# MPLP++: Fast, Parallel Dual Block-Coordinate Ascent for Dense Graphical Models 

Siddharth Tourani ${ }^{1}$, Alexander Shekhovtsov ${ }^{2}$, Carsten Rother ${ }^{1}$, Bogdan<br>Savchynskyy ${ }^{1}$<br>${ }^{1}$ Visual Learning Lab, IWR, Uni. Heidelberg, Germany<br>${ }^{2}$ Centre for Machine Perception,Czech Technical University,Prague, Czech Republic


#### Abstract

Dense, discrete Graphical Models with pairwise potentials are a powerful class of models which are employed in state-of-the-art computer vision and bio-imaging applications. This work introduces a new MAP-solver, based on the popular Dual Block-Coordinate Ascent principle. Surprisingly, by making a small change to the low-performing solver, the Max Product Linear Programming (MPLP) algorithm [1], we derive the new solver MPLP++ that significantly outperforms all existing solvers by a large margin, including the state-of-the-art solver Tree-Reweighted Sequential (TRW-S) message-passing algorithm [2]. Additionally, our solver is highly parallel, in contrast to TRW-S, which gives a further boost in performance with the proposed GPU and multi-thread CPU implementations. We verify the superiority of our algorithm on dense problems from publicly available benchmarks, as well, as a new benchmark for 6D Object Pose estimation. We also provide an ablation study with respect to graph density.


Keywords: Graphical models, MAP-Inference, Block-Coordinate-Ascent, Dense Graphs, Message Passing Algorithms

## 1 Introduction

Undirected discrete graphical models with dense neighbourhood structure are known to be much more expressive than their sparse counterparts. A striking example is the fully-connected Conditional Random Field (CRF) model with Gaussian pairwise potentials [3], significantly improving the image segmentation field, once an efficient solver for the model was proposed. More recently, various applications in computer vision and bio-imaging have successfully used fully-connected or densely-connected, pairwise models with non-Gaussian potentials. Non-Gaussian potentials naturally arise from application-specific modelling or the necessity of robustness potentials. A prominent application of the non-Gaussian fully-connected CRF case achieved state-of-theart performance in 6D object pose estimation problem [4], with an efficient applicationspecific solver. Other examples of densely connected models were proposed in the area of stereo-reconstruction [5], body pose estimation [6, 7, 8], bio-informatics [9] etc.

An efficient solver is a key condition to make such expressive models efficient in practice. This work introduces such a solver, which outperforms all existing methods for a class of dense and semi-dense problems with non-Gaussian potentials. This includes Tree-Reweighted Sequential (TRW-S) message passing, which is typically used
for general pairwise models. We would like to emphasize that efficient solvers for this class are highly desirable, even in the age of deep learning. The main reason is that such expressive graphical models can encode information which is often hard to learn from data, since very large training datasets are needed to learn the application specific prior knowledge. In the the above mentioned 6D object pose estimation task, the pairwise potentials encode length-consistency between the observed data and the known 3D model and the unary potentials are learned from data. Other forms of combining graphical models with CNNs such as Deep-Structured-Models [10] can also benefit from the proposed solver.

Linear Programming (LP) relaxation is a powerful technique that can solve exactly all known tractable maximum a posteriori (MAP) inference problems for undirected graphical models (those known to be polynomially solvable) [11]. Although there are multiple algorithms addressing the MAP inference, which we discuss in $\S 2$, the linear programs obtained by relaxing the MAP-inference problem are not any simpler than general linear programs [12]. This implies that algorithms solving it exactly are bounded by the computational complexity of the general LP and do not scale well to problems of large size. Since LP relaxations need not be tight, solving it optimally is often impractical. On the other hand, block coordinate ascent (BCA) algorithms for the LP dual problem form an approach delivering fast and practically useful approximate solutions. TRW-S [2] is probably the most well-known and efficient solver of this class, as shown in [13]. Since due to the graph density the model size grows quadratically with the number of variables, a scalable solver must inevitably be highly parallelizable to be of practical use. Our work improves another well-known BCA algorithm of this type, MPLP [1] (Max Product Linear Programming algorithm) and proposes a parallel implementation as explained next.
Contribution We present a new state-of-the-art parallel BCA algorithm of the same structure as MPLP, i.e. an elementary step of our algorithm updates all dual variables related to a single graph edge. We explore the space of such edge-wise updates and propose that a different update rule inspired by [14] can be employed, which significantly improves the practical performance of MPLP. Our method with the new update is termed MPLP++. The difference in the updates stems from the fact that the optimization in the selected block of variables is non-unique and the way the update utilizes the degrees of freedom to which the block objective is not sensitive significantly affects subsequent updates and thus the whole performance.

We propose the following theoretical analysis. We show that MPLP++ converges towards arc-consistency, similarly to the convergence result of [15] for min-sum diffusion. We further show, that given any starting point, an iteration of the MPLP++ algorithm, which processes all edges of the graph in a specified order always results in a better objective than the same iteration of MPLP. For multiple iterations this is not theoretically guaranteed, but empirically observed in all our test cases. All proofs relating to the main paper are given in the supplement.

Another important aspect that we address is parallelization. TRW-S is known as a "sequential" algorithm. However it admits parallelization, especially for bipartite graphs [2], which is exploited in specialized implementations [16, 17] for 4-connected grid graphs. The parallelization there gives a speed-up factor of $O(n)$, where $n$ is the number of nodes in the graph. A parallel implementation for dense graphs has not been
proposed. We observe that in MPLP a group of non-incident edges can be updated in parallel. We pre-compute a schedule maximizing the number of edges that can be processed in parallel using an exact or a greedy maximum matching. The obtainable theoretical speed-up is at least $n / 2$ for any graph, including dense ones. We consider two parallel implementations, suitable for CPU and GPU architectures respectively. A further speed-up is possible by utilizing parallel algorithms for lower envelopes in message passing (see $\S 4$ and [18]).

The new MPLP++ method consistently outperforms all its competitors, including TRW-S [2], in the case of densely (not necessarily fully) connected graphs, even in the sequential setting: In our experiments it is 2 to 10 times faster than TRW-S and 5 to 150 times faster than MPLP depending on the dataset and the required solution precision. The empirical comparison is conducted on several datasets. As there are only few publicly available ones, we have created a new dataset related to the 6 D pose estimation problem [4]. Our code and the new benchmark dataset will be made publicly available together with the paper.

## 2 Related work

The general MAP inference problem for discrete graphical models (formally defined in $\S 3$ ) is NP-hard and is also hard to approximate [19]. A natural linear programming relaxation is obtained by formulating it as a $0-1$ integer linear program (ILP) and relaxing the integrality constraints [20] (see also the recent review [21]). A hierarchy of relaxations is known [22], from which the so-called Base LP relaxation, also considered in our work, is the simplest one. It was shown [11,23] that this relaxation is tight for all tractable subclasses of the problem. For many other classes of problems it provides approximation guarantees, reviewed [19]. A large number of specialized algorithms were developed for this relaxation.

Apart from general LP solvers, a number of specialized algorithms exist that take advantage of the problem structure and guarantee convergence to an optimal solution of the LP relaxation. This includes proximal [24, 25, 26, 27], dual sub-gradient [28, 29, 30], bundle [31], mirror-descent [32] smoothing-based [33, 34, 35] and (quasi-) Newton [36] methods.

