
Partial Optimality of Dual Decomposition for MAP Inference in Pairwise MRFs

Alexander Bauer
TU Berlin

Shinichi Nakajima
TU Berlin
AIP RIKEN

Nico Görnitz
TU Berlin

Klaus-Robert Müller
TU Berlin
Korea University
MPI for Informatics

Abstract

Markov random fields (MRFs) are a powerful tool for modelling statistical dependencies for a set of random variables using a graphical representation. An important computational problem related to MRFs, called maximum a posteriori (MAP) inference, is finding a joint variable assignment with the maximal probability. It is well known that the two popular optimisation techniques for this task, linear programming (LP) relaxation and dual decomposition (DD), have a strong connection both providing an optimal solution to the MAP problem when a corresponding LP relaxation is tight. However, less is known about their relationship in the opposite and more realistic case. In this paper, we explain how the fully integral assignments obtained via DD partially agree with the optimal fractional assignments via LP relaxation when the latter is not tight. In particular, for binary pairwise MRFs the corresponding result suggests that both methods share the partial optimality property of their solutions.

1 Introduction

The framework of graphical models such as Markov random fields (MRFs) [8, 15, 21] provides a powerful tool for modelling statistical dependencies for a set of random variables using a graphical representation. It is of a fundamental importance for many practical application areas including natural language processing, information retrieval, computational biology, and computer vision. A related computational problem, called

maximum a posteriori (MAP) inference, is finding a joint variable assignment with the maximal probability. The vast amount of existing methods for solving the discrete MAP problem (see [7] for an overview) can be divided roughly into three different groups: methods based on graph-cuts, methods based on message passing, and polyhedral methods. The latter group is tightly connected to a popular approach of linear programming (LP) relaxation [17, 21, 2, 5], which is based on reformulating the original combinatorial problem as an integer linear problem (ILP) and then relaxing the integrality constraints on the variables. Besides providing a lower bound on the optimal value its popularity is partially¹ due to the fact that for binary pairwise MRFs with submodular energies it is guaranteed to find an optimal solution [8, 22]. Unfortunately, for bigger problems with a high number of variables and constraints it becomes impractical due to the expensive consumption in memory and computation.

As an alternative to the LP relaxation, dual decomposition (DD) [12, 18, 16, 4, 13] provides an effective parallelisation framework for solving the MAP problem. Furthermore, it has an appealing property that any found solution comes with a certificate of optimality which allows for efficient evaluation whether a corresponding variable assignment is primal optimal. Finally, it is well known that the two optimisation techniques have a strong connection both providing an optimal solution if the LP relaxation is tight. However, less is known about their relationship in the opposite and more realistic case. Instead the main focus of the existing literature is on tightening the standard LP relaxation [12, 19, 17]. In contrast, the aim of this paper is to investigate the connections between the two techniques (in the original formulation) if the LP relaxation is not tight. More precisely, we focus on the following issue. The main idea of DD is to decompose a given MRF into different trees (or other subgraphs) on

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

¹It also provides an optimal solution for any tree-structured MRF.

which inference can be performed efficiently and trying to enforce an agreement on the overlapping variables between different trees to obtain global consistency. If the LP relaxation is not tight, some trees will provide inconsistent assignments which disagree on the overlapping parts. Here we analyse the nature of this disagreement and get the following main results:

- given an optimal (fractional) solution of the LP relaxation, there always exists an optimal variable assignment via DD which agrees with the integral part of the LP solution
- for binary pairwise MRFs in a non degenerate case², the unambiguous part among all optimal assignments from different trees in a decomposition coincides with the integral part of the (fractional) optimal assignment via LP relaxation.

We note that the first result holds also for non binary MRFs with arbitrary higher order potentials. On the other hand, for binary pairwise MRFs it implies that a corresponding assignment contains the strongly persistent part³ of the LP solution [21, 6]. The second result suggests a strategy how to extract the strongly persistent part by looking at the intersection of all optimal assignments to the overlapping trees. In this sense, both methods LP relaxation and DD share a partial optimality property of their solutions.

Note that a corresponding fractional solution obtained via LP relaxation (for binary pairwise MRFs) is half integral. That is, every fractional node is equal to 0.5. This provides no information about the preferences of the fractional variables in that case preventing use of rounding techniques. In contrast, DD always provides a fully integral assignment which inherits the strongly persistent part of the LP relaxation making DD an appealing optimisation method.

Finally, we argue that depending on the final goal there is a decision to make with respect to the degree of a corresponding decomposition. Usually, the main goal is to find an accurate (integral) assignment to the variables in a MRF. In that case a decomposition over spanning trees is more beneficial. It significantly speeds up the convergence but most importantly it is straightforward how to extract an optimal assignment. On the other hand if we want to extract the strongly

²By a non degenerate case we mean the case where the LP relaxation has a unique (fractional) solution. That is, the optimum is attained at a corner and not at an edge or a facet of a corresponding polytope.

