

# Running MAP Inference on Million Node Graphical Models: A High Performance Computing Perspective

Chen Jin<sup>†</sup>, Qiang Fu<sup>#</sup>, Huahua Wang<sup>#</sup>, William Hendrix<sup>†</sup>, Zhengzhang Chen<sup>†</sup>, Ankit Agrawal<sup>†</sup>,  
Arindam Banerjee<sup>#</sup>, Alok Choudhary<sup>†</sup>

<sup>#</sup>University of Minnesota, Twin Cities

<sup>†</sup>Northwestern University

Email:<sup>†</sup>{chen.jin, whendrix, zzc472, ankitag, choudhar}@eecs.northwestern.edu

Email:<sup>#</sup>{qifu, huwang, banerjee}@cs.umn.edu

*Abstract—*

An important problem in discrete graphical models is the maximum a posteriori (MAP) inference problem. Recent research has been focusing on the development of parallel MAP inference algorithm, which scales to graphical models of millions of nodes. In this paper, we introduce a parallel implementation of the recently proposed Bethe-ADMM algorithm using Message Passing Interface (MPI), which allows us to fully utilize the computing power provided by the modern supercomputers with thousands of cores. Experimental results demonstrate that for a broad class of problems, our parallel implementation of Bethe-ADMM scales almost linearly even with thousands of cores.

*Keywords*-Alternating Direction Method of Multipliers; Markov Random Field; Maximum a Posteriori Inference; Message Passing Interface

## I. INTRODUCTION

Discrete graphical models have found applications in a wide variety of problems, including image analysis [13], speech recognition [14], bioinformatics [9] and error correcting codes. Hidden Markov models (HMMs) [5], Markov random fields (MRFs) [29], and conditional random fields (CRFs) are popular examples of discrete graphical models which have found widespread usage in data analysis. Given a discrete graphical model with known structure and parameters, the problem of finding the most likely configuration of the states is known as the *maximum a posteriori* (MAP) inference problem. For a tree-structured graph, the problem can be solved efficiently using a suitable dynamic programming algorithm such as the max-product algorithm [63]. For example, for HMMs, the most likely sequence of latent states can be found efficiently using the Viterbi decoding algorithm [10]. For general graphs, however, the MAP inference problem is a computationally intractable integer program and is NP-hard [29]. In practice, these problems can be approached by considering a linear programming (LP) relaxation of the integer program. Over the past few years, several algorithms have been proposed to solve such graph-structured LPs [29], [18], [22], [21]. Such approaches can be broadly classified into two groups:

primal methods which work with the original variables [25] and dual methods, which works on the dual variables [27].

In the context of large scale data analysis, the ability to apply such methods efficiently to graphs of over millions or hundreds of millions of nodes is important and necessary. Consider the problem of detecting droughts from precipitation data of the past 100 years at a temporal resolution of a month and spatial resolution of  $0.5^\circ \times 0.5^\circ$  over land. A 3-dimensional MRF (latitude  $\times$  longitude  $\times$  time) with neighborhood dependencies is a suitable model for such analysis since droughts have both spatial and temporal continuity. Assuming a boolean indicator variable of drought at each space-time location, the graph-structured LP relaxation of the MAP inference problem in this context has to work with approximately 7 million variables and about double that many constraints. The key bottleneck, even in the advanced algorithms for solving graph-structured LPs, is that they are inherently sequential [26]. Given that climate datasets are available at much higher resolutions, especially from climate model outputs used by the Intergovernmental Panel on Climate Change (IPCC) for future climate projections, we need algorithms for solving graph structured LPs which efficiently scale to problem sizes of millions or hundreds of millions of nodes. Further, due to the generality of the framework, the algorithms can find applications in other domains, such as community detection in large scale social networks, which can have millions to hundreds of millions of users (nodes) [28].

Driven by the emerging need for scalable MAP inference algorithms, in this paper, we adopt the recently proposed Bethe-ADMM algorithm [12], [15] for solving graph-structured LPs. The overall structure of the algorithm is based on two ideas: tree-based decomposition of a graph-structured LP [29] and the alternating direction method of multipliers (ADMM) [6]. The tree decomposition breaks the problem into small but overlapping parts, each involving small number of variables and constraints. The algorithm iterates between doing updates to variables in individual parts in parallel followed by suitable aggregation, all within the framework of ADMM. However, unlike standard ADMM,

Bethe-ADMM is a novel inexact ADMM augmented by a Bregman divergence induced by the Bethe entropy. The unusual modification in Bethe-ADMM leads to an efficient projection of partial solutions to subsets of constraints, leading to highly efficient iterations and avoids double-loop algorithm.

To illustrate the efficiency of the Bethe-ADMM algorithm, we implement it using Message Passing Interface (MPI), which is a natural fit for the parallel algorithm given its flexible message passing mechanism, along with its portability and wide adoption in distributed and high performance clusters. Another advantage of using MPI is that its I/O interface is optimized for a wide variety of underlying parallel file systems (PFS) and sustains high I/O bandwidth. We evaluate our algorithms on both large-scale simulation and real precipitation data. The empirical results show that we manage to obtain almost linear speedup in the number of cores used.

The rest of the paper is organized as follows: We briefly review the MAP inference problem in Section II. We introduce the Bethe-ADMM algorithm in Section III, and discuss its MPI implementation in detail in Section IV. We present the experimental results in Section V and conclude in Section VI.

## II. PROBLEM DEFINITION

A pairwise Markov random field (MRF) is defined on an undirected graph  $G = (V, E)$ , where  $V$  is the vertex set and  $E$  is the edge set. Each node  $u \in V$  has a random variable  $X_u$  associated with it, which can take value  $x_u$  in some discrete space  $\mathcal{X} = \{1, \dots, k\}$ . Concatenating all the random variables  $X_u, \forall u \in V$ , we obtain an  $n$  dimensional random vector  $\mathbf{X} = \{X_u | u \in V\} \in \mathcal{X}^n$ . We assume that the distribution  $P$  of  $\mathbf{X}$  is a Markov Random Field [30], meaning that it factors according to the structure of the undirected graph  $G$  as follows: With  $f_u : \mathcal{X} \mapsto \mathbb{R}, \forall u \in V$  and  $f_{uv} : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}, \forall (u, v) \in E$  denoting nodewise and edgewise potential functions respectively, the distribution takes the form  $P(\mathbf{x}) \propto \exp \left\{ \sum_{u \in V} f_u(x_u) + \sum_{(u, v) \in E} f_{uv}(x_u, x_v) \right\}$ .