However, as it was shown in [37], the linear programs arising from the relaxation of the MAP inference have the same computational complexity as general LPs. At the same time, if the problem is not known to belong to a tractable class, solving the relaxation to optimality may be of low practical utility. A comparative study [13] notes that TRW-S [2] is the most efficient solver for the relaxation in practice. It belongs to the class of BCA methods for the LP dual that includes also MPLP [1], min-sum diffusion [38, 15] and DualMM [14] algorithms. BCA methods are not guaranteed to solve the LP dual to optimality. They may get stuck in a suboptimal point and be unable to compute primal LP solutions unless integer solutions are found (which is not always the case). However, they scale extremely well, take advantage of fast dynamic programming techniques, solve exactly all submodular problems [39] and provide good approximate solutions in general vision benchmarks. Suboptimal solutions of the relaxation can also be employed in exact solvers [40, 41] using cutting plane or branch-and-cut techniques and in methods identifying a part of optimal solution [42, 43].

## 3 Preliminaries

Notation $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ denotes an undirected graph, with vertex set $\mathcal{V}$ (we assume $\mathcal{V}=$ $\{1, \ldots|\mathcal{V}|\})$ and edge set $\mathcal{E}$. The notation $u v \in \mathcal{E}$ will mean that $\{u, v\} \in \mathcal{E}$ and $u<v$ with respect to the order of $\mathcal{V}$. Each node $u \in \mathcal{V}$ is associated with a label from a finite set of labels $\mathcal{Y}$ (for brevity w.l.o.g. we will assume that it is the same set for all nodes). The label space for a pair of nodes $u v \in \mathcal{E}$ is $\mathcal{Y}^{2}$ and for all nodes it is $\mathcal{Y}^{\mathcal{V}}$.

For each node and edge the unary $\theta_{u}: \mathcal{Y} \rightarrow \mathbb{R}, u \in \mathcal{V}$ and pairwise cost functions $\theta_{u v}: \mathcal{Y}^{2} \rightarrow \mathbb{R}, u v \in \mathcal{E}$ assign a cost to a label or label pair respectively. Let $\mathcal{I}=$ $(\mathcal{V} \times \mathcal{Y}) \cup\left(\mathcal{E} \times \mathcal{Y}^{2}\right)$ be the index set enumerating all labels and label pairs in neighbouring graph nodes. Let the cost vector $\theta \in \mathbb{R}^{\mathcal{I}}$ contain all values of the functions $\theta_{u}$ and $\theta_{u v}$ as its coordinates.

The MAP-inference problem for the graphical model defined by a triple $\left(\mathcal{G}, \mathcal{Y}^{\mathcal{V}}, \theta\right)$ consists in finding the labelling with the smallest total cost, i.e.:

$$
\begin{equation*}
y^{*}=\underset{y \in \mathcal{Y}^{\mathcal{V}}}{\arg \min }\left[E(y \mid \theta):=\sum_{v \in \mathcal{V}} \theta_{v}\left(y_{v}\right)+\sum_{u v \in \mathcal{E}} \theta_{u v}\left(y_{u v}\right)\right] \tag{1}
\end{equation*}
$$

This problem is also known as energy minimization for graphical models and is closely related to weighted and valued constraint satisfaction problems. The total cost $E$ is also often called energy.

The problem (1) is in general NP-hard and is also hard to approximate [19]. A number of approaches to tackle it in different practical scenarios is reviewed in [13, 44]. One of the widely applicable techniques is based on (approximately) solving its linear programming (LP) relaxation as discussed in $\S 2$.
Dual Problem Most existing solvers for the LP relaxation tackle its dual form, which we introduce now. This is because the LP dual has much fewer variables than the primal and can also be written in the form of unconstrained concave (piecewise-linear) maximization. It is based on the fact that the representation of the energy function $E(y \mid \theta)$ using unary $\theta_{u}$ and pairwise $\theta_{u v}$ costs is not unique. There exist other costs $\hat{\theta} \in \mathbb{R}^{\mathcal{I}}$ such that $E(y \mid \hat{\theta})=E(y \mid \theta)$ for all labelings $y \in \mathcal{Y}^{\mathcal{V}}$.

It is known (see e.g. [21]) and straightforward to check that such equivalent costs can be obtained with an arbitrary vector $\phi:=\left(\phi_{v \rightarrow u}(s) \in \mathbb{R} \mid u \in \mathcal{V}, v \in \mathrm{Nb}(u), s \in\right.$ $\mathcal{Y}$ ), where $\mathrm{Nb}(u)$ is the set of neighbours of $u$ in $\mathcal{G}$, as follows:

$$
\begin{align*}
\hat{\theta}_{u}(s) \equiv \theta_{u}^{\phi}(s) & :=\theta_{u}(s)+\sum_{v \in \mathrm{Nb}(u)} \phi_{v \rightarrow u}(s)  \tag{2}\\
\hat{\theta}_{u v}(s, t) \equiv \theta_{u v}^{\phi}(s, t) & :=\theta_{u v}(s, t)-\phi_{v \rightarrow u}(s)-\phi_{u \rightarrow v}(t) .
\end{align*}
$$

The cost vector $\theta^{\phi}$ is called reparametrized and the vector $\phi$ is known as reparametrization. Costs related by (2) are also called equivalent. Other established terms for reparametrization are equivalence preserving [45] or equivalent transformations [20].

By swapping min and $\sum$ operations in (1) one obtains a lower bound on the energy $D\left(\theta^{\phi}\right) \leq E(y \mid \theta)$ for all $y$, which reads

$$
\begin{equation*}
D\left(\theta^{\phi}\right):=\sum_{u \in \mathcal{V}} \min _{s \in \mathcal{Y}} \theta_{u}^{\phi}(s)+\sum_{u v \in \mathcal{E}} \min _{(s, t) \in \mathcal{Y}^{2}} \theta_{u v}^{\phi}(s, t) \tag{3}
\end{equation*}
$$

Although the energy $E(y \mid \theta)$ remains the same for all equivalent cost vectors (e.g. $E(y \mid \theta)=E\left(y \mid \theta^{\phi}\right)$ ), the lower bound $D\left(\theta^{\phi}\right)$ depends on reparametrization, which is $D(\theta) \neq D\left(\theta^{\phi}\right)$. Therefore, a natural maximization problem arises as maximization of the lower bound over all equivalent costs: $\max _{\phi} D\left(\theta^{\phi}\right)$. It is known (e.g. [21]) that this maximization problem can be seen as a dual formulation of the LP relaxation of (1). We will write $D(\phi)$ to denote $D\left(\theta^{\phi}\right)$, since the cost vector $\theta$ is clear from the context.The function $D(\phi)$ is concave, piece-wise linear and therefore non-smooth. In many applications the dimensionality of $\phi$ often exceeds $10^{5}$ to $10^{6}$ and the respective dual problem $\max _{\phi} D(\phi)$ is large scale.

## 4 Dual Block-Coordinate Ascent

As we discussed in $\S 2$, BCA methods, although not guaranteed to solve the dual to the optimality, provide good solutions for many practical instances and scale very well to large problems. The fastest such methods are represented by methods working with chain subproblems [2], [14] or their generalizations [46].

The TRW-S algorithm can be seen as updating a block of dual variables $\left(\phi_{v \rightarrow u}, v \in\right.$ $\mathrm{Nb}(u))$ "attached" to a node $u$, during each elementary step. The same block of variables is also used in the min-sum diffusion algorithm [15], as well as in the convex message passing [47], and [46] gives a generalization of such methods.

