriented edges is odd. Algorithm 1 produces such an orientation (Karpinski and Rytter, 1998). It receives an undirected graph  $\mathcal{G}_{ext}$  and constructs a copy  $\mathcal{G}'_{ext} := (\mathcal{V}_{\mathcal{G}'_{ext}}, \mathcal{E}_{\mathcal{G}'_{ext}})$  with properly oriented edges  $\mathcal{E}_{\mathcal{G}'_{ext}}$ .

It is convenient that  $\mathcal{G}_{ext}$  is bi-connected, i.e. it has no articulation points. If needed, we add dummy edges with zero weight which do not alter the partition function or the original model.

---

**Algorithm 1** Pfaffian orientation

---

**Arguments:** undirected bi-connected extended graph  $\mathcal{G}_{ext}$ .

- 1: Construct a planar embedding  $\bar{\mathcal{G}}_{ext}$  of  $\mathcal{G}_{ext}$ .
  - 2: Construct a spanning tree  $T$  of  $\bar{\mathcal{G}}_{ext}$ .
  - 3: Construct a graph  $H$  having vertices corresponding to the faces of  $\bar{\mathcal{G}}_{ext}$ :  
Connect two vertices in  $H$  if the respective face boundaries share an edge not in  $T$ .  
 $H$  is a tree. Root  $H$  to the external face.
  - 4:  $\mathcal{G}'_{ext} := T$ .
  - 5: Orient all edges in  $\mathcal{G}'_{ext}$  arbitrarily.
  - 6: **for all** face (vertex in  $H$ ) traversed in post-order **do**
  - 7:   Add to  $\mathcal{G}'_{ext}$  the unique edge not in  $\mathcal{G}'_{ext}$ .
  - 8:   Orient it such that the number of clock-wise oriented edges is odd.
  - 9: **end for**
  - 10: **RETURN** directed bi-connected extended graph  $\mathcal{G}'_{ext}$ .
- 

Finally, denote  $\mu_{ij}$  the weight of the edge between nodes  $i$  and  $j$  in  $\mathcal{G}'_{ext}$ . We create the following skew-symmetric matrix  $\hat{A} = -\hat{A}^t$ :

$$\hat{A}_{ij} = \begin{cases} +\mu_{ij} & \text{if } (i, j) \in \mathcal{E}_{\mathcal{G}'_{ext}} \\ -\mu_{ij} & \text{if } (j, i) \in \mathcal{E}_{\mathcal{G}'_{ext}} \\ 0 & \text{otherwise} \end{cases}.$$

This matrix is known as the Tutte matrix of  $\mathcal{G}'_{ext}$  and the Pfaffian of  $\hat{A}$  gives the desired sum up to the overall sign. The Pfaffian of  $\hat{A} = \pm\sqrt{\text{Det}(\hat{A})}$ . However,  $z_\emptyset$  can be either positive or negative, and computing the value of the Pfaffian with the sign yet uncertain

is not sufficient. Furthermore, since each element  $\hat{A}_{ij}$  can be negative not only due to the Pfaffian orientation but also if  $\mu_{ij}$  is negative, the sign of the Pfaffian needs to be *corrected*. This problem is fixed with the help of the original Kasteleyn's binary matrix:

$$\hat{B}_{ij} = \begin{cases} +1 & \text{if } (i, j) \in \mathcal{E}_{\mathcal{G}'_{ext}} \\ -1 & \text{if } (j, i) \in \mathcal{E}_{\mathcal{G}'_{ext}} \\ 0 & \text{otherwise} \end{cases}.$$

If the sign of  $\text{Pf}(\hat{B})$  is negative then the sign of  $\text{Pf}(\hat{A})$  is changed. Notice that the absolute value of  $\text{Pf}(\hat{B})$  coincides with the number of perfect matchings or the number of loops included in the sum if no additional edges have been added. The sign of  $\text{Pf}(\hat{B})$  represents the correction. Therefore, the corrected value of  $z_\emptyset$  is:

$$z_\emptyset = \text{sign} \left( \text{Pf} \left( \hat{B} \right) \right) \cdot \text{Pf} \left( \hat{A} \right).$$

Calculation of  $z_\emptyset$  can therefore be performed in time  $\mathcal{O}(N^3)$  where  $N$  is the number of nodes of  $\mathcal{G}_{ext}$  (Galbiati and Maffioli, 1994). For the special case of binary planar graphs with zero local fields the 2-regular partition function coincides with the exact partition function  $Z = Z_\emptyset = Z^{BP} \cdot z_\emptyset$  since the other terms in the loop series vanish.

## 2.2 Computing the Full Loop Series Using Perfect Matching

Chertkov et al. (2008) established that  $z_\emptyset$  is just the first term of a finite sum involving Pfaffians. We briefly reproduce their results here and provide an algorithm for computing the full loop series as a Pfaffian series.

Consider  $\mathcal{T}$  defined as the set of all possible triplets (vertices with degree three in the original graph  $\mathcal{G}$ ). For each possible subset  $\Psi \in \mathcal{T}$ , including an *even* number of triplets, there exists a unique correspondence between loops in  $\mathcal{G}$  including the triplets in  $\Psi$  and perfect matchings in another extended graph  $\mathcal{G}_{ext_\Psi}$  constructed after removal of the triplets  $\Psi$  in  $\mathcal{G}$ . Using this representation the full loop series can be represented as a Pfaffian series, where each term  $Z_\Psi$  is tractable and is a product of the respective Pfaffian and the  $\mu_{a;\bar{a}}$  terms associated with each triplet of  $\Psi$ :<sup>2</sup>

$$z = \sum_{\Psi} Z_\Psi \qquad Z_\Psi = z_\Psi \prod_{a \in \Psi} \mu_{a;\bar{a}} \qquad (6)$$

$$z_\Psi = \text{sign} \left( \text{Pf} \left( \hat{B}_\Psi \right) \right) \cdot \text{Pf} \left( \hat{A}_\Psi \right).$$

The 2-regular partition function thus corresponds to  $\Psi = \emptyset$ . We refer to the remaining terms of the series as higher order Pfaffian terms. Notice that matrices  $\hat{A}_\Psi$  and  $\hat{B}_\Psi$  depend on the removed triplets and therefore each  $z_\Psi$  requires different matrices and different edge orientations. In addition, after removal of vertices in  $\mathcal{G}$  the resulting graph may be disconnected. As before, in these cases we add dummy edges to  $\mathcal{G}_{ext}$  with zero weight to make the graph bi-connected again.