An important problem in the context of MRF is that of *maximum a posteriori* (MAP) inference,

which is to compute the configuration  $\mathbf{x}^*$  with the largest probability:

$$\mathbf{x}^* \in \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}^n} \exp \left\{ \sum_{u \in V} f_u(x_u) + \sum_{(u, v) \in E} f_{uv}(x_u, x_v) \right\}. \quad (1)$$

The above optimization problem is equivalent to the following integer programming (IP) problem:

$$\mathbf{x}^* \in \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}^n} \left\{ \sum_{u \in V} f_u(x_u) + \sum_{(u, v) \in E} f_{uv}(x_u, x_v) \right\}. \quad (2)$$

The complexity of (2) depends critically on the structure of the underlying graph. When  $G$  is a tree structured graph, the MAP inference problem can be solved efficiently via the max-product algorithm [19]. However, for an arbitrary graph  $G$ , the MAP inference algorithm is usually computationally intractable. The intractability motivates the development of algorithms to solve the MAP inference problem approximately. In this paper, we focus on the Linear Programming (LP) relaxation method [29], [8]. The LP relaxation of MAP inference problem is defined on a set of pseudomarginals  $\mu_u$  and  $\mu_{uv}$ , which are non-negative, normalized and locally consistent [29], [8]:

$$\begin{aligned} \mu_u(x_u) &\geq 0, \quad \forall u \in V, \\ \sum_{x_u \in \mathcal{X}_u} \mu_u(x_u) &= 1, \quad \forall u \in V, \\ \mu_{uv}(x_u, x_v) &\geq 0, \quad \forall (u, v) \in E, \\ \sum_{x_u \in \mathcal{X}_u} \mu_{uv}(x_u, x_v) &= \mu_v(x_v), \quad \forall (u, v) \in E. \end{aligned} \quad (3)$$

We denote the polytope defined by (3) as  $L(G)$  and the LP relaxation of MAP inference problem (2) becomes solving the following LP:

$$\begin{aligned} \max_{\boldsymbol{\mu} \in L(G)} \langle \boldsymbol{\mu}, \mathbf{f} \rangle &= \sum_{u \in V} \sum_{x_u} \mu_u(x_u) f_u(x_u) \\ &+ \sum_{(u, v) \in E} \sum_{x_u, x_v} \mu_{uv}(x_u, x_v) f_{uv}(x_u, x_v), \end{aligned} \quad (4)$$

subject to the constraint that  $\boldsymbol{\mu} \in L(G)$ . If the solution  $\boldsymbol{\mu}$  to (4) is an integer solution, it is guaranteed to be the optimal solution of (2). Otherwise, one can apply rounding schemes [24], [25] to round the fractional solution to an integer solution.

Although standard LP solvers can be used to solve the optimization problem (4), they are usually inefficient for the MAP inference problem [33], mainly because they fail to take advantage of the underlying graph structure. Specially designed MAP inference algorithms usually exploit the structures of the dependency graphs and solve the problem efficiently. Broadly speaking, LP-based approximate MAP inference algorithms can be categorized into primal and dual algorithms which solve the primal (4) and the dual of (2) respectively.

## III. ALGORITHM

In this section, we first show how to solve (4) by the ADMM based on tree decomposition. The resulting algorithm can be a double loop algorithm since some updates do not have closed-form solution. We then introduce the Bethe-ADMM algorithm where every update can be computed efficiently.

### A. ADMM for MAP Inference

To facilitate the discussion in the sequel, we first outline the basic ADMM [6] updates. Suppose our optimization problem is as follows:

$$\min_{\mathbf{x} \in \mathbb{R}^n} \sum_{i=1}^N f_i(\mathbf{x}), \quad (5)$$

and we assume that  $f_i, i = 1, \dots, N$  are convex functions. We can rewrite (5) with local variables  $\mathbf{x}_i \in \mathbb{R}^n$  and a global variable  $\mathbf{z} \in \mathbb{R}^n$ :

$$\begin{aligned} \min_{\mathbf{x}_i, \mathbf{z}} \quad & \sum_{i=1}^N f_i(\mathbf{x}_i) \\ \text{subject to} \quad & \mathbf{x}_i = \mathbf{z}, \quad i = 1, \dots, N. \end{aligned} \quad (6)$$

Note that (7) is a consensus constraint and we use it to make sure that the local variables and the global variable are of the same value.

Let  $\lambda_i \in \mathbb{R}^n, i = 1, \dots, N$  be the Lagrangian multipliers. The Lagrangian function of (6) is:

$$L(\mathbf{x}_i, \lambda_i, \mathbf{z}) = \sum_{i=1}^N \left( f_i(\mathbf{x}_i) + \lambda_i^T (\mathbf{x}_i - \mathbf{z}) \right). \quad (8)$$

The ADMM algorithm imposes a quadratic penalty on (8) and the augmented Lagrangian then becomes:

$$L(\mathbf{x}_i, \lambda_i, \mathbf{z}) = \sum_{i=1}^N \left( f_i(\mathbf{x}_i) + \langle \lambda_i, \mathbf{x}_i - \mathbf{z} \rangle + \underbrace{\frac{\beta}{2} \|\mathbf{x}_i - \mathbf{z}\|_2^2}_{\text{quadratic penalty}} \right), \quad (9)$$

where  $\beta > 0$  is the positive penalty parameter. The problem (9) and (8) are equivalent because for any feasible  $\mathbf{x}_i$  and  $\mathbf{z}$ , the penalty term is zero. The ADMM algorithm minimizes (9) iteratively and consists of the following updates:

$$\mathbf{x}_i^{t+1} = \underset{\mathbf{x}_i}{\operatorname{argmin}} \left( f_i(\mathbf{x}_i) + \langle \lambda_i^t, \mathbf{x}_i - \mathbf{z}^t \rangle + \frac{\beta}{2} \|\mathbf{x}_i - \mathbf{z}^t\|_2^2 \right), \quad (10)$$

$$\mathbf{z}^{t+1} = \underset{\mathbf{z}}{\operatorname{argmin}} \left( - \sum_{i=1}^N \langle \lambda_i^{t+1}, \mathbf{z} \rangle + \frac{\beta}{2} \|\mathbf{x}_i^{t+1} - \mathbf{z}\|_2^2 \right), \quad (11)$$

$$\lambda_i^{t+1} = \lambda_i^t + \beta (\mathbf{x}_i^{t+1} - \mathbf{z}^{t+1}). \quad (12)$$

We first show how to decompose (4) into a series of subproblems. We can decompose the graph  $G$  into overlapping subgraphs and rewrite the optimization problem with consensus constraints to enforce the pseudomarginals on subgraphs (local variables) to agree with  $\mu$  (global variable). Throughout the paper, we focus on tree-structured decompositions. To be more specific, let  $\mathbb{T} = \{(V_1, E_1), \dots, (V_{|\mathbb{T}|}, E_{|\mathbb{T}|})\}$  be a collection of subgraphs of  $G$  which satisfies two criteria: (i) Each subgraph  $\tau =$

$(V_\tau, E_\tau)$  is a tree-structured graph and (ii) Each node  $u \in V$  and each edge  $(u, v) \in E$  is included in at least one subgraph  $\tau \in \mathbb{T}$ . We also introduce local variable  $\mathbf{m}_\tau \in L(\tau)$  which is the pseudomarginal [29], [8] defined on each subgraph  $\tau$ . We use  $\theta_\tau$  to denote the potentials on subgraph  $\tau$ . We denote  $\mu_\tau$  as the components of global variable  $\mu$  that belong to subgraph  $\tau$ . Note that since  $\mu \in L(G)$  and  $\tau$  is a tree-structured subgraph of  $G$ ,  $\mu_\tau$  always lies in  $L(\tau)$ . In the newly formulated optimization problem, we will impose consensus constraints for sharing nodes and edges. For the ease of exposition, we simply use the equality constraint  $\mu_\tau = \mathbf{m}_\tau$  to enforce the consensus.