However, the update coefficients in TRW-S are related to the density of the graph and its advantage diminishes when the graph becomes dense. We show that for dense graphs updating a block of dual variables $\phi_{u \leftrightarrow v}=\left(\phi_{v \rightarrow u}, \phi_{u \rightarrow v}\right)$ associated to an edge $u v \in \mathcal{E}$ can be more efficient. Such updates were previously used in the MPLP algorithm [1]. We show that our MPLP++ updates differ in detail but bring a significant improvement in performance.

Block Optimality Condition For further analysis of BCA algorithms we will require a sufficient condition of optimality w.r.t. the selected block of dual variables. The restriction of the dual to the block of variables $\phi_{u \leftrightarrow v}$ is given by the function:

$$
\begin{equation*}
D_{u v}\left(\phi_{u \leftrightarrow v}\right):=\min _{s t \in \mathcal{Y}^{2}} \theta_{u v}^{\phi}(s, t)+\min _{s \in \mathcal{Y}} \theta_{u}^{\phi}(s)+\min _{t \in \mathcal{Y}} \theta_{v}^{\phi}(t), \tag{4}
\end{equation*}
$$

Maximizing $D_{u v}$ is equivalent to performing a BCA w.r.t. $\phi_{u \leftrightarrow v}$ for $D(\phi)$. The necessary and sufficient condition of maximum of $D_{u v}$ are given by the following.

Proposition 1. Reparametrization $\phi_{u \leftrightarrow v}$ maximizes $D_{u v}(\cdot)$ iff there exist $(s, t) \in \mathcal{Y}^{2}$ such that s minimizes $\theta_{u}^{\phi}(\cdot)$, t minimizes $\theta_{v}^{\phi}(\cdot)$ and $(s, t)$ minimizes $\theta_{u v}^{\phi}(\cdot, \cdot)$.

This condition is trivial to check, it is a special case of arc consistency [21] or weak tree-agreement [2] when considering a simple graph with one edge. It is also clear that $\phi_{u \leftrightarrow v}$ satisfying this condition is not unique. For BCA algorithms it means that there are degrees of freedom in the block, which do not directly affect the objective. By moving on a plateau, they can nevertheless affect subsequent BCA updates. Therefore the performance of the algorithm will be very much dependent on the particular BCA update rule satisfying Proposition 1.


Fig. 1. Illustration of the considered BCA-updates. The gray boxes in the figure represent graph nodes. The black dots in them the labels. The edges connecting the black dots make up the pairwise costs. The numbers adjacent to the edges and labels are the pairwise and unary costs respectively.

Block Coordinate Ascent Updates Given the form of the restricted dual (4) on the edge $u v$, all BCA-updates can be described as follows. Assume $\theta$ is the current reparametrized cost vector, i.e. $\theta=\bar{\theta}^{\bar{\phi}}$ for some initial $\bar{\theta}$ and the current reparametrization $\bar{\phi}$.

Definition 1. A BCA update takes on the input an edge $u v \in \mathcal{E}$ and costs $\theta_{u v}(s, t)$, $\theta_{u}(s)$ and $\theta_{v}(t)$ and outputs a reparametrization $\phi_{u \leftrightarrow v}$ satisfying Proposition 1. W.l.o.g., we assume that it will also satisfy $\min _{s, t} \theta_{u v}^{\phi}(s, t)=0 .{ }^{1}$

According to (2) a BCA-update results in the following reparametrized potentials:

$$
\begin{equation*}
\theta_{u}^{\phi}=\theta_{u}+\phi_{v \rightarrow u}, \quad \theta_{v}^{\phi}=\theta_{v}+\phi_{u \rightarrow v}, \quad, \theta_{u v}^{\phi}(s, t)=\theta_{u v}-\phi_{v \rightarrow u}-\phi_{u \rightarrow v} . \tag{5}
\end{equation*}
$$

Note that since all reparametrizations constitute a vector space, after a BCA-update $\phi$ we can update the current total reparametrization as $\bar{\phi}:=\bar{\phi}+\phi$.

We will consider BCA-updates of the following form: first construct the aggregated $\operatorname{cost} g_{u v}(s, t)=\theta_{u v}(s, t)+\theta_{u}(s)+\theta_{v}(t)$. This corresponds to applying a reparametrization $\dot{\phi}_{u \leftrightarrow v}=\left(-\theta_{u},-\theta_{v}\right)$, which gives $\theta_{u v}^{\dot{\phi}}=g_{u v}, \theta_{u}^{\circ}=\theta_{v}^{\circ}=0$. After that, a BCA update forms a new reparametrization $\phi_{u \leftrightarrow v}$ such that $\theta^{\dot{\phi}+\phi}$ satisfies Proposition 1.

Such BCA-updates can be represented in the following form, which will be simpler for defining and analysing the algorithms.

Definition 2. Consider a BCA-update using a composite reparametrization $\dot{\phi}_{u \leftrightarrow v}+$ $\phi_{u \leftrightarrow v}$. It can be then fully described by the reparametrization mapping $\gamma: g_{u v} \rightarrow$ $\left(\theta_{u}^{\gamma}, \theta_{v}^{\gamma}\right)$, where $g_{u v} \in \mathbb{R}^{\mathcal{Y}_{u v}}, \theta_{u}^{\gamma}=\phi_{v \rightarrow u}$ and $\theta_{v}^{\gamma}=\phi_{u \rightarrow v}$.

By construction, $\theta_{u}^{\gamma}$ matches the reparametrized unary term $\theta_{u}^{\dot{\phi}+\phi}, \theta_{v}^{\gamma}$ is alike and the reparametrized pairwise term is given by $\theta_{u v}^{\gamma}=g_{u v}-\phi_{v \rightarrow u}-\phi_{u \rightarrow v}=g_{u v}-\theta_{u}^{\gamma}-\theta_{v}^{\gamma}$. In what follows, BCA-update will mean specifically the reparametrization mapping $\gamma$. We define now several BCA-updates that will be studied further.

- The uniform BCA-update is given by the following reparametrization mapping $\mathcal{U}$ :

$$
\begin{equation*}
\theta_{u}^{\mathcal{U}}(s)=\theta_{v}^{\mathcal{U}}(t):=\frac{1}{2} \min _{s^{\prime}, t^{\prime} \in \mathcal{Y}} g_{u v}\left(s^{\prime}, t^{\prime}\right), \forall s, t \in \mathcal{Y} . \tag{U}
\end{equation*}
$$

${ }^{1}$ This fixes the ambiguity w.r.t. a constant that can be otherwise added to the edge potential and subtracted from one of the unary potentials. This constant does not affect the performance of algorithms.

This is indeed just an example, to illustrate the problem of non-uniqueness of the minimizer. It is easy to see that this update satisfies Proposition 1 since both $\theta_{u}^{\mathcal{U}}(\cdot)$ and $\theta_{v}^{\mathcal{U}}(\cdot)$ are constant and therefore any pairwise minimizer of $\theta_{u v}^{\mathcal{U}}$ is consistent with them.

- The MPLP BCA-update is given by the following reparametrization mapping $\mathcal{M}$ :

$$
\begin{align*}
\theta_{u}^{\mathcal{M}}(s) & :=\frac{1}{2} \min _{t \in \mathcal{Y}} g_{u v}(s, t), \forall s \in \mathcal{Y}, \\
\theta_{v}^{\mathcal{M}}(t) & :=\frac{1}{2} \min _{s \in \mathcal{Y}} g_{u v}(s, t), \forall t \in \mathcal{Y} . \tag{M}
\end{align*}
$$

The MPLP algorithm [1] can now be described as performing iterations by applying BCA-update $\mathcal{M}$ to all edges of the graph in a sequence.

