

Belief-Propagation for Weighted b-Matchings on Arbitrary Graphs and its Relation to Linear Programs with Integer Solutions

Mohsen Bayati* Christian Borgs* Jennifer Chayes* Riccardo Zecchina†

Abstract

We consider the general problem of finding the minimum weight b-matching on arbitrary graphs. We prove that, whenever the linear programming (LP) relaxation of the problem has no fractional solutions, then the belief propagation (BP) algorithm converges to the correct solution. We also show that when the LP relaxation has fractional solution then BP algorithm can be used to solve the LP. These results are notable in several regards: (1) It is one of a very small number of proofs showing correctness of BP without any constraint on the graph structure. (2) Instead of showing that BP leads to a PTAS, we give a finite bound for the number of iterations after which BP has converged to the exact solution. (3) Variants of the proof work for both synchronous and asynchronous BP; to the best of our knowledge, it is the first proof of convergence and correctness of an asynchronous BP algorithm for a combinatorial optimization problem. (4) It works for both ordinary b-matchings and the more difficult case of perfect b-matchings. (5) Together with the recent work of Sanghavi, Malioutov and Wiskly [41] they are the first complete proofs showing that tightness of LP implies correctness of BP.

1 Introduction

Motivated by the cavity method in statistical physics, very fast distributed heuristic algorithms have recently been developed for the solution of random constraint satisfaction problems [28], [12], [15], [1]. Similar heuristic methods have been known for many years [20] in the context of coding theory. And a variety of specific examples of such algorithms have been developed in artificial intelligence, signal processing, and digital communications. Well-known examples include the Viterbi algorithm, the iterative decoding algorithm in turbo codes and in low-density parity-check codes [38], Pearl’s belief propagation algorithm for Bayesian networks [35], the Kalman filter, and certain fast Fourier transform (FFT) algorithms. Very recent applications can also be found in systems biology [19], [22], [57], computer vision [44], and data clustering [17].

In some cases, the algorithms generated by the cavity method are exactly of the form of a classic belief propagation (max-product or min-sum) i.e., a message-passing algorithm for efficiently computing marginal probabilities or finding the assignment with highest probability of a joint discrete probability distribution defined on a graph. The belief propagation (BP) algorithm converges to a correct solution if the associated graph is a tree, and may be also a good heuristic for some graphs with cycles. In other cases, the cavity method may lead to a more involved survey propagation (SP) algorithm [28], in which some form of correlation among variables is controlled.

In this paper, we study the problem of finding the minimum weight b-matchings in *arbitrary* graphs via the min-sum version of BP algorithm¹.

Our Results. Let $G = (V, E)$ be an undirected graph with edge weights w_{ij} for each edge $\{i, j\} \in E$ and node capacities b_i for each node $i \in V$. The iterative message-passing algorithm based on synchronous

*Microsoft Research; {mohsenb , borgs , jchayes}@microsoft.com

†Politecnico Di Torino; riccardo.zecchina@polito.it

¹Throughout this paper, the term BP algorithm refers to min-sum version of the BP algorithm.

BP for solving the weighted perfect b-matching problem (see our Section 2 for the precise definition) is the following *simple* procedure: At each time, every vertex of the graph sends (real valued) messages to each of its neighbors. The message transmitted at time t from i to j is w_{ij} minus the b_i^{th} minimum of the messages previously received by i at time $t - 1$ from all of its neighbors except j . At the end of each iteration, every vertex i selects b_i of its adjacent edges that correspond to the b_i smallest received messages.

We will show the following result: For arbitrary graphs G , and all sets of weights $\{w_{ij}\}$, after $O(n)$ iterations, the set of selected edges converges to the correct solution, i.e., to the minimum weight perfect b-matching of G , provided that the LP relaxation of the problem (see Section 2 for definitions) has no fractional solutions. Additionally we introduce a new construction, a *generalized computation tree*, which allows us to analyze the more complicated case of BP with an asynchronous updating scheme, and prove convergence and correctness of it when each edge of the graph transmits at least $\theta(n)$ messages. To the best of our knowledge, this technique is new and can be applied in the analysis of asynchronous BP in other problems as well. These are extensions of the previous results of [5] and [24] which showed convergence and correctness of the above algorithm for bipartite graphs.² Moreover, our proof gives a better understanding of the often-noted but poorly understood connection between BP and LP through *graph covers*. We also modify our BP algorithm and its analysis to include the problem of finding the *non-perfect* weighted b-matchings. Recently and independently from our work a similar result for the scenario of using synchronous BP for non-perfect 1-matchings was shown by Sanghavi, Malioutov and Willsky [41]. Finally we will show that when LP has a solution that is partially fractional then a slightly modified BP converges to correct values for the integer parts and oscillates for the fractional parts of the LP solution, hence solving the LP.

Related Works. The weighted b-matching problem is an important problem in combinatorial optimization. It belongs to a family of integer linear programs which have been well-studied and can be solved in strongly polynomial time [13], [14]. For extensive surveys see [23] and [36]. In physics, the study of the random 1-matching problem goes back to the work of Mèzard and Parisi [27] who made a celebrated conjecture for the expected optimum weight ($\pi^2/6$) that was proven to be exact a decade later by Aldous [2].

BP algorithms have been the subject of extensive study in several communities. The general BP algorithm is known to be correct on graphs with no cycles [35]. For graphs with a single cycle, convergence and correctness of BP have also been rigorously analyzed [3], [52]. For arbitrary graphs, relatively little is known about the correctness of BP, although some interesting progress has been made in [54], [49], [45], [58]. Performance of the BP algorithm usually depends on the length of cycles in graphs; most analytical results require that the graphs have no short cycles (i.e., that they are large-girth graphs) [38], [4], [21]. For the case of weighted matchings and a few other problems, there were initially surprising results that BP works correctly on graphs with many short cycles ([53], [39], [5], [29], [34]).

Recent works have also suggested a connection between the BP algorithm and linear programming (LP) in particular problems. A relationship between iterative decoding of channel codes and LP decoding was studied in [16], [47], [46]. In fact our proof is based on the notion of graph-cover that is used in [47] and [46] as well. Other relationships were noted in the context of BP algorithms with convex free energies [51], [50], [55], and in the case of BP algorithms for resource allocations [31]. For weighted 1-matchings, the connection was studied [6] in the context of similarities between BP equations and the primal-dual auction algorithm of Bertsekas [8]. And it was further clarified recently for non-perfect 1-matchings in [40] and [41] where it was shown that BP does not converge to the correct solution if the LP relaxation has fractional solutions. Another recent result studies this connection for the weighted independent set problem [42]. We will compare our work with some of these results in the “Technical Contribution” section below.

Finally, we note that the BP equations for solving the weighted matching problem which we use in this paper have been previously studied in [6], [24]. These equations are also very similar to equations for weighted matching problems and traveling salesman problems given in [27], [48], [2], [21], and to equations for various other problems given in [59], [37], [26].

²Both of these results were assuming that the minimum weight matching is unique. Note that if there is more than one solution, then one can construct a fractional solution to the LP relaxation.

Technical Contribution. The main contributions of our results and techniques can be summarized as follows:

1. *BP for the weighted matching* was first used in [5] and its correctness and convergence was shown for bipartite graphs with unique optimum solution. That proof relied heavily on the fact that the minimum weight matching of a bipartite graph is locally optimal on any cycle since the cycles of a bipartite graph have even length. The same technique was used in [24] to extend the result to b-matchings in bipartite graphs. But this technique fails for graphs containing cycles with odd length. In order to bypass this difficulty we use graph covers for connecting the bipartite graphs to general graphs and use proof techniques introduced in [5], and [24]. In the preliminary version of this paper [7] we give a completely different and independent proof using complementary slackness conditions of the LP relaxation and its dual. This proof is included in the appendix of this paper.
2. *Connection between LP and BP* has been suggested and analyzed by various groups (as we discussed above), but our result together with [41], to the best of our knowledge, are the first ones which show both convergence and correctness of the BP algorithm when LP relaxation has no fractional solutions. We also use the notion of graph-covers to solve the LP relaxation via BP. One related result, [42], studies only properties of the BP fixed points and their relation to the LP, *conditioned* on the convergence of the BP algorithm. Similarly in another recent work, [55], which generalizes methods of [50] and [51], the connection of the BP algorithm and LP relaxation is studied in the converged case of the BP. The authors also study interesting variations of the BP which have convex free energies.
3. *The asynchronous BP*, which includes the synchronous version as a special case, has been a more popular version for practical purposes. But, due to its more complicated structure, it has not been the subject of much rigorous study. To the best of our knowledge, our work is the first correctness and convergence proof of asynchronous BP for a combinatorial optimization problem. Another advantage of our proof is the construction of a new tool, the generalized computation tree, which can be used for the analysis of the both convergence and correctness of asynchronous message-passing algorithms including BP. Without the notion of a suitable computation tree the existing methods, free energy analysis [58][49] or Lipschitz functions [21][4], do not give correctness and convergence at the same time.

Organization of the Paper. The rest of the paper is organized as follows. In Section 2, we provide the setup, define the weighted b-matching problem, and describe the LP relaxation. In Section 3, we describe our algorithm for the minimum weighted perfect b-matching problem, and state our main result. The analysis of our algorithm is given in Section 4. The extension of our algorithm and results to the *non-perfect* minimum weighted b-matching problem are presented in Section 5. In Section 6, we state the asynchronous version of the BP algorithm and present its analysis, and finally in Section 7 we show how BP can be used to solve LP. The appendix is dedicated to the alternative proofs for the results of Sections 3-6 using dual variables of the LP relaxation.

2 Definitions and Problem Statement

Consider an undirected simple graph $G = (V, E)$, with vertices $V = \{1, \dots, n\}$, and edges E . Let each edge $\{i, j\}$ have weight $w_{ij} \in \mathbb{R}$. Denote the set of neighbors of each vertex i in G by $N(i)$. Let $\mathbf{b} = (b_1, \dots, b_n)$ be a sequence of positive integers such that $b_i \leq \deg_G(i)$. A subgraph M of G is called a *b-matching* (*perfect b-matching*) if the degree of each vertex i in M is at most b_i (equal to b_i). Denote the set of b-matchings (perfect b-matchings) of G by $M_G(\mathbf{b})$ ($PM_G(\mathbf{b})$), and assume that it is non-empty. Clearly $PM_G(\mathbf{b}) \subset M_G(\mathbf{b})$.