The new optimization problem we formulate based on graph decomposition is then as follows:

$$\min_{\mathbf{m}_\tau, \mu} \sum_{\tau=1}^{|\mathbb{T}|} \rho_\tau \langle \mathbf{m}_\tau, \theta_\tau \rangle \quad (13)$$

$$\text{subject to} \quad \mathbf{m}_\tau - \mu_\tau = 0, \quad \tau = 1, \dots, |\mathbb{T}| \quad (14)$$

$$\mathbf{m}_\tau \in L(\tau), \quad \tau = 1, \dots, |\mathbb{T}| \quad (15)$$

where  $\rho_\tau$  is a positive constant associated with each subgraph. We use the consensus constraints (14) to make sure that the pseudomarginals agree with each other in the sharing components across all the tree-structured subgraphs. Besides the consensus constraints, we also impose feasibility constraints (15), which guarantee that, for each subgraph, the local variable  $\mathbf{m}_\tau$  lies in  $L(\tau)$ . When the constraints (14) and (15) are satisfied, the global variable  $\mu$  is guaranteed to lie in  $L(G)$ .

To make sure that problem (4) and (13) are equivalent, we also need to guarantee that

$$\min_{\mathbf{m}_\tau} \sum_{\tau=1}^{|\mathbb{T}|} \rho_\tau \langle \mathbf{m}_\tau, \theta_\tau \rangle = \max_{\mu} \langle \mu, \mathbf{f} \rangle, \quad (16)$$

assuming the constraints (14) and (15) are satisfied. It is easy to verify that, as long as (16) is satisfied, the choice of  $\rho_\tau$  and  $\theta_\tau$  do not change the problem. Let  $\mathbf{1}[\cdot]$  be a binary indicator function and  $\mathbf{l} = -\mathbf{f}$ . A straightforward approach to obtain the potential  $\theta_\tau$  can be:

$$\theta_{\tau,u}(x_u) = \frac{l_u(x_u)}{\sum_{\tau'} \rho_{\tau'} \mathbf{1}[u \in V_{\tau'}]}, \quad u \in V_\tau,$$

$$\theta_{\tau,uv}(x_u, x_v) = \frac{l_{uv}(x_u, x_v)}{\sum_{\tau'} \rho_{\tau'} \mathbf{1}[(u, v) \in E_{\tau'}]}, \quad (u, v) \in E(\tau).$$

Plugging in the equality constraints, we then have the augmented Lagrangian of (13) as:

$$L(\mathbf{m}_\tau, \mu_\tau, \lambda_\tau) = \sum_{\tau=1}^{|\mathbb{T}|} \left( \rho_\tau \langle \mathbf{m}_\tau, \theta_\tau \rangle + \langle \lambda_\tau, \mathbf{m}_\tau - \mu_\tau \rangle + \frac{\beta}{2} \|\mathbf{m}_\tau - \mu_\tau\|_2^2 \right), \quad (17)$$

where  $\lambda_\tau$  is the dual variable and  $\beta > 0$  is the penalty parameter. The following updates constitute a single iteration

of the ADMM [6]:

$$\mathbf{m}_\tau^{t+1} = \operatorname{argmin}_{\mathbf{m}_\tau \in L(\tau)} \langle \mathbf{m}_\tau, \rho_\tau \boldsymbol{\theta}_\tau + \boldsymbol{\lambda}_\tau^t \rangle + \frac{\beta}{2} \|\mathbf{m}_\tau - \boldsymbol{\mu}_\tau^t\|_2^2, \quad (18)$$

$$\boldsymbol{\mu}_\tau^{t+1} = \operatorname{argmin}_{\boldsymbol{\mu}} \sum_{\tau=1}^{|\mathbb{T}|} \left( -\langle \boldsymbol{\mu}_\tau, \boldsymbol{\lambda}_\tau^t \rangle + \frac{\beta}{2} \|\mathbf{m}_\tau^{t+1} - \boldsymbol{\mu}_\tau\|_2^2 \right), \quad (19)$$

$$\boldsymbol{\lambda}_\tau^{t+1} = \boldsymbol{\lambda}_\tau^t + \beta(\mathbf{m}_\tau^{t+1} - \boldsymbol{\mu}_\tau^{t+1}). \quad (20)$$

Now, the problem turns to whether the updates (18) and (19) can be solved efficiently which we analyze as follows:

**Updating  $\boldsymbol{\mu}$ :** Since we have an unconstrained optimization problem (19) and the objective function decomposes component-wise, taking the derivatives and setting them to zero yield the solution. In particular, let  $S_u$  be the set of subgraphs which contain node  $u$ , for the node components, we have:

$$\mu_u^{t+1}(x_u) = \frac{1}{|S_u| \beta} \sum_{\tau \in S_u} (\beta m_{\tau,u}^{t+1}(x_u) + \lambda_{\tau,u}^t(x_u)). \quad (21)$$

(21) can be further simplified by observing that  $\sum_{\tau \in S_u} \lambda_{\tau,u}^t(x_u) = 0$ :

$$\mu_u^{t+1}(x_u) = \frac{1}{|S_u|} \sum_{\tau=1}^T m_{\tau,u}^{t+1}(x_u). \quad (22)$$

Similarly, let  $S_{uv}$  be the subgraphs which contain edge  $(u, v)$  and the update for the edge components is:

$$\mu_{u,v}^{t+1}(x_u, x_v) = \frac{1}{|S_{uv}|} \sum_{\tau \in S_{uv}} m_{\tau,uv}^{t+1}(x_u, x_v). \quad (23)$$

**Updating  $\mathbf{m}_\tau$ :** We need to solve a quadratic optimization problem for each tree-structured subgraph. Unfortunately, we do not have a closed-form solution for (18) in general. One possible approach, similar to the proximal algorithm, is to first obtain the solution  $\tilde{\mathbf{m}}_\tau$  to the unconstrained problem of (18) and then project  $\tilde{\mathbf{m}}_\tau$  to  $L(\tau)$ :

$$\mathbf{m}_\tau = \operatorname{argmin}_{\mathbf{m} \in L(\tau)} \|\mathbf{m} - \tilde{\mathbf{m}}_\tau\|_2^2. \quad (24)$$

If we adopt the cyclic Bregman projection algorithm [7] to solve (24), the algorithm becomes a double-loop algorithm, i.e., the cyclic projection algorithm projects the solution to each individual constraint of  $L(\tau)$  until convergence and the projection algorithm itself is iterative.