³For binary pairwise MRFs, the integral part of an optimal solution of the LP relaxation is known to be strongly persistent. That is, any optimal solution of a corresponding MAP problem must agree with this partial integral assignment.

persistent part, then a decomposition over edges is more appropriate. In that case it is straightforward to extract the unambiguous part, but at the same time it is NP-hard to construct an optimal (and globally consistent) assignment from the individual edges in a decomposition.

2 Notation and Background

2.1 MAP Inference as an Optimisation Problem

For a set of n discrete variables $\mathbf{x} = \{x_1, \dots, x_n\}$ taking values from a finite set S we define the energy of a pairwise MRF factorising over a graph $G = (\mathcal{V}, \mathcal{E})$ according to:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \theta_i(x_i) + \sum_{(i,j) \in \mathcal{E}} \theta_{i,j}(x_i, x_j), \quad (1)$$

where the functions $\theta_i(\cdot): S \rightarrow \mathbb{R}$, $\theta_{i,j}(\cdot, \cdot): S \times S \rightarrow \mathbb{R}$ denote the corresponding unary and pairwise potentials, respectively. The *maximum a posteriori* (MAP) problem, that is, computing an assignment with the highest probability is equivalent to the problem of finding an assignment which minimises the energy.

Probably the most popular method for solving this problem is based on the linear programming (LP) relaxation technique. For this purpose, the MAP problem is first represented as an (equivalent) integer linear problem (ILP):

$$\underset{\boldsymbol{\mu} \in \mathcal{X}_G}{\text{minimise}} \quad \boldsymbol{\theta}^\top \boldsymbol{\mu} \quad (2)$$

where $\boldsymbol{\mu}$ corresponds to a joint variable assignment \mathbf{x} in the standard overcomplete representation [21]. That is, $\boldsymbol{\theta}$ is a vector with entries $\theta_i(x_i)$ for all $i \in \mathcal{V}$, $x_i \in S$ and $\theta_{i,j}(x_i, x_j)$ for all $(i, j) \in \mathcal{E}$, $x_i, x_j \in S$, and $\boldsymbol{\mu}$ is a binary vector of indicator functions for nodes and edges $\mu_i(x_i), \mu_{i,j}(x_i, x_j) \in \{0, 1\}$, where $\mu_i(s) = 1 \Leftrightarrow x_i = s$ and $\mu_{i,j}(s_1, s_2) = 1 \Leftrightarrow x_i = s_1 \wedge x_j = s_2$. The set \mathcal{X}_G corresponds to all valid assignments of a pairwise MRF over a graph G and has the following compact representation:

$$\mathcal{X}_G := \left\{ \boldsymbol{\mu} \in \mathbb{R}^d \left| \begin{array}{l} \sum_{x_i} \mu_i(x_i) = 1 \\ \forall i \in \mathcal{V} \\ \sum_{x_i} \mu_{i,j}(x_i, x_j) = \mu_j(x_j) \\ \forall (i, j) \in \mathcal{E}, \forall x_j \in S \\ \sum_{x_j} \mu_{i,j}(x_i, x_j) = \mu_i(x_i) \\ \forall (i, j) \in \mathcal{E}, \forall x_i \in S \\ \mu_i(x_i) \in \{0, 1\} \\ \forall i \in \mathcal{V}, \forall x_i \in S \\ \mu_{i,j}(x_i, x_j) \in \{0, 1\} \\ \forall (i, j) \in \mathcal{E}, \forall x_i, x_j \in S \end{array} \right. \right\} \quad (3)$$

A convex hull of this set, which we denote by $\mathcal{M}_G := \text{conv } \mathcal{X}_G$ plays a special role in the optimisation and is known as the *marginal polytope* of a corresponding MRF. Namely, the problem (2) is equivalent to the one where we replace the set \mathcal{X}_G by its convex hull \mathcal{M}_G , that is

$$\min_{\boldsymbol{\mu} \in \mathcal{X}_G} \boldsymbol{\theta}^\top \boldsymbol{\mu} = \min_{\boldsymbol{\mu} \in \mathcal{M}_G} \boldsymbol{\theta}^\top \boldsymbol{\mu}. \quad (4)$$