The weight of a (perfect or non-perfect) b-matching M , denoted by W_M , is defined by $W_M = \sum_{ij} w_{ij} 1_{\{i,j\} \in M}$. In the next two sections, we will restrict ourselves to the case of *perfect* b-matchings. We will extend the analysis to (possibly non-perfect) b-matchings in Section 5. The minimum weight perfect b-Matching (b-MWPM), M^* , is defined by $M^* = \operatorname{argmin}_{M \in PM_G(\mathbf{b})} W_M$. The goal of this paper is to find M^* via a

min-sum belief propagation algorithm. Throughout the paper, we will assume that M^* is unique. Let ϵ be the difference between the weights of M^* and the second minimum weight b-matching; i.e.

$$\epsilon = \min_{M \neq M^*} (W_M) - W_{M^*}.$$

due to the uniqueness of the M^* , $\epsilon > 0$. Also define w^* to be $w^* = \max_{\{i,j\} \in E} (w_{ij})$

Linear Programming Relaxation. Assigning variables $x_{ij} \in \{0, 1\}$ to the edges in E , we can express the weighted perfect b-matching problem as the problem of finding a vector $x \in \{0, 1\}^{|E|}$ that minimizes the total weight $\sum_{ij \in E} x_{ij} w_{ij}$, subject to the constraints $\sum_{j \in N(i)} x_{ij} = b_i$ for all $i \in V$. Relaxing the constraint that x_{ij} is integer, this leads to the following linear program:

$$\begin{aligned} \min & \quad \sum_{\{i,j\} \in E} x_{ij} w_{ij} \\ \text{subject to} & \quad \sum_{j \in N(i)} x_{ij} = b_i \quad \forall i \\ & \quad 0 \leq x_{ij} \leq 1 \quad \forall \{i, j\} \in E \end{aligned} \tag{1}$$

We say the LP relaxation (1) has *no fractional solution* if, every optimal solution x of LP satisfies $x \in \{0, 1\}^{|E|}$. Note that absence of fractional solutions implies uniqueness of integer solutions, since any convex combination of two integer solutions is a solution to the LP as well. We want to show that the BP algorithm for our problem converges to the correct solution, provided the LP relaxation (1) has no fractional solution.

3 Algorithm and Main Results

The following algorithm is a synchronous implementation of BP for finding the minimum weight perfect b-matching (b-MWPM). The main intuition behind this algorithm (and, indeed, all BP algorithms) is that each vertex of the graph assumes the graph has no cycles, and makes the best (greedy) decision based on this assumption. This is shown in more detail in Section 4.1.

Before applying the BP algorithm, we remove all *trivial* vertices from the graph. A vertex i is called trivial if $\deg_G(i) = b_i$. This is because all of the edges adjacent to i should be in every perfect b-matching. Therefore the graph can be simplified by removal of all trivial vertices and their adjacent edges.

Algorithm Sync-BP.

- (1) At times $t = 0, 1, \dots$, each vertex sends real-valued messages to each of its neighbors. The message of i to j at time t is denoted by $m_{i \rightarrow j}(t)$.
- (2) Messages are initialized³ by $m_{i \rightarrow j}(0) = w_{ij}$ for all $\{i, j\} \in E$.
- (3) For $t \geq 1$, messages in iteration t are obtained from messages in iteration $t - 1$ recursively as follows:

$$\forall \{i, j\} \in E: \quad m_{i \rightarrow j}(t) = w_{ij} - b_i^{t^{\text{th}}\text{-min}_{\ell \in N(i) \setminus \{j\}}} \left[m_{\ell \rightarrow i}(t-1) \right] \tag{2}$$

where $k^{\text{th}}\text{-min}[A]$ denotes the k^{th} minimum⁴ of set A .

- (4) The estimated b-MWPM at the end of iteration t is $M(t) = \cup_{i=1}^n E_i(t)$ where $E_i(t) = \{\{i, j_1\}, \dots, \{i, j_{b_i}\}\}$ is such that $N(i) = \{j_1, j_2, \dots, j_{\deg_G(i)}\}$ and $m_{j_1 \rightarrow i}(t) \leq m_{j_2 \rightarrow i}(t) \leq \dots \leq m_{j_{\deg_G(i)} \rightarrow i}(t)$. i.e., among all i 's neighbors, choose edges to the b_i neighbors that transfer the smallest incoming messages to i .

³We show in Section 4.4 that the messages can be initialized by any arbitrary values.

⁴Note that the $b_i^{t^{\text{th}}\text{-min}_{\ell \in N(i) \setminus \{j\}}}$ is well defined since we assumed that all trivial vertices are removed and thus there are at least $b_i + 1$ elements in the set $N(i)$ for each i .

(5) Repeat (3)-(4) until $M(t)$ converges⁵.

In Corollary 1, we will show the main intuition behind the equation (2) and how it is derived. But we note that one can also use the graphical model representations of [5], [24], [40] to obtain the standard BP equations for this problem, which, after some algebraic calculations, yield the recursive equation (2).

The main result of the paper is rather surprising: it says that the above algorithm, which is designed for graphs with no cycle (i.e., for trees), works correctly for a much larger family of graphs including those with many short cycles.

Theorem 1 *Assume that the LP relaxation (1) has no fractional solution. Then the algorithm Sync-BP converges to M^* after at most $\lceil \frac{2nw^*}{\epsilon} \rceil$ iterations.*

Note 1. We will show in Section 7 that if BP is used properly then it can solve all cases of the LP relaxation 1.

If the LP relaxation (1) has a fractional solution whose cost is strictly less than W_{M^*} , then [40], [41] have shown for the case of 1-matching that BP does not converge to M^* . It is straightforward to generalize this to perfect b-matching as well. But for the case in which the LP relaxation has a fractional solution whose cost is equal to W_{M^*} , BP fails in general. This is because the b_i^{th} minimum in equation (2) is not unique, and one needs an oracle to make the right decision. If such an oracle exists, then BP converges to M^* .

4 Analysis of the Synchronous BP

This section contains the analysis of the synchronous BP algorithm for perfect b-matchings. First, in Section 4.1 we show one derivation of the equations for Sync-BP and its representation in term of the so-called computation tree. Next, in Section 4.2 we introduce the notion of graph-covers which connects the graph G to a bipartite graph \tilde{G} that has all the information for finding the minimum weight perfect b-matching in the graph G . This connection is used in Section A.3 to prove that, when the LP relaxation has no fractional solutions, then solutions on the computation tree are the same as the solutions on the original graph G .

4.1 Computation Tree and Derivation of Sync-BP

The main idea behind the algorithm Sync-BP is that it assumes the graph G has no cycle. In other words, it finds the b-MWPM of a graph G' that has the same local structure as G but no cycles. In this section we rigorously define such graph G' (computation tree) and show its connection with the Sync-BP algorithm.

Computation Tree. For any $i \in V$, let T_i^t be the t -level computation tree corresponding to i , defined as follows: T_i^t is a weighted tree of height $t + 1$, rooted at i . All tree-nodes have labels from the set $\{1, \dots, n\}$ according to the following recursive rules:

- (a) The root has label i .
- (b) The set of labels of the $deg_G(i)$ children of the root is equal to $N(i)$.
- (c) If s is a non-leaf node whose parent has label r , then the set of labels of its children is $N(s) \setminus \{r\}$.

Note 2. T_i^t is often called the *unwrapped tree* at node i . The computation tree is constructed by replicating the local connectivity of the original graph. The messages received by node i in the belief propagation algorithm after t iterations in graph G are equivalent to those that would have been received by the root i in

⁵ The subgraph $M(t)$ is not necessarily a perfect b-matching of G but we will show that after $O(n)$ iterations it will be the minimum weight perfect b-matching.

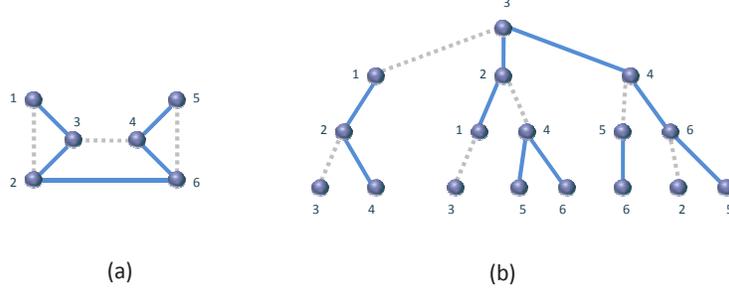


Figure 1: Part (a) shows a graph G where dashed and gray edges represent a 1-matching. Part (b) shows the computation tree T_3^2 corresponding to G where the set of dashed and gray edges form a 1-TMWPM.

the computation tree, if the messages were passed up along the tree from the leaves to the root. Computation trees have been used in most of the previous analyses of BP algorithms; see e.g. [20, 5, 52, 54, 53, 18].

A subtree \mathcal{M} of edges in the computation tree T_i^t is called a perfect *tree-b-matching* if for each *non-leaf* vertex with label i we have $\deg_{\mathcal{M}}(i) = b_i$. Now denote the minimum weight perfect tree-b-matching (b-TMWPM) of the computation tree T_i^t by $\mathcal{N}^*(T_i^t)$. We will show that Sync-BP can be seen as a dynamic programming procedure that finds the minimum weight perfect tree-b-matching over the computation tree. Figure 1 shows a graph G and one of its corresponding computation tree.

Sync-BP Equations. Consider the computation tree T_i^t . Let us assume that $\deg_G(i) = k$, and that i_1, \dots, i_k are neighbors of i in G which are children of the root i as well. Let us denote the subtree of T_i^t that consists of the root edge (i, i_j) and all descendants of i_j by $T_{i_j \rightarrow i}^t$. Given this, we define the following weights and weight differences:

$$\begin{aligned} W_{i_j \rightarrow i}^+(t) &= \text{Weight of b-TMWPM in } T_{i_j \rightarrow i}^t \text{ that contains the root edge } (i, i_j). \\ W_{i_j \rightarrow i}^-(t) &= \text{Weight of b-TMWPM in } T_{i_j \rightarrow i}^t \text{ that does not contain the root edge } (i, i_j). \\ n_{i_j \rightarrow i}(t) &= W_{i_j \rightarrow i}^+(t) - W_{i_j \rightarrow i}^-(t). \end{aligned}$$

Clearly, for any edge $\{i, j\}$ of graph G the real number $n_{j \rightarrow i}(t)$ is well-defined; the next lemma shows its relation with the messages passed in Sync-BP.

