Belief propagation and loop calculus for the permanent of a non-negative matrix

2010 J. Phys. A: Math. Theor. 43 242002

(http://iopscience.iop.org/1751-8121/43/24/242002)

IOP FTC
IOP Publishing — Journal of Physics A: Mathematical and Theoretical
J. Phys. A: Math. Theor. **43** (2010) 242002 (11pp) doi:10.1088/1751-8113/43/24/242002

FAST TRACK COMMUNICATION

Belief propagation and loop calculus for the permanent of a non-negative matrix

Yusuke Watanabe[1] **and Michael Chertkov**[2,3]

[1] Institute of Statistical Mathematics, 10-3 Midori-cho, Tachikawa, Tokyo 190-8562 Japan
[2] Center for Nonlinear Studies and Theoretical Division, LANL, NM 87545, USA
[3] New Mexico Consortium, Los Alamos, NM 87544, USA

E-mail: watay@ism.ac.jp and chertkov@lanl.gov

Received 18 February 2010, in final form 2 May 2010
Published 25 May 2010
Online at stacks.iop.org/JPhysA/43/242002

Abstract
We consider computation of the permanent of a positive $(N \times N)$ non-negative matrix, $P = \left(P_i^j \middle| i, j = 1, \ldots, N\right)$, or equivalently the problem of weighted counting of the perfect matchings over the complete bipartite graph $K_{N,N}$. The problem is known to be of likely exponential complexity. Stated as the partition function Z of a graphical model, the problem allows for exact loop calculus representation (Chertkov M and Chernyak V 2006 *Phys. Rev.* E **72** 065102) in terms of an interior minimum of the Bethe free energy functional over non-integer doubly stochastic matrix of marginal beliefs, $\beta = \left(\beta_i^j \middle| i, j = 1, \ldots, N\right)$, also correspondent to a fixed point of the iterative message-passing algorithm of the belief propagation (BP) type. Our main result is an explicit expression of the exact partition function (permanent) in terms of the matrix of BP marginals, β, as $Z = \mathrm{Perm}(P) = Z_{\mathrm{BP}}\,\mathrm{Perm}\left(\beta_i^j\left(1-\beta_i^j\right)\right) / \prod_{i,j}\left(1-\beta_i^j\right)$, where Z_{BP} is the BP expression for the permanent stated explicitly in terms of β. We give two derivations of the formula, a direct one based on the Bethe free energy and an alternative one combining the Ihara graph-ζ function and the loop calculus approaches. Assuming that the matrix β of the BP marginals is calculated, we provide two lower bounds and one upper bound to estimate the multiplicative term. Two complementary lower bounds are based on the Gurvits–van der Waerden theorem and on a relation between the modified permanent and determinant, respectively.

PACS numbers: 02.60.−x, 47.11−j, 89.20.−a

1. Introduction

The problem of calculating the permanent of a non-negative matrix arises in many contexts in statistics, data analysis and physics. For example, it is intrinsic to the parameter learning of a flow used to follow particles in turbulence and to cross-correlate two subsequent images

[1]. However, the problem is $\#P$-hard [2], meaning that solving it in a time polynomial in the system size, *N*, is unlikely. Therefore, when the size of the matrix is sufficiently large, one naturally looks for ways to approximate the permanent. A very significant breakthrough was achieved with the invention of a so-called fully polynomial randomized algorithmic schemes (FPRAS) for the permanent problem [3]: the permanent is approximated in a polynomial time, with high probability and within an arbitrarily small relative error. However, the complexity of this FPRAS is $O(N^{11})$, making it impractical for the majority of realistic applications. This motivates the task of finding a lighter deterministic or probabilistic algorithm capable of evaluating the permanent more efficiently.

This communication continues the thread of [1, 4] and [5], where the belief propagation (BP) algorithm was suggested as an efficient heuristic of good (but not absolute) quality to approximate the permanent. The BP family of algorithms, originally introduced in the context of error-correction codes [6] and artificial intelligence [7], can generally be stated for any graphical model [8]. The exactness of the BP on any graph without loops suggests that the algorithm can be an efficient heuristic for evaluating the partition function or for finding a maximum likelihood (ML) solution for the graphical model (GM) defined on sparse graphs. However, in the general loopy cases, one would normally not expect BP to work well, thus making the heuristic results of [1, 4, 5] somehow surprising, even though not completely unexpected in view of the existence of polynomially efficient algorithms for the ML version of the problem [9, 10], also realized in [11] via an iterative BP algorithm. This raises the questions of understanding the performance of BP: what does it do well and what does it miss? It also motivates the challenge of improving the BP heuristics.

An approach potentially capable of handling the question and the challenge was recently suggested in the general framework of GM. The loop series/calculus (LS) of [12, 13] expresses the ratio between the partition function (PF) of a binary GM and its BP estimate in terms of a finite series, in which each term is associated with the so-called generalized loop (a subgraph with all vertices of degree larger than 1) of the graph. Each term in the series, as well as the BP estimate of the partition function, is expressed in terms of a doubly stochastic matrix of marginal probabilities, $\beta = \left(\beta_i^j \middle| i, j = 1, \ldots, N\right)$, for matching pairs to contribute a perfect matching. This matrix β describes a minimum of the so-called Bethe free energy, and it can also be understood as a fixed point of an iterative BP algorithm. The first term in the resulting LS is equal to 1. Accounting for all the loop corrections, one recovers the exact expression for the PF. In other words, the LS holds the key to understanding the gap between the approximate BP estimate for the PF and the exact result. In sections 2 and 4, we will give a technical introduction to the variational Bethe free energy (BFE) formulation of BP and a brief overview of the LS approach for the permanent problem, respectively.