---

2. We omit the loop index in the triplet term  $\mu_{a;\bar{a}}$  because nodes have at most degree three and therefore the set  $\bar{a}$  always coincide in all loops which contain that triplet.

Figure 1d shows loops corresponding to the higher order Pfaffian terms on our illustrative example. The first and second subsets of triplets (d.1 and d.2) include summation over two loops whereas the remaining Pfaffian terms include uniquely one loop.

Exhaustive enumeration of all the subsets of triplets leads to a  $2^{|\mathcal{T}|}$  time algorithm, which is prohibitive. However, many triplet combinations may lead to forbidden configurations. Experimentally, we found that a principled way to look for higher order Pfaffian terms with large contribution is to search first for pairs of triplets, then groups of four, and so on. For large graphs, this also becomes intractable. Actually, the problem is very similar to the problem of selecting loop terms  $r_C$  with largest contribution. The advantage of the Pfaffian representation, however, is that  $Z_\emptyset$  is always the Pfaffian term that accounts for the largest number of loop terms and is the most contributing term in the series. In this work we do not derive any heuristic for searching Pfaffian terms with larger contributions. Instead, in Section 3.1 we study the full Pfaffian series and subsequently we restrict ourselves on the accuracy of  $Z_\emptyset$ .

Algorithm 2 describes the full procedure to compute all terms using the representation of expression (6). The main loop of the algorithm can be interrupted at any time, thus leading to a sequence of algorithms producing corrections incrementally.

---

**Algorithm 2** Pfaffian series

---

**Arguments:** Forney graph  $\mathcal{G}$

- 1:  $z := 0$ .
  - 2: **for all** ( $\Psi \in \mathcal{T}$ ) **do**
  - 3:   Build extended graph  $\mathcal{G}_{ext_\Psi}$  applying rules of Figure 2.
  - 4:   Set Pfaffian orientation in  $\mathcal{G}_{ext_\Psi}$  according to Algorithm 1
  - 5:   Build matrices  $\hat{A}$  and  $\hat{B}$ .
  - 6:   Compute Pfaffian with sign correction  $z_\Psi$  according to Equation (3).
  - 7:    $z := z + z_\Psi \prod_{a \in \Psi} \mu_{a;\bar{a}}$ .
  - 8: **end for**
  - 9: **RETURN**  $Z^{BP} \cdot z$
- 

### 3. Experiments

In this Section we study numerically the proposed algorithm. To facilitate the evaluation and the comparison with other algorithms we focus on the binary Ising model, a particular case of the model (1) where factors only depend on the disagreement between two variables and take the form  $f_a(\sigma_{ab}, \sigma_{ac}) = \exp(J_{a;\{ab,ac\}}\sigma_{ab}\sigma_{ac})$ . We consider also nonzero local potentials parametrized by  $f_a(\sigma_{ab}) = \exp(J_{a;\{ab\}}\sigma_{ab})$  in all variables so that the model becomes planar-intractable.

We create different inference problems by choosing different interactions  $\{J_{a;\{ab,ac\}}\}$  and local field parameters  $\{J_{a;\{ab\}}\}$ . To generate them we draw independent samples from a Normal distribution  $\{J_{a;\{ab,ac\}}\} \sim \mathcal{N}(0, \beta/2)$  and  $\{J_{a;\{ab\}}\} \sim \mathcal{N}(0, \beta\Theta)$ , where  $\Theta$  and  $\beta$  determine how difficult the inference problem is. Generally, for  $\Theta = 0$  the planar problem is tractable. For  $\Theta > 0$ , small values of  $\beta$  result in weakly coupled variables (easy problems)

and large values of  $\beta$  in strongly coupled variables (hard problems). Larger values of  $\Theta$  result in easier inference problems.

In the next Subsection we analyze the full Pfaffian series using a small example and compare it with the original representation based on the loop series. Next, we compare our algorithm with the following ones: <sup>3</sup>

**Truncated Loop-Series for BP (TLSBP)** (Gómez et al., 2007), which accounts for a certain number of loops by performing depth-first-search on the factor graph and then merging the found loops iteratively. We adapted TSLBP as an any-time algorithm (**anyTLSBP**) such that the length of the loop is used as the only parameter instead of the two parameters  $S$  and  $M$  (see Gómez et al. (2007) for details). This is equivalent to setting  $M = 0$  and discard  $S$ . In this way, anyTLSBP does not compute all possible loops of a certain length (in particular, complex loops <sup>4</sup> are not included), but is more efficient than TLSBP.

**Cluster Variation Method (CVM-Loopk)** A double-loop implementation of CVM (Heskes et al., 2003). This algorithm is a special case of generalized belief propagation (Yedidia et al., 2005) with convergence guarantees. We use as outer clusters all (maximal) factors together with loops of four ( $k=4$ ) or six ( $k=6$ ) variables in the factor graph.

**Tree-Structured Expectation Propagation (TreeEP)** (Minka and Qi, 2004). This method performs exact inference on a base tree of the graphical model and approximates the other interactions. The method yields good results if the graphical model is very sparse.

When possible, we also compare with the following two variational methods which provide upper bounds on the partition function:

**Tree Reweighting (TRW)** (Wainwright et al., 2005) which decomposes the parametrization of a probabilistic graphical model as a mixture of spanning trees of the model, and then uses the convexity of the partition function to get an upper bound.

**Planar graph decomposition (PDC)** (Globerson and Jaakkola, 2007) which decomposes the parametrization of a probabilistic graphical model as a mixture of tractable planar graphs (with zero local field).