## B. Bethe-ADMM

Instead of solving (18) exactly, a common way in inexact ADMMs [32], [17] is to linearize the objective function in

(18), i.e., the first order Taylor expansion at  $\mathbf{m}_\tau^t$ , and add a new quadratic penalty term such that

$$\mathbf{m}_\tau^{t+1} = \operatorname{argmin}_{\mathbf{m}_\tau \in L(\tau)} \langle \mathbf{y}_\tau^t, \mathbf{m}_\tau - \mathbf{m}_\tau^t \rangle + \frac{\alpha}{2} \|\mathbf{m}_\tau - \mathbf{m}_\tau^t\|_2^2, \quad (25)$$

where  $\alpha$  is a positive constant and

$$\mathbf{y}_\tau^t = \rho_\tau \boldsymbol{\theta}_\tau + \boldsymbol{\lambda}_\tau^t + \beta(\mathbf{m}_\tau^t - \boldsymbol{\mu}_\tau^t). \quad (26)$$

However, as discussed in the previous section, the quadratic problem (25) is generally difficult for a tree-structured graph and thus the conventional inexact ADMM does not lead to an efficient update for  $\mathbf{m}_\tau$ . Next we show that, by taking the tree structure into account, an inexact minimization of (18) augmented with a Bregman divergence induced by Bethe entropy leads to efficient update of  $\mathbf{m}_\tau$ .

The basic idea in the new algorithm is that we replace the quadratic term in (25) with a Bregman-divergence term  $d_\phi(\mathbf{m}_\tau || \mathbf{m}_\tau^t)$  such that

$$\mathbf{m}_\tau^{t+1} = \operatorname{argmin}_{\mathbf{m}_\tau \in L(\tau)} \langle \mathbf{y}_\tau^t, \mathbf{m}_\tau - \mathbf{m}_\tau^t \rangle + \alpha d_\phi(\mathbf{m}_\tau || \mathbf{m}_\tau^t), \quad (27)$$

is efficient to solve for tree  $\tau$ . Expanding the Bregman divergence and removing the constants, we can rewrite (27) as

$$\mathbf{m}_\tau^{t+1} = \operatorname{argmin}_{\mathbf{m}_\tau \in L(\tau)} \langle \mathbf{y}_\tau^t / \alpha - \nabla \phi(\mathbf{m}_\tau^t), \mathbf{m}_\tau \rangle + \phi(\mathbf{m}_\tau). \quad (28)$$

For a tree-structured problem, what convex function  $\phi(\mathbf{m}_\tau)$  should we choose? Recall  $\mathbf{m}_\tau$  defines the marginal distributions of a tree-structured distribution  $p_{\mathbf{m}_\tau}$  over the nodes and edges:

$$\begin{aligned} \mathbf{m}_{\tau,u}(x_u) &= \sum_{\neg x_u} p_{\mathbf{m}_\tau}(x_1, \dots, x_u, \dots, x_n), \quad \forall u \in V_\tau, \\ \mathbf{m}_{\tau,uv}(x_u, x_v) &= \sum_{\neg x_u, \neg x_v} p_{\mathbf{m}_\tau}(x_1, \dots, x_u, x_v, \dots, x_n), \quad \forall (uv) \in E_\tau. \end{aligned}$$

It is well known that the sum-product algorithm [19] efficiently computes the marginal distributions for a tree structured graph. It can also be shown that the sum-product algorithm solves the following optimization problem [30] for tree  $\tau$ :

$$\max_{\mathbf{m}_\tau \in L(\tau)} \langle \mathbf{m}_\tau, \boldsymbol{\eta}_\tau \rangle + H_{\text{Bethe}}(\mathbf{m}_\tau), \quad (29)$$

where  $H_{\text{Bethe}}(\mathbf{m}_\tau)$  is the Bethe entropy of  $\mathbf{m}_\tau$ . The Bethe entropy on tree  $\tau$  is defined as:

$$H_{\text{Bethe}}(\mathbf{m}_\tau) = \sum_{u \in V_\tau} H_u(\mathbf{m}_{\tau,u}) - \sum_{(u,v) \in E_\tau} I_{uv}(\mathbf{m}_{\tau,uv}), \quad (30)$$

where  $H_u(\mathbf{m}_{\tau,u})$  is the entropy function on each node  $u \in V_\tau$  and  $I_{uv}(\mathbf{m}_{\tau,uv})$  is the mutual information on each edge  $(u, v) \in E_\tau$ .

Combing (28) and (29), we set  $\boldsymbol{\eta}_\tau = \nabla \phi(\mathbf{m}_\tau^t) - \mathbf{y}_\tau^t / \alpha$  and choose  $\phi$  to be the negative Bethe entropy of  $\mathbf{m}_\tau$  so

that (28) can be solved efficiently in linear time via the sum-product algorithm.

For the sake of completeness, we summarize Bethe-ADMM algorithm as follows :

$$\mathbf{m}_\tau^{t+1} = \underset{\mathbf{m}_\tau \in L(\tau)}{\operatorname{argmin}} \langle \mathbf{y}_\tau^t / \alpha - \nabla \phi(\mathbf{m}_\tau^t), \mathbf{m}_\tau \rangle + \phi(\mathbf{m}_\tau), \quad (31)$$

$$\boldsymbol{\mu}^{t+1} = \underset{\boldsymbol{\mu}}{\operatorname{argmin}} \sum_{\tau=1}^T \left( -\langle \boldsymbol{\lambda}_\tau^t, \boldsymbol{\mu}_\tau \rangle + \frac{\beta}{2} \|\mathbf{m}_\tau^{t+1} - \boldsymbol{\mu}_\tau\|_2^2 \right), \quad (32)$$

$$\boldsymbol{\lambda}_\tau^{t+1} = \boldsymbol{\lambda}_\tau^t + \beta(\mathbf{m}_\tau^{t+1} - \boldsymbol{\mu}_\tau^{t+1}), \quad (33)$$

where  $\mathbf{y}_\tau^t$  is defined in (26) and  $\phi$  is defined in (30). Due to the space constraint, we refer the readers to [12] for the detailed convergence analysis of the Bethe-ADMM algorithm.