Lemma 1 *For all $1 \leq i, j \leq n$ such that $\{i, j\}$ is an edge of G and all $t = 0, 1, \dots$, the following is true: $n_{j \rightarrow i}(t) = m_{j \rightarrow i}(t)$.*

Proof We proceed by induction on t . For $t = 0$ by definition the computation tree T_i^0 has height 1. Therefore each branch $T_{i_j \rightarrow i}^0$ consists of a single root edge (i, i_j) . Thus $W_{i_j \rightarrow i}^+(0) = w_{ii_j}$ and $W_{i_j \rightarrow i}^-(0) = 0$ which gives: $n_{i_j \rightarrow i}(0) = w_{ii_j}$, and by definition this is equal to $m_{i_j \rightarrow i}(0)$. Now for the general case consider the computation tree T_i^t and fix a branch $T_{i_j \rightarrow i}^t$. Denote the children of i_j in this branch by j_1, \dots, j_ℓ with $\ell = \deg_G(i_j) - 1$ (by rule (c) from the construction of the computation tree described above). For simplicity of notation let $a = b_{i_j}$. Without loss of generality assume that the children j_1, \dots, j_ℓ are ordered so that

$$W_{j_1 \rightarrow i_j}^+(t-1) \leq W_{j_2 \rightarrow i_j}^+(t-1) \leq \dots \leq W_{j_\ell \rightarrow i_j}^+(t-1).$$

Now it is not hard to see that

$$\begin{aligned} W_{i_j \rightarrow i}^+(t) &= w_{ii_j} + \sum_{r=1}^{a-1} W_{j_r \rightarrow i_j}^+(t-1) + \sum_{r=a}^{\ell} W_{j_r \rightarrow i_j}^-(t-1) \\ W_{i_j \rightarrow i}^-(t) &= \sum_{r=1}^a W_{j_r \rightarrow i_j}^+(t-1) + \sum_{r=a+1}^{\ell} W_{j_r \rightarrow i_j}^-(t-1), \end{aligned}$$

so that

$$\begin{aligned}
n_{i_j \rightarrow i}(t) &= W_{i_j \rightarrow i}^+(t) - W_{i_j \rightarrow i}^-(t) \\
&= w_{ii_j} + W_{j_a \rightarrow i_j}^-(t-1) - W_{j_{a+1} \rightarrow i_j}^-(t-1) \\
&= w_{ii_j} - n_{j_a \rightarrow i_j}(t-1) \\
&= w_{ii_j} - a^{\text{th-min}}_{r \in N(i_j) \setminus \{i\}} \left(n_{j_r \rightarrow i_j}(t-1) \right).
\end{aligned}$$

Therefore we have shown that variables $n_{j \rightarrow i}(t)$ satisfy the same recursive relation as variables $m_{j \rightarrow i}(t)$, equation (2), and satisfy the same initial conditions. Thus they are equal. \blacksquare

It follows immediately from the above lemma that the set of edges $E_i(t)$ which is selected in iteration t of the algorithm Sync-BP consists of exactly the same edges which are adjacent to root i in $\mathcal{M}^*(T_i^t)$. This is formalized in the following corollary.

Corollary 1 *The algorithm Sync-BP solves the b-TMWPM problem on the computation tree. In particular, for each vertex i of G , the set of $E_i(t)$ which was chosen at the end of iteration t by Sync-BP is exactly the set of b_i edges which are attached to the root in b-TMWPM of T_i^t .*

Corollary 1 characterizes the estimated b-MWPM, $M(t)$, and will be used in the proof of the main result in Subsection 4.3. In the next subsection we present a very useful connection between graph G and its double-cover graph that is crucial for the proofs of Subsection 4.3.

4.2 Graph Covers

For a simple graph $G = (V, E)$, define the *double-cover* of G as a bipartite graph $\tilde{G} = (V_1 \cup V_2, E(\tilde{G}))$ where V_1, V_2 are exact copies of V (i.e. $v \in V$ if and only if $v_1 \in V_1, v_2 \in V_2$) and each edge of G has exactly two copies in \tilde{G} (i.e. $\{u, v\} \in E$ if and only if $\{u_1, v_2\} \in E(\tilde{G}), \{u_2, v_1\} \in E(\tilde{G})$). For each vertex (edge) of G there are exactly two vertices (edges) in \tilde{G} . Note that G and \tilde{G} have the same local structure that is for any vertex v in G the vertices v_1, v_2 in \tilde{G} have the same neighborhood as v . For weighted graphs we assign the weight w_{ij} of the edge $\{i, j\}$ to both edges $\{i_1, j_2\}, \{i_2, j_1\}$ of \tilde{G} . Figure 2 shows a graph G and its double-cover \tilde{G} . We can now write an analogues LP relaxation to (1) for \tilde{G} .

$$\begin{aligned}
&\min && \sum_{\{i,j\} \in E} (x_{i_1 j_2} + x_{i_2 j_1}) w_{ij} \\
&\text{subject to} && \sum_{j_2 \in N(i_1)} x_{i_1 j_2} = \sum_{j_1 \in N(i_2)} x_{i_2 j_1} = b_i \quad \forall i \\
&&& 0 \leq x_{i_r j_s} \leq 1 \quad \forall \{i_r, j_s\} \in E(\tilde{G})
\end{aligned} \tag{3}$$

Now we investigate the relationship between the LP relaxations given in (1) and (3). The following lemma characterizes this relationship.

Lemma 2 *Let $\mathbf{x}^* = (x_{ij}^*)$ and $\tilde{\mathbf{x}}^* = (x_{i_r j_s}^*)$ be the optimum solutions for the LP relaxations (1) and (3) respectively.*

- (a) $\sum_{\{i,j\} \in E} (x_{i_1 j_2}^* + x_{i_2 j_1}^*) w_{ij} = 2 \sum_{\{i,j\} \in E} x_{ij}^* w_{ij}$.
- (b) *If the optimum \mathbf{x}^* of LP (1) is integer and is unique then the optimum $\tilde{\mathbf{x}}^*$ of LP (3) is also integer and is unique. Moreover for all $\{i, j\} \in E$ the following holds $x_{i_1 j_2}^* = x_{i_2 j_1}^* = x_{ij}^*$. We call such optimum of LP (3), a symmetric integer solution.*
- (c) *If the optimum \mathbf{x}^* of LP (1) is integer and is fractional then there exists an optimum $\tilde{\mathbf{x}}^*$ of LP (3) which is integer and satisfies the following constraint. For all $\{i, j\} \in E$, either $x_{i_1 j_2}^* = x_{i_2 j_1}^* = x_{ij}^*$ holds or $x_{i_1 j_2}^* = 1 - x_{i_2 j_1}^*, x_{ij}^* = .5$ hold. We call such optimum of LP (3), a non-symmetric integer solution.*

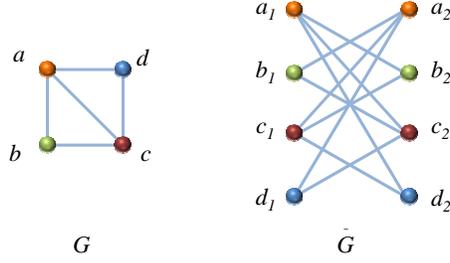


Figure 2: A graph G with four vertices and its double-cover \tilde{G} on the right.

Proof Starting from an optimum \mathbf{x}^* for the LP (1), it is easy to see that $f(\mathbf{x}^*) = (y_{i_r, j_s})$ where $y_{i_1 j_2} = y_{j_1 i_2} = x_{ij}^*$ is a feasible solution for the LP (3) and hence one side of the equality in (a) is proven as follows:

$$\sum_{\{i,j\} \in E} (x_{i_1 j_2}^* + x_{i_2 j_1}^*) w_{ij} \leq \sum_{\{i,j\} \in E} (y_{i_1 j_2} + y_{i_2 j_1}) w_{ij} = 2 \sum_{\{i,j\} \in E} x_{ij}^* w_{ij}.$$

On the other hand starting from an optimum $\tilde{\mathbf{x}}^*$ of the LP (3) one can define a feasible solution for the LP (1) by $g(\tilde{\mathbf{x}}^*) = (z_{ij})$ where $z_{ij} = (x_{i_1 j_2}^* + x_{i_2 j_1}^*)/2$. Now the other side of the equality in part (a) is proven by the following:

$$\sum_{\{i,j\} \in E} (x_{i_1 j_2}^* + x_{i_2 j_1}^*) w_{ij} = 2 \sum_{\{i,j\} \in E} z_{ij} w_{ij} \geq 2 \sum_{\{i,j\} \in E} x_{ij}^* w_{ij}.$$

As a corollary of this discussion one can see that the mappings f, g provide a correspondence (not necessarily one to one) between the optimums of LPs (1) and (3).

Now to prove (b), if the optimum \mathbf{x}^* of the LP (1) is integer then by the above discussion $f(\mathbf{x}^*)$ is an integer optimum of the LP (3). And since by assumption, \mathbf{x}^* is the unique optimum of (1) then $g(f(\mathbf{x}^*))$ which is an optimum of (1), is equal to \mathbf{x}^* . This proves (b).

For (c) we use the well-known fact that vertices of the b-matching polytope on bipartite graphs are all integer solutions [43]. So There exist an optimum $\tilde{\mathbf{x}}^*$ for (3) which is integer. Since the optimum of (1) is unique therefore it has to be equal to $g(\tilde{\mathbf{x}}^*)$. Thus $\mathbf{x}^* = g(\tilde{\mathbf{x}}^*)$. But $\tilde{\mathbf{x}}^*$ being integer means that for all $\{i, j\} \in E$ we either have $x_{i_1 j_2}^* = x_{i_2 j_1}^* = x_{ij}^*$ or $x_{i_1 j_2}^* = 1 - x_{i_2 j_1}^*$, $x_{ij}^* = .5$ which proves (c). ■

Note that the above analysis reproves a well-known fact about coordinates of the vertices of the polytope for LP (1). It shows that those coordinates are from the set $\{0, .5, 1\}$.

Next we will show that the algorithm Sync-BP for graphs G and \tilde{G} is the same. In particular both graphs G and \tilde{G} have similar local structure and therefore they have similar computation trees. The only difference is that the vertices with a fixed label i in the computation tree of G now have labels i_1 or i_2 in the computation tree of \tilde{G} depending on the parity of their distance from the root. More specifically the following lemma is straightforward.

Lemma 3 *There exists an isomorphism $\phi : T_i^t \rightarrow T_{i_1}^t$ between the computation trees of graphs G and \tilde{G} that preserves the roots and for any vertex v with label j in T_i^t , the label of $\phi(j)$ is j_1 (or j_2) if v has even (or odd) distance from the root of T_i^t . Similar isomorphism exists between T_i^t and $T_{i_2}^t$.*

Therefore the minimum weight perfect tree-b-matchings for T_i^t , $T_{i_1}^t$, and $T_{i_2}^t$ choose the same edges at the roots. This shows the following corollary.

Corollary 2 *The output of the algorithm Sync-BP for a vertex i of graph G is the same as the output of the algorithm Sync-BP, started with the same initial conditions, for vertices i_1 and i_2 of the double-cover \tilde{G} .*

The notion of graph-covers and its relation with LP and BP has been studied before in the context of coding theory [47], [46] and similar notions in the combinatorial optimization context exist in [43].

4.3 Proof of Theorem 1

We will prove Theorem 1, namely that if the LP relaxation (1) has no fractional solution and hence M^* is unique, then Sync-BP converges to the correct b-MWPM. We will do this by using Corollary 2 to reduce the problem to bipartite graphs and then use Lemma 2 and results of [5] and [24] to prove correctness of the BP. Alternatively, we provide a different proof of Theorem 1 which is also independent from the previous results [5] and [24] in the appendix. This second proof was also presented in the preliminary version of this paper [7].