To evaluate the accuracy of the approximations we consider errors in  $Z$  and, when possible, computational cost as well. As shown in Gómez et al. (2007), errors in  $Z$ , obtained from a truncated form of the loop series, are very similar to errors in single variable marginal probabilities, which can be obtained by conditioning over the variables under interest. We only consider tractable instances for which  $Z$  can be computed via the junction tree algorithm (Lauritzen and Spiegelhalter, 1988) using 8GB of memory. When studying the

3. We use the libDAI library (Mooij, 2008) for algorithms **CVM-Loopk**, **TreeEP** and **TRW**.

4. A complex loop is defined as a loop which can not be expressed as the union of two or more circuits or simple loops.

scalability of the approaches, we Given an approximation  $Z'$  of  $Z$ , the error measure used in this manuscript is:

$$\text{error}Z' = \frac{|\log Z - \log Z'|}{\log Z}.$$

As in Gómez et al. (2007), we use four different message updates for BP: fixed and random sequential updates, parallel (or synchronous) updates, and residual belief propagation (RBP), a method proposed by Elidan et al. (2006) which selects the next message to be updated which has maximum *residual*, a quantity defined as an upper bound on the distance of the current messages from the fixed point. We report non-convergence when none of the previous methods converged. We report convergence at iteration  $t$  when the maximum absolute value of the updates of all the messages from iteration  $t - 1$  to  $t$  is smaller than a threshold  $\vartheta = 10^{-14}$ .

### 3.1 Full Pfaffian Series

In the previous Section we have described two equivalent representations for the exact partition function in terms of the loop series and the Pfaffian series. Here we analyze numerically how these two representations differ using an example, shown in Figure 3 as a factor graph, for which all terms of both series can be computed. We analyze a single instance, parametrized using  $\Theta = 0.1$  and different pairwise interactions  $\beta \in \{0.1, 0.5, 1.5\}$ .




Figure 3: Planar bipartite factor graph used to compare the full Pfaffian series with the loop series. Circles and black squares denote variables and factors respectively.

We use TLSBP to retrieve all loops, 8123 for this example, and Algorithm 2 to compute all Pfaffian terms. To compare the two approximations we sort all contributions, either loops or Pfaffians, by their absolute values in descending order, and then analyze how the errors are corrected as more terms are included in the approximation. We define partition




Figure 4: Comparison between the full loop series and the full Pfaffian series. Each row corresponds to a different value of the interaction strength  $\beta$ . **Left column** shows the error, considering loop terms  $Z^{TLSBP}(l)$  in log-log scale. Shaded regions include all loop terms (not necessarily 2-regular loops) required to reach the same (or better) accuracy than the accuracy of the 2-regular partition function  $Z_\emptyset$ . **Middle column** shows the error considering Pfaffian terms  $Z^{Pf}(p)$  also in log-log scale. The first Pfaffian term corresponds to  $Z_\emptyset$ , marked by a circle. **Right column** shows the values of the first 100 Pfaffian terms sorted in descending order in  $|Z_\Psi|$  and excluding  $z_\emptyset$ .

functions for the truncated series in the following way:

$$Z^{TLSBP}(l) = Z^{BP} \left( 1 + \sum_{i=1 \dots l} r_{C_i} \right), \quad Z^{Pf}(p) = Z^{BP} \left( \sum_{i=1 \dots p} Z_{\Psi_i} \right).$$

Then  $Z^{TLSBP}(l)$  accounts for the  $l$  most contributing loops and  $Z^{Pf}(p)$  accounts for the  $p$  most contributing Pfaffian terms. In all cases, the Pfaffian term with largest absolute value  $Z_{\Psi_1}$  corresponds to  $z_\emptyset$ .

Figure 4 shows the error  $Z^{TLSBP}$  (first column) and  $Z^{Pf}$  (second column) for both representations. For weak interactions ( $\beta = 0.1$ ) BP converges fast and provides an accurate approximation with an error of order  $10^{-4}$ . Summation of less than 50 loop terms (top-left panel) is enough to obtain machine precision accuracy. Notice that the error is almost

reduced totally with the  $z_\emptyset$  correction (top-middle panel). In this scenario, higher order terms of the Pfaffian series are negligible (top-right panel).

As we increase  $\beta$ , the quality of the BP approximation decreases. The number of loop corrections required to correct the BP error then increases. In this example, for intermediate interactions ( $\beta = 0.5$ ) the first Pfaffian term  $z_\emptyset$  improves considerably, more than five orders of magnitude, on the BP error, whereas approximately 100 loop terms are required to achieve a similar correction (gray region of middle-left panel).

For strong interactions ( $\beta = 1.5$ ) BP converges after many iterations and gives a poor approximation. In this scenario also a larger proportion of loop terms (bottom-left panel) is necessary to correct the BP result up to machine precision. Looking at the bottom-left panel we find that approximately 200 loop terms are required to achieve the same correction as the one obtained by  $z_\emptyset$ . The  $z_\emptyset$  is quite accurate (bottom-middle panel).

As the right panels show, higher order Pfaffian contributions change progressively from a flat sequence of small terms to an alternating sequence of positive and negative terms which grow in absolute value as  $\beta$  increases. These oscillations are also present in the loop series expansion.

In general, we conclude that the  $z_\emptyset$  correction to the BP approximation can give a significant improvement even in hard problems for which BP converges after many iterations. Notice again that calculating  $z_\emptyset$ , the most contributing term in the Pfaffian series, does not require explicit search of loop or Pfaffian terms.

### 3.2 Grids

After analyzing the full Pfaffian series on a small random example we now restrict our attention to the  $Z_\emptyset$  approximation using Ising grids (nearest neighbor connectivity). First, we compare that approximation with other inference methods for different types of interactions (attractive or mixed) and then study the scalability of the method in the size of the graphs.

#### 3.2.1 ATTRACTIVE INTERACTIONS

We first focus on binary models with interactions that tend to align the neighboring variables to the same value,  $J_{a;\{ab,ac\}} > 0$ . If local fields are also positive  $J_{a;\{ab\}} > 0, \forall a \in \mathcal{V}$ , Sudderth et al. (2008) showed that, under some additional condition, the BP approximation is a *lower-bound* of the exact partition function and all loops (and therefore Pfaffian terms too) have the same sign<sup>5</sup>. Although this is not formally proved for general models with attractive interactions regardless of the sign of the local fields, numerical results suggest that this property holds as well for this type of models.

We generate grids with positive interactions and local fields, that is  $|\{J_{a;bc}\}| \sim \mathcal{N}(0, \beta/2)$  and  $|\{J_{a;b}\}| \sim \mathcal{N}(0, \beta\Theta)$ , and study the performance for various values of  $\beta$ , as well as for strong  $\Theta = 1$  and weak  $\Theta = 0.1$  local fields.