Our results. In this communication, we develop an LS-based approach to describe the quality of the BP approximation for the permanent of a non-negative matrix. (i) Our natural starting point is the analysis of the BP solution itself conducted in section 3. Evaluating the permanent of the non-negative matrix, $P = \left(\left(p_i^j\right)^{1/T} \middle| i, j = 1, \ldots, N\right)$, dependent on the temperature parameter, $T \in [0, \infty]$, we find that a non-integer BP solution is observed only at $T > T_c$, where T_c is defined by (15). (ii) At $T > T_c$, we derive an alternative representation for the LS in section 5. The entire LS is collapsed to a product of two terms: the first term is an easy-to-calculate function of β, and the second term is the permanent of the matrix $\beta .* (1-\beta) = \left(\beta_i^j\left(1-\beta_i^j\right)\right)$. (The binary operator $.*$ denotes the element-wise multiplication of matrices.) This is our main result stated in theorem 3, and the majority of the consecutive statements of our communication follows from it. We also present yet another, alternative, derivation of theorem 3 using the multivariate Ihara–Bass formula for the graph zeta-function in subsection 5.2. (iii) section 6 presents two easy-to-calculate lower bounds for the LS. The lower bound stated in corollary 7 is based on the Gurvits–van der Waerden theorem applied to $\mathrm{Perm}(\beta.*(1-\beta))$. Interestingly enough, this lower bound is invariant with respect to the BP transformation, i.e. it is exactly equivalent to the lower bound derived via application of the van der Waerden–Gurvits theorem to the original permanent. Another lower bound is stated in theorem 8. Note that as follows from an example discussed in the text, the two lower bounds are complementary: the latter is stronger at sufficiently small temperatures, while the former dominates the large T region. (iv) Section 7 discusses an upper bound on the transformed permanent based on the application of the Godzil–Gutman formula and the Hadamard inequality. Possible future extensions of the approach are discussed in section 8.

2. Background (I): graphical models, Bethe free energy and belief propagation

The permanent of a non-negative matrix, $P = \left(\left(p_i^j\right)^{1/T} | i, j = 1, \ldots, N\right)\left(0 \leqslant p_i^j, 0 \leqslant T \leqslant \infty\right)$, is a sum over the set of permutations on $\{1, \ldots, N\}$, which can be parameterized via binary-component vectors, σ, corresponding to perfect matchings (PM) on the complete bipartite graph $K_{N,N}$:

$$\left\{\sigma = \left(\sigma_i^j\right) \in \{0,1\}^{N\times N} | \forall i : \sum_{j=1}^N \sigma_i^j = 1,\ \forall j : \sum_{i=1}^N \sigma_i^j = 1\right\}. \tag{1}$$

This binary interpretation allows us to represent the permanent as the partition function (PF), Z, of a probabilistic model over the set of perfect matchings. Each perfect matching, σ, is realized with the probability

$$\mathcal{P}(\sigma) = \frac{1}{Z} P^\sigma; \qquad P^\sigma \equiv \prod_{(i,j)\in E} \left(p_i^j\right)^{\sigma_i^j/T}, \qquad Z \equiv \sum_{\sigma:PM} \left(p_i^j\right)^{\sigma_i^j/T} = \mathrm{Perm}(P), \tag{2}$$

where $E = \{(i,j) | i, j = 1, \ldots, N\}$ are the edges of $K_{N,N}$. In the zero-temperature limit, $T \to 0$, (2) selects one special ML solution, $\sigma_* = \arg\max_\sigma P^\sigma$. (Here and below we assume that P is non-degenerate, in the sense that at $T \to 0$, $\mathcal{P}(\sigma) \to 0$ for $\forall\, \sigma \neq \sigma_*$.)

For a generic GM, assigning (un-normalized) weight P^σ to a state σ, one defines the exact variational (called Gibbs, in statistical physics, and Kullback–Leibler in statistics) functional

$$\mathcal{F}\{b(\sigma)\} \equiv T \sum_\sigma b(\sigma) \ln \frac{b(\sigma)}{P^\sigma}. \tag{3}$$

One finds that under the condition that the belief, $b(\sigma)$, understood as a proxy to the probability $\mathcal{P}(\sigma)$, is normalized to unity, $\sum_{\sigma\in PM} b(\sigma) = 1$, the Gibbs functional is convex and it achieves its only minimum at $b(\sigma) = \mathcal{P}(\sigma)$ and $\mathcal{F}\{\mathcal{P}\} = -T \ln Z$.