Let us summarize results of [5] and [24] by the following theorem:

Theorem 2 *For any weighted bipartite graph with unique minimum weight b-matching (and gap ϵ), the algorithm Sync-BP, converges to the correct solution within $\lceil \frac{2nw^*}{\epsilon} \rceil$.*

Note that ϵ, w^* for bipartite graphs are defined the same way as in Section 3 for general graphs. Now we are ready to prove Theorem 1.

First note that since LP (1) has no fractional solution then by Lemma 2(a) the LP (3) has unique integer solution which is also symmetric. And therefore proofs given in [5] and [24] show correctness of Sync-BP in the double-cover \tilde{G} after $\lceil \frac{2nw^*}{\epsilon} \rceil$ iterations. Now using Corollary 2, we obtain correctness of Sync-BP for the original graph G after $\lceil \frac{2nw^*}{\epsilon} \rceil$ iterations. ■

4.4 Independence from Initial Conditions.

We would like to point out that changing the initial condition for the messages in step (2) of Sync-BP to any arbitrary values does not change the convergence and correctness of algorithm Sync-BP. The only effect of initial condition is on the number of iterations needed for convergence. Theorem 1 remains true by *re-defining* w^* according to: $w^* = \max_{\{i,j\} \in E} |w_{ij}| + \max_{\{i,j\} \in E} |m_{i \rightarrow j}(0)|$. This follows because, by changing the initial condition, the algorithm Sync-BP runs over a slightly modified computation tree. The new computation tree is almost the same computation tree as T_i^t , except that the leaf edges of the tree have arbitrary weights and not w_{ij} 's from G . As it appears in proof of Lemma 3 in [5], the only place where the weight of leaf edges appears is the inequality (16) of that paper, which will be satisfied by new definition of w^* (similarly in proof of Lemma 1.3 of [24]).

5 Extension to Possibly Non-Perfect b-Matchings

In this section we show that the algorithm and the results of the previous sections can be easily generalized to the case of b-matchings (subgraphs H of G such that degree of each vertex i in H is *at most* b_i). Let $U(H) \subset V$ be the set of *unsaturated* vertices of G (vertices $i \in V$ such that $deg_H(i) < b_i$). Similar to Section 2, the minimum weight b-Matching (b-MWM), H^* , is the b-Matching such that

$$H^* = \operatorname{argmax}_{H \in \mathcal{M}_G(b)} W_H.$$

Note that H^* does not include any edge with positive weight because removing such edges from H^* reduces its weight while keeping it a b-matching. Therefore in this section we assume that for all $\{i, j\} \in E$: $w_{ij} \leq 0$. The LP relaxation is slightly different from before:

$$\begin{aligned} & \min && \sum_{\{i,j\} \in E} x_{ij} w_{ij} \\ & \text{subject to} && \\ & && \sum_{j \in N(i)} x_{ij} \leq b_i \quad \forall i \\ & && 0 \leq x_{ij} \leq 1 \quad \forall \{i, j\} \in E \end{aligned} \tag{4}$$

Similar to Section 2 we define $\epsilon' > 0$ to be the difference between W_{H^*} and weight of the second best b-matching.

Now we can present the modified algorithm Sync-BP for finding b-MWM in G :

Algorithm Sync-BP(2).

- (1) At times $t = 0, 1, \dots$, each vertex sends real-valued messages to each of its neighbors. The message of i to j at time t is denoted by $m_{i \rightarrow j}(t)$.
- (2) Messages are initialized by $m_{i \rightarrow j}(0) = w_{ij}$ for all $\{i, j\} \in E$.
- (3) For $t \geq 1$, messages in iteration t are obtained from messages in iteration $t - 1$ recursively as follows:

$$\forall \{i, j\} \in E: \quad m_{i \rightarrow j}(t) = w_{ij} - \min \left(0, b_i^{t^{\text{th}}} - \min_{\ell \in N(i) \setminus \{j\}} \left[m_{\ell \rightarrow i}(t-1) \right] \right) \quad (5)$$

where $k^{\text{th}}\text{-min}(A)$ denotes the k^{th} minimum⁶ of set A .

- (4) The estimated b-MWM at the end of iteration t is $H(t) = \cup_{i=1}^n F_i(t)$ where $F_i(t) = \{\{i, j_1\}, \dots, \{i, j_{c_i}\}\}$ is such that $m_{j_\ell \rightarrow i}(t) < 0$ for all $1 \leq \ell \leq c_i$, i.e., choose edges that transfer negative messages to i .
 - (5) Repeat (3)-(4) until $H(t)$ converges.
-

The results for b-matchings generalize as follows:

Theorem 3 *Assume that the LP relaxation (10) has no fractional solution. Then the algorithm Sync-BP(2) converges to H^* after at most $\lceil \frac{4nw^*}{\epsilon'} \rceil$ iterations.*

The proof of Theorem 3 is similar to the one of Section 4 with the following modifications:

1. The computation tree T_i^t and b-TMWM are defined as before, while Lemma 1 is slightly modified. A careful analysis of W^+ and W^- for the tree-b-matchings yields equations (5) for finding b-TMWM in the computation tree. This is how the new equations are obtained.
2. The LP (3) that is defined on the double-cover \tilde{G} should be modified to address possibly non-perfect b-matching. Then Lemmas 2-4 and Corollary 2 hold for the modified LP.
3. The proof of Lemmas 2, 3 in [5] and Lemma 1.3 should be slightly modified. In particular, the alternating path that is constructed in those papers can be different: One can show that if the b-TMWM $\mathcal{N}^*(T_i^t)$ and the tree-b-matching \mathcal{H}^* choose different sets of edges at the root i , then an alternating path can be constructed in T_i^t (as in [5], [24]) which includes the root i . But endpoints of this alternating path are either leaves of T_i^t or vertices inside T_i^t which have labels from $U(H^*)$ (are un-saturated in G by H^*). In the case in which there is at least one leaf as an endpoint, the same argument as in [5], [24] can be used since length of the path grows with the depth of computation tree. But in the case in which both endpoints are non-leaf vertices of the computation tree one can show by switching the edges on the path, a better b-TMWM can be achieved. In particular using the proof of Lemma 3 in [5] we can partition the path to many simple alternating cycles and an alternating simple path. For the cycles, as in [5], switching the edges improves the weight by at least ϵ' . For the simple path, situation is better here compared to [5]. We can show switching yields a gain of at least ϵ' . This is because both endpoints of the path belong to $U(H^*)$ which means switching the edges of the path (in graph G) yields a valid b-matching which should have larger weight than W_{H^*} by at least ϵ' .

⁶Here $b_i^{t^{\text{th}}}\text{-min}_{\ell \in N(i) \setminus \{j\}}$ is defined to be 0 if $\text{deg}_G(i) = b_i$.

6 Analysis of the Asynchronous BP

In this section we study the asynchronous version of the BP algorithm. The update equations are exactly analogous to the synchronous version, but at each time only a subset of the edges are updated in an arbitrary order. Consider the set \vec{E} of all directed edges in the G ; i.e., $\vec{E} = \{(i \rightarrow j) \text{ s.t. } i \neq j \in V\}$. Let A be a sequence $\vec{E}(1), \vec{E}(2), \dots$ of subsets of the set \vec{E} . Then the asynchronous BP algorithm corresponding to the sequence A can be obtained by modifying only the step (3) of the algorithm Sync-BP for the perfect b-matchings:

(3) For $t \geq 1$, messages in iteration t are obtained from messages in iteration $t - 1$ recursively as follows:

$$m_{i \rightarrow j}(t) = \begin{cases} w_{ij} - b_i^{t^{\text{th}} - \min_{\ell \in N(i) \setminus \{j\}} \left[m_{\ell \rightarrow i}(t-1) \right]} & \text{if } (i \rightarrow j) \in \vec{E}(t) \\ m_{i \rightarrow j}(t-1) & \text{otherwise} \end{cases}$$

Note 3. This is the most general form of the asynchronous BP and it includes the synchronous version ($\vec{E}(t) = \vec{E}$ for all $t = 1, 2, \dots$) as a special case. In many applications, a special case of the asynchronous BP is used for which each set $\vec{E}(t)$ consists of a single element.

We assume that the sequence A of the updates does not have *redundancies*. That is, no edge direction $(i \rightarrow j) \in \vec{E}$ is re-updated before at least one of its incoming edge directions ($(\ell \rightarrow i)$ for $\ell \in N(i) \setminus \{j\}$) is updated. More formally, if $(i \rightarrow j) \in \vec{E}(t) \cap \vec{E}(t+s)$ and $(i \rightarrow j) \notin \cup_{r=1}^{s-1} \vec{E}(t+r)$, then at least for one $\ell \in N(i) \setminus \{j\}$, we should have $(\ell \rightarrow i) \in \cup_{r=1}^{s-1} \vec{E}(t+r)$.

Let us denote the above algorithm by Async-BP. We claim that, if each edge direction $(i \rightarrow j) \in \vec{E}$ is updated $\theta(n)$ times, then the same result as Theorem 1 can be proved here. That is, let $u(t)$ be the minimum number of times that an edge direction of the graph G appears in the sequence $\vec{E}(1), \dots, \vec{E}(t)$; i.e.,

$$u(t) = \min_{(i \rightarrow j) \in \vec{E}} \left(\left| \left\{ \ell : \text{ s.t. } 1 \leq \ell \leq t \text{ and } (i \rightarrow j) \in \vec{E}(\ell) \right\} \right| \right).$$

From the definition, $u(t)$ is a non-decreasing function of t . We claim that the following result holds:

Theorem 4 *Assume that the LP relaxation (1) has no fractional solution. Then the algorithm Async-BP converges to M^* after at most t iterations, provided $u(t) > \frac{2nw^*}{\epsilon}$.*

Before proving the above theorem let us define the notion of generalized computation tree for the asynchronous version of the BP algorithm.

6.1 Generalized Computation Tree for the Asynchronous BP

In order to define the generalized computation tree (GCT) for the asynchronous BP, we will begin with some definitions. For any $(i \rightarrow j) \in \vec{E}(t)$, define $R_{i \rightarrow j}^t$ to be the *computation branch* of i to j at time t which is a weighted rooted tree (not necessarily a balanced rooted tree) and recursively defined according to the following rules:

- (a) The root has label j .
- (b) The root has only one child which has label i .
- (c) If $t = 0$, then the child i has no child ($R_{i \rightarrow j}^0$ is just a single edge $\{i, j\}$).
- (d) For $t > 0$, if $(i \rightarrow j) \notin \vec{E}(t)$ then $R_{i \rightarrow j}^t = R_{i \rightarrow j}^{t-1}$. Otherwise the child i has $\deg_G(i) - 1$ children which have all of labels in the set $N(i) \setminus \{j\}$ and for any child r of i the subtree that consists of all descendants of r and the edge $\{r, i\}$ to i is $R_{r \rightarrow i}^{t-1}$.