Figure 5 shows the average error over 50 instances reported by different methods. Using this setup, BP converged in all instances using sequential updates of the messages. The error curves of all methods show an initial growth and a subsequent decrease, a fact explained by the phase transition occurring in this model for  $\Theta = 0$  and  $\beta \approx 1$  (Mooij and Kappen,

---

5. The condition is that all single variable beliefs at the BP fixed point must satisfy  $m_{ab} = b_{ab}(+1) - b_{ab}(-1) > 0, \forall(a, b) \in \mathcal{E}$

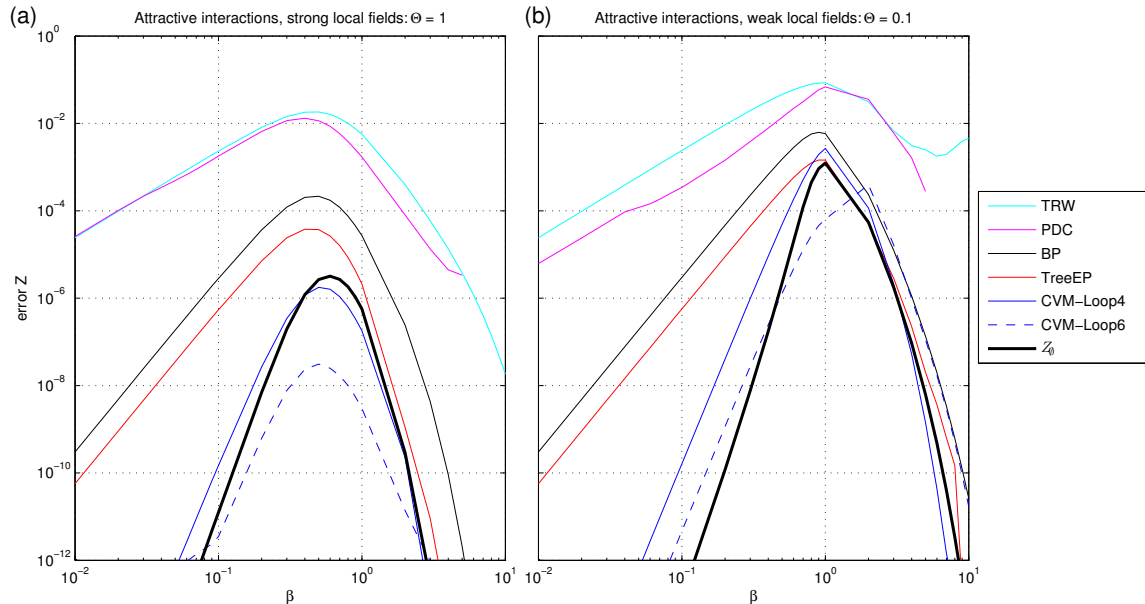


Figure 5: 7x7 grid attractive interactions and positive local fields. Error averages over 50 random instances in function of the difficulty of the problem. **(a)** Strong local fields. **(b)** Weak local fields.

2005). As the difference between the two plots suggest, errors are larger as  $\Theta$  approaches zero. Notice, that  $Z_\Theta = Z$  for the limit case of  $\Theta = 0$ .

We observe that in all instances  $Z_\Theta$  *always improves* over the BP approximation. Corrections are most significant for weak interactions  $\beta < 1$  and strong local fields. For strong interactions  $\beta > 1$  and weak local fields the improvement is less significant.

It appears that the  $Z_\Theta$  approximation performs better than TreeEP in all cases except for very strong couplings, where they show very similar results. Interestingly,  $Z_\Theta$  performs very similar to CVM-Loop4 which is known to be a very accurate approximation for this type of model, see Yedidia et al. (2000) for instance. We observe that in order to obtain better average results than  $Z_\Theta$  using CVM, we need to select larger outer-clusters such as loops of length 6, which increases dramatically the computational cost.

The methods which provide upper bounds on  $Z$  (PDC and TRW) report the largest average error. PDC performs slightly better than TRW, as was shown in Globerson and Jaakkola (2007) for the case of mixed interactions. We remark that the worse performance of PDC for stronger couplings and weak local fields might be attributed to implementation artifacts, since for  $\beta > 4$  we have numerical precision errors. In general, both upper bounds are significantly less tight than the lower bounds provided by BP and  $Z_\Theta$ .




Figure 6: 7x7 grid mixed interactions. Error averaged over 50 random instances as a function of the problem difficulty for (a) strong local fields and (b) weak local fields. Bottom panels show percentage of cases when BP converges for (c) strong local fields and (d) weak local fields.

### 3.2.2 MIXED INTERACTIONS

We now analyze a more general Ising grid model where interactions and local fields can have mixed signs. In that case,  $Z^{BP}$  and  $Z_\theta$  are no longer lower bounds on  $Z$  and loop terms can be positive or negative. Figure 6 shows results using this setup. Top panels show average errors and bottom panels show percent of instances in which BP converged using at least one of the methods described above.

For strong local fields (subplots a,c), we observe that  $Z_\theta$  improvements over BP results become less significant as  $\beta$  increases. It is important to note that  $Z_\theta$  always improves on the BP result, even when the couplings are very strong ( $\beta = 10$ ) and BP fails to converge for a small percentage of instances.  $Z_\theta$  performs slightly better than CVM-Loop4 and substantially better than TreeEP for small and intermediate  $\beta$ . All three methods show similar results for strong couplings  $\beta > 2$ . As in the case of attractive interactions, the best results are attained using CVM-loop6.

For the case of weak local fields (subplots b,d), BP fails to converge near the transition to the spin-glass phase. For  $\beta = 10$ , BP converges only in less than 25% of the instances. In the most difficult domain,  $\beta > 22$ , all methods under consideration give similar results (all comparable to BP). Moreover, it may happen that  $Z_\theta$  degrades the  $Z^{BP}$  approximation because loops of alternating signs have major influence in the series. This result was