The BP method offers an approximation which is exact when the underlying GM is a tree. As shown in [8], the BP approach can also be stated for a general GM as a relaxation of the Gibbs functional (3). In this paragraph we briefly review the concept of [8] with application to the permanent problem. For the GM (2), the BP approximation for the state beliefs becomes

$$b(\sigma) \approx b_{BP}(\sigma) = \frac{\prod_i b_i(\sigma_i) \prod_j b^j(\sigma^j)}{\prod_{(i,j)\in E} b_i^j\left(\sigma_i^j\right)}, \tag{4}$$

where $\forall\, i, j$: $\sigma_i = \left(\sigma_i^j \in \{0,1\} \middle| j = 1, \ldots, N\right)$ s.t. $\sum_j \sigma_i^j = 1$ and $\sigma^j = \left(\sigma_i^j \in \{0,1\} \middle| i = 1, \ldots, N\right)$ s.t. $\sum_i \sigma_i^j = 1$, i.e. σ_i and σ^j each has only N allowed states corresponding to allowed local perfect matchings for the vertices i and j, respectively. The vertex and edge beliefs are related to each other according to

$$\forall (i,j) \in E: \quad b_i^j\left(\sigma_i^j\right) = \sum_{\sigma_i \backslash \sigma_i^j} b_i(\sigma_i) = \sum_{\sigma^j \backslash \sigma_i^j} b^j(\sigma^j), \tag{5}$$

and the beliefs, as probabilities, should also satisfy the normalization conditions:

$$\forall (i,j) \in E: \quad b_i^j(1) + b_i^j(0) = 1. \tag{6}$$

Note that our notations for beliefs are not identical to the ones used in [8]: the multi-variable beliefs, b_i, are associated with vertices of $K_{N,N}$, and the single-variable beliefs, b_i^j, are associated with the edges of the graph. Substituting (4) into (3) and approximating $\sum_{\sigma \in \mathrm{PM}} b(\sigma) f\left(\sigma_i^j\right)$ with $\sum_{\sigma_i^j} b_i^j\left(\sigma_i^j\right) f\left(\sigma_i^j\right)$, etc, one arrives at the BFE functional

$$\mathcal{F}_{BP}\left\{b_i^j\left(\sigma_i^j\right); b_i(\sigma_i); b^j(\sigma^j)\right\} \equiv E - TS, \qquad E \equiv \sum_{(i,j)} b_i^j(1) \log\left(p_i^j\right), \tag{7}$$

$$S \equiv \sum_{(i,j)} \sum_{\sigma_i^j} b_i^j\left(\sigma_i^j\right) \ln b_i^j\left(\sigma_i^j\right) - \sum_i \sum_{\sigma_i} b_i(\sigma_i) \ln b_i(\sigma_i) - \sum_j \sum_{\sigma^j} b^j(\sigma^j) \ln b^j(\sigma^j). \tag{8}$$

Note that the BFE functional is bounded from below and generally non-convex, and thus finding the absolute minimum of the BFE is the main task of the BFE approximation. The BP approximation Z_{BP} of the partition function is given by $\mathcal{F}_{\mathrm{BP}} = -T \ln Z_{\mathrm{BP}}$ at a minimum of the BFE.

Moreover, the variational formulation of (5)–(8) can be significantly simplified in our case; one can utilize (5), (6) and express $b_i(\sigma_i)$, $b^j(\sigma^j)$ and $b_i^j\left(\sigma_i^j\right)$ solely in terms of the $\beta_i^j \equiv b_i^j(1)$ variables, satisfying doubly stochastic constraints

$$\forall (i,j) \in E: 0 \leqslant \beta_i^j \leqslant 1; \qquad \forall\, i: \sum_j \beta_i^j = 1; \qquad \forall\, j: \sum_i \beta_i^j = 1. \tag{9}$$

The entropy (8) becomes

$$\begin{aligned} S\{\beta_i^j\} &= \sum_{(i,j)} \left(\beta_i^j \log \beta_i^j + \left(1-\beta_i^j\right) \log\left(1-\beta_i^j\right)\right) - \sum_i \sum_j \beta_i^j \log \beta_i^j - \sum_j \sum_i \beta_i^j \log \beta_i^j \\ &= \sum_{(i,j)} \left(\left(1-\beta_i^j\right) \ln\left(1-\beta_i^j\right) - \beta_i^j \ln \beta_i^j\right). \end{aligned} \tag{10}$$

Therefore, the Bethe-free energy approach applied to the GM (2) results in minimization of the Bethe-free energy (BFE) functional

$$\mathcal{F}_{\mathrm{BP}}\{\beta\} = T \sum_{(i,j)\in E} \left(\beta_i^j \ln \frac{\beta_i^j}{\left(p_i^j\right)^{1/T}} - \left(1-\beta_i^j\right) \ln\left(1-\beta_i^j\right)\right), \tag{11}$$

over $\beta = \left(\beta_i^j\right)$ under the constraints (9).