Since finding an optimal solution of the above ILP or equivalently minimising its linear objective over the marginal polytope is in general intractable, we usually consider the following relaxation:

$$\text{minimise}_{\boldsymbol{\mu} \in L_G} \boldsymbol{\theta}^\top \boldsymbol{\mu} \quad (5)$$

where we optimise over a bigger set $L_G \supseteq \mathcal{M}_G \supseteq \mathcal{X}_G$ called the *local consistency polytope* of a MRF over a graph G , which results from relaxing the integrality constraints $\mu_i(x_i), \mu_{i,j}(x_i, x_j) \in \{0, 1\}$ in the definition of \mathcal{X}_G by allowing the corresponding variables to take all real values in the interval $[0, 1]$. That is,

$$L_G := \left\{ \boldsymbol{\mu} \in \mathbb{R}^d \left| \begin{array}{l} \sum_{x_i} \mu_i(x_i) = 1 \\ \forall i \in \mathcal{V} \\ \sum_{x_i, x_j} \mu_{i,j}(x_i, x_j) = \mu_j(x_j) \\ \forall (i, j) \in \mathcal{E}, \forall x_j \\ \sum_{x_j} \mu_{i,j}(x_i, x_j) = \mu_i(x_i) \\ \forall (i, j) \in \mathcal{E}, \forall x_i \\ \mu_{i,j}(x_i, x_j) \geq 0 \\ \forall (i, j) \in \mathcal{E}, \forall x_i, x_j \in S \end{array} \right. \right\} \quad (6)$$

Note that the non-negativity of the unary variables $\mu(x_i) \geq 0$ implicitly follows from the combination of the agreement constraints between node and edge variables and the non-negativity of the latter.

2.2 Optimisation via Dual Decomposition

We now briefly review the DD framework for MAP inference in (pairwise) MRFs [12]. The main idea is to decompose the original intractable optimisation problem (OP) in (2) over a graph G into a set of tractable inference problems over subtrees $\{\mathcal{T}_j\}_{j=1}^m$, $\mathcal{T}_j \subseteq G$, which are coupled by a set of agreement constraints to ensure the consistency. That is, each of the individual subproblems $j \in \{1, \dots, m\}$ corresponds to the MAP inference on a subtree $\mathcal{T}_j = (\mathcal{V}_j, \mathcal{E}_j)$ of the original MRF. More precisely, we define the following OP

$$\text{minimise}_{\boldsymbol{\mu}^1 \in \mathcal{X}_{\mathcal{T}_1}, \dots, \boldsymbol{\mu}^m \in \mathcal{X}_{\mathcal{T}_m}, \boldsymbol{\nu}} \sum_{j=1}^m \boldsymbol{\theta}^j \top \boldsymbol{\mu}^j \quad (7)$$

$$\text{subject to} \quad \begin{array}{l} \mu_i^j(x_i) = \nu_i(x_i) \\ \forall j \in \{1, \dots, m\}, \forall i \in \mathcal{V}_j, \forall x_i \in S \end{array}$$

where each vector $\boldsymbol{\mu}^j$ denotes the variables of a local subproblem with respect to \mathcal{T}_j , and $\boldsymbol{\nu}$ is a set of global

variables $\nu_i(x_i)$ on which the variables $\mu_i^j(x_i)$ of the (overlapping) subproblems must agree. We can choose any decomposition with the only condition that the corresponding trees together must cover all the nodes and edges of G , that is, $\mathcal{V} = \bigcup_{j=1}^m \mathcal{V}_j$ and $\mathcal{E} = \bigcup_{j=1}^m \mathcal{E}_j$, as well as $\boldsymbol{\theta}^\top \boldsymbol{\mu} = \sum_{j=1}^m \boldsymbol{\theta}^j \top \boldsymbol{\mu}^j$. Note that OP in (7) is equivalent to the ILP in (2). A corresponding LP relaxation given by

$$\text{minimise}_{\boldsymbol{\mu}^1 \in L_{\mathcal{T}_1}, \dots, \boldsymbol{\mu}^m \in L_{\mathcal{T}_m}, \boldsymbol{\nu}} \sum_{j=1}^m \boldsymbol{\theta}^j \top \boldsymbol{\mu}^j \quad (8)$$

$$\text{subject to} \quad \begin{array}{l} \mu_i^j(x_i) = \nu_i(x_i) \\ \forall j \in \{1, \dots, m\}, \forall i \in \mathcal{V}_j, \forall x_i \in S \end{array}$$

is equivalent to the OP in (5) in the sense that both have the same optimal value and the same optimal solution set.