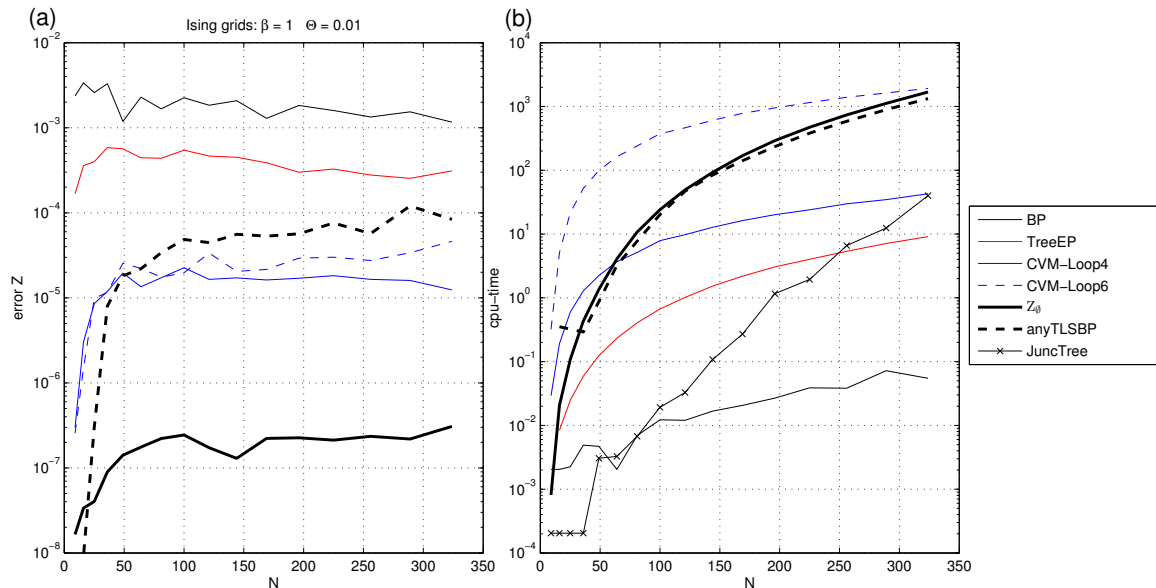


Figure 7: Results on regular grids: scaling with grid size for strong interactions  $\beta = 1$  and very weak local fields  $\Theta = 0.01$ . BP converged in all cases. **(a)** Error medians over 50 instances. **(b)** Cpu time (log-scale).

also reported in Gómez et al. (2007) when loop terms, instead of Pfaffian terms, were considered.

### 3.2.3 SCALING WITH GRAPH SIZE

We now study how the accuracy of the  $Z_\Theta$  approximation changes as we increase the size of the grid. We generate random grids with mixed couplings for  $\sqrt{N} = \{4, \dots, 18\}$  and focus on a regime of very weak local fields  $\Theta = 0.01$  and strong couplings  $\beta = 1$ , a difficult configuration according to the previous results. We compare  $Z_\Theta$  also with anyTLSBP, a variant of our previous algorithm for truncating the loop series. We run anyTLSBP by selecting loops shorter than a given length, and the length is increased progressively. To provide a fair comparison between both methods, we run anyTLSBP for the same amount of cpu time as the one required to obtain  $Z_\Theta$ .

Figure 7a shows the errors of different methods. Since variability in the errors is larger than before, we take the median for comparison. In order of increasing accuracy we get BP, TreeEP, anyTLSBP, CVM-Loop6, CVM-Loop4 and  $Z_\Theta$ . We note that larger clusters in CVM does not necessarily result in better performance.

Overall, we can see that results are roughly independent of the network size  $N$  in almost all methods that we compare. The error of anyTLSBP starts being the smallest but soon increases because the proportion of loops captured decreases very fast. For  $N > 64$ , anyTLSBP performs worse than CVM. The  $Z_\Theta$  correction, on the other hand, stays flat and we can conclude that it scales reasonably well. Interestingly, although  $Z_\Theta$  and TLSBP use




Figure 8: Two examples of planar graphs used for comparison between methods. We fix the number of concentric polygons to 9 and change the degree  $d$  of the central node within the range  $[3, \dots, 25]$ . **(left)** Graph for  $d = 3$ . **(right)** Graph for  $d = 25$ . Here nodes represent variables and edges pairwise interactions. We also add external fields which depend on the state of each nodes (not drawn).

different ways to truncate the loop series, both methods show similar scaling behavior for large graphs.

Figure 7b shows the cpu time for all the tested approaches averaged over all cases. Concerning the approximate inference methods, in order of increasing cost, we have BP, TreeEP, CVM-Loop4,  $Z_\theta$  with anyTLSBP, and CVM-Loop6. Although the cpu time required to compute  $Z_\theta$  scales with  $O(N_{Ext}^3)$ , its curve shows the steepest growth. We discuss how to correct this caveat in Section 4. The cpu time of the junction tree method quickly increases with the tree-width of the graphs. For large enough  $N$ , exact solution via the junction tree method is no longer feasible because of its memory requirements. In contrast, for all approximate inference methods, memory demands do not represent a limitation.

### 3.3 Radial grid graphs

In the previous subsection we analyzed the quality of the  $Z_\theta$  correction for graphs with a regular grid structure. Here, we carry over the analysis of the  $Z_\theta$  correction using planar graphs which consist of concentric polygons with a variable number of sides. Figure 8 illustrates these spider-web like graphs. We generate them as factor graphs with pairwise interactions which we subsequently convert to an equivalent Forney graph. (See Appendix A for details). Again, local field potentials are parametrized using  $\Theta = 0.01$  and interactions using  $\beta = 1$ . We analyze the error in  $Z$  as a function of the degree  $d$  of the central node.

Figure 9a shows the median of errors in  $Z$  of 50 random instances. First, we see that the variability of all methods, in particular anyTLSBP, is larger than in the regular grid