To analyze the minima of the BFE, we incorporate Lagrange multipliers μ_i, μ^j enforcing the constraints in (9). Looking for a stationary point of the Lagrange function over the β variables, one arrives at the following set of quadratic equations for each (of N^2) variables, β_i^j:

$$\forall (i,j) \in E: \quad \beta_i^j\left(1-\beta_i^j\right) = \left(p_i^j\right)^{1/T} \exp\left(\mu_i + \mu^j\right). \tag{12}$$

One observes that any solution of (9), (12) at $T > 0$ that contains at least one β_i^j which is not integer does not contain any integers among all β_i^j. In fact, our main focus will be on these non-integer (interior) solutions of (9), (12). To find a solution of BP (9), (12) one relies on an iterative procedure. For a description of a set of iterative BP algorithms convergent to a minimum of the BFE for the perfect matching problem we refer the interested reader to [1, 4, 5].

Remark 1. Note that just derived BP approximation differs from the so-called mean-field (MF) approximation corresponding to the following ansatz:

$$b(\sigma) \approx b_{\mathrm{MF}}(\sigma) = \prod_{(i,j)\in E} b_i^j\left(\sigma_i^j\right), \tag{13}$$

enforcing statistical independence of the edge beliefs. If one substitutes $b(\sigma)$ by $b_{\mathrm{MF}}(\sigma)$ in (3) and also accounts for the normalization condition (6), which may be understood here as one enforcing the 'Fermi exclusion principle' for an edge (i, j) to contribute a perfect matching, $\sigma_i^j = 1$, the resulting expression for the MF free energy will turn into the BP expression (11) with the first term there changing the sign to $-$. One expects that BP approximation outperforms MF approximation in accuracy. Consider, for example, $N = 10$ and $\beta_i^j = 1/N$; then the exact BP and MF entropies are $\ln(10!) \approx 15.10$, $100(.9\ln(.9) - .1\ln(.1)) \approx 13.54$ and $100(-.9\ln(.9) - .1\ln(.1)) \approx 32.50$, respectively. An intuitive explanation for MF overestimating the entropy term is related to the fact that MF ignores correlations related to competitions between neighboring edges for contributing a perfect matching.

3. Threshold behavior of BP at low temperatures

As discovered in [11], at $T = 0$, properly scheduled iterative version of BP converges efficiently to the ML solution of the problem. In this context it is natural to ask the question of how a non-integer solution of BP emerges with a temperature increase. To address this question, we first consider the following homogeneous example.

Example 1. Define a homogeneous weight model biased toward a perfect matching solution, $\sigma_i^j = \delta_i^j : p_i^j = 1$ if $i \neq j$ and $p_i^i = W$ $(W > 1)$. Looking for β in the homogeneous form

$$\beta_i^j(T) = \begin{cases} 1 - \epsilon(N-1) & \text{if} \quad i = j \\ \epsilon & \text{otherwise}, \end{cases} \tag{14}$$

one observes that this ansatz for β solves the BP (9), (12) at ϵ equal to $\epsilon_{\min} = (N - 1 - W^{1/T})/((N-1)^2 - W^{1/T})$. At $T = \infty$, the probabilities are uniform, i.e. β from (14) with $\epsilon = \epsilon_{\min}$ is $\beta_i^j = 1/N$ for all $(i, j) \in E$. Now consider lowering the temperature and observe that at $T_c = \ln W/\ln(N-1)$ the nontrivial solution, with $\beta_i^j \neq 0, 1$ for all $(i, j) \in E$, turns exactly into the isolated/trivial ML one, $\beta_i^j = \delta_i^j$. Obviously one finds that the BFE, $\mathcal{F}_{\mathrm{BF}}$, considered as a function of ϵ, achieves its minimum at $\epsilon = \epsilon_{\min}$ if $T > T_c$. Exactly at $T = T_c$, this $\epsilon_{\min} = 0$ and the nontrivial solution merges into the isolated ML solution. The dependence of the BFE on ϵ for different T (at some exemplary values of N and W) is shown in figure 1(*a*). The partition function can be calculated efficiently. Counting the configurations straightforwardly (in a brute force combinatorial manner), one derives $Z = \sum_{k=0}^{N} W^{(N-k)/T}\binom{N}{k} D_k$. The following recursion is used to evaluate the number of permutation coefficient, D_k: $\forall k \geqslant 2$, $D_k = (k-1)(D_{k-1} + D_{k-2})$, $D_0 = 1$, $D_1 = 0$. A comparison of $T\ln Z$ and $T\ln Z_{\mathrm{BP}}$ as functions of T is shown in figure 1(*b*).

Returning to the case of an arbitrary non-negative P, we discover that this phenomenon of the nontrivial solution splitting at some finite nonzero (!!) temperature from the ML configuration is generic.

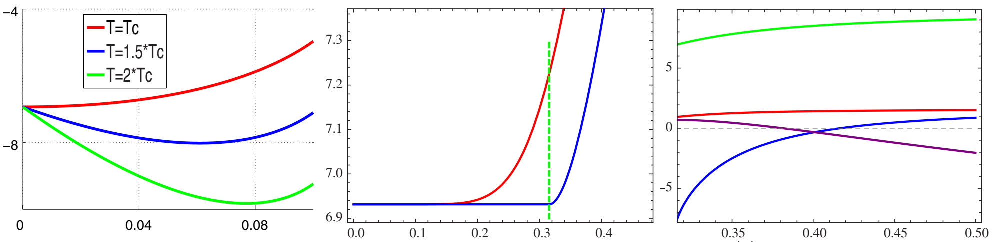

(a) $\mathcal{F}_{BP}$ vs ϵ. (b) $T \ln Z$ vs T. (c) $\ln(Z/Z_{BP})$ vs T for different estimators.