It is easy to see that the update of  $\mathbf{m}_\tau$  in (31) is independent for each tree  $\tau$ , which motivates the parallel implementation of the Bethe-ADMM algorithm. We describe our implementation in detail in the next section.

#### IV. PARALLEL IMPLEMENTATION

In this section, we explain the key components of our MPI implementation in detail. Our goal is to run the Bethe-ADMM algorithm on modern high performance computers with thousands of cores and it requires us to adopt the best parallelization practice. To achieve this goal, we carefully design our MPI implementation so that the underlying parallel computing architecture can be fully utilized.

##### A. Existing Frameworks

Various computational frameworks have been proposed and successfully applied to various inference problems in large scale graphical models, such as GraphLab [1] and Pregel [20].

GraphLab [1] provides a high level parallel abstraction through a *data graph*, which encodes the computational structure and the dependencies of the subproblems. It also provides a set of *consistency models* which insulates the users from the complexity of synchronization, data races and deadlocks. However, to our best knowledge, there are no experimental results to show that graphLab runs over thousands of computer cores.

On the other side, Pregel, a large-scale graph processing system proposed in [20], and its open source counterpart Giraph [2], adopt the bulk synchronous parallel (BSP) computing model, by which the computation is carried out in the sequence of supersteps. The same user-defined function is applied to each vertex in parallel within each superstep, while the global synchronization barrier is enforced at inter-supersteps. Different from our focus on HPC systems, these two frameworks dedicate efforts on the distributed commodity hardware and relies on the distributed file system.

Nevertheless, we all endorse using a pure messaging passing model for data exchange rather than either remote procedure call or shared memory model.

Even though we only construct first order Markov Random Fields in the previous problem formulation, Bethe-ADMM algorithm is general enough to apply to high order MRFs. Thus, our goal is to design a MPI framework to solve the MAP inference problems in arbitrary large graphs.

##### B. Bethe-ADMM using MPI

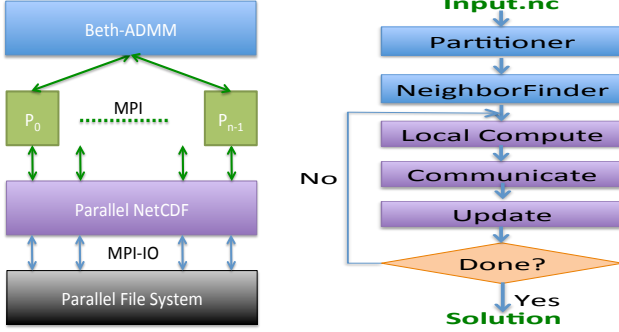
Since the update of  $\mathbf{m}_\tau$  in (31) for each tree is independent, the Bethe-ADMM algorithm is inherently parallel. In the parallel Bethe-ADMM algorithm, each process only maintains the information of a subset of trees in  $\mathbb{T}$  and  $\mathbf{m}_\tau$  is updated simultaneously. According to (32), the update of variable  $\boldsymbol{\mu}$  involves averaging over  $\mathbf{m}_\tau$  from the relevant trees. If these trees belong to different processes, the value of  $\mathbf{m}_\tau$  needs to be exchanged among the processes so that  $\boldsymbol{\mu}$  can be computed correctly. Because of the communication among the processes, the message passing framework is a good fit for our parallel implementation. Hence, we implement the Bethe-ADMM algorithm using MPI. We also make the following implementation assumptions: (i) The MRF dependency graph is a regular grid shaped graph, e.g., two dimensional four nearest neighbor grid. (ii) Each tree structured subgraph is simply an edge of  $G$ . (iii) The input to the MAP inference algorithm is some data file, which has the potential and graph structure information.

An efficient parallel implementation is more challenging than an efficient sequential implementation. To fully utilize the computing power provided by the underlying parallel architecture, we need to address the following issues:

- How to design an efficient I/O scheme to load the data files, i.e., node potentials, edge potentials and graph structure?
- How to decompose the graph so that the work load on each process is balanced?
- How to efficiently figure out, for each process, what ‘messages’ it needs to exchange with other processes?

We depict in Fig 1(a) how the entire application is deployed in the high performance computing environment. We take advantage of the underlying parallel file system (PFS) in high performance computing systems along with the parallel I/O library, PnetCDF in our case, so that the data files can be accessed in parallel among all the processes.

We also design a simple heuristic to partition the graph to achieve load balancing. With the help of PnetCDF’s metadata APIs to access the graph structure information, we deploy a decentralized partitioning algorithm to figure out, for each MPI process, the information it needs to exchange with other processes. After each process reads the data file in parallel to fetch the relevant nodewise and edgewise potentials, it computes the local variables  $\mathbf{m}_\tau$ ,



(a) Beth-ADMM deployment. (b) Beth-ADMM parallel implementation.

communicates with other processes and update the global variable  $\mu$ .

Once the graph is decomposed among processes, Beth-ADMM algorithm can be executed iteration by iteration until the termination condition is met.

The key components of our MPI implementation is illustrated in Fig 1(b), each of which is described in detail below.

### C. Parallel File Loading

The data file used as input to the MAP inference algorithm contains the nodewise and edgewise potentials and the graph structure information. We represent the graph as a set of edges with two node ids. (Figure 1(c) shows an example on a simple grid graph.) A naive way to load the data file is to have a master process read the entire data file and send to other slave processes the information they need. This approach is clearly not efficient because a slave process remains idle when other slave processes receive data from the master process. Our approach is to take advantage of the PFS, which stripes a file across multiple storage devices and enables parallel access to the data file.

To be more specific, we adopt the PnetCDF [3] file format for parallel data file loading. The PnetCDF is suitable for our implementation because the potential data and graph structure information can be easily stored as PnetCDF multi-dimensional arrays. A PnetCDF file also provides a rich suite of APIs that allow users to define metadata which describe datasets in detail, such as the number of nodes and edges of a given graph, the type of graphs and the dimensions of the datasets. Moreover, it integrates tightly with MPI-IO and the underlying PFS so that our algorithm can achieve a high degree of parallelism in terms of I/O operations. Given the number of edges, the number of processes and its own MPI rank, each process could easily determine its own set of edges to be processed.

### D. Graph Partitioning

To take advantage of the parallel architecture, the workload should be split evenly among the processes and the par-

tion should also minimize the intercommunication among the processes. This problem is usually NP hard and most practical solutions are based on heuristics. For example, in Pregel [20] and Giraph [2], the solution is to use a node-centric partition, where assignment of a node to a partition depends solely on the node id. The simplest implementation is to calculate the hash value of each node id and modulus by  $N$ , where  $N$  is the number of partitions. However this simple heuristic comes with a cost that neighboring nodes are likely to be distributed on different processes and thus incur high communication overhead.