- The new MPLP++ BCA-update, that we propose, is based on the handshake operation [14]. It is given by the following procedure defining the reparametrization mapping $\mathcal{H}$ :

$$
\begin{align*}
\theta_{u}^{\mathcal{H}}(s) & :=\theta_{u}^{\mathcal{M}}(s), \quad \theta_{v}^{\mathcal{H}}(s):=\theta_{v}^{\mathcal{M}}(s), \forall s \in \mathcal{Y}, \\
\theta_{v}^{\mathcal{H}}(t) & :=\theta_{v}^{\mathcal{H}}(t)+\min _{s \in \mathcal{Y}}\left[g_{u v}(s, t)-\theta_{v}^{\mathcal{H}}(t)-\theta_{u}^{\mathcal{H}}(s)\right], \forall t \in \mathcal{Y},  \tag{H}\\
\theta_{u}^{\mathcal{H}}(s) & :=\theta_{u}^{\mathcal{H}}(s)+\min _{t \in \mathcal{Y}}\left[g_{u v}(s, t)-\theta_{v}^{\mathcal{H}}(t)-\theta_{u}^{\mathcal{H}}(s)\right], \forall s \in \mathcal{Y} .
\end{align*}
$$

In other words, the MPLP++ update first performs the MPLP update and then pushes as much cost from the pairwise factor to the unary ones, as needed to fulfill $\min _{t} \theta_{u v}^{\mathcal{M}}(s, t)=$ $\min _{s} \theta_{u v}^{\mathcal{M}}(s, t)=0$ for all labels $s$ and $t$ in the nodes $u$ and $v$ respectively. It is also easy to see that the assignment $\theta_{v}^{\mathcal{H}}(s):=\theta_{v}^{\mathcal{M}}(s)$ together with the second line in $\mathcal{H}$ can be equivalently substituted by $\theta_{v}^{\mathcal{H}}(t):=\min _{s \in \mathcal{Y}}\left[g_{u v}(s, t)-\theta_{u}^{\mathcal{H}}(s)\right], \forall t \in \mathcal{Y}$. This allows to perform the MPLP++ update with 3 minimizations over $\mathcal{Y}^{2}$ instead of 4. Fig. 1 shows the result of applying the three BCA-updates on a simple two-node graph.

It is straightforward to show that $\mathcal{M}$ and $\mathcal{H}$ also satisfy Definition 1 (see supplement) and therefore are liable BCA-updates. In spite of that, the behavior of all three updates is notably different, as it is shown by example in Fig. 2. Therefore, proving only that some algorithm is a BCA, does not imply its efficiency.

Message Passing Importantly for performance, updates $\mathcal{U}, \mathcal{M}$ and $\mathcal{H}$ can be computed using a subroutine computing $\min _{t}\left[\theta_{u v}(s, t)+a(t)\right]$ for all $s$, where $\theta$ is a fixed initial pairwise potential and $a$ is an arbitrary input unary function. This operation occurring in dynamic programming and all of the discussed BCA algorithms, is known as message passing. In many cases of practical interest it can be implemented in time $O(|\mathcal{Y}|)$ (e.g. for Potts, absolute difference and quadratic costs) rather than $O\left(|\mathcal{Y}|^{2}\right)$ in the general case, using efficient sequential algorithms, e.g., [48]. Using $|\mathcal{Y}|$ processors, the computation time can be further reduced to $O(\log |\mathcal{Y}|)$ with appropriate parallel algorithms, e.g., [18].

Primal Rounding BCA-algorithms iterating BCA-updates give only a lower bound to the MAP-inference problem (1). To obtain a primal solution, we use a sequential rounding procedure similar to the one proposed in [2]. Assuming we have already computed a primal integer solution $x_{v}^{*}$ for all $v<u$, we want to compute $x_{u}^{*}$. To do so, we use the following equation for the assignment

$$
\begin{equation*}
x_{u}^{*} \in \arg \min _{x_{u} \in \mathcal{Y}}\left[\theta_{u}\left(x_{u}\right)+\sum_{v<u \mid u v \in \mathcal{E}} \theta_{u v}\left(x_{u}, x_{v}^{*}\right)\right], \tag{6}
\end{equation*}
$$



Fig. 2. Choosing the right BCA-update is important. Notation has the same meaning as in Fig. 1, black circles denote locally optimal labels, white circles - the non-optimal ones. Solid lines correspond to the locally optimal pairwise costs connected to the locally optimal labels. Omitted lines in pairwise interactions denote infinite pairwise costs. $e_{i}$ denotes edge indexes, edge processing order is according to the subscript $i$. (a) Uniform update gets stuck and is unable to optimize the dual further. (b) MPLP and MPLP ++ attain the dual optimum in one iteration.
where $\theta$ is the reparametrized potential produced by the algorithm.

## 5 Theoretical Analysis

As we prove below, MPLP++ in the limit guarantees to fulfill a necessary optimality condition, related to the arc-consistency [21] and the weak tree-agreement [2].
Arc-Consistency Let $\llbracket \cdot \rrbracket$ be the Iverson bracket, i.e. $\llbracket A \rrbracket=1$ if $A$ holds. Otherwise $\llbracket A \rrbracket=0$. Let $\bar{\theta}_{u}(s):=\llbracket \theta_{u}(s)=\min _{s^{\prime}} \theta_{u}\left(s^{\prime}\right) \rrbracket$ and $\bar{\theta}_{u v}(s, t):=\llbracket \theta_{u v}(s, t)=$ $\min _{s^{\prime}, t^{\prime}} \theta_{u v}\left(s^{\prime}, t^{\prime}\right) \rrbracket$ be binary vectors with values 1 assigned to the locally minimal labels and label pairs. Let also logical and and or operations be denoted as $\wedge$ and $\vee$. To the binary vectors they apply coordinate-wise.

Definition 3. A vector $\bar{\theta} \in\{0,1\}^{\mathcal{I}}$ is called arc-consistent, if $\bigvee_{t \in \mathcal{Y}} \bar{\theta}_{u v}(s, t)=\bar{\theta}_{u}(s)$ for all $\{u, v\} \in \mathcal{E}$ and $s \in \mathcal{Y}$.

However, arc-consistency itself is not necessary for dual optimality. The necessary condition is existence of the node-edge agreement, which is a special case of the weak tree agreement [2] when individual nodes and edges are considered as the trees in a problem decomposition. This condition is also known as a non-empty kernel [21] / arcconsistent closure [45] of the cost vector $\theta$.

Definition 4. We will say that the costs $\theta \in \mathbb{R}^{\mathcal{I}}$ fulfill the node-edge agreement, if there is an arc-consistent vector $\xi \in\{0,1\}^{\mathcal{I}}$ such that $\xi \wedge \bar{\theta}=\xi$.

Convergence of MPLP++ It is clear that all BCA algorithms are monotonous and converge in the dual objective value as soon as the dual is bounded (i.e., primal is feasible). However, such convergence is weaker than convergence in the reparametrization than is desired. To this end we were able to show something in between the two: the convergence of MPLP++ in a measure quantifying violation of the node-edge agreement, a result analogous to [15].