Figure 1. This figure contains a set of illustrations based on the homogeneous example 1 discussed in the text. $N = 10$ and $W = 2$ are chosen for these illustrations. (b) $T \ln Z$ for the homogeneous model (red) and respective BP expression, $T \ln Z_{BP}$ (blue) as functions of the temperature, T. Green dashed line mark T_c. (c) comparison of different estimations of $\ln(\mathrm{Perm}(\beta.*(1-\beta))/\prod_{(i,j)}(1-\beta_i^j))$ versus the temperature parameter T, where β is the matrix of marginal beliefs evaluated at a fixed point of BP equations. Red, blue, purple, green and dashed-gray lines show the exact expression, the lower bound of corollary 7, the lower bound of theorem 8, the upper bound of proposition 9 and the BP expression, respectively. (a) $\mathcal{F}_{\mathrm{BP}}$ versus ϵ. (b) $T \ln Z$ versus T. (c) $\ln(Z/Z_{\mathrm{BP}})$ versus T for different estimators.

(This figure is in colour only in the electronic version)

Proposition 1. *For any non-negative matrix* $P = \left((p_i^j)^{1/T} \middle| i, j = 1, \ldots, N\right)$ *one finds a special (we call it critical) temperature,* T_c*, such that for* $T > T_c + \varepsilon$ *a nontrivial solution of BP, corresponding to a local non-saturated minimum of* $\mathcal{F}_{\mathrm{BP}}$*, dominating the respective value corresponding to the maximum likelihood solution, is realized for at least a sufficiently small positive* ε*. This special solution coincides with the best perfect matching solution at* $T = T_c$ *and it does not exist for* $T < T_c$*. The critical temperature* T_c *solves*

$$\det\left(P_i^j - 2\sigma_{*i}^j P_i^j\right) = 0, \tag{15}$$

where σ_* *is the ML configuration.*

Proof. Our proof of the proposition is constructive. Let us look for a solution of the BP equations weakly deviating from the ML configuration σ_*. Without loss of generality we assume that $\sigma_{*i}^j = \delta_i^j$. We introduce $v_i^j = \beta_i^j\left(1-\beta_i^j\right) \ll 1$ and observe that a nontrivial solution, approaching the ML one at $v \to 0$, is $\beta_i^j = \left(1 - \left(1 - 2\delta_i^j\right)\left[1 - 4v_i^j\right]^{1/2}\right)/2$. Linearizing the normalization condition, over v one derives $\forall i : v_i^i = \sum_{j \neq i} v_i^j; \forall\, j : v_j^j = \sum_{i \neq j} v_i^j$. On the other hand, the BP equation (12), complemented by the set of linear constraints on v, translates into $\forall\, i : P_i^i U^i = \sum_{j \neq i} P_i^j U^j; \forall\, j : P_j^j U_j = \sum_{i \neq j} P_i^j U_i$, where $U_i = \exp(\mu_i)$ and $U^j = \exp(\mu^j)$. Requiring that the later equations have a nontrivial solution (with nonzero v), one arrives at the critical temperature condition (15). It is then straightforward to verify that the extension of the nontrivial solution into the $T < T_c$ domain is unphysical (as some elements of the respective small v solution are negative), while the BFE associated with the nontrivial solution for $T > T_c$ is smaller than the one corresponding to the ML perfect matching. □

Conjecture 2. *We conjecture that the non-integer solution of BP equations discussed in proposition 1 extends beyond the small* $T_c + \varepsilon$ *vicinity of* T_c*, and this solution transitions smoothly at* $T \to \infty$ *into the obvious fully homogeneous solution,* $\beta_i^j = 1/N$ *for all* $(i, j) \in E$*. Another plausible conjecture is that no other non-integer solutions exist at* $T < T_c$*; therefore,* *when the non-integer solution discussed in the proposition emerges at $T = T_c$ it, in fact, gives a global minimum of the BFE.*

4. Background (II): loop calculus and series

Here we consider $T > T_c$ where, according to the main result of the previous section, there exists a solution of (9), (12) lying in the interior of the doubly stochastic matrix polytope. We assume that such a nontrivial solution of the BP equations is found.

As shown in [12, 13], the exact partition function of a generic GM can be expressed in terms of a LS, where each term is computed explicitly using the BP solution. Adapting this general result to the permanent, bulky yet straightforward algebra leads to the following exact expression for the partition function Z from (2):

$$Z/Z_{\rm BP} = z_{\rm LS}; \qquad z_{\rm LS} \equiv 1 + \sum_{C\neq\emptyset} r_C,$$
$$r_C \equiv \left(\prod_{i\in C}(1-q_i)\right)\left(\prod_{j\in C}(1-q^j)\right)\prod_{(i,j)\in C}\frac{\beta_i^j}{1-\beta_i^j}. \tag{16}$$

The variables β are in accordance with (9), (12) and C stands for an arbitrary generalized loop, defined as a subgraph of the complete bipartite graph with all its vertices having a degree larger than 1. The q_i (or q^j) in (16) are the C-dependent degrees, i.e. $q_i = \sum_{j|(i,j)\in C} 1$ and $q^j = \sum_{i|(i,j)\in C} 1$. According to (16), those loops with an even/odd number of vertices give positive/negative contributions r_C.