The edge between nodes labeled i, j in the tree is assigned weight w_{ij} for $1 \leq i < j \leq n$. Now for any vertex $i \in V$ and any t , the GCT R_i^t is a weighted rooted tree with root i such that all its branches starting from the root are the computation branches $R_{r \rightarrow i}^t$ for all $r \in N(i)$. Since the GCT R_i^t is not necessarily balanced, we will define its *depth* to be the length of the shortest path from the root i to a leaf and denote it by $d(R_i^t)$.

Similarly to the Section 4.1, we can define the minimum weight perfect tree-b-matching (b-TMWPM) for GCT R_i^t and denote it by $\mathcal{M}^*(R_i^t)$. Moreover, arguments similar to the ones in the Section 4.1 show that the algorithm Async-BP is solving the b-TMWPM for GCTs R_i^t . In other words, the following corollary holds:

Corollary 3 *The algorithm Async-BP solves the b-TMWPM problem on the GCT. In particular, for each vertex i of the G , the set $E_i(t)$ that was chosen at the end of iteration t by Async-BP is exactly the set of b_i edges that are attached to the root in b-TMWPM of R_i^t .*

6.2 Technical Analysis of the Asynchronous BP

Now we can use the same analysis as in Section A.3 to show that if the depth of the generalized computation tree (GCT) is large enough, then for any vertex i , its neighbors in M^* (b-MWPM of G) are exactly those children that are selected in $\mathcal{N}^*(R_i^t)$ (b-TMWPM of R_i^t). We will show this by relating the function $u(t)$ to the depth of the GCT. Here is the main lemma which summarizes the above claim:

Lemma 4 *If the LP relaxation (1) has no fractional solution, then for any vertex i of G and for any t such that $u(t) > \frac{2nw^*}{\epsilon}$, the set of edges that are adjacent to root i in $\mathcal{N}^*(R_i^t)$ are exactly those edges that are connected to i in M^* .*

The proof of Lemma 4 is similar to the proof of Lemma 3 from [5] or Lemma 1.3 from [24], with the following slight modifications:

- (i) One can construct alternating paths P_ℓ in the same way as before for $1 \leq \ell \leq d(R_i^t)$.
- (ii) The depth of the GCT R_i^t is related to $u(t)$ according to the following lemma:

Lemma 5 *For any vertex $i \in V$ and any t , the depth of any computation branch at time t is at least $u(t)$; i.e., $d(R_{i \rightarrow j}^t) \geq u(t)$.*

This tells us $d(R_i^t) \geq u(t) > \frac{2nw^*}{\epsilon}$. Applying this to the path $P_{d(R_i^t)}$, analogously to the use of Lemma 3 in [5], gives us the proof of Lemma 4. Therefore all that is needed is a proof of Lemma 5.

Proof of Lemma 5 The proof follows easily by looking at the construction of the computation branch. Each computation branch $R_{i \rightarrow j}^t$ grows at time t if $(i \rightarrow j) \in \vec{E}(t)$. And if this is the case, the depth increases by at least one due to the “no redundancy condition” on the updating sequence. So if each edge is updated at least $u(t)$ times then the depth of its computation branch grows by at least $u(t)$. ■

Finally we note that the same algorithm as Async-BP and the same result as Theorem 4 can be stated and proved for the (possibly non-perfect) b-matchings as well.

7 Solving LP with BP

So far we have been proving correctness of the BP algorithm provided that the LP relaxation has no fractional solution. In this section we show that with a slight perturbation of the weights, one can solve the LP relaxation with BP. We will do this for all possibilities for the LP solutions (i.e. when LP has fractional solutions or when the integer optimum is non-unique). Note that we can first assume that the LP relaxation has a unique optimum which is a corner of its polytope. If that is not the case then we can add a small and

random perturbation to the weights which breaks the ties and forces the uniqueness of LP optimum with high probability. If we solve this new LP with BP then we have found an optimum of the original LP.

Now we will focus on the double cover \tilde{G} of graph G . We saw in Section 4 that if LP on G has no fractional solution then LP on \tilde{G} has a unique symmetric integer solution which can be found with BP. But if the LP on G has fractional solution then LP on \tilde{G} has many optimums including two non-symmetric integer solutions. We would like to perturb the weights, in order to break the tie and force the LP of \tilde{G} to have a unique non-symmetric integer solution with high probability. This can be done by adding a small and random perturbation to the edges of \tilde{G} (the edges $\{i_1, j_2\}$ and $\{i_2, j_1\}$ will be perturbed differently). Now that LP on \tilde{G} has unique integer solution we can find it with BP and then we can construct an LP optimum for G by assigning to x_{ij} values 0, .5, or 1 depending on whether none, one, or both of the edges $\{i_1, j_2\}$ and $\{i_2, j_1\}$ in \tilde{G} are part of the non-symmetric optimum found by BP.

The above discussion describes a often observed behavior of the BP algorithm in practice: BP converges to the correct values for the integer parts and oscillates for the fractional parts of the LP solution. This statement is correct if the above analysis was true for the original weights in the graph G without any perturbation. However one can construct counterexamples in which BP gives ambiguous answers when there is no perturbation. Although these special cases have very low probability of occurrence in practice.

8 Acknowledgements

We would like to thank László Lovász, Andrea Montanari, Elchannan Mossel and Amin Saberi for useful discussions. This work was done while Riccardo Zecchina was a Visiting Researcher in the Theory Group at Microsoft Research, and was supported by the Microsoft Technical Computing Initiative.

References

- [1] D. Achlioptas, F. Ricci-Tersenghi, “On the solution-space geometry of random constraint satisfaction problems,” *STOC* 2006.
- [2] D. Aldous, “The zeta (2) Limit in the Random Assignment Problem,” *Random Structures and Algorithms*, Vol. 18, pp. 381-418, 2001.
- [3] S. M. Aji, G. B. Horn and R. J. McEliece, “On the Convergence of Iterative Decoding on Graphs with a Single Cycle,” in *Proc. IEEE Int. Symp. Information Theory*, 1998, p. 276.
- [4] M. Bayati and C. Nair, “A rigorous proof of the cavity method for counting matchings”, *Allerton conference on communication, control and computing*, 2006.
- [5] M. Bayati, D. Shah, and M. Sharma, “Maximum weight matching via max-product belief propagation,” *Preliminary version appeared at IEEE ISIT 2005. Longer version to appear in IEEE Trans. Information Theory*, 2007.
- [6] M. Bayati, D. Shah, and M. Sharma, “Max-product for maximum weight matching: convergence, correctness and LP duality,” in *IEEE Int. Symp. Information Theory*, 2006.
- [7] M. Bayati, C. Borgs, J. Chayes and R. Zecchina, “Belief-Propagation for Weighted b-Matchings on Arbitrary Graphs and its Relation to Linear Programs with Integer Solutions”, in *arXiv*, <http://www.arxiv.org/abs/0709.1190v1>, September 8, 2007.
- [8] D. P. Bertsekas, “The auction algorithm: A distributed relaxation method for the assignment problem,” *Annals of Operations Research*, vol. 14, 1988.
- [9] D. P. Bertsekas, “Auction algorithms for network flow problems: A tutorial introduction,” *Computational Optimization and Applications*, vol. 1, pp. 7–66, 1992.

- [10] S. Boyd and L. Vandenberghe, “Convex Optimization”, *Cambridge University Press*, 2004.
- [11] A. Braunstein, M. Mezard, and R. Zecchina, “Survey propagation: an algorithm for satisfiability,” *Random Structures and Algorithms*, vol. 27, pp. 201–226, 2005.
- [12] A. Braunstein, R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, “Polynomial iterative algorithms for coloring and analyzing random graphs,” *Phys. Rev. E* 68, 036702, 2003.
- [13] J. Edmonds and E. Johnson, “Matching: A well-solved class of integer linear programmes,” *Combinatorial Structures and their Applications, Calgary International Conference, Gordon and Breach*, 89-92, 1970.
- [14] J. Edmonds and R. Karp, “Theoretical improvements in algorithmic efficiency for network flow problems,” *Journal of ACM*, vol. 18, pp. 264–284, 1972.
- [15] U. Feige, E. Mossel and D. Vilenchik, “Complete convergence of message passing algorithms for some satisfiability problems”, *In Proceedings of Random 2006*, LNCS 4110 Springer, 339–350, 2006.
- [16] J. Feldman, M. Wainwright and D. Karger, “Using linear programming to decode binary linear codes”, *IEEE Transactions on Information Theory*, vol. 51, pp. 954-972, 2005.
- [17] B. Frey and D. Dueck, “Clustering by Passing Messages Between Data Points”, *Science* 315, 972, 2007.
- [18] B.J. Frey, R. Koetter, “Exact inference using the attenuated max-product algorithm”, *Advanced Mean Field Methods: Theory and Practice*, ed. Manfred Opper and David Saad, MIT Press, 2000.
- [19] N. Friedman, “Inferring Cellular Networks Using Probabilistic Graphical Models,” *Science*; 303(5659):799-805, Feb 6 2004.
- [20] R. G. Gallager, “Low Density Parity Check Codes,” Cambridge, MA: MIT Press, 1963.
- [21] D. Gamarnik, T. Nowicki and G. Swirszcz, “Maximum Weight Independent Sets and Matchings in Sparse Random Graphs. Exact Results using the Local Weak Convergence Method”, *Random Structures and Algorithm*, Vol.28, No. 1, pp. 76-106, 2005.
- [22] I. Gat-Viks, A. Tanay, D. Rajjman and R. Shamir, “Factor graph network models for biological systems,” *Proc. of RECOMB 2005*, pp. 31-47, Lecture Notes in Bioinformatics 3500, Springer, Berlin, 2005.
- [23] A. Gerards, “Matching. Volume 7 of ,” *Hand book of Operation Research and Management Science*, Chapter 3, pp. 135-224. North-Holland, 1995.
- [24] B. Huang, T. Jebara, “Loopy belief propagation for bipartite maximum weight b-matching”, *Artificial Intelligence and Statistics (AISTATS)*, March, 2007.
- [25] E. Maneva, E. Mossel and M. J. Wainwright., “A New Look at Survey Propagation and its Generalizations,” *SODA*, 2005.
- [26] E. Marinari, G. Semerjian and V. Van Kerrebroeck, “Finding long cycles in graphs”, *Phys. Rev.*, 75, 066708, 2007.
- [27] M. Mezard and G. Parisi “Mean-field equations for the matching and travelling Salesman problems,” *Eurhophysics letters*, Vol. 2, pp. 913-918, 1986.
- [28] M. Mezard and R. Zecchina “Random K-satisfiability: from an analytic solution to a new efficient algorithm,” *Phys.Rev. E E*, 66, 056126, 2002.
- [29] C. Moallemi and B. Van Roy, “Consensus Propagation,” *IEEE Transactions on Information Theory*, Vol. 52, No. 11, pp. 4753-4766, 2006.
- [30] C. Moallemi and B. Van Roy, “Convergence of the Min-Sum Message Passing Algorithm for Quadratic Optimization,” preprint 2006.