Theorem 1. The MPLP ++ algorithm converges to node-edge agreement.
Another important question is the comparison of different BCA methods. When comparing different algorithms, an ultimate goal is to prove faster convergence of one compared to the other one. We cannot show that the new MPLP++ algorithm has a better theoretical convergence rate. First, such rates are generally unknown for BCA algorithms for non-smooth functions. Second, considered algorithms are all of the same family and it is likely that their asymptotic rates are the same. Instead, we study the dominance, the condition that allows to rule out BCA updates which are always inferior to others. Towards this end we show that given the same starting dual point, one iteration of MPLP++ always results in a better objective value than that of MPLP and uniform BCA. While this argument does not extend theoretically to multiple iterations of each method, we show that it is still true in practice for all used datasets and a significant speed-up (up to two orders) is observed. The experimental comparison in $\S 7$ gives results in wall-clock time as well as in a machine-independent count of the message passing updates performed.

### 5.1 Analysis of BCA-updates

Definition 5. A BCA-iteration $\alpha$ is defined by the BCA-update $\gamma$ applied to all edges $\mathcal{E}$ in some chosen order. Let also $D(\alpha)$ represent the dual objective value with the reparametrization defined by the iteration $\alpha$ on the input costs $\theta$.

We will analyze BCA-iterations of different BCA-updates w.r.t. the same sequence of edges. The goal is to show that an iteration of the new update dominates the baselines in the dual objective. This property is formally captured by the following definition.

Definition 6. We will say that a BCA-iteration $\alpha$ dominates a BCA-iteration $\beta$, if for any input costs it holds $D(\alpha) \geq D(\beta)$.

In order to show it, we introduce now and prove later the dominance relations of individual BCA-updates. They are defined not on the dual objective but on all unary components.
Definition 7. Let $\gamma$ and $\delta$ be two BCA-updates. We will say that update $\gamma$ dominates $\delta$ (denoted as $\gamma \geq \delta$ ) if for any $g_{u v} \in \mathbb{R}^{\mathcal{Y}^{2}}$ it holds that $\gamma\left[g_{u v}\right] \geq \delta\left[g_{u v}\right]$, where the inequality is understood as component-wise inequalities $\theta_{u}^{\gamma}\left[g_{u v}\right] \geq \theta_{u}^{\delta}\left[g_{u v}\right]$ and $\theta_{v}^{\gamma}\left[g_{u v}\right] \geq \theta_{v}^{\delta}\left[g_{u v}\right]$.

We can show the following dominance results. Recall that $\mathcal{U}, \mathcal{H}$ and $\mathcal{M}$ are the uniform, MPLP and MPLP++ BCA-updates, respectively.

Proposition 2. The following BCA-dominances hold: $\mathcal{H} \geq \mathcal{M} \geq \mathcal{U}$.
It is easy to see that the dominance Definition 7 is transitive and so also $\mathcal{H} \geq \mathcal{U}$. We will prove that such coordinate-wise dominance of BCA-updates implies also the dominance in the dual objective whenever the following monotonicity property holds:

Definition 8. A BCA-update $\gamma$ is called monotonous if $\left(\theta_{u} \geq \theta_{u}^{\prime}, \theta_{v} \geq \theta_{v}^{\prime}\right)$ implies $\gamma\left[\theta_{u}+\theta_{u v}+\theta_{v}\right] \geq \gamma\left[\theta_{u}^{\prime}+\theta_{u v}+\theta_{v}^{\prime}\right]$ for all $\theta, \theta^{\prime}$.

Proposition 3. Updates $\mathcal{U}$ and $\mathcal{M}$ are monotonous. The update $\mathcal{H}$ is not monotonous.
With these results we can formulate our main claim about domination in the objective value for the whole iteration.

Theorem 2. Let BCA-update $\gamma$ dominate BCA-update $\mu$ and let $\mu$ be monotonous. Then a BCA-iteration with $\gamma$ dominates a BCA-iteration with $\mu$.

From Proposition 2, Proposition 3 and Theorem 2 it follows now that BCA-iteration of MPLP++ dominates that of MPLP, which in its turn dominates uniform.

## 6 Parallelization

To optimize $D\left(\theta^{\phi}\right)$, we have to perform local operations on a graph, that per reparametrization influence only one edge $u v \in \mathcal{E}$ and it's incident vertices $u$ and $v$. The remaining graph $\mathcal{G}^{\prime}=\left(\mathcal{V}-\{u, v\}, \mathcal{E}-\left\{I_{u} \cup I_{v}\right\}\right)$ remains unchanged. This gives rise to opportunities for parallelization. However, special care has to be taken to prevent race conditions which occurs when two or more threads access shared data and they try to change it at the same time.

Consider the case of Fig. 3, choosing edges 1 and 2 or 1 and 6 to process in parallel. These edges have vertex A in common, which would lead to race conditions. Processing edges 1 and 3 in parallel would lead to more parallelization as there are no conflicting nodes. Thus, for maximal parallelization we have to come up with an ordering of edges where threads working in parallel do not have common vertices.

Finding such edges without intersecting vertices is a well-studied problem in combinatorial optimization [49]. A matching $\mathcal{M} \subset \mathcal{E}$ in graph $\mathcal{G}$ is a set of edges such that no two edges in $\mathcal{M}$ share a common vertex. A matching is maximum if it includes the largest number of edges, $|\mathcal{M}|$. Every edge in a matching can be processed in parallel without race conditions ensuing. There exist efficient greedy algorithms to find a maximum matching which we use. This gives rise to Algorithm 1 for covering all edges of the graph while ensuring good parallelization. To cover the entire graph, we call a matching algorithm repeatedly, until all edges are exhausted.

Initially, in line 1 the edge queue $\mathcal{Q}_{\mathcal{E}}$ is empty. In line 3 , a maximum matching $\mathcal{E}_{\mathcal{M}}$ is found. This is added to $\mathcal{Q}_{\mathcal{E}}$ in line 4 . This continues until all edges have been exhausted, i.e. the edges remaining $\mathcal{E}_{\mathcal{R}}$ is empty. The queue thus has a structure $\mathcal{Q}_{\mathcal{E}}=\left(\mathcal{E}_{\mathcal{M}}^{1}, \mathcal{E}_{\mathcal{M}}^{2}, \ldots, \mathcal{E}_{\mathcal{M}}^{n}\right)$, ordered left to right. $\mathcal{E}_{\mathcal{M}}^{i}$ being the $i^{t h}$ matching computed. The threads running in parallel can keep popping edges from $\mathcal{Q}_{\mathcal{E}}$ and processing them without much need for mutex locking.

We have different implementation algorithms for GPUs and multi-core CPUs.
CPU Implementation: Modern CPUs consist of multiple cores with each core having one hardware threads. Hyper-threading allows for an additional thread per core, but with lesser performance in the second thread compared to the first.