5. Loop series as a permanent

This section, explaining the main result of the communication, is split into two parts. In subsection 5.1 we give a simple derivation of a very compact representation for the LS (16) following directly from the BFE formulation. Subsection 5.2 contains an alternative derivation of this main formula from LS using the concept of the Ihara–Bass graph ζ-function [14, 15].

We also find it appropriate here to make the following general remark. Even though discussion of the manuscript is limited to permanents, counting perfect matchings over $K_{N,N}$, all the results reported in this section allows for straightforward generalizations to weighted counting of perfect matchings over arbitrary (and not necessarily bipartite) graphs.

5.1. Permanent representation for $Z/Z_{\rm BP}$

Theorem 3. *For any non-integer solution of the BP equations (9), (12), the following is true:*

$$\mathrm{Perm}(P)/Z_{BP} = \mathrm{Perm}(\beta.*(1-\beta))\prod_{(i,j)\in E}\left(1-\beta_i^j\right)^{-1}, \tag{17}$$

where $A.*B$ is the element-by-element multiplication of the A and B matrices.

Proof. From the definition of the BFE, $\mathcal{F}_{\rm BP} = -T\ln Z_{\rm BP}$ and (9), (12), one derives

$$Z_{\rm BP} = \prod_{(i,j)\in E}\left[\left(1-\beta_i^j\right)\left(\frac{\left(p_i^j\right)^{1/T}}{\beta_i^j\left(1-\beta_i^j\right)}\right)^{\beta_i^j}\right] = \prod_{(i,j)\in E}\left(1-\beta_i^j\right)\prod_i \mathrm{e}^{-\mu_i}\prod_j \mathrm{e}^{-\mu^j}.$$

On the other hand (12) results in $\mathrm{Perm}(P) = \mathrm{Perm}(\beta.*(1-\beta))\prod_i \exp(-\mu_i)\prod_j \exp(-\mu^j)$. Combining the two formulas we arrive at (17). □

Remark 2. Note that if one considers expanding the permanent on the rhs of (17) over the elements of the matrix $\beta.*(1-\beta)$, each element of the expansion will be positive, in contrast with the LS of (16). Moreover, the number of terms in the Perm-expansion is significantly smaller than those in the original LS.

5.2. From LS to the permanent representations for Z/Z_{BP}

Here we discuss the relation between the two complementary representations of Z/Z_{BP}, i.e. between the LS expression (16) and the permanent formula (17). We do this in two steps, stated in the two theorems presented consequently, one relating the LS to an average of a determinant, and another expressing it via the permanent of $\beta.*(1-\beta)$.

Theorem 4 (LS as an average of the determinant)**.** *Let $\vec{E}$ be the set of directed edges obtained by duplicating undirected edges E of $K_{N,N}$. Define the edge-adjacency matrix $\mathcal{M}$ of the complete bipartite graph $K_{N,N}$ according to $\mathcal{M}_{i\to j,k\to l} = \delta_{l,i}(1-\delta_{j,k})$. Let $x = (x_{i\to j})_{(i\to j)\in\vec{E}}$ be the set of random variables that satisfies $\langle x_{i\to j}\rangle = 0$, $\langle x_{i\to j}x_{j\to i}\rangle = 1$ and $\langle x_{i\to j}x_{k\to l}\rangle = 0$ $(\{i,j\}\neq\{k,l\})$. (Here and below $\langle\cdots\rangle_x$ stands for the mathematical expectation over the random variables x.) Then, the following relation holds: $z_{\mathrm{LS}} = \langle\det[I - \mathrm{i}\mathcal{B}\mathcal{M}]\rangle_x$, where $\mathcal{B} = \mathrm{diag}\left(\sqrt{\beta_i^j/\left(1-\beta_i^j\right)}x_{i\to j}\right)$.*

Proof. For a general undirected graph G, the Ihara–Bass formula [14, 15] states that

$$\zeta_G^{-1}(u) = \det[I - u\mathcal{M}] = (1-u)^{|E|-|V|}\det[I + u^2(\mathcal{D}-I) - u\mathcal{A}], \tag{18}$$

where $\mathcal{A}$ is the adjacency matrix and $\mathcal{D} = \mathrm{diag}(q_i; i\in V)$ is the degree matrix of G. If we take the limit $u\to\infty$, this formula implies $\det\mathcal{M} = (-1)^{|E|}\prod_{i\in V}(1-q_i)$. Expanding the determinant, one derives

$$\det[I - \mathrm{i}\mathcal{B}\mathcal{M}] = \sum_{\{e_1,\ldots,e_n\}\subset\vec{E}} \det\mathcal{M}|_{\{e_1,\ldots,e_n\}}(-i)^k\prod_{l=1}^n(\mathcal{B})_{e_l,e_l}. \tag{19}$$