In our implementation, we adopt edge-centric partition, where we evenly divide the edges among all the processes. With no assumption on the shape of the given graph models, edge-centric decomposition can be applied to any arbitrary graphs. If the underlying dependency graph is a simple graph such as a regular shaped grid graph, edge partition is empirically a good choice, as shown by the experimental results in Section V.

Since the graph structure such as the list of edges and the number of edges are recorded in the PnetCDF file as metadata, we can retrieve this information using PnetCDF APIs and evenly divide the edges among all the processes. The edge list needs to be sorted based on the end points such that the close vertices are more likely assigned to the same process together. Even though this wouldn't give us the optimal cut of the graph, we can still preserve the local connectivity on one process. Once each process calculates its own workload, it only reads its own portion of data from the PnetCDF file. The `nc_get_vara` API provided by PnetCDF library employs MPI-IO techniques and retrieves subarrays from datasets with high read bandwidth.

One widely used program, METIS[16], uses multilevel algorithms to recursively partition graphs and it exhibits high efficiency for large graphs with millions of vertices. We could also add an extra partitioner hook for libraries like METIS, such that its partitioning results can be incorporated in our framework when solving more generalized graphs.

Even though later in our experiments, we mainly focus on MAP inference on pairwise MRFs, in principle, there is no loss of generality in restricting to pairwise interactions, since any higher-order MRF can be converted to pairwise MRFs by introducing auxiliary random variables[31]. Moreover, by applying edge-centric decomposition without making assumptions about the graph type, the techniques described in this paper can all be generalized to apply directly to general MRFs.

### E. Inter-process communication

After the graph decomposition step, each process reads from the input PnetCDF file, retrieves the nodewise and edgewise potentials and computes  $m_\tau$ . To compute  $\mu$ , a simple solution is to have a master process collect the value of  $m_\tau$  from the slave processes and compute  $\mu$  according

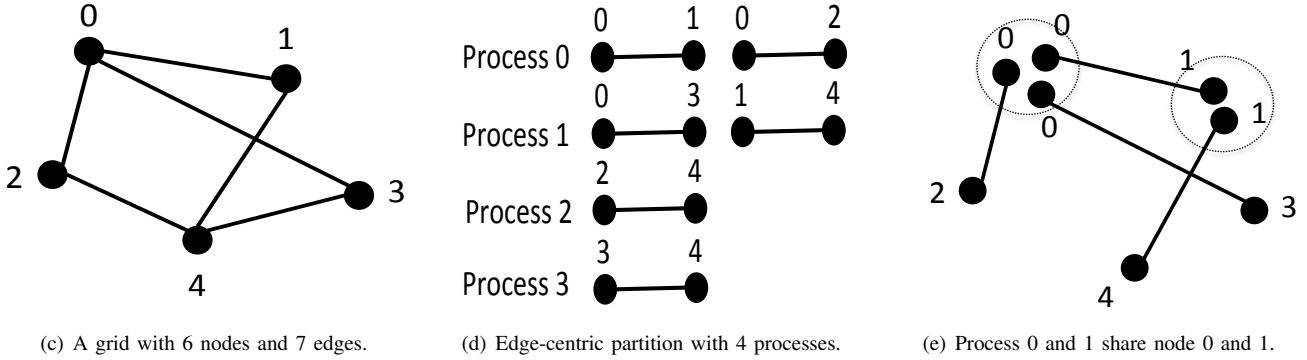


Figure 1. 1(c): We label the nodes row by row and represent the graph structure as a set of edges: (0, 1), (0, 2), (0, 3), (1, 4), (2, 4), (3, 4). 1(d): We use 4 MPI processes and apply edge-centric partition. 1(e): The node list of process 0 can be represented as:  $\{\{0, 1, 2\}\}$  and the node list of process 1 can be represented as:  $\{\{0, 1, 3, 4\}\}$ . These two processes share node 1 and 0. Node 0 has the degree of 3, its partial degree is 2 on process 0 and 1 on process 1.

to (32). After  $\mu$  is updated, the master process has to send  $\mu$  back to each slave process so that  $m_\tau$  can be computed in the next iteration. This approach is clearly not efficient and we adopt a fully distributed algorithm: each process maintains a copy of the relevant elements of  $\mu$ , receives  $m_\tau$  from other processes and updates  $\mu$  according to (32).

---

#### Algorithm 1 NeighborFinder

---

- 1: **procedure** NEIGHBORFINDER
  - 2:   idList = getNodeId()
  - 3:   pairCount = idList.size()
  - 4:   MPI\_Allgather(  
5:     *pairCount*, 1, MPI\_INT,  
6:     *countArr*, 1, MPI\_INT, *comm*)
  - 7:   Copy idList to sendBuf
  - 8:   Construct *displacementArr* from *countArr*
  - 9:   MPI\_Allgatherv(  
10:     *sendBuf*, 2 \* *pairCount*, MPI\_INT,  
11:     *recvBuf*, 2 \* *countArr*, *displacementArr*,  
12:     MPI\_INT, *comm*)
  - 13:   Compute neighbor processes by comparing idLists
  - 14:   Count partial degree of sharing nodes
  - 15:   Exchange partial degree with neighbor processes
  - 16:   Compute full degree of sharing nodes
  - 17: **end procedure**
- 

To apply the above distributed algorithm, each process needs to figure out the neighbor processes with which it exchanges the value of  $m_\tau$ . This can be done by comparing the node ids of each process and a pair of processes need to communicate with each other if they have sharing nodes. To be more specific, we compactly represent the node list of a process as a list of pairs  $\{v_i, l_i\}$ , where  $l_i$  is the length of continuous ids starting from  $v_i$ . (Figure 1(e) illustrates the compact representation of node lists on two processes.) Each process then gathers  $\{(v_i, l_i)\}$  from all other processes,


compares the lists with its own node list and decides what processes it communicates with. Beside deciding the neighbor processes, each process also needs to figure out the degrees of the sharing nodes. The degree (count) information will be used when the averaging operation is performed according to (32). As a result, the neighbor process also exchanges the local partial degree of the sharing nodes and compute the full degree accordingly. Algorithm 1 summaries the above procedure.

Algorithm 2 shows the details on how MPI APIs are employed in our framework. We reduce the communication cost by exchanging the partial sum of  $m_\tau$  rather than individual  $m_\tau$  one by one. Further we use asynchronous MPI APIs, which allow messages to be sent or received asynchronously without blocking the following operations.

## V. EXPERIMENTAL RESULTS

In this section, we present experimental results on a simulation dataset and a precipitation dataset. Our experiments are conducted on Hopper [4], the Cray XE6 parallel machine at the National Energy Research Scientific Computing Center.