Processing an edge is a short-lived task and launching a separate thread for each edge would have excessive overhead. To process lightweight tasks we use the threadpool design pattern. A thread-pool keeps multiple threads waiting for tasks to be executed concurrently by a supervising program. The threads are launched only once and


Fig. 3. Figure shows parallelization details and performance comparison for CPU and GPU. (a) shows the details of how the edge schedule is computed for maximizing throughput. The first row shows how edge selection is carried out by finding matchings and adding the edges in these matchings to the queue. The second row shows how the threads are launched for the CPU and GPU. For the CPU, threads are launched dynamically at different time instances, and no synchronization is carried out across all threads. This is due to a local memory locking mechanism (mutex) for the CPU. For the GPU, many threads are launched simultaneously and synchronized simultaneously via a memory barrier. This barrier is shown as the vertical line with Synch. (b) shows the running time comparison between single-threaded, multi-threaded and GPU versions of the MPLP++ algorithm. The GPU takes some time to load memory from the host (CPU) to device (GPU). This is why it takes longer to get started than the CPU versions.

```
Algorithm 1 Compute Edge Schedule
Require: \(\mathcal{G}=(\mathcal{V}, \mathcal{E})\)
    Initial: \(\mathcal{Q}_{\mathcal{E}}:=\emptyset \quad\) (Empty edge queue),
    \(\mathcal{E}_{\mathcal{R}}:=\mathcal{E} \quad\) (Initial pool of edges)
    while \(\mathcal{E}_{\mathcal{R}}!=\emptyset\) do
        \(\mathcal{E}_{\mathcal{M}}:=\) Maximum_Matching \(\left(\mathcal{V}, \mathcal{E}_{\mathcal{R}}\right)\)
        \(\mathcal{Q}_{\mathcal{E}} \cdot \operatorname{push}\left(\mathcal{E}_{\mathcal{M}}\right) \quad\) (Push maximum-matching \(\mathcal{E}_{\mathcal{M}}\) to the queue)
        \(\mathcal{E}_{\mathcal{R}}:=\mathcal{E}_{\mathcal{R}}-\mathcal{E}_{\mathcal{M}} \quad\) (Remove matched edges from \(\mathcal{E}_{\mathcal{R}}\) )
    end while
```

continuously process tasks from a task-queue. As they are launched only once, the latency in execution due to overhead in thread creation and destruction is avoided.

In the case of our algorithm the task queue is $\mathcal{Q}_{\mathcal{E}}$. The thread picks up the index of the edge to process and performs the MPLP ++ operation $(\mathcal{H})$. During the MPLP ++ operation the node and edge structures are locked by mutexes. One iteration of the algorithm is complete when all edges have been processed. Since multiple iterations maybe required to reach convergence, our task queue is circular, letting the threads restart the reparameterization process from the first element of $\mathcal{Q}_{\mathcal{E}}$. The ordering of $\mathcal{Q}_{\mathcal{E}}$ prevents heavy lock contention of mutexes.

GPU Implementation: Unlike CPUs, GPUs do not use mutexes for synchronization. The threads in each GPU processor are synchronized via a hardware barrier synchronization. A barrier for a group of threads stops the threads at this point and prevents them from proceeding until all other threads/processes reach this barrier.

This is where the ordering of $\mathcal{Q}_{\mathcal{E}}$ comes handy. Recall the structure of $\mathcal{Q}_{\mathcal{E}}=$ $\left(\mathcal{E}_{\mathcal{M}}^{1}, \mathcal{E}_{\mathcal{M}}^{2}, \ldots, \mathcal{E}_{\mathcal{M}}^{n}\right)$. Barrier synchronization can be used between the matchings $\mathcal{E}_{\mathcal{M}}^{i}$ and $\mathcal{E}_{\mathcal{M}}^{i+1}$, allowing for the completion of the MPLP ++ BCA update operations for $\mathcal{E}_{\mathcal{M}}^{i}$, before beginning the processing of $\mathcal{E}_{\mathcal{M}}^{i+1}$. This minimizes the time threads spend waiting while the other threads complete, as they have no overlapping areas to write to in the memory.

## 7 Experimental Evaluation

In our experiments, we use a 4-core Intel i7-4790K CPU @ 4.00GHz, with hyperthreading, giving 8 logical cores. For GPU experiments, we used the NVIDIA Tesla K80 GPU with 4992 cores.
Compared Algorithms We compare different algorithms in two regimes: the wallclock time and the machine-independent regime, where we count the number of operations performed. For the latter one we use the notion of the oracle call, which is an operation like $\min _{t \in \mathcal{Y}} g_{u v}(s, t), \forall s \in \mathcal{Y}$ involving single evaluation of all costs of a pairwise factor. When speaking about oracle complexity we mean the number of oracle calls per single iteration of the algorithm. As different algorithms have different oracle complexities we define a normalized iteration as exactly $|\mathcal{E}|$ messages for an even comparison across algorithms. We compare the following BCA schemes:

- The Tree-Reweighted Sequential (TRWS) message processing [2] algorithm has consistently been the best performing method on several benchmarks like [13]. Its oracle complexity is $2|\mathcal{E}|$. We use the multicore implementation introduced in [42] for comparison, which is denoted as TRWS (MT) when run on multiple cores.
- The Max-Product Linear Programming MP LP [1] algorithm with BCA-updates $(\mathcal{M})$. It's oracle complexity is thus $2|\mathcal{E}|$. For MP LP we have our own multi-threaded implementation that is faster than the original one by a factor of 4
- The Min-Sum Diffusion Algorithm MSD [15], is one of the earliest (in the 70s) proposed BCA algorithms for graphical models. The oracle complexity of the method is $4|\mathcal{E}|$.
- The MPLP++ algorithm with BCA-updates $(\mathcal{H})$ has the oracle complexity of $3|\mathcal{E}|$. The algorithm is parallelized as described in $\S 6$ for both CPU and GPU. The corresponding legends are MPLP++ (MT) and MPLP++ (GPU).

Dense datasets To show the strength of our method we consider the following datasets with underlying densely connected graphs:

- worms dataset [50] consists of 30 problem instances coming from the field of bioimaging. The problems' graphs are dense but not fully connected with about $0.1 \cdot \mathcal{V}^{2}$ of edges, up to 600 nodes and up to 1500 labels.
- protein-folding dataset [51] taken from the OpenGM benchmark [13]. This dataset has 21 problem instances with $33-1972$ variables. The models are fully connected and have $81-503$ labels per node.
- pose is the dataset inspired by the recent work [4] showing state-of-the-art performance on the 6 D object pose estimation problem. In [4] this problem is formulated as MAP-inference in a fully connected graphical model. The set of nodes of the
underlying graph coincides with the set of pixels of the input image (up to some downscale factor), which requires specialized heuristics to obtain practically useful solutions in reasonable time due to the large problem size. Contrary to the original work, we assume that the position of the object is given by its mask (we used ground-truth data from the validation set, but assume the mask could be provided by some segmentation method) and treat only the pixels inside the mask as the nodes of the corresponding graphical model. Otherwise, the unary and pairwise costs are constructed in the same way as in [4]. We have got the learned unary costs from the authors of [4] and merely tuned the hyper-parameters on the validation set. This dataset has 32 problem instances with $600-4800$ variables each and 13 labels per node. The models are all fully connected.

Sparse datasets Although our method is not the best performing one on sparse graphical models, we include the comparison on the following four-connected grid-graph based benchmark datasets for fairness:

- color-seg from [13] has Potts pairwise costs. The nodes contain up to 12 labels.
- stereo from the Middlebury MRF benchmark [52]. This dataset consists of 3 models with truncated linear pairwise costs and 16,20 and 60 labels respectively.