- [31] C. Moallemi and B. Van Roy, “A Message-Passing Paradigm for Resource Allocation,” preprint June 2007.
- [32] C. Moallemi and B. Van Roy, “Convergence of the Min-Sum Algorithm for Convex Optimization,” preprint May 2007.
- [33] M. Mezard, A. Montanari, “Reconstruction on trees and spin glass transition,” *J. Stat. Phys.*, 124, 1317-1350, 2006.
- [34] A. Montanari, B. Prabhakar, and D. Tse., “Belief Propagation Based Multi-User Detection”, *Allerton Conference on Communication, Control, and Computing*, 2006.
- [35] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco, CA: Morgan Kaufmann, 1988.
- [36] W. Pulleyblank, “Matchings and extensions.” Volume 1 of *Handbook of Combinatorics*, Chapter 3, pp. 179-232, North Holland, 1995.
- [37] M. Pretti and M. Weigt, “Sudden emergence of q-regular subgraphs in random graphs”, *Europhys. Lett.*, 75 8, 2006.
- [38] T. Richardson and R. Urbanke, “The Capacity of Low-Density Parity Check Codes under Message-Passing Decoding,” *IEEE Trans. Info. Theory*, Vol. 47, pp 599-618, 2001.
- [39] P. Rasmussen and B. Van Roy, “ An Analysis of Belief Propagation on the Turbo Decoding Graph with Gaussian Densities,” *IEEE Transactions on Information Theory*, Vol. 47, No. 2, pp. 745-765, 2001.
- [40] S. Sanghavi, “Equivalence of LP Relaxation and Max-Product for Weighted Matching in General Graphs”, *IEEE Information Theory Workshop*, September 2007.
- [41] S. Sanghavi, D. Malioutov, A. Willsky “Linear programming analysis of loopy belief propagation for weighted matching”, to appear in *NIPS*, 2007.
- [42] S. Sanghavi, D. Shah and A. Willsky, “Message Passing for Max-weight Independent Set”, to appear in *NIPS* 2007.
- [43] A. Schrijver, “Combinatorial Optimization”, Springer-Verlag, 2003.
- [44] M. Tappen and W. Freemand, “Graph cuts and belief propagation for stereo, using identical MRF parameters”, *ICCV*, 2003.
- [45] S. Tatikonda and M. I. Jordan, “Loopy belief propagation and Gibbs measures,” *In D. Koller and A. Darwiche (Eds.), Uncertainty in Artificial Intelligence (UAI), Proceedings of the Eighteenth Conference*, 2002.
- [46] P.O. Vontobel and R. Koetter, “Graph-cover decoding and finite-length analysis of message-passing iterative decoding of LDPC codes,” to appear in *IEEE Trans. Inform. Theory*, <http://www.arxiv.org/abs/cs.IT/0512078>.
- [47] P.O. Vontobel and R. Koetter, “On the relationship between linear programming decoding and min-sum algorithm decoding,” *Proc. ISITA* 2004, Parma, Italy, pp. 991–996, Oct. 10-13, 2004.
- [48] J. Wastlund, “The traveling salesman problem in the stochastic mean field model”, <http://www.mai.liu.se/~jowas>, 2006.
- [49] M. Wainwright, T. Jaakkola, and A. Willsky, “Tree Consistency and Bounds on the Performance of the Max-Product Algorithm and its Generalizations”, *Statistics and Computing*, 14, 2004.
- [50] 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, 2005.

- [51] M. Wainwright, T. Jaakkola, and A. Willsky, “MAP estimation via agreement on trees: message-passing and linear programming”, *IEEE Transactions on Information Theory*, 51(11):3697:3711, 2005.
- [52] Y. Weiss, “Correctness of local probability propagation in graphical models with loops,” *Neural Comput.*, Vol. 12, pp. 1-42, 2000.
- [53] Y. Weiss and W. Freeman, “Correctness of belief propagation in Gaussian graphical models of arbitrary topology,” *Neural Comput.*, Vol. 13, Issue 10, pp 2173-2200, 2001.
- [54] Y. Weiss and W. T. Freeman, ”On the Optimality of Solutions of the Max–Product Belief–Propagation Algorithm in Arbitrary Graphs”, *IEEE Trans. Info. Theory*, 47: 2, 2001.
- [55] Y. Weiss, C. Yanover and T. Meltzer “MAP Estimation, Linear Programming and Belief Propagation with Convex Free Energies,” *UAI*, 2007.
- [56] N. Wiberg, “Codes and Decoding on General Graphs”, *Ph.D. thesis*, Linköping University, Sweden, 1996.
- [57] C. Yanover and Y. Weiss, “Approximate inference and protein folding,” *Advances in Neural Processing Systems*, 2002.
- [58] J. Yedidia, W. Freeman and Y. Weiss, “Understanding Belief Propagation and its Generalizations,” Mitsubishi Elect. Res. Lab., TR-2001-22, 2000.
- [59] L. Zdeborová and M. Mézard, “The number of matchings in random graphs”, *J. Stat. Mech.*, 2006.

Appendix

A Alternative analysis of the algorithms in Sections 3-6

For the proofs of this section first we define the dual of LP relaxation (1).

$$\begin{array}{ccc|ccc}
 \min & \sum_{\{i,j\} \in E} x_{ij} w_{ij} & & \max & \sum_{i=1}^n b_i y_i - \sum_{\{i,j\} \in E} \lambda_{ij} & \\
 \text{subject to} & & & \text{subject to} & & \\
 & \sum_{j \in N(i)} x_{ij} = b_i & \forall i & & w_{ij} + \lambda_{ij} \geq y_i + y_j & \forall \{i, j\} \in E \\
 & 0 \leq x_{ij} \leq 1 & \forall \{i, j\} \in E & & \lambda_{ij} \geq 0 & \forall \{i, j\} \in E \\
 & \text{Primal LP} & & & \text{Dual LP} & \\
 & & & & & (6)
 \end{array}$$

Next, we need to define the complementary slackness conditions.

A.1 Complementary Slackness Conditions.

Complementary slackness conditions for the LP and its dual state that the variables $x^* = (x_{ij}^*)$ and $y^* = (y_i^*)$, $\lambda^* = (\lambda_{ij}^*)$ are optimum solutions to the LP relaxation and its dual (6), respectively, if and only if for all edges $\{i, j\}$ of G both $x_{ij}^*(w_{ij} + \lambda_{ij}^* - y_i^* - y_j^*) = 0$. and $(x_{ij}^* - 1)\lambda_{ij}^* = 0$ hold. See [10], [43] for more information about LP, dual LP and complementary slackness conditions.

Using the fact that the LP has no fractional solution, one can deduce the following modified complementary slackness conditions: For all $\{i, j\} \in M^*$; $w_{ij} + \lambda_{ij}^* = y_i^* + y_j^*$ and for all $\{i, j\} \notin M^*$; $\lambda_{ij}^* = 0$.

By these conditions and the fact that $\lambda_{ij}^* \geq 0$, we have that $w_{ij} \leq y_i^* + y_j^*$ for all $\{i, j\} \in M^*$, and $w_{ij} \geq y_i^* + y_j^*$ for all $\{i, j\} \notin M^*$. However, it is in general not true that these inequalities are strict even when the LP has no fractional solution. Let S be the set of those edges in G for which $|w_{ij} - y_i^* - y_j^*| > 0$. We will assume the minimum gap is ϵ'' . i.e. $\epsilon'' = \min_{\{i,j\} \in S} |w_{ij} - y_i^* - y_j^*| > 0$. Throughout this paper we assume that there exist an edge in G for which the strict inequality $|w_{ij} - y_i^* - y_j^*| > 0$ holds and therefore $\epsilon'' > 0$ is well defined. The other cases, where for each $\{i, j\} \in E$ the equality $w_{ij} = y_i^* + y_j^*$ holds, happen only for special cases and are discussed in Section A.4. Let also $L = \max_{1 \leq i \leq n} |y_i^*|$.

Now we are ready to state the analogous version of Theorem 1 and its proof.

Theorem 5 *Assume that the LP relaxation (1) has no fractional solution. Then the algorithm Sync-BP converges to M^* after at most $\lceil \frac{2nL}{\epsilon''} \rceil$ iterations.*

A crucial part of all the proofs in the appendix is the following lemma.

A.2 Main Technical Lemma

In this section we state our main technical lemma which connects the complementary slackness conditions from Section A.1 to paths on the graph G and on the computation tree. This lemma is a key step in our proof. Its proof is quite delicate, and provides the connection between the absence of fractional solutions and the correctness of BP.

Definition A path $P = (i_1, i_2, \dots, i_k)$ in G is called *alternating path* if:

- (a) There exist a partition of edges of P into two sets A, B such that either $(A \subset M^*, B \cap M^* = \emptyset)$ or $(A \cap M^* = \emptyset, B \subset M^*)$. Moreover A (B) consists of all odd (even) edges; i.e. $A = \{(i_1, i_2), (i_3, i_4), \dots\}$ ($B = \{(i_2, i_3), (i_4, i_5), \dots\}$).
- (b) The path P might intersect itself or even repeat its own edges but no edge is repeated immediately. That is, for any $1 \leq r \leq k - 2$: $i_r \neq i_{r+1}$ and $i_r \neq i_{r+2}$.

P is called an *alternating cycle* if $i_1 = i_k$.

Lemma 6 *Assume that the LP relaxation (1) has no fractional solution. Then for any alternating path P of length at least $2n$, there exists an edge $\{i, j\} \in P$ such that the inequality $|w_{ij} - y_i^* - y_j^*| > 0$ holds. That is, $P \cap S \neq \emptyset$.*

Proof We will consider two cases:

Case I) Existence of an even simple cycle in P .

Consider the subgraph of G that is generated by edges and vertices of P . If this subgraph contains an alternating cycle C that does not intersect itself (simple cycle) and has even length, then we will show that $C \cap S \neq \emptyset$. Let $C = (j_1, \dots, j_{2\ell}, j_1)$. Without loss of generality assume that odd edges belong to M^* and even edges do not. That is, for all $1 \leq r \leq \ell$:

$$\{i_{2r-1}, i_{2r}\} \in M^* \quad , \quad \{i_{2r}, i_{2r+1}\} \notin M^*$$

where $j_{2\ell+1} = j_1$. To prove $C \cap S \neq \emptyset$, assume the contrary; that is, assume for all edges $\{i, j\}$ of C : $w_{ij} = y_i^* + y_j^*$. The weight of M^* -edges of C is equal to weight of their complement in C , due to the fact that

$$\sum_{r=1}^{\ell} w_{j_{2r}j_{2r+1}} = \sum_{s=1}^{2\ell} y_r^* = \sum_{r=1}^{\ell} w_{j_{2r-1}j_{2r}}.$$

Now one can obtain a perfect b-matching M' in G which is different from M^* and has the same weight as M^* . This can be done by defining $M' = M^*$ outside cycle C , and $M' = C \setminus M^*$ on cycle C . However, this contradicts the uniqueness assumption for b-MWPM in G which holds due to the fact that the LP relaxation has no fractional solution. Hence we are done.