Hopper is a 6384 compute node cluster where each compute node consists of two twelve-core AMD Magny-Cours processors with a theoretical peak performance of 8.4 GFlop/sec per core. 6000 compute nodes have 32 GB DD3 memory each and the rest have 64 GB memory each. Hopper runs “Cray Linux Environment” (CLE) which is a restricted low-overhead operating system that has been optimized for high performance computing. The PFS is Lustre with 156 I/O servers (OSTs). The measured peak write performance on Hopper is 35 GB per second. To maximize the possible read bandwidth, we stripe our input file across 128 stripes and the file stripe size is set to 1 MB.



(a) Time spent on the I/O phase and the Bethe-ADMM optimization. The I/O cost is low. (b) Time spent on the three steps of the Bethe-ADMM optimization. The communication overhead processes can be negligible. (c) Almost linear speedup in the number of MPI

Figure 2. Results on the simulation dataset with 10 million nodes and 20 million edges using 8-1024 MPI processes. The I/O and communication cost is relatively low. Overall, the MPI implementation achieves almost linear speedup in the number of processes.

### Algorithm 2 Exchange $m_\tau$ among neighbor processes

```

1: procedure EXCHMSG
2:   for Each node  $u$  do
3:      $partial\_sum[u] = 0$ 
4:   end for
5:   for Each edge  $\tau(u, v)$  do
6:      $partial\_sum[u] += m_\tau[u]$ 
7:      $partial\_sum[v] += m_\tau[v]$ 
8:   end for
9:    $idx = 0$ 
10:  MPI_Request request[neighbors.size() * 2]
11:  for  $i$  in neighbors do
12:     $sharing\_node = getSharingNode(i)$ 
13:    copy partial_sum[sharing_node] to sendBuf
14:    MPI_Isend(sendBuf, k*sharing_node.size(),
15:             MPI_FLOAT, i, rank,
16:             comm, &request[idx++])
17:    MPI_Irecv(recvBuf, k*sharing_node.size(),
18:             MPI_FLOAT, i, i,
19:             comm, &requests[idx++])
20:  end for
21: end procedure

```

#### A. Simulation Dataset

We show experimental results on a simulation dataset. The underlying graph is a 2 dimensional  $1,000 \times 10,000$  grid with  $k = 3$  and the potentials are random numbers in  $[0, 1]$ . The resulting MRF has 10 million nodes and approximately 20 million edges. We apply the edge-centric partitioning and run the Bethe-ADMM algorithm for 100 iterations.

Figure 2(a) shows the run time performance using 8 to 1024 MPI processes. The algorithm runs about half an hour on 8 process, and dramatically reduces to 16 seconds on 1024 processes. The input file size is close to 1GB and data loading only takes 1.2 seconds. We attribute the speedup to

our adoption of PNetCDF as well as stripping the input file across 128 OSTs.

Figure 2(b) illustrates the average time it takes per process to compute  $m_\tau$ , update  $\mu$  and communicate with neighbor processes respectively. The error bars show the minimum and maximum time spent on these three steps across all the processes. Since we evenly distribute the edges to the processes, the time spent on computing  $m_\tau$  has little fluctuation among the processes. The time to update  $\mu$ , however, also depends on the number of neighbor processes and the number of shared nodes, hence the fluctuation between the min and max time among all the processes becomes more obvious as the number of processes increases. The plot shows the communication cost incurred by the edge-centric partitioning is negligible. The main reason that the communication cost is so small is because when we partition the grid, we sweep row edges and column edges from top to bottom, which essentially behaves as row partitioning where each process has at most 2 neighbors and only the boundary data are exchanged.


Figure 2(c) shows that the Bethe-ADMM algorithm implementation achieves almost linear speedup while the speedup of the entire implementation (I/O phase + Bethe-ADMM optimization) starts to deviate from the ideal case after 256 processes. This happens because communication costs dominate once the work is spread out enough.

#### B. CRU Precipitation Dataset

The dataset used in this section is the Climate Research Unit (CRU) precipitation dataset [23], which has monthly precipitation from the years 1901-2006. The dataset is of high gridded spatial resolution ( $360 \times 720$ , i.e., 0.5 degree latitude  $\times$  0.5 degree longitude) and only includes the precipitation over land.

Our goal is to detect major droughts of the last century based on precipitation. We formulate the drought detection





(a) Time spent on the I/O phase and Bethe-ADMM optimization. The I/O cost is low. (b) Time spent on the three steps of the Bethe-ADMM optimization. The communication overhead is low. (c) Almost linear speedup in the number of MPI processes

Figure 3. Results on the CRU dataset with 7,146,520 nodes and 20,777,480 edges using 8-1024 MPI processes. The I/O and communication cost is relatively low. Overall, the MPI implementation achieves almost linear speedup in the number of processes.

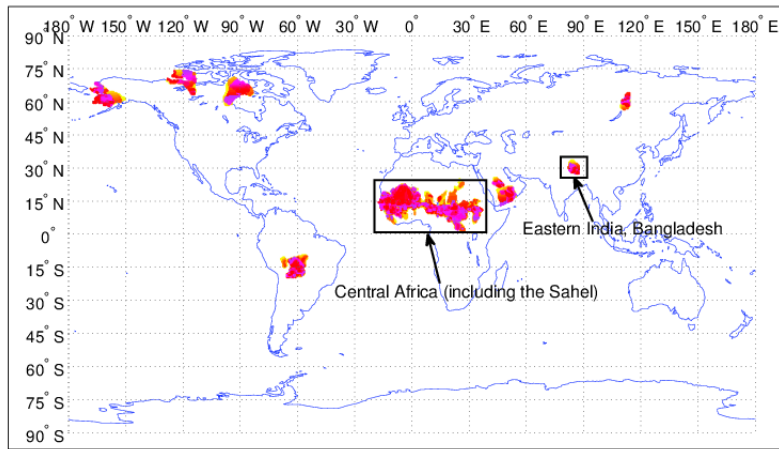


Figure 4. Major droughts starting within the period 1961-1970, which include the three decade long Sahel drought and the drought in eastern India in the 1960s.

problem as the one of estimating the most likely configuration of a binary hidden MRF. In the underlying graph, each node represents a location and it can be in two possible states: dry and normal. We use a four nearest neighbor grid ( $m = 360, n = 720$ ) to model the global dependency and replicate it 106 times. The resulting graph is similar to the ones used in the previous section and the structure respects the CRU dataset, i.e., it only has the nodes that correspond to the locations with precipitation record. Overall, the three dimensional grid has 7,146,520 nodes and 20,777,480 edges.

We design the potential functions carefully from the CRU datasets to enforce label consistency, i.e., neighboring nodes should take same values. We refer the readers to [11] regarding the details on designing potential functions. We obtain the integer solution after rounding the node pseudomarginals and we can detect droughts based on it. Figure 4 shows the detected droughts in the 1960s.