Algorithm Convergence Fig. 4 shows convergence of the considered algorithms in a sequential setting for the protein and stereo datasets as representatives of the dense and sparse problems. For other datasets the behavior is similar, therefore we moved the corresponding plots to the supplement $\S$ A. The proposed MPLP++ method outperforms all its competitors on all dense problem instances and is inferior to TRWS only on sparse ones. This holds for both comparisons: the implementation-independent normalized iteration and the implementation-dependent running time ones. Fig. 5 shows relative time speed-ups of the considered methods as a function of the attained dual precision. The speed of MPLP, as typically the slowest one, is taken to be 1 , i.e., for other algorithms the speed-up compared to MPLP is plotted. This includes also the CPU-parallel versions of MPLP++, TRWS and the GPU-parallel MPLP++. Fig. 5 also shows that for dense models MPLP ++ is $2-10$ faster than TRWS in the sequential setting and $7-40$ times in the parallel setting. The speed-up w.r.t. MP LP is $5-150$ times depending on the dataset and the required precision.
Performance Degradation with Graph Sparsification In this test we gradually and randomly removed edges from the graphs of the pose dataset and measures performance of all algorithms. Fig. 6 shows that up to at least $10 \%$ of all possible edges MPLP + + leads the board. Only when the number of edges drops to $5 \%$ and less, TRWS starts to outperform MPLP++. Note, the density of edges in grid graphs considered above does not exceed the $0.007 \%$ level.

## 8 Conclusions and Outlook

Block-coordinate ascent methods remain a perspective research direction for creating efficient parallelizable dual solvers for MAP-inference in graphical models. We have presented one such solver beating the state-of-the-art on dense graphical models with arbitrary potentials. The method is directly generalizable to higher order models, which we plan to investigate in the future.


Fig. 4. Improvement in dual as a function of time and iterations for the protein-folding and stereo dataset. The algorithms we compare have different message-passing schemes and end up doing different amounts of work per iteration. Hence, for a fair comparison across algorithms we define a normalized iteration as exactly $|\mathcal{E}|$ messages. This is also equal to number of messages passed in a normal iteration divided by its oracle complexity. (a) and (b) are for the dense proteinfolding dataset, where both per unit time and per iteration, MPLP++ outperforms TRWS and all other algorithms. In the sparse stereo dataset (c,d) TRWS beats all other algorithms. Results are averaged over the entire dataset. The dual is normalized to 1 for equal weighing of every instance in the dataset.


Fig. 5. Relative speed w.r.t the MPLP algorithm in converging to within $\varepsilon$ of the best attained dual optima $D^{*}\left(\theta^{\phi}\right)$, i.e. $D^{*}\left(\theta^{\phi}\right)-\varepsilon$. The plot shows the speedups of all the algorithms relative to MPLP for $\varepsilon$ 's $0.001,0.01,0.1,1$ and 10 of $D^{*}\left(\theta^{\phi}\right)$. Figure (a) shows MPLP ++ is $50 \times$ faster than MPLP in converging to within $10 \%$ of $D^{*}\left(\theta^{\phi}\right)$. Figure (b) and (c) likewise show an order of magnitude speedup. Figure (d) shows that for the stereo dataset consisting of sparse graphs TRWS dominates MPLP++. Convergence only till $0.1 \%$ of $D^{*}\left(\theta^{\phi}\right)$ is shown for stereo as only TRWS converges to the required precision.

## 9 Acknowledgements

This project has received funding from the ERC under the European Unions Horizon 2020 research and innovation program (grant agreement No 647769). A. Shekhovtsov was supported by Czech Science Foundation grant 18-25383S.


Fig. 6. Degradation with Sparsity (Dual vs Time): (a)-(d) show graphs with decreasing average connectivity given as percentage of possible edges in figure subcaption. In (a)-(c) MPLP ++ outperforms TRWS. MPLP++ is resilient to graph sparsification even when $90 \%$ of the edges have been removed. Only when more than $95 \%$ of the edges have been removed as in (d) TRWS outperforms MP LP ++.