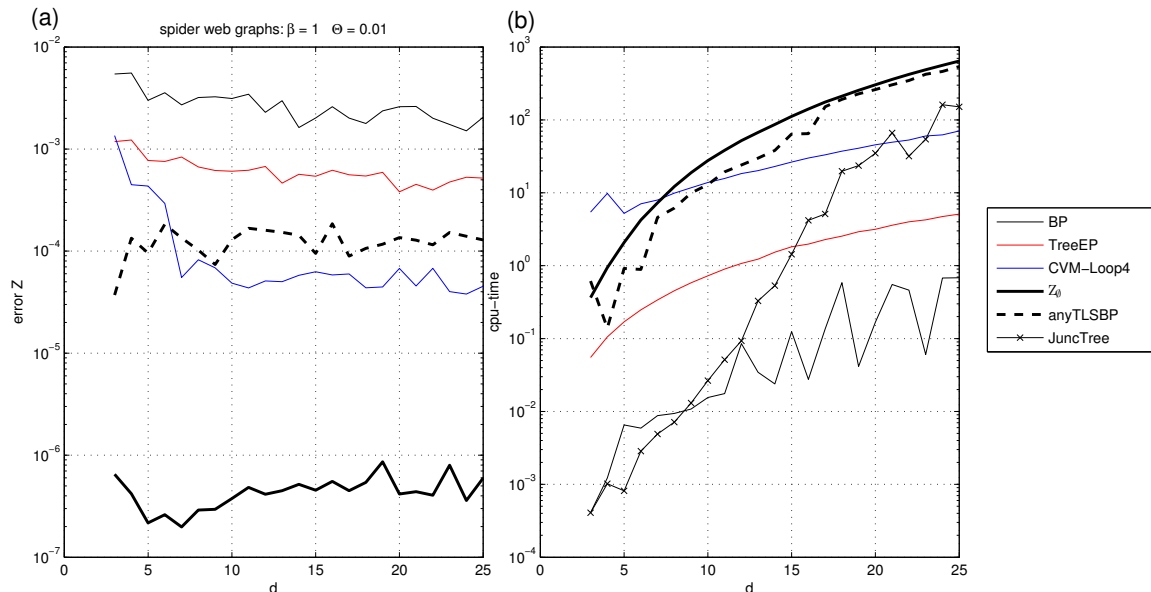


Figure 9: Results on spider-web like graphs: scaling with the degree of the central node for  $\beta = 1$  and  $\Theta = 0.01$ . BP converged in all cases. **(a)** Error medians over 50 instances. **(b)** Cpu time (log-scale).

scenario. Also, the improvement of CVM-Loop4 over BP is slightly less significant, possibly caused by the existence of the central node with a large degree. CVM-Loop6 does not converge for instances with  $d > 4$  before  $10^4$  seconds and is not included in the analysis. We can say that all approaches present results independent of the degree  $d$ .

The  $Z_0$  approximation is the best method compared to the other tested approaches. The improvements of  $Z_0$  on CVM-Loop4 (the second best method) can be of more than two orders of magnitude and more than three orders of magnitude compared to BP.

Computational costs are shown in 9b. The best performance of  $Z_0$  comes at the cost of being the most expensive approximate inference approach for which we obtain results. Again, for larger graphs, exact solution via the junction tree is not feasible due to the large tree-width.

#### 4. Discussion

We have presented an approximate algorithm based on the exact loop calculus framework for inference on planar graphical models defined in terms of binary variables. The proposed approach improves the estimate for the partition function provided by BP without an explicit search of loops.

The algorithm is illustrated on the example of ordered and disordered Ising model on a planar graph. Performance of the method is analyzed in terms of its dependence on the system size. The complexity of the partition function computation is exponential in the

general case, unless the local fields are zero, when it becomes polynomial. We tested our algorithm on regular grids and planar graphs with different structures. Our experiments on regular grids show that significant improvements over BP are always obtained if single variable potentials (local magnetic fields) are sufficiently large. The quality of this correction degrades with decrease in the amplitude of external field, to become exact at zero external fields. This suggests that the difficulty of the inference task changes abruptly from very easy, with no local fields, to very hard, with small local fields, and then decays again as external fields become larger.

The  $Z_\theta$  correction turns out to be competitive with other state of the art methods for approximate inference of the partition function. First of all, we showed that  $Z_\theta$  is much more accurate than upper bounds based methods such as TRW or PDC. This illustrates that such methods come at the cost of less accurate approximations. We have also shown that for the case of grids with attractive interactions, the lower bound provided by  $Z_\theta$  is the most accurate.

Secondly, we found that  $Z_\theta$  performs much better than treeEP for weak and intermediate couplings and shows competitive results for strong interactions. Concerning CVM, we showed that using larger outer clusters does not necessarily lead to better approximations. In general, the  $Z_\theta$  correction presented better results than CVM for our choice of regions.

Finally, we have presented a comparison of  $Z_\theta$  with TLSBP, which is another algorithm for the BP-based loop series using the loop length as truncation parameter. On the one hand, the calculation of  $Z_\theta$  involves a *re-summation* of many loop terms which in the case of TLSBP are summed individually. This consideration favors the  $Z_\theta$  approach. On the other hand,  $Z_\theta$  is restricted to the class of 2-regular loops whereas TLSBP also accounts for terms corresponding to more complex loop structures in which nodes can have degree larger than two. Overall, for planar graphs, we have shown evidence that the  $Z_\theta$  approach is better than TLSBP when the size of the graphs is not very small. We emphasize, however, that TLSBP can be applied to non-planar binary graphical models too.

Currently, the shortcoming of the presented approach is in its relatively costly implementation. However, since the bottleneck of the algorithm is the Pfaffian calculation and not the algorithm itself (used to obtain the extended graphs and the associated matrices), it is easy to devise more efficient methods than the one used here. Thus, one may substitute brute-force evaluation of the Pfaffians by a smarter one available for planar graphs. This reduces the cost from  $\mathcal{O}(N^3)$  to  $\mathcal{O}(N^{3/2})$  (Galluccio et al., 2000; Loh and Carlson, 2006). Besides, the Pfaffian of  $\hat{B}$  is binary, see Eq. (6), making it possible to improve using a bit-matrix representation (Schraudolph and Kamenetsky, 2008). Alternatively one could think of a strategy which does not require the Pfaffian of  $\hat{B}$ . All these technical issues are the focus of our continuing investigation.

In this manuscript we have focused on inference problems defined on planar graphs with symmetric pairwise interactions and, to make the problems difficult, we have introduced local field potentials. Notice however, that the algorithm can also be used to solve models with more complex interactions, i.e. more than pairwise as in the case of the Ising model (see Chertkov et al., 2008, for a discussion of possible generalizations). This makes our approach more powerful than other approaches, namely, (Globerson and Jaakkola, 2007; Schraudolph and Kamenetsky, 2008), designed specifically for the pairwise interaction case.

Although planarity is a severe restriction, we emphasize that planar graphs appear in many contexts such as computer vision and image processing, magnetic and optical recording, or network routing and logistics. It would also be interesting (and possible) to consider extensions of the algorithm developed in the manuscript for approximate inference of some class of non-planar graphs. Thus, following the approach of Globerson and Jaakkola (2007), one can think of other types of spanning subgraphs more general than "easy" planar graphs for which exact computation can be performed using perfect matching. The correction  $Z_\theta$  can be an accurate approximation for this spanning subgraphs and the resulting approximation method would also provide bounds on the exact result.