We run the Bethe-ADMM algorithm on the CRU dataset for 500 iterations with edge-centric partitioning. The input

PNetcdf file is around 530 MB. The runtime performance, as shown in Figure 3(a) exhibits the nice decreasing trend as it does on the simulation data. The algorithm takes less than 2 minutes to complete with 1024 MPI processes which would run more than two hours with 8 processes. The amount of time saved by our implementation is tremendous.

Figure 3(b) illustrates the average time per process to compute  $m_\tau$ , communicate with neighbors and update  $\mu$  respectively. The error bars mark the minimum and maximum time spent on these three steps across all the processes. The communication cost on the CRU dataset is no longer negligible anymore. This is because the underlying 3 dimensional grid has missing nodes (CRU only has precipitation over land) and when we apply edge-centric partitioning, each process may have more than two neighbors. Hence as the number of processes increases, the number of neighbors for each process is more dynamic and the communication pattern becomes more complicated. Figure 3(c) plots the almost linear speedup on the CRU dataset. It also shows

the trend that our implementation is scalable beyond 1024 processes. This is understandable because I/O time is only 2% of the total execution time, even at 1024 MPI processes.

## VI. CONCLUSIONS

In this paper, we proposed a parallel MAP inference algorithm for large-scale MRFs. Based on the ‘tree decomposition’ idea from the MAP inference literature and the alternating direction method from the optimization literature, the problem is divided into a set of tree-structured subproblems, each of which can be solved independently and efficiently via the sum-product algorithm.

The entire framework is implemented with the standard MPI library in C++ and can be easily ported on any HPC systems or even distributed clusters. We evaluated our algorithm on both synthetic and real-world datasets and the experimental results verify our parallel design can scale almost linearly with the number of MPI processes up to thousands of computer cores for grid-structured graphs.

## ACKNOWLEDGMENT

This work is supported in part by the following grants: NSF awards CCF-1029166, ACI-1144061, IIS-1343639, and CCF-1409601; DOE award DESC0007456; AFOSR award FA9550-12-1-0458; NIST award 70NANB14H012.

## REFERENCES

- [1] <http://graphlab.org>.
- [2] <http://giraph.apache.org>.
- [3] <http://cucis.ece.northwestern.edu/projects/PnetCDF>.
- [4] <http://www.nersc.gov/users/computational-systems/hopper>.
- [5] J. Bilmes. A gentle tutorial on the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden markov models. Technical Report ICSI-TR-97-02, University of Berkeley, 1997.
- [6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [7] Y. Censor and S. Zenios. *Parallel optimization: theory, algorithms, and applications*. Oxford University Press, 1998.
- [8] C. Chekuri, S. Khanna, J. Naor, and L. Zosin. A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM Journal on Discrete Mathematics*, 18(3):608–625, Mar. 2005.
- [9] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison. *Biological sequence analysis*. Cambridge University Press, 2006.
- [10] G. D. Fornay. The viterbi algorithm. *Proceedings of the IEEE*, 61(3):268–278, 1973.
- [11] Q. Fu, A. Banerjee, S. Liess, and P. K. Snyder. Drought detection of the last century: An MRF-based approach. In *Proceedings of the SIAM International Conference on Data Mining*, 2012.
- [12] Q. Fu, H. Wang, and A. Banerjee. Bethe-ADMM for tree based parallel MAP inference. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, 2013.
- [13] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.
- [14] X. Huang, A. Acero, and H.-W. Hon. *Spoken language processing: A guide to theory, algorithm, and system development*. Prentice Hall, 2001.
- [15] C. Jin, Q. Fu, H. Wang, A. Agrawal, W. Hendrix, W.-k. Liao, M. M. A. Patwary, A. Banerjee, and A. Choudhary. Solving combinatorial optimization problems using relaxed linear programming: a high performance computing perspective. In *Proceedings of the 2nd International Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications, BigMine ’13*, pages 39–46, New York, NY, USA, 2013. ACM.
- [16] G. Karypis and V. Kumar. Metis – unstructured graph partitioning and sparse matrix ordering system, version 2.0. Technical report, 1995.
- [17] S. P. Kasiviswanathan, P. Melville, A. Banerjee, and V. Sindhvani. Emerging topic detection using dictionary learning. In *Proceedings of the Twentieth ACM international conference on Information and knowledge management*, 2011.
- [18] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1568–1583, oct. 2006.
- [19] F. R. Kschischang, B. J. Frey, and H. A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, 47(2):498–519, 2001.
- [20] G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM International Conference on Management of Data*, 2010.
- [21] A. F. Martins, P. M. Aguiar, M. A. Figueiredo, N. A. Smith, and E. P. Xing. An augmented Lagrangian approach to constrained MAP inference. In *Proceedings of the Twenty-Eighth International Conference on Machine Learning*, 2011.
- [22] O. Meshi and A. Globerson. An alternating direction method for dual MAP LP relaxation. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2011.
- [23] T. D. Mitchell, T. R. Carter, P. D. Jones, M. Hulme, and M. New. *A comprehensive set of high-resolution grids of monthly climate for Europe and the globe: the observed record (1901-2000) and 16 scenarios (2001-2100)*. Tyndall Centre for Climate Change Research, 2004.

- [24] P. Raghavan and C. D. Thompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. *Combinatorica*, 7(4):365–374, 1987.
- [25] P. Ravikumar, A. Agarwal, and M. J. Wainwright. Message-passing for graph-structured linear programs: Proximal methods and rounding schemes. *Journal of Machine Learning Research*, 11:1043–1080, 2010.
- [26] P. Ravikumar and J. Lafferty. Quadratic programming relaxations for metric labeling and markov random field map estimation. In *ICML*, pages 737–744, 2006.
- [27] D. Sontag, A. Globerson, and T. Jaakkola. Introduction to dual decomposition for inference. In S. Sra, S. Nowozin, and S. J. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2011.
- [28] M. Szell, R. Lambiotte, and S. Thurner. Multirelational organization of large-scale social networks. *Proceedings of the National Academy of Sciences USA*, 107(31):13636–13641, 2010.
- [29] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. MAP estimation via agreement on (hyper)trees: Message-passing and linear-programming approaches. *IEEE Transactions of Information Theory*, 51(11):3697–3717, 2005.
- [30] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, 2008.
- [31] M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(12):1–305, 2008.
- [32] J. Yang and Y. Zhang. Alternating direction algorithms for l1-problems in compressive sensing. *SIAM Journal on Scientific Computing*, 33(1):250–278, 2011.
- [33] C. Yanover, T. Meltzer, and Y. Weiss. Linear programming relaxations and belief propagation: an empirical study. *Journal of Machine Learning Research*, 7:1887–1907, 2006.