Case II) There is no even simple cycle in P .

Let $P = \{i_1, i_2, \dots, i_k\}$. Since P has length at least $2n$, it must repeat a vertex. We also add a natural direction to each edge $\{i_j, i_{j+1}\}$ that is from i_j to i_{j+1} . Consider the first vertex that is revisited by starting from i_1 and walking along P . That is, consider the smallest numbers r, s such that $1 \leq r < s \leq n + 1$ and $i_r = i_s$. Now we break P into three connected pieces as follows:

- (i) Simple path $P_0 = (i_1, i_2, \dots, i_r)$ (this part will be ignored).
- (ii) Simple cycle $C_1 = (i_r, i_{r+1}, \dots, i_s)$.
- (iii) Path $P_1 = (i_{s+1}, i_{s+2}, \dots, i_k)$.

From now on we are going to assume that path P_0 does not even exist. Basically we will show that there is one edge from S which is in $C_1 \cup P_1$. Since we assumed that P has no even simple cycle, it follows that C_1 has odd length ($s - r$ is odd). Since the length of P is at least $2n$, it follows that P has to intersect itself at least twice and there must be another vertex that is revisited after i_r . Consider the smallest numbers r', s' such that $r \leq r' < s' \leq k$ and $i_{r'} = i_{s'}$. Denote this new simple cycle by C_2 ; i.e., $C_2 = (i_{r'}, i_{r'+1}, \dots, i_{s'})$. Again since C_2 is an alternating path, it has to have odd length ($s' - r'$ is odd).

Now we claim that $s \leq r'$. Again assume the contrary, that $r < r' < s$. We obtain a contradiction by finding an even simple cycle in P . Break path C_1 in two simple paths $Q_1 = (i_r, i_{r+1}, \dots, i_{r'})$ and $Q_2 = (i_{r'}, i_{r'+1}, \dots, i_s)$, and define the simple path $Q_3 = (i_s, i_{s+1}, \dots, i_{s'})$. Now consider the simple cycle $C_3 = Q_1 \cup Q_3$. The length of C_3 is equal to $r' - r + s' - s$, which has the same parity as $s - r + s' - r'$, which is even. Therefore C_3 is an even cycle. Moreover, the fact that the parities of r' and s' are different guarantees the alternation of adjacent edges $\{i_{r'-1}, i_{r'}\}$ and $\{i_{s'-1}, i_{s'}\}$ in cycle C_3 . Similarly the difference in parity between r and s implies alternation of adjacent edges $\{i_r, i_{r+1}\}$ and $\{i_s, i_{s+1}\}$ in cycle C_3 . Thus C_3 is an even length alternating simple cycle, which is a contradiction. So the claim $s \leq r'$ is proved.

Now we are left with a final possibility which uses the integrality of the LP optimum solution. Consider the following three pieces of path P :

- (i) Simple odd cycle C_1 .
- (ii) Simple path $P_2 = (i_{s+1}, i_{s+2}, \dots, i_{r'})$ (could be only a point).
- (iii) Simple odd cycle C_2 .

If $(C_1 \cup P_2 \cup C_2) \cap S = \emptyset$, this means that for all edges $\{i, j\} \in C_1 \cup P_2 \cup C_2$, the equality $w_{ij} = y_i^* + y_j^*$ holds. We will reach a contradiction by showing the existence of an optimum fractional solution for LP relaxation (1). This is done by defining x' as follows:

$$\forall \{i, j\} \in E : \quad x'_{ij} = \begin{cases} x_{ij}^* & \text{if } \{i, j\} \notin C_1 \cup P_2 \cup C_2 \\ 1 - x_{ij}^* & \text{if } \{i, j\} \in P_2 \\ 0.5 & \text{if } \{i, j\} \in C_1 \cap C_2. \end{cases}$$

First we need to show that x' is a feasible solutions for the LP. For this, all we need to show is that x' satisfies the same local constraints as x^* on vertices of $C_1 \cup P_2 \cup C_2$. Since all $C_1 \cup P_2 \cup C_2$ is a connected alternating path, then for all vertices $i_\ell \in C_1 \cup P_2 \cup C_2$ ($\ell \notin \{r, s, r', s'\}$) it is clear that $x'_{(\ell-1)\ell} + x'_{\ell(\ell+1)} = x_{(\ell-1)\ell}^* + x_{\ell(\ell+1)}^* = 1$. For $\ell = r$, using the fact that length of C_1 is odd and path $C_1 \cup P_2$ is an alternating sub-path of P , either $x_{r(r+1)}^* = x_{(s-1)s}^* = 1$, $x_{s(s+1)}^* = 0$ or $x_{r(r+1)}^* = x_{(s-1)s}^* = 0$, $x_{s(s+1)}^* = 1$, which leads to $x'_{r(r+1)} = x'_{(s-1)s} = 0.5$, $x'_{s(s+1)} = 1$ or $x'_{r(r+1)} = x'_{(s-1)s} = 0.5$, $x'_{s(s+1)} = 0$, respectively. In both cases, x' satisfies same local constraint as x^* at i_r . A similar argument holds at $i_{r'}$.

Next we show that x' has the same cost as x^* . This is done by applying the equality $w_{ij} = y_i^* + y_j^*$ to all edges of $C_1 \cup P_2 \cup C_2$ as follows:

$$\begin{aligned}
\sum_{\{i,j\} \in C_1 \cup P_2 \cup C_2} w_{ij} x_{ij}^* &= \sum_{i \in C_1 \cup P_2 \cup C_2} y_i^* + x_{i_r, i_{r+1}}^* y_{i_r}^* + x_{i_{r'}, i_{r'+1}}^* y_{i_{r'}}^* \\
&= \sum_{i \in C_1 \cup C_2} y_i^* + \sum_{\{i,j\} \in P_2} w_{ij} x'_{ij} \\
&= \sum_{\{i,j\} \in C_1 \cup C_2 \cup P_2} w_{ij} x'_{ij}.
\end{aligned}$$

This completes the proof of Lemma 6. ■

A.3 Proof of Theorem 5

We will prove Theorem 5, using a similar argument to the proof in Section 4.3. In other words we show that if the LP relaxation (1) has no fractional solution and hence M^* is unique, then Sync-BP converges to the correct b-MWPM. We will do this by showing that if the depth of computation tree is large enough, then for any vertex i , its neighbors in M^* (b-MWPM of G) are exactly those children that are selected in $\mathcal{N}^*(T_i^t)$ (b-TMWPM of T_i^t). Here is the main lemma that summarizes the above claim:

Lemma 7 *If the LP relaxation (1) has no fractional solution, then for any vertex i of G and for any $t > \frac{2nL}{\epsilon r}$, the set of edges that are adjacent to root i in $\mathcal{N}^*(T_i^t)$ are exactly those edges that are connected to i in M^* .*

Before entering into the details of the proof here is a high level overview of the underlying argument. Consider the computation tree (T_i^t) rooted at vertex i and look at $\mathcal{N}^*(T_i^t)$. We will assume that the claim of the lemma does not hold. That is, we assume that at the root, $\mathcal{N}^*(T_i^t)$ does not choose the same edges as M^* -edges adjacent to i . Then we use the property of perfect tree-b-matchings, namely that each non-leaf vertex j is connected to exactly b_j of its neighbors, to construct a new perfect tree-b-matching on the computation tree. This new perfect tree-b-matching is going to have less total weight if the depth of the computation tree is large enough. This last step uses an alternating path argument which is a highly non-trivial generalization of the technique of [5] for the case of perfect 1-matching in bipartite graphs. For this part we will use the solutions to the dual LP (6).

Proof of Lemma 7 Let us denote the lifting of a perfect b-matching M^* to a perfect tree-b-matching on T_i^t by \mathcal{M}^* . That is, \mathcal{M}^* consists of all edge of the computation tree with endpoint labels i, j such that $\{i, j\} \in M^*$ as an edge in G . The goal is to show that $\mathcal{N}^*(T_i^t)$ and \mathcal{M}^* have the same set of edges at the root of the computation tree. To lighten the notation, we denote the b-TMWPM of T_i^t by \mathcal{N}^* .

Assume the contrary, that there exist children i_{-1}, i_1 of root i such that $\{i, i_1\} \in \mathcal{M}^* \setminus \mathcal{N}^*$ and $\{i, i_{-1}\} \in \mathcal{N}^* \setminus \mathcal{M}^*$. Since both $\mathcal{M}^*, \mathcal{N}^*$ are perfect tree-b-matchings, they have b_{i_1} edges connected to i_1 . Therefore there exist a child i_2 of i_1 such that $\{i_1, i_2\} \in \mathcal{N}^* \setminus \mathcal{M}^*$. Similarly there is a child i_{-2} of i_{-1} such that $\{i_{-1}, i_{-2}\} \in \mathcal{M}^* \setminus \mathcal{N}^*$. Therefore we can construct a set of alternating paths P_ℓ , $\ell \geq 0$, in the computation tree, that contain edges from \mathcal{M}^* and \mathcal{N}^* alternatively defined as follows. Let $i_0 = \text{root } i$ and $P_0 = (i_0)$ be a single vertex path. Let $P_1 = (i_{-1}, i_0, i_1)$, $P_2 = (i_{-2}, i_{-1}, i_0, i_1, i_2)$ and similarly for $r \geq 1$, define P_{2r+1} and P_{2r+2} recursively as follows:

$$P_{2r+1} = (i_{-(2r+1)}, P_{2r}, i_{2r+1}) \quad , \quad P_{2r+2} = (i_{-(2r+2)}, P_{2r+1}, i_{2r+2})$$

where $i_{-(2r+1)}, i_{2r+1}$ are nodes at level $2r+1$ such that $\{i_{2r}, i_{2r+1}\} \in \mathcal{M}^* \setminus \mathcal{N}^*$ and $\{i_{-2r}, i_{-(2r+1)}\} \in \mathcal{N}^* \setminus \mathcal{M}^*$. Similarly $i_{-(2r+2)}, i_{2r+2}$ are nodes at level $2r+2$ such that $\{i_{2r+1}, i_{2r+2}\} \in \mathcal{N}^* \setminus \mathcal{M}^*$ and $\{i_{-(2r+1)}, i_{-(2r+2)}\} \in \mathcal{M}^* \setminus \mathcal{N}^*$. Note that, by definition, such paths P_ℓ for $0 \leq \ell \leq t$ exist since the tree T_i^t has $t+1$ levels and can support a path of length at most $2t$ as defined above. Now consider the path P_t of length $2t$. It is an alternating path on the computation tree with edges from \mathcal{M}^* and \mathcal{N}^* . Let us refer to the edges of \mathcal{M}^* (\mathcal{N}^*) as the \mathcal{M}^* -edges (\mathcal{N}^* -edges) of P_t .