## Acknowledgments

We acknowledge J. M. Mooij for providing the libDAI framework and A. Windsor for the planar graph functions of the boost graph library. We also thank V. Y. Chernyak, J. K. Johnson and N. Schraudolph for interesting discussions and A. Globerson for providing the Matlab sources of PDC. This research is part of the Interactive Collaborative Information Systems (ICIS) project, supported by the Dutch Ministry of Economic Affairs, grant BSIK03024. The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396.

## Appendix A: Converting a factor graph to a Forney Graph.

A probabilistic model is usually represented as a Bayesian Network or a Markov Random Field. Since bipartite factor graphs subsume both models, we show here how to convert a factor graph model defined in terms of binary variables to a more general Forney graph representation, for which the presented algorithm can be directly applied to.

On a bipartite factor graph  $\mathcal{G}_{\mathcal{F}} = (\mathcal{V}_{\mathcal{F}}, \mathcal{E}_{\mathcal{F}})$  the set  $\mathcal{V}_{\mathcal{F}}$  is composed of a set of variable nodes  $\mathcal{I}$  and a set of factor nodes  $\mathcal{J}$ . Each variable node  $i \in \mathcal{I}, i := \{1, 2, \dots\}$  represents a variable which takes values  $\sigma_i = \{\pm 1\}$ . We label factor nodes using capital letters so that  $a = \{A, B, \dots\}, a \in \mathcal{J}$  denotes a factor node which has an associated function  $f_a(\sigma_a)$  defined on a subset of variables  $\bar{a} \in \mathcal{I}$ . An (undirected) edge exists between two nodes  $(a, i) \in \mathcal{E}_{\mathcal{F}}$  if  $i \in \bar{a}$ .

Given  $\mathcal{G}_{\mathcal{F}}$ , a direct way to obtain an equivalent Forney graph  $\mathcal{G}$  is: first, to create a node  $\delta_i \in \mathcal{V}$  for each variable node  $i \in \mathcal{V}_{\mathcal{F}}$ , and second, to associate a new binary variable  $\delta_i a$  with values  $\sigma_{\delta_i a} = \{\pm 1\}$  to edges  $(\delta_i, a) \in \mathcal{E}$ . Nodes  $\delta_i \in \mathcal{V}$  are *equivalent factor nodes* denoting the characteristic function:  $\delta_i(\sigma_a) = 1$  if  $\sigma_{\delta_i a} = \sigma_{\delta_i b}, \forall a, b \in \bar{\delta}_i$  and zero otherwise. Finally, factor nodes  $c \in \mathcal{V}_{\mathcal{F}}$  correspond to the same factor nodes  $c$  in  $\mathcal{V}$  but defined in terms of the new variables  $\delta_i c, \forall i \in \bar{c}$ .

Figure 10 shows an example of this transformation. Notice that, although we impose an direction in the edge labels, they remain undirected:  $(\delta_i, a) = (a, \delta_i), \forall \delta_i, a \in \mathcal{V}$ . For variables  $i \in \mathcal{V}_{\mathcal{F}}$  which only appear in two factors, such as variable 3, the corresponding  $\delta_3$  node is redundant and can be removed. The joint distribution of  $\mathcal{G}_{\mathcal{F}}$  is related to the joint

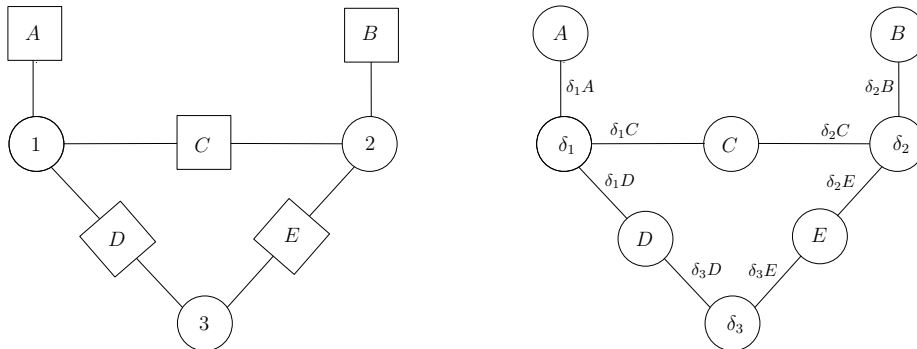


Figure 10: **(a)** An factor graph  $\mathcal{G}_{\mathcal{F}}$  and **(b)** an equivalent Forney graph  $\mathcal{G}$ .

distribution of  $\mathcal{G}$  by:

$$\begin{aligned} & \frac{1}{Z} f_A(\sigma_1) f_B(\sigma_2) f_C(\sigma_1, \sigma_2) f_D(\sigma_1, \sigma_3) f_E(\sigma_2, \sigma_3) \\ & \equiv \frac{1}{Z} f_A(\sigma_{\delta_1 A}) f_B(\sigma_{\delta_2 B}) f_C(\sigma_{\delta_1 C}, \sigma_{\delta_2 C}) f_D(\sigma_{\delta_1 D}, \sigma_{\delta_3 D}) f_E(\sigma_{\delta_2 E}, \sigma_{\delta_3 E}) \\ & \quad f_{\delta_1}(\sigma_{\delta_1 A}, \sigma_{\delta_1 C}, \sigma_{\delta_1 D}) f_{\delta_2}(\sigma_{\delta_2 B}, \sigma_{\delta_2 C}, \sigma_{\delta_2 E}) f_{\delta_3}(\sigma_{\delta_3 D}, \sigma_{\delta_3 E}). \end{aligned} \quad (7)$$

Once  $\mathcal{G}$  has been generated following the previous procedure it may be the case that the nodes  $\delta_i \in \mathcal{V}$  have degree three or larger. This happens if a variable  $i$  appears in more than 3 factor nodes on  $\mathcal{G}_{\mathcal{F}}$ . It is easy to convert  $\mathcal{G}$  to a graph were all  $\delta_i$  nodes have maximum degree three by introducing new auxiliary variables  $\delta_{i_1}, \delta_{i_2}, \dots$  and new equivalent nodes. For instance, if variable  $i \in \mathcal{V}_{\mathcal{F}}$  appears in 4 factors  $A, B, C, D$ :

$$f_{\delta_i}(\sigma_{\delta_i A}, \sigma_{\delta_i B}, \sigma_{\delta_i C}, \sigma_{\delta_i D}) \equiv f_{\delta_{i_1}}(\sigma_{\delta_{i_1} A}, \sigma_{\delta_{i_1} B}, \sigma_{\delta_{i_1}}) f_{\delta_{i_2}}(\sigma_{\delta_{i_2} C}, \sigma_{\delta_{i_2} D}).$$

Notice that although the models are equivalent, the number of loops in  $\mathcal{G}$  may be larger than in  $\mathcal{G}_{\mathcal{F}}$ . In the case that a factor in  $\mathcal{G}_{\mathcal{F}}$  involves more than three variables, as sketched in Chertkov et al. (2008), one could split the node of degree  $N$  into auxiliary nodes of degree  $N-1$  and compute  $Z_{\emptyset}$  on the transformed model. Alternatively, one can reduce the number of variables that enter a factor by clamping.