Evaluating the expectation of each summand in (19), one observes that it is nonzero only if $(i\to j)\in\{e_1,\ldots,e_n\}$ implies $(j\to i)\in\{e_1,\ldots,e_n\}$, thus arriving at

$$\langle\det[I - \mathrm{i}\mathcal{B}\mathcal{M}]\rangle_x = \sum_{C\subset E}(-1)^{|C|}\det\mathcal{M}|_C\prod_{(i,j)\in C}\frac{\beta_i^j}{1-\beta_i^j} = 1 + \sum_{\emptyset\neq C\subset E} r_C.$$

□

Theorem 5 (from LS to permanent)**.** *For the doubly stochastic matrix of BP beliefs β and LS defined in (16), one derives*

$$z_{LS} = \mathrm{Perm}(\beta.*(1-\beta))\prod_{(i,j)\in E}\left(1-\beta_i^j\right)^{-1}.$$

Proof. We use theorem 4, choosing the random variables $x_i^j = x_{i\to j} = x_{j\to i}$ that take ± 1 values with probability 1/2. We also utilize a multivariate version of the Ihara–Bass formula from [16] to derive the following expression for z_{LS} proving the theorem:

$$\det[I - \mathrm{i}\mathcal{B}\mathcal{M}] = \det\begin{bmatrix} 0 & \sqrt{\beta.*(1-\beta)}.*x \\ (\sqrt{\beta.*(1-\beta)}.*x)^T & 0\end{bmatrix}\prod_{(i,j)\in E}\left(1-\beta_i^j\right)^{-1},$$

$$z_{LS} = \langle\det(\sqrt{\beta.*(1-\beta)}.*x)^2\rangle_x\prod_{(i,j)}\left(1-\beta_i^j\right)^{-1} = \mathrm{Perm}(\beta.*(1-\beta))\prod_{(i,j)}\left(1-\beta_i^j\right)^{-1}.$$

□

6. Invariance of the Gurvits–van der Waerden lower bound and new lower bounds for the permanent

Van der Waerden [17] conjectured that the minimum of the permanent over the doubly stochastic matrices is $N^N/N!$, and it is only attained when all entries of the matrix are $1/N$. Though the conjecture appears to be simple, it remained open for over 50 years before Falikman [18] and Egorychev [19] finally proved it. Recently Gurvits [20] found an alternative, surprisingly short and elegant proof that also allowed for a number of unexpected extensions of the Van der Waerden conjecture. We call it the Gurvits–van der Waerden theorem. (See e.g. [21].) A simplified form of this theorem is as follows.

Theorem 6 (Gurvits–van der Waerden theorem [20, 21])**.** *For an arbitrary non-negative $N \times N$ matrix A,*

$$\mathrm{Perm}(A) \geqslant \mathrm{cap}(p_A)\frac{N^N}{N!}, \qquad \textit{where} \quad p_A(x) \equiv \prod_i \sum_j a_{i,j} x_j, \qquad \mathrm{cap}(p_A) \equiv \inf_{x\in\mathbb{R}^N_{>0}} \frac{p_A(x)}{\prod_j x_j}.$$

We have found that the lower bound of theorem 6 has a 'good' property with respect to the BP transformation. As stated in theorem 3, BP transforms the permanent to another permanent. Therefore, applying theorem 6 to both sides of (17), one naturally asks how do the two lower bounds compare? A somewhat surprising result is that the Gurvits–van der Waerden theorem is invariant with respect to the BP transformation. Namely, $\mathrm{cap}(p_P) = Z_{BP} * \mathrm{cap}(p_{\beta.*(1-\beta)}) \prod_{(i,j)\in E} \left(1-\beta_i^j\right)^{-1}$. The lower bound for $\mathrm{Perm}(\beta.*(1-\beta))$ based on theorem 6 is

Corollary 7.

$$\mathrm{Perm}(\beta.*(1-\beta)) \geqslant \frac{N!}{N^N} \prod_{(i,j)\in E} \left(1-\beta_i^j\right)^{\beta_i^j}.$$

Proof. This bound is the result of a direct application of the inequality $\sum_j \beta_i^j \left(1-\beta_i^j\right) x_j \geqslant \prod_j \left[\left(1-\beta_i^j\right) x_j\right]^{\beta_i^j}$ to theorem 6. □

We also obtain another lower bound which improves the bound of corollary 7 at sufficiently low values of the temperature. See figure 1(*c*) for an illustration.

Theorem 8. *For an arbitrary perfect matching Π (permutation of $\{1, \ldots, N\}$),*

$$\mathrm{Perm}(\beta.*(1-\beta)) \geqslant 2\prod_i \beta_i^{\Pi(i)}\left(1-\beta_i^{\Pi(i)}\right).$$

Proof. Without loss of generality, we assume that Π is the identity permutation. From the positivity of entries and (9), we have $\mathrm{Perm}(\beta.*(1-\beta)) \geqslant \prod_i \beta_i^i \,\mathrm{Perm}(X)$, where $X_{ij} = \delta_{i,j} + (1-2\delta_{i,j})\beta_i^j$. Since β is a stochastic matrix, $\det X = 0$, and thus $\mathrm{Perm}(X) \geqslant 2\prod_i \left(1-\beta_i^i\right)$. □

Note, for the sake of completeness, that a comprehensive review of other bounds on permanents of specialized matrices (for example 0, 1 matrices) can be found in [22].