We will now modify the perfect tree-b-matching \mathcal{N}^* by replacing all \mathcal{N}^* -edges of P_t with their complement in P_t (\mathcal{M}^* -edges of P_t). It is straightforward that this process produces a new perfect tree-b-matching \mathcal{N}' in T_i^t .

Let us assume, for the moment, the following lemma:

Lemma 8 *The weight of the perfect tree-b-matching \mathcal{N}' is strictly less than that of \mathcal{N}^* on T_i^t .*

This completes the proof of Lemma 7 since Lemma 8 shows that \mathcal{N}^* is not the minimum weight perfect tree-b-matching on T_i^t , leading to a contradiction. \blacksquare

Now, we provide the proof of Lemma 8.

Proof of Lemma 8 It suffices to show that the total weight of the \mathcal{N}^* -edges of P_t is more than the total weight of \mathcal{M}^* -edges of P_t . For each vertex $i_r \in P_t$ consider the value $y_{i_r}^*$ from the optimum solution to the dual LP (6). Using the inequality $w_{ij} \leq y_i^* + y_j^*$ for edges of \mathcal{M}^* , we obtain:

$$\sum_{\{i,j\} \in P_t \cap \mathcal{M}^*} w_{ij} \leq \left(\sum_{r=-t}^t y_{i_r}^* \right) - y_{i_{(-1)t}}^* - k_1 \epsilon'' \quad (7)$$

where k_1 is the number of \mathcal{M}^* -edges of P_t that belong to S , i.e., the number of \mathcal{M}^* -edges of P_t endowed with the strict inequality $w_{ij} \leq y_i^* + y_j^*$, with a gap of at least ϵ'' . On the other hand, using the inequality $w_{ij} \geq y_i^* + y_j^*$ for edges of \mathcal{N}^* we have:

$$\sum_{\{i,j\} \in P_t \cap \mathcal{N}^*} w_{ij} \geq \left(\sum_{r=-t}^t y_{i_r}^* \right) - y_{i_{(-1)t+1}}^* + k_2 \epsilon'' \quad (8)$$

where now k_2 is number of \mathcal{N}^* -edges of P_t that belong to S , or equivalently the number of times the inequality $w_{ij} \geq y_i^* + y_j^*$ is strict with a gap of at least ϵ'' . One finds

$$\begin{aligned} \sum_{\{i,j\} \in P_t \cap \mathcal{N}^*} w_{ij} - \sum_{\{i,j\} \in P_t \cap \mathcal{M}^*} w_{ij} &= y_{i_{(-1)t}}^* - y_{i_{(-1)t+1}}^* + (k_1 + k_2) \epsilon'' \\ &\stackrel{(a)}{\geq} (k_1 + k_2) \epsilon'' - 2L \stackrel{(b)}{\geq} (k_1 + k_2) \epsilon'' - 2L \stackrel{(c)}{>} 0 \end{aligned} \quad (9)$$

where (a) uses definition of L from Section A.1 and (b) uses the fact that for all i, j : $\lambda_{ij}^* \geq 0$. The main step is (c), which uses Lemma 6 as follows. Path P_t has length $2t$, and each continuous piece of it with length $2n$ has a projection to the graph G which satisfies the conditions of Lemma 6. This means the path has at least one edge from the set S . Thus $(k_1 + k_2) \geq \frac{2t}{2n} > \frac{2L}{\epsilon''}$. This completes the proof of Lemma 8. \blacksquare

Note 4. The proof of Section 4.4 on the independence of Sync-BP's convergence and correctness from the initial conditions carries over here as well. This is done by *re-defining* L according to: $L = \max_{1 \leq i \leq n} |y_i^*| + \max_{\{i,j\} \in E} |m_{i \rightarrow j}(0)|$ and noting that, in the proof of Lemma 7, the only place where the weight of leaf edges appears is the inequality (a) in equation (9), which will be satisfied by new definition of L .

A.4 Sync-BP is Correct When ϵ'' is Not Well-Defined

Recall from the discussion given above about the complementary slackness conditions that, if for all edges $\{i, j\} \in E$ the equality $w_{ij} = y_i^* + y_j^*$ holds, then ϵ'' is not well defined. In this section we show that these rare cases do not cause any trouble. We will show that the condition $t > \frac{2nL}{\epsilon''}$ in the main theorem can be replaced by $t > n$. This is shown by proving the following lemma instead of Lemma 7.

Lemma 9 *If the LP relaxation (1) has no fractional solution, then, for any vertex i of G and for any $t > n$, the set of edges that are adjacent to root i in $\mathcal{N}^*(T_i^t)$ are exactly those edges that are connected to i in M^* .*

Proof The proof is similar to the proof of Lemma 7. If after iteration t , the claim of the Lemma 9 does not hold, then the alternating path P_t can be constructed as before. Now since the length of P_t is greater than $2n$, one can use the technical Lemma 6 for the projection of the path P_t onto G to show that the strict inequality $|w_{ij} - y_i^* - y_j^*| > 0$ happens for at least one edge. This contradicts the above assumption at the beginning of the Section. Therefore Lemma 9 is true. \blacksquare

A.5 Modifications for the possibly non-perfect matchings

Similar to the perfect matching case the LP relaxation and its dual are:

$$\begin{array}{ccc|ccc}
\min & \sum_{\{i,j\} \in E} x_{ij} w_{ij} & & \max & \sum_{i=1}^n -b_i y_i - \sum_{\{i,j\} \in E} \lambda_{ij} & \\
\text{subject to} & & & \text{subject to} & & \\
& \sum_{j \in N(i)} x_{ij} \leq b_i & \forall i & & w_{ij} + \lambda_{ij} \geq -y_i - y_j & \forall \{i,j\} \in E \\
& 0 \leq x_{ij} \leq 1 & \forall \{i,j\} \in E & & \lambda_{ij} \geq 0 & \forall \{i,j\} \in E \\
& \text{Primal LP} & & & \text{Dual LP.} & \\
& & & & & (10)
\end{array}$$

Complementary slackness now reads, for all $\{i,j\} \in E$: $x_{ij}^*(w_{ij} + \lambda_{ij}^* + y_i^* + y_j^*) = 0$, $(x_{ij}^* - 1)\lambda_{ij}^* = 0$ and for all $i \in V$: $(\sum_{j \in N(i)} x_{ij} - b_i)y_i^* = 0$.

Similarly to the perfect matching case, we can write the following modified complementary slackness condition using the fact that the LP relaxation has no fractional solution:

- (CS'-i) For all $\{i,j\} \in H^*$; $w_{ij} + \lambda_{ij}^* + y_i^* + y_j^* = 0$.
- (CS'-ii) For all $\{i,j\} \notin H^*$; $\lambda_{ij}^* = 0$.
- (CS'-iii) For all $i \in U(H^*)$; $y_i^* = 0$.

Let S' be set of those edges in G for which $|w_{ij} + y_i^* + y_j^*| > 0$. We will assume the minimum gap is ϵ''' . That is

$$0 < \epsilon''' = \min_{\{i,j\} \in S'} |w_{ij} + y_i^* + y_j^*|.$$

The quantity L' is defined similarly to L by $L' = \max_{1 \leq i \leq n} |y_i^*|$.

The algorithm Sync-BP(2) will remain unchanged and the modified theorem for its convergence and correctness is:

Theorem 6 *Assume that the LP relaxation (10) has no fractional solution. Then the algorithm Sync-BP(2) converges to H^* after at most $\lceil \frac{4nL'}{\epsilon'''} \rceil$ iterations.*

The proof of Theorem 6 is similar to the one of Section A.3, with the following modifications:

1. The computation tree T_i^t and b-TMWm are defined as before, while Lemma 1 is slightly modified. A careful analysis of W^+ and W^- for the tree-b-matchings yields equations (5) for finding b-TMWm in the computation tree. This is how the new equations are obtained.
2. The technical lemma from Section A.2 is still true and its proof does not change because the definition of alternating paths is preserved and because all cycles involved in the proof turn out to be adjacent to exactly one edge of H^* .

3. The proof of Lemmas 7 and 8 should be slightly modified. In particular, the alternating path P_t can be different: One can show that if the b-TMWM $\mathcal{N}^*(T_i^t)$ and the tree-b-matching \mathcal{H}^* choose different sets of edges at the root i , then an alternating path can be constructed as before in T_i^t which includes the root i . But endpoints of this alternating path P_t are either leaves of T_i^t or vertices inside T_i^t which have labels from $U(H^*)$ (are un-saturated in G by H^*). In the case in which there is at least one leaf as an endpoint of P_t , the same argument as equation (9) in Section A.3 can be used since length of P_t is at least t . This shows $(k_1 + k_2) \geq \frac{t}{2n} > \frac{2L'}{\epsilon''}$. But in the case in which both endpoints of P_t are non-leaf vertices of the computation tree, then using condition (CS'-iii), the analogous version of equation (9) is as follows:

$$\sum_{\{i,j\} \in P_t \cap \mathcal{N}^*} w_{ij} - \sum_{\{i,j\} \in P_t \cap \mathcal{H}^*} w_{ij} = (k_1 + k_2)\epsilon'''. \quad (11)$$

Now all that is needed is to show $k_1 + k_2 > 0$. We will show this by the following extension of technical Lemma 6.

Lemma 10 *Assume that the LP relaxation (10) has no fractional solution. Then for any alternating path P with endpoints from the set $U(H^*)$, there exists an edge $\{i, j\} \in P$ such that the inequality $|w_{ij} + y_i^* + y_j^*| > 0$ holds. That is, $P \cap S' \neq \emptyset$.*

Proof For paths P with length at least $2n$, we can use Lemma 6, so there is nothing to do. If a subgraph generated by P includes at least two cycles, then the same argument as in the proof of Lemma 6 can be used. Therefore we can assume P intersects itself at most once. So P can be written as a union $C \cup P_1$ where C is an odd simple alternating cycle and P_1 is a simple alternating path (either C or P_1 can be empty, but not at the same time). Next, one can define a different solution x' to the LP (10) which has the same cost as x^* by defining $x' = 1 - x^*$ on path P_1 and setting x' equal to 0.5 on C . x' will still be a feasible solution since the endpoints of P_1 are elements of $U(H^*)$ and the edge adjacent to them in path P_1 is not in H^* . This contradicts the no fractional solution assumption on the LP. ■

A.6 Modifications for the Asynchronous BP

The algorithm Async-BP and its analysis will be exactly similar to Section 6 except that the main theorem is slightly modified to:

Theorem 7 *Assume that the LP relaxation (1) has no fractional solution. Then the algorithm Async-BP converges to M^* after at most t iterations, provided $u(t) > \frac{2nL}{\epsilon''}$.*

And all the references to Lemma 3 of [24] should be replaced to Lemma 7 of this paper.