In the corresponding dual problems the goal is to maximise the dual function of the OPs in (7) and (8) according to

$$\text{maximise}_{\mathbf{u} \in \mathcal{U}} g_7(\mathbf{u}), \quad (9)$$

where

$$g_7(\mathbf{u}) = \inf_{\boldsymbol{\mu}^1 \in \mathcal{X}_{\mathcal{T}_1}, \dots, \boldsymbol{\mu}^m \in \mathcal{X}_{\mathcal{T}_m}} \left\{ \sum_{j=1}^m (\boldsymbol{\theta}^j + \mathbf{u}^j)^\top \boldsymbol{\mu}^j \right\}$$

and

$$\text{maximise}_{\mathbf{u} \in \mathcal{U}} g_8(\mathbf{u}), \quad (10)$$

where

$$g_8(\mathbf{u}) = \inf_{\boldsymbol{\mu}^1 \in L_{\mathcal{T}_1}, \dots, \boldsymbol{\mu}^m \in L_{\mathcal{T}_m}} \left\{ \sum_{j=1}^m (\boldsymbol{\theta}^j + \mathbf{u}^j)^\top \boldsymbol{\mu}^j \right\}$$

respectively, over a restricted set of dual values

$$\mathcal{U} := \left\{ \mathbf{u}: \sum_{j: i \in \mathcal{V}_j} \mathbf{u}_i^j(x_i) = 0, i \in \{1, \dots, n\}, x_i \in S \right\} \quad (11)$$

We here overload the notation in the following sense. The dual variables have the following form $\mathbf{u} = (\mathbf{u}^1, \dots, \mathbf{u}^m)$, $\mathbf{u}^j = (\dots, u_i^j(x_i), \dots)$. Therefore, since we ignore the edges, the number of dual variables \mathbf{u}^j is smaller than the dimensionality of $\boldsymbol{\mu}^j$ (or $\boldsymbol{\theta}^j$). However, for the algebraic operations (e.g. inner product) to make sense we implicitly assume that the vector \mathbf{u}^j (if required) is appropriately filled with zeros to get the same dimensionality as $\boldsymbol{\mu}^j$. A derivation of the above dual problems can be found in [12]. There are different ways to solve a corresponding dual problem. The most popular is a subgradient method for convex non differentiable objectives. Alternatively, we could use a variant of block coordinate descent or cutting plane algorithm.

3 Connections between LP Relaxation and DD

The facts summarised in Subsection 3.1 are mainly known. We provide them for the sake of completeness. In Subsection 3.2 we present new insights in the connections between the two optimisation techniques.

3.1 Case 1: LP Relaxation yields an Integral Solution

It is well known that if the LP relaxation is tight, both the LP relaxation and DD provide an integral optimal solution to the MAP problem. From a different perspective, this means that strong duality holds for the OP in (7). That is, there is zero duality gap between optimal values of OPs in (7) and (9). Equivalently, it implies the existence of consistent optimal assignments to subtrees according to a chosen decomposition. We summarise these insights in the following lemma.

Lemma 1 *The following claims are equivalent:*

- (i) LP relaxation in (5) has an integral solution
- (ii) strong duality holds for problem (7)
- (iii) $\bar{\mu}_i^{j_1}(x_i) = \bar{\mu}_i^{j_2}(x_i) \quad \forall j_1, j_2 \in \{1, \dots, m\}, i \in \mathcal{V}_{j_1} \cap \mathcal{V}_{j_2}, x_i \in S$

where $\bar{\mu} = (\bar{\mu}^1, \dots, \bar{\mu}^m)$ is a (not necessarily unique) minimiser of the Lagrangian $\mathcal{L}(\cdot, \dots, \cdot, \mathbf{u}^*)$ for OP in (7) and \mathbf{u}^* is a dual optimal.

3.2 Case 2: LP Relaxation yields a Fractional Solution

If the LP relaxation is not tight, a corresponding optimal solution μ^* will have fractional components. Given such a fractional solution, we denote by $\mathcal{I} \subseteq \{1, \dots, n\}$ the indices of the variables x_1, \dots, x_n , which have been assigned an integral value in μ^* , and by $\mathcal{F} \subseteq \{1, \dots, n\}$ the remaining set of indices corresponding to the fractional part. Formally, $i \in \mathcal{I} \Leftrightarrow \forall x_i \in S: \mu_i^*(x_i) \in \{0, 1\}$ and $\mathcal{F} = \{1, \dots, n\} \setminus \mathcal{I}$. In contrast to LP relaxation, assignments produced via DD are fully integral. Given a tree decomposition of a corresponding MRF, the optimal assignments to different subtrees, however, will partially disagree on the overlapping parts. Here we denote by $\mathcal{A} \subseteq \{1, \dots, n\}$ the indices of the variables x_1, \dots, x_n with unique assignment. Formally, $i \in \mathcal{A}$, if for every tree which contains x_i and every optimal assignment to that tree, x_i has the same value. Similarly we denote by $\mathcal{D} \subseteq \{1, \dots, n\}$ the set of indices for which at least two different trees disagree on their optimal assignments, that is, $\mathcal{D} = \{1, \dots, n\} \setminus \mathcal{A}$. We now provide a formal analysis of the relationship between the sets \mathcal{I} and \mathcal{A} , or equivalently between \mathcal{F} and \mathcal{D} .