7. New upper bound for permanent

Proposition 9.

$$\mathrm{Perm}(\beta.*(1-\beta)) \leqslant \prod_j \left(1-\sum_i \left(\beta_i^j\right)^2\right).$$

Proof. We use the Godzil–Gutman representation for permanents [23]

$$\mathrm{Perm}(\beta.*(1-\beta)) = \langle \det(\sqrt{\beta.*(1-\beta)}.*\sigma)^2\rangle_\sigma, \tag{20}$$

where $\sigma_i^j = \pm 1$, with $i, j = 1, \ldots, N$, are independent random variables taking values ± 1 of equal probability. Each row of the matrix $\sqrt{\beta.*(1-\beta)}.*\sigma$ has the squared Euclid norm $\sum_i \beta_i^j\left(1-\beta_i^j\right) = 1-\sum_i\left(\beta_i^j\right)^2$. Therefore, the upper bound is obtained from the Hadamard inequality $|\det(a_1,\ldots,a_n)| \leqslant \|a_1\|\cdots\|a_n\|$. □

8. Path forward

We consider this study to be the beginning of further research along the following lines: (1) more detailed analysis of the BP solution, in particular, study of T_c, e.g. concerning its dependence on the matrix size, analysis of the BP solution dependence on temperature, and the construction of an iterative algorithm provably convergent to a nontrivial BP solution for $T > T_c$; (2) explanation of the BP invariance with respect to the Gurvits–van der Warden lower bound; (3) development of a deterministic and/or randomized polynomial algorithm for estimating the permanent with provable guarantees based on the loop calculus expression; and (4) numerical tests of the lower and upper bounds for realistic large-scale problems.

Acknowledgments

We are thankful to Leonid Gurvits for educating us, through his course of lectures given at CNLS/LANL, about existing approaches in the 'mathematics of the permanent'. YW acknowledges support of the Students Visit Abroad Program of the Graduate University for Advanced Studies which allowed him to spend two months at LANL and he is also grateful to CNLS at LANL for its hospitality. Research at LANL was carried out under the auspices of the National Nuclear Security Administration of the US Department of Energy at Los Alamos National Laboratory under Contract no DE C52-06NA25396. MC also acknowledges partial support of NMC via the NSF collaborative grant, CCF-0829945, on 'Harnessing Statistical Physics for Computing and Communications'.

References

[1] Chertkov M, Kroc L, Krzakala F, Vergassola M and Zdeborova L 2010 *Proc. Natl Acad. Sci.* **107** 7663–8 (arXiv:0909.4256)
[2] Valiant L 1979 *Theor. Comput. Sci.* **8** 189–201
[3] Jerrum M, Sinclair A and Vigoda E 2004 *J. ACM* **51** 671–97
[4] Chertkov M, Kroc L and Vergassola M 2008 Belief propagation and beyond for particle tracking arXiv:0806.1199
[5] Huang B and Jebara T 2009 Approximating the permanent with belief propagation arXiv:0908.1769
[6] Gallager R 1963 *Low Density Parity Check Codes* (Cambridge, MA: MIT Press)
[7] Pearl J 1988 *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (San Francisco, CA: Morgan Kaufmann)

[8] Yedidia J S, Freeman W T and Weiss Y 2005 Information theory *IEEE Trans. Inf. Theory* **51** 2282–312
[9] Kuhn H W 1955 The Hungarian method for the assignment algorithm *Nav. Res. Logist. Q.* **1** 83–97
[10] Bertsekas D 1992 Auction algorithms for network flow problems: a tutorial introduction *Comput. Optim. Appl.* **1** 7–66
[11] Bayati M, Shah D and Sharma M 2008 *IEEE Trans. Inf. Theory* **54** 1241–51 (*Proc. IEEE Int. Symp. Information Theory, 2006*)
[12] Chertkov M and Chernyak V 2006 *Phys. Rev.* E **73** 065102
[13] Chertkov M and Chernyak V Y 2006 *J. Stat. Mech.* P06009 (arXiv:cond-mat/0603189)
[14] Ihara Y 1966 *J Math. Soc. Japan* **18** 219–35
[15] Bass H 1992 *Int. J. Math.* **3** 717–97
[16] Watanabe Y and Fukumizu K 2009 Graph Zeta Function in the Bethe Free Energy and Loopy Belief Propagation *Adv. Neural Inf. Process. Syst.* **22** 2017–25
[17] van der Waerden B 1926 *[Aufgabe] 45, Jahresbericht der Deutschen Mathematiker-Vereinigung* **35** 117
[18] Falikman D 1981 *Math. Notes* **29** 475–9
[19] Egorychev G 1981 *Siberian Math. J.* **22** 854–9
[20] Gurvits L 2008 *Electron. J. Comb.* **15** R66
[21] Laurent M and Schrijver A 2009 On Leonid Gurvits' proof for permanents
[22] Lovász L and Plummer M 1986 *Matching Theory (North-Holland Mathematics Studies* vol 121*) (Annals of Discrete Mathematics* vol 29*)* (Amsterdam: Elsevier)
[23] Godsil C D and Gutman I 1981 *J. Graph Theory* **5** 137–44