## References

- F. Barahona. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241–3253, 1982. URL <http://stacks.iop.org/0305-4470/15/3241>.
- M. Chertkov and V. Y. Chernyak. Loop series for discrete statistical models on graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(06):P06009, 2006a.
- M. Chertkov and V. Y. Chernyak. Loop calculus helps to improve Belief Propagation and linear programming decodings of LDPC codes. In *invited talk at 44th Allerton Conference*, September 2006b.

- M. Chertkov, V. Y. Chernyak, and R. Teodorescu. Belief propagation and loop series on planar graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(05): P05003 (19pp), 2008. URL <http://stacks.iop.org/1742-5468/2008/P05003>.
- G. Elidan, I. McGraw, and D. Koller. Residual belief propagation: Informed scheduling for asynchronous message passing. In *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence (UAI-06)*, Boston, Massachusetts, July 2006. AUAI Press.
- M. Fisher. On the dimer solution of the planar Ising model. *Journal of Mathematical Physics*, 7(10):1776–1781, 1966.
- Jr. Forney, G.D. Codes on graphs: normal realizations. *IEEE Transactions on Information Theory*, 47(2):520–548, Feb 2001. ISSN 0018-9448. doi: 10.1109/18.910573.
- B. J. Frey and D. J. C. MacKay. A revolution: belief propagation in graphs with cycles. In *Advances in Neural Information Processing Systems 10*, pages 479–486, Cambridge, MA, 1998. MIT Press.
- G. Galbiati and F. Maffioli. On the computation of Pfaffians. *Discrete Applied Mathematics*, 51(3):269–275, 1994. ISSN 0166-218X. doi: [http://dx.doi.org/10.1016/0166-218X\(92\)00034-J](http://dx.doi.org/10.1016/0166-218X(92)00034-J).
- A. Galluccio, M. Loeb, and J. Vondrák. New algorithm for the Ising problem: partition function for finite lattice graphs. *Physical Review Letters*, 84(26):5924–5927, Jun 2000. doi: 10.1103/PhysRevLett.84.5924.
- A. Globerson and T. S. Jaakkola. Approximate inference using planar graph decomposition. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 473–480. MIT Press, Cambridge, MA, 2007.
- V. Gómez, J. M. Mooij, and H. J. Kappen. Truncating the loop series expansion for belief propagation. *Journal of Machine Learning Research*, 8:1987–2016, 2007. ISSN 1533-7928.
- T. Heskes, K. Albers, and H. J. Kappen. Approximate inference and constrained optimization. In *Proceedings of the 19th Annual conference on Uncertainty in Artificial Intelligence (UAI-03)*, pages 313–320, San Francisco, CA, 2003. Morgan Kaufmann Publishers.
- M. Karpinski and W. Rytter. Fast parallel algorithms for graph matching problems. pages 164–170. Oxford University Press, USA, 1998.
- P. W. Kasteleyn. Dimer statistics and phase transitions. *Journal of Mathematical Physics*, 4(2):287–293, 1963. doi: 10.1063/1.1703953. URL <http://link.aip.org/link/?JMP/4/287/1>.
- S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society. Series B-Methodological*, 50(2):154–227, 1988.

- H.-A. Loeliger. An introduction to factor graphs. *Signal Processing Magazine, IEEE*, 21(1):28–41, Jan. 2004. ISSN 1053-5888. doi: 10.1109/MSP.2004.1267047.
- Y. L. Loh and E. W. Carlson. Efficient algorithm for random-bond Ising models in 2d. *Physical Review Letters*, 97(22):227205, 2006. doi: 10.1103/PhysRevLett.97.227205. URL <http://link.aps.org/abstract/PRL/v97/e227205>.
- T. Minka and Y. Qi. Tree-structured approximations by expectation propagation. In Sebastian Thrun, Lawrence Saul, and Bernhard Schölkopf, editors, *Advances in Neural Information Processing Systems 16*. MIT Press, Cambridge, MA, 2004.
- J. M. Mooij. libDAI: A free/open source C++ library for discrete approximate inference methods, 2008. <http://mloss.org/software/view/77/>.
- J. M. Mooij and H. J. Kappen. On the properties of the Bethe approximation and loopy belief propagation on binary networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2005(11):P11012, 2005.
- K. P. Murphy, Y. Weiss, and M. I. Jordan. Loopy Belief Propagation for approximate inference: An empirical study. In *Proceedings of the 15th Annual Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pages 467–475, San Francisco, CA, 1999. Morgan Kaufmann Publishers.
- J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, San Francisco, CA, 1988. ISBN 1558604790.
- N. Schraudolph and D. Kamenetsky. Efficient exact inference in planar Ising models. In *Advances in Neural Information Processing Systems 22*. MIT Press, Cambridge, MA, 2008.
- E. Sudderth, M. Wainwright, and A. Willsky. Loop series and Bethe variational bounds in attractive graphical models. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1425–1432. MIT Press, Cambridge, MA, 2008.
- M. Wainwright, T. Jaakkola, and A. Willsky. A new class of upper bounds on the log partition function. *IEEE Transactions on Information Theory*, 51(7):2313–2335, July 2005.
- J. S. Yedidia, W. T. Freeman, and Y. Weiss. Generalized belief propagation. In T.K. Leen, T.G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 689–695, December 2000.
- J. S. Yedidia, W. T. Freeman, Y. Weiss, and A. L. Yuille. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Transactions on Information Theory*, 51(7):2282–2312, July 2005.