Theorem 1 *Let μ^* be an optimal fractional solution of the LP relaxation (5), $\bar{\mathbf{u}}$ a dual optimal for OP in (7), and $\mathcal{L}: \mathcal{X}_{\mathcal{T}_1} \times \dots \times \mathcal{X}_{\mathcal{T}_m} \times \mathcal{U} \rightarrow \mathbb{R}$ a corresponding Lagrangian. There always exists a set of minimisers $\bar{\mu}^1 \in \mathcal{X}_{\mathcal{T}_1}, \dots, \bar{\mu}^m \in \mathcal{X}_{\mathcal{T}_m}$ of the Lagrangian $\mathcal{L}(\cdot, \dots, \cdot, \bar{\mathbf{u}})$ which agree with the integral part \mathcal{I} of μ^* , that is,*

$$\forall j \in \{1, \dots, m\}, i \in \mathcal{I} \cap \mathcal{V}_j, x_i \in S: \quad \bar{\mu}_i^j(x_i) = \mu_i^*(x_i),$$

for short $\bar{\mu}_{\mathcal{I}}^j = \mu_{\mathcal{I}}^*$.

The result in the above theorem has the most intuitive interpretation in the case of a decomposition into spanning trees, that is, when every tree \mathcal{T}_j covers all the nodes ($\mathcal{V}_j = \mathcal{V}$) of the original graph. In that case Theorem 1 implies that for any optimal solution of the LP relaxation and a corresponding assignment \mathbf{x}^* with an integral part \mathcal{I} , there exist optimal assignments $\mathbf{x}^1, \dots, \mathbf{x}^m$ (from a dual solution $\bar{\mathbf{u}}$) for the different spanning trees which agree on a set of nodes \mathcal{A} with $x_i^j = x_i^*$ for all $i \in \mathcal{I} \subseteq \mathcal{A}$. An immediate question arising is whether the two sets \mathcal{I} and \mathcal{A} are equal. The answer is no. In general, the two sets are not the same and \mathcal{I} will usually be a proper subset of \mathcal{A} .

Theorem 1 motivates the following simple heuristic for getting an approximate integral solution, which is especially suitable for a decomposition over spanning trees when using subgradient optimisation. Namely, we can consider optimal assignments for every spanning tree and choose the best according to the value of the primal objective. Increasing the number of trees in a decomposition also increases the chance of finding a good assignment. Furthermore, during the optimisation we can repeat this for the intermediate results after each iteration of a corresponding optimisation algorithm saving the currently best solution. Obviously, this only can improve the quality of the resulting assignment. Note that it is the usual praxis with subgradient methods to save the intermediate results since the objective is not guaranteed to improve in every step but can even get worse. In that sense, the above heuristic does not impose additional computational cost.

Finally, it turns out that in a non degenerate case, where LP relaxation has a unique solution, the relationship $\mathcal{I} \subseteq \mathcal{A}$ (and therefore $\mathcal{D} \subseteq \mathcal{F}$) holds for all minimiser of a corresponding Lagrangian $\mathcal{L}(\cdot, \dots, \cdot, \bar{\mathbf{u}})$ supported by the following theorem.

Theorem 2 *Let μ^* be a unique optimal solution of the LP relaxation (5), $\bar{\mathbf{u}}$ a dual optimal for OP in (7), and $\mathcal{L}: \mathcal{X}_{\mathcal{T}_1} \times \dots \times \mathcal{X}_{\mathcal{T}_m} \times \mathcal{U} \rightarrow \mathbb{R}$ a corresponding Lagrangian. Each set of minimisers $\bar{\mu}^1 \in \mathcal{X}_{\mathcal{T}_1}, \dots, \bar{\mu}^m \in \mathcal{X}_{\mathcal{T}_m}$ of the Lagrangian $\mathcal{L}(\cdot, \dots, \cdot, \bar{\mathbf{u}})$ agrees with the integral part \mathcal{I} of μ^* , that is,*

$$\forall j \in \{1, \dots, m\}, i \in \mathcal{I} \cap \mathcal{V}_j, x_i \in S: \quad \bar{\mu}_i^j(x_i) = \mu_i^*(x_i)$$

for short $\bar{\mu}_{\mathcal{I}}^j = \mu_{\mathcal{I}}^*$.

In particular, for binary pairwise MRFs this result implies that each assignment obtained via DD is partially optimal in a sense that it always contains the strongly persistent part of the (fractional) solution of the LP relaxation. Provided the fractional part is small, it suggests that the obtained assignments will often have a low energy close to the optimum even if the LP relaxation is not tight. This fact and the possibility of a parallel computation renders the dual decomposition a practical tool for the MAP inference upon the LP relaxation. We also note that the property $\mathcal{I} \subseteq \mathcal{A}$ in Theorem 1 still holds for non binary MRFs with arbitrary higher order cliques, but \mathcal{I} is not guaranteed to be strongly persistent anymore. In the following we build on an additional lemma.

Lemma 2 *Assume the setting of Theorem 1. For every subproblem $j \in \{1, \dots, m\}$ over a tree \mathcal{T}_j there are minimisers $\bar{\mu}^j, \hat{\mu}^j \in \mathcal{X}_{\mathcal{T}_j}$, where*

$$\mu_i^*(x_i) = \frac{1}{2}(\bar{\mu}_i^j(x_i) + \hat{\mu}_i^j(x_i)) \quad (12)$$

holds for all $i \in \mathcal{V}_j, x_i \in S$.

The above lemma ensures that for each optimal solution μ^* of the LP relaxation, for each tree in a given decomposition there always exist two different assignments which agree exactly on the nodes corresponding to the integral (strongly persistent) part \mathcal{I} of μ^* . It is also worth noting that when LP relaxation is not tight the sets of minimising assignments to the different trees are disjoint on the overlapping parts due to Lemma 1.

Lemma 2 and Theorem 2 together imply that the unambiguous part among all optimal assignments from all trees in a decomposition coincides with the integral part of the (fractional) optimal assignment via LP relaxation giving rise to the following theorem.

Theorem 3 *Let μ^* be a unique optimal solution for the LP relaxation (5) and $\bar{\mathbf{u}}$ a dual optimal for OP in (7). Then the unambiguous part \mathcal{A} of optimal assignments among all the overlapping subproblems in a decomposition coincides with the integral part \mathcal{I} of μ^* . That is, $\mathcal{A} = \mathcal{I}$.*

That is, we can extract the strongly persistent part from DD by considering the intersection of optimal assignments for individual trees. This is in particular convenient for a decomposition over single edges, since computing the set of optimal assignments for an edge is straightforward.

4 Related Work

Several previous works [20, 9, 11, 16] including the original paper [12] on DD for MAP inference comment on the optimality of DD and LP relaxation when the latter is tight. While in [12] the authors make an explicit statement, in [20] the same question is approached by introducing a notion of a tree agreement, which is generalised to a weak tree agreement in [9]. Both, however, can be seen as a special case of [12] in the sense that the corresponding lower bound on the optimal value of the MAP problem is equal to the optimal value of a corresponding dual problem in case of tree agreement of all the involved subproblems.

Concerning the case where the LP relaxation is not tight, to the best of our knowledge, there are no previous works which directly address the question how the optimal solutions of the LP relaxation are related to the assignments produced via DD. Instead the existing works proceed with a discussion of tightening a corresponding relaxation [12, 19, 17]. In [10] the authors provide a related discussion for tree-reweighted (TRW) message passing and show that the partial assignment corresponding to the unambiguous part in the intersection of all optimal assignments for individual trees is strongly persistent. The TRW algorithms are less accurate than DD, however, in case of binary pairwise MRFs both achieve dual optimal value. Our results are based on a different proof than in [10], which supports a similar statement that the unambiguous part in a corresponding assignment is strongly persistent, but additionally implies that it is exactly the integral part of the LP relaxation when the corresponding solution is unique. Furthermore, our result in Theorem 1 extends also to the case of arbitrary graphs.

Another related question to the discussion in the present paper is the problem of recovering optimal primal solutions of an LP relaxation from dual optimal solutions via DD [12, 18, 14, 1]. In particular, [1] shows that the strongly persistent part of an LP solution (for a binary pairwise MRF) can be recovered from the DD based on subgradient optimisation. However, it is not the case with other optimisation methods. In contrast, our corresponding result holds independently of the chosen optimisation technique.

In the case of binary pairwise MRFs each optimal solution of the LP relaxation is known to be strongly persistent. In this paper we show that (in a non-grease case) this property is inherited by every solution provided via DD in the sense that persistent part is a subset of the corresponding assignment. Since the fractional part of a solution of the LP relaxation conveys no useful information about the preference of the variables in the fractional area to be in specific state, the pre-

sented result further supports the practical usefulness of the DD in that case.

5 Numerical Validation

Here we present a numerical experiment summarised in Figure 1 to validate the theoretical statements in the paper. For this purpose we considered a 5×5 Ising grid model corresponding to 25 pixels and defined its energy according to the following procedure. The unary potentials have been all set to zero. The values of the corresponding edge potentials have been selected uniformly at random from an interval $[-0.5, 0.5]$. This process has been repeated until the corresponding LP relaxation yielded a fractional solution μ^* (see the fourth plot in the first row).

Given such an energy, we then considered a decomposition of the above Ising model into two spanning trees \mathcal{T}_1 (all the vertical edges) and \mathcal{T}_2 (all the horizontal edges). We used the subgradient method to solve a corresponding dual problem and got optimal assignments $\bar{\mu}^1$ and $\bar{\mu}^2$ corresponding to subproblems \mathcal{T}_1 and \mathcal{T}_2 , respectively (see the first two plots in the first row).

The last two plots in the upper row in Figure 1 can serve as a validation of Theorem 1. Namely, the red area in these two plots corresponds to the variables on which each assignment disagrees with the LP solution μ^* . As we can see it happens only for the fractional area, where each variable in μ^* has the value 0.5. In other words, $\bar{\mu}^1$ and $\bar{\mu}^2$ agree with the integral part of μ^* .

To support Theorem 2 we computed further optimal assignments $\bar{\mu}_2^1, \dots, \bar{\mu}_7^1$ for the subproblem \mathcal{T}_1 additionally to $\bar{\mu}^1$. These are visualised in the second row in Figure 1. We overlay these assignments with the assignment μ^* in a transparent way to emphasise the difference to the LP solution. We see that all these optimal assignments agree with the integral part of μ^* .

The last plot in the third row validates the statement in Lemma 2. It visualises the average of the the two plots above corresponding to the optimal assignments $\bar{\mu}_6^1$ and $\bar{\mu}_7^1$ for the subproblem \mathcal{T}_1 .

Finally, the first plot in the third row supports the claim in Theorem 3. Namely, the blue area corresponds to the unambiguous part \mathcal{A} where each variable has a unique value in all the optimal assignments among the two subproblems. The red area marks the ambiguous part, where each variable has different values in different assignments. We see that the equality $\mathcal{A} = \mathcal{I}$ holds.

6 Conclusion

We presented the established frameworks of linear programming (LP) relaxation and dual decomposition (DD) for the task of MAP inference in discrete MRFs. In the case when a corresponding LP relaxation is tight, these two methods are known to be equivalent both providing an optimal MAP assignment. However, less is known about their relationship in the opposite and more realistic case. While it is known that both methods have the same optimal objective value also in the non tight regime, it is an interesting question if there are other properties they share. In particular, the connection between the solutions of LP relaxation and the assignments which can be extracted by DD has not been clarified. For example, even if the solution of LP relaxation is unique, there might be multiple optimal (but disagreeing) assignments via DD. What is the nature of this ambiguity? Are all these assignments equivalent or can we even extract additional information about the optimal solutions from analysing the disagreement behaviour? These and other questions were the main motivation for our paper. Here we successfully provided a few novel findings explaining how the fully integral assignments obtained via DD agree with the optimal fractional assignments via LP relaxation when the latter is not tight. More specifically, we have proved:

- given an optimal (fractional) solution of the LP relaxation, there always exists an optimal variable assignment via DD which agrees with the integral part of the LP solution; this also holds for non binary models with arbitrary higher order potentials
- for binary pairwise MRFs (in a non degenerate case) the first result holds for every optimal assignment which can be extracted via DD
- for binary pairwise MRFs (in a non degenerate case) the unambiguous part among all optimal assignments from different trees in a decomposition coincides with the integral part of the (fractional) optimal assignment via LP relaxation

In particular, for binary pairwise MRFs the integral part of an optimal solution provided via LP relaxation is known to be strongly persistent. Therefore, due to the properties listed above, we can conclude that (for this case) both methods LP relaxation and DD share the partial optimality property of their solutions.

Practically, it has the following implications. 1) If the goal is to find an accurate MAP assignment we can use LP relaxation first to fix the integral part and then apply an approximation algorithm (e.g. loopy belief




Figure 1: Illustration of numerical validation of the theoretical results including Theorem 1, 2, 3 and Lemma 2. We consider a decomposition of an 5×5 Ising grid model into two spanning trees (more precisely, disconnected forests): \mathcal{T}_1 consisting of all vertical edges and \mathcal{T}_2 consisting of all horizontal edges. The corresponding optimal assignments for both trees are denoted by $\bar{\mu}^1$ and $\bar{\mu}^2$, respectively. μ^* is a corresponding optimal assignment from the LP relaxation. Here, the red area corresponds to label 1, blue area to label 0, and green area to label 0.5. The individual plots in the first row visualise a corresponding assignment or a disagreement between two assignments. The second row provides (additionally to $\bar{\mu}^1$) further optimal assignments for the subproblem \mathcal{T}_1 . The first plot in the third row illustrates the unambiguous part among all optimal and overlapping assignments in blue and the ambiguous part in red. The last plot illustrates an average of the two assignments above.

propagation) to set the remaining variables in the fractional part – each fractional node has the value 0.5 showing no preference of a corresponding variable to be in a specific state. On the other hand, DD provides fully integral assignments, which agree with the integral part of LP relaxation. 2) If we are only interested in the persistent part, we can extract the corresponding partial assignment from DD by considering an intersection of all the optimal assignments to the individual trees in a decomposition. The unambiguous part then coincides with the strongly persistent part of LP relaxation, provided a corresponding solution is unique, or to a subset of it in the opposite case. In particular, considering a decomposition over individual edges, is more beneficial in that case. Namely, finding all optimal assignments to trees becomes a trivial task.

To summarise, the LP relaxation is a popular method for discrete MAP inference in pairwise graphical models because of its appealing (partial) optimality properties. However, this method does not scale nicely due to the extensive memory requirements restricting its practical use for bigger problems. Here, DD provides an effective alternative via distributed optimisation, which scales to problems of arbitrary size — we can always consider

a decomposition over the individual edges. Finally, the results presented in this paper suggest, that when using DD instead of the LP relaxation we do not lose any of the nice properties of the latter. Both methods provide (for binary pairwise models) exactly the same information about their solutions.

An interesting question is which findings in the paper can be extended beyond the pairwise models involving higher order potentials. We will continue this investigation in the future works.

Acknowledgments

This research was supported by the Federal Ministry of Education and Research under the Berlin Big Data Center 2 project (FKz 01IS18025A), and by the World Class University Program through the National Research Foundation of Korea funded by the Ministry of Education, Science, and Technology, under Grant R31-10008.

References

- [1] K. M. Anstreicher and L. A. Wolsey. Two "well-known" properties of subgradient optimization. *Math. Pro-*

- gram.*, 120(1):213–220, 2009.
- [2] D. Bertsimas and J. N. Tsitsiklis. *Introduction to linear optimization*. Athena scientific series in optimization and neural computation. Athena Scientific, Belmont (Mass.), 1997.
- [3] M. M. Deza and M. Laurent. *Geometry of Cuts and Metrics*. Algorithms and Combinatorics, vol. 15. Springer, 1997.
- [4] H. Everett. Generalized lagrange multiplier method for solving problems of optimum allocation of resources. *Operations Research*, 11(3):399–417, may 1963.
- [5] M. Grötschel, L. Lovasz, and A. Schrijver. *Geometric algorithms and combinatorial optimization*. Algorithms and combinatorics. Springer-Verlag, Berlin, New York, 1988.
- [6] P. L. Hammer, P. Hansen, and B. Simeone. Roof duality, complementation and persistency in quadratic 0-1 optimization. *Math. Program.*, 28(2):121–155, 1984.
- [7] J. H. Kappes, B. Andres, F. A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, T. Kröger, J. Lellmann, N. Komodakis, B. Savchynskyy, and C. Rother. A comparative study of modern inference techniques for structured discrete energy minimization problems. *International Journal of Computer Vision*, 115(2):155–184, 2015.
- [8] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques - Adaptive Computation and Machine Learning*. The MIT Press, 2009.
- [9] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(10):1568–1583, 2006.
- [10] V. Kolmogorov and M. J. Wainwright. On the optimality of tree-reweighted max-product message-passing. In *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence, Edinburgh, Scotland, July 26-29, 2005*, pages 316–323, 2005.
- [11] V. Kolmogorov and M. J. Wainwright. On the optimality of tree-reweighted max-product message-passing. *CoRR*, abs/1207.1395, 2012.
- [12] N. Komodakis, N. Paragios, and G. Tziritas. MRF energy minimization and beyond via dual decomposition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(3):531–552, 2011.
- [13] D. G. Luenberger. *Introduction to Linear and Nonlinear Programming*. Addison-Wesley, 1973.
- [14] A. Nedic and A. E. Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, 2009.
- [15] J. Pearl. *Probabilistic reasoning in intelligent systems - networks of plausible inference*. Morgan Kaufmann series in representation and reasoning. Morgan Kaufmann, 1989.
- [16] A. M. Rush and M. Collins. A tutorial on dual decomposition and lagrangian relaxation for inference in natural language processing. *CoRR*, abs/1405.5208, 2014.
- [17] D. Sontag. *Approximate Inference in Graphical Models using LP Relaxations*. PhD thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, 2010.
- [18] 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.
- [19] D. Sontag, T. Meltzer, A. Globerson, T. S. Jaakkola, and Y. Weiss. Tightening LP relaxations for MAP using message passing. *CoRR*, abs/1206.3288, 2012.
- [20] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. MAP estimation via agreement on trees: message-passing and linear programming. *IEEE Trans. Information Theory*, 51(11):3697–3717, 2005.
- [21] 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.
- [22] J. Wang and S. Yeung. A compact linear programming relaxation for binary sub-modular MRF. In *Energy Minimization Methods in Computer Vision and Pattern Recognition - 10th International Conference, EMMCVPR 2015, Hong Kong, China, January 13-16, 2015. Proceedings*, pages 29–42, 2014.