## References

[1] Globerson, A., Jaakkola, T.S.: Fixing Max-Product: Convergent Message Passing Algorithms for MAP LP-Relaxations. In: Advances in Neural Information Processing Systems 20. (2008)
[2] Kolmogorov, V.: Convergent tree-reweighted message passing for energy minimization. IEEE transactions on pattern analysis and machine intelligence 28(10) (2006) 1568-1583
[3] Krähenbühl, P., Koltun, V.: Efficient inference in fully connected CRFs with gaussian edge potentials. In: Advances in neural information processing systems. (2011) 109-117
[4] Michel, F., Kirillov, A., Brachmann, E., Krull, A., Gumhold, S., Savchynskyy, B., Rother, C.: Global hypothesis generation for 6D object pose estimation. arXiv preprint (2017)
[5] Kolmogorov, V., Rother, C.: Comparison of energy minimization algorithms for highly connected graphs. In: European Conference on Computer Vision, Springer (2006) 1-15
[6] Bergtholdt, M., Kappes, J., Schmidt, S., Schnörr, C.: A study of parts-based object class detection using complete graphs. International journal of computer vision 87(1-2) (2010) 93
[7] Nowozin, S., Rother, C., Bagon, S., Sharp, T., Yao, B., Kohli, P.: Decision Tree Fields. In: Computer Vision (ICCV), 2011 IEEE International Conference on, IEEE (2011) 1668-1675
[8] Kirillov, A., Schlesinger, D., Zheng, S., Savchynskyy, B., Torr, P.H., Rother, C.: Joint training of generic CNN-CRF models with stochastic optimization. In: Asian Conference on Computer Vision, Springer (2016) 221-236
[9] Kainmueller, D., Jug, F., Rother, C., Myers, G.: Active graph matching for automatic joint segmentation and annotation of C. elegans. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer (2014) 81-88
[10] Chen, L.C., Schwing, A., Yuille, A., Urtasun, R.: Learning deep structured models. In: International Conference on Machine Learning. (2015) 1785-1794
[11] Kolmogorov, V., Thapper, J., Zivny, S.: The power of linear programming for generalvalued CSPs. SIAM Journal on Computing 44(1) (2015) 1-36
[12] Průša, D., Werner, T.: LP relaxation of the potts labeling problem is as hard as any linear program. IEEE Trans. Pattern Anal. Mach. Intell. 39(7) (2017) 1469-1475
[13] Kappes, J.H., Andres, B., Hamprecht, F.A., Schnörr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B.X., Kröger, T., Lellmann, J., Komodakis, N., Savchynskyy, B., Rother, C.: A comparative study of modern inference techniques for structured discrete energy minimization problems. International Journal of Computer Vision (2015) 1-30
[14] Shekhovtsov, A., Reinbacher, C., Graber, G., Pock, T.: Solving dense image matching in real-time using discrete-continuous optimization. In: CVWW. (2016) 13
[15] Schlesinger, M., Antoniuk, K.: Diffusion algorithms and structural recognition optimization problems. Cybernetics and Systems Analysis 47(2) (2011) 175-192
[16] Choi, J., Rutenbar, R.A.: Hardware implementation of mrf map inference on an fpga platform. In: Field Programmable Logic and Applications (FPL), 2012 22nd International Conference on, IEEE (2012) 209-216
[17] Hurkat, S., Choi, J., Nurvitadhi, E., Martínez, J.F., Rutenbar, R.A.: Fast hierarchical implementation of sequential tree-reweighted belief propagation for probabilistic inference. In: Field Programmable Logic and Applications (FPL), 2015 25th International Conference on, IEEE (2015) 1-8
[18] Chen, W., Wada, K.: On computing the upper envelope of segments in parallel. In: Proceedings. 1998 International Conference on Parallel Processing (Cat. No.98EX205). (Aug 1998) 253-260
[19] Li, M., Shekhovtsov, A., Huber, D.: Complexity of discrete energy minimization problems. In: European Conference on Computer Vision. (2016) 834-852
[20] Schlesinger, M.I.: Syntactic analysis of two-dimensional visual signals in noisy conditions. Kibernetika 4(113-130) (1976) 1
[21] Werner, T.: A linear programming approach to max-sum problem: A review. IEEE transactions on pattern analysis and machine intelligence 29(7) (2007)
[22] Živný, S., Werner, T., Průša, D.a. In: The Power of LP Relaxation for MAP Inference. The MIT Press, Cambridge, USA (December 2014) 19-42
[23] Thapper, J., Živný, S.: The power of linear programming for valued CSPs. In: Symposium on Foundations of Computer Science (FOCS). (2012) 669-678
[24] Ravikumar, P., Agarwal, A., Wainwright, M.: Message-passing for Graph-structured Linear Programs: Proximal Methods and Rounding Schemes. JMLR 11 (2010) 1043-1080
[25] Martins, A.F.T., Figueiredo, M.A.T., Aguiar, P.M.Q., Smith, N.A., Xing, E.P.: An Augmented Lagrangian Approach to Constrained MAP Inference. In: ICML. (2011)
[26] Meshi, O., Globerson, A.: An Alternating Direction method for Dual MAP LP Relaxation. In: ECML/PKDD (2). (2011) 470-483
[27] Schmidt, S., Savchynskyy, B., Kappes, J., Schnörr, C.: Evaluation of a first-order primaldual algorithm for mrf energy minimization. In: EMMCVPR 2011. (2011)
[28] Storvik, G., Dahl, G.: Lagrangian-based methods for finding MAP solutions for MRF models. IEEE Transactions on Image Processing 9(3) (2000) 469-479
[29] Schlesinger, M., Giginyak, V.: Solution to structural recognition (max,+)-problems by their equivalent transformations. in 2 Parts. Control Systems and Computers (1-2) (2007)
[30] Komodakis, N., Paragios, N., Tziritas, G.: MRF optimization via dual decomposition: Message-passing revisited. In: Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on, IEEE (2007) 1-8
[31] Kappes, J.H., Savchynskyy, B., Schnörr, C.: A bundle approach to efficient MAP-inference by lagrangian relaxation. In: Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on, IEEE (2012) 1688-1695
[32] Luong, D.V., Parpas, P., Rueckert, D., Rustem, B.: Solving MRF minimization by mirror descent. In: International Symposium on Visual Computing, Springer (2012) 587-598
[33] Savchynskyy, B., Kappes, J., Schmidt, S., Schnörr, C.: A study of Nesterov's scheme for Lagrangian decomposition and MAP labeling. In: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, IEEE (2011) 1817-1823
[34] Savchynskyy, B., Schmidt, S., Kappes, J., Schnörr, C.: Efficient MRF energy minimization via adaptive diminishing smoothing. arXiv preprint arXiv:1210.4906 (2012)
[35] Meshi, O., Globerson, A., Jaakkola, T.S.: Convergence rate analysis of MAP coordinate minimization algorithms. In: Advances in Neural Information Processing Systems. (2012) 3014-3022
[36] Kannan, H., Komodakis, N., Paragios, N.: Newton-type methods for inference in higherorder markov random fields. In: IEEE International Conference on Computer Vision and Pattern Recognition. (2017)
[37] Prusa, D., Werner, T.: Universality of the local marginal polytope. PAMI 37(4) (April 2015)
[38] Kovalevsky, V., Koval, V.: A diffusion algorithm for decreasing energy of max-sum labeling problem. Glushkov Institute of Cybernetics, Kiev, USSR (1975) Unpublished.
[39] Schlesinger, M.I., Flach, B.: Some solvable subclasses of structural recognition problems. In: Czech Pattern Recognition Workshop. Volume 2000. (2000) 55-62
[40] Cooper, M.C., De Givry, S., Sánchez, M., Schiex, T., Zytnicki, M., Werner, T.: Soft arc consistency revisited. Artificial Intelligence 174(7-8) (2010) 449-478
[41] Savchynskyy, B., Kappes, J.H., Swoboda, P., Schnörr, C.: Global MAP-optimality by shrinking the combinatorial search area with convex relaxation. In: Advances in Neural Information Processing Systems. (2013) 1950-1958
[42] Shekhovtsov, A., Swoboda, P., Savchynskyy, B.: Maximum persistency via iterative relaxed inference with graphical models. PAMI (2017)
[43] Swoboda, P., Shekhovtsov, A., Kappes, J.H., Schnorr, C., Savchynskyy, B.: Partial optimality by pruning for MAP-inference with general graphical models. PAMI 38(7) (2016) 1370-1382
[44] Hurley, B., OSullivan, B., Allouche, D., Katsirelos, G., Schiex, T., Zytnicki, M., De Givry, S.: Multi-language evaluation of exact solvers in graphical model discrete optimization. Constraints 21(3) (2016) 413-434
[45] Cooper, M., Schiex, T.: Arc consistency for soft constraints. Artificial Intelligence 154(1-2) (2004) 199-227
[46] Kolmogorov, V.: A new look at reweighted message passing. IEEE transactions on pattern analysis and machine intelligence 37(5) (2015) 919-930
[47] Hazan, T., Shashua, A.: Norm-Product Belief Propagation: Primal-Dual Message-Passing for approximate inference (2008)
[48] Aggarwal, A., Klawe, M.M., Moran, S., Shor, P., Wilber, R.: Geometric applications of a matrix-searching algorithm. Algorithmica 2(1) (Nov 1987) 195-208
[49] Schrijver, A.: Combinatorial Optimization: Polyhedra and Efficiency. Volume 24. Springer Science \& Business Media (2003)
[50] Kainmueller, D., Jug, F., Rother, C., Meyers, G.: Graph matching problems for annotating c. elegans. http://dx.doi.org/10.15479/AT:ISTA:57(2017) Accessed: 2017-09-10.
[51] Yanover, C., Schueler-Furman, O., Weiss, Y.: Minimizing and learning energy functions for side-chain prediction. Journal of Computational Biology 15(7) (2008) 899-911
[52] Szeliski, R., Zabih, R., Scharstein, D., Veksler, O., Kolmogorov, V., Agarwala, A., Tappen, M., Rother, C.: A comparative study of energy minimization methods for markov random fields with smoothness-based priors. IEEE transactions on pattern analysis and machine intelligence 30(6) (2008) 1068-1080
[53] Boyd, S., Vandenberghe, L.: Convex optimization. Cambridge university press (2004)
[54] Bartle, R., Sherbert, D.: Introduction to real analysis. John Wiley \& Sons Canada, Limited (2000)

# MPLP++: Fast, Parallel Dual Block-Coordinate Ascent for Dense Graphical Models (ECCV'18 Appendix) 

## A Additional Experimental Results



Fig. A.1. Dual vs Normalized Iterations on three other datasets. The experimental setup and notation is the same as in Fig. 4. Problems in cases (a) and (b) have dense graphs where MPLP++ outperforms all other algorithms by a substantial margin. In the case (c) graphs are sparse and TRWS is dominant.


Fig. A.2. Dual vs Time (Single Threaded): Fig. shows dual as a function of time for the single threaded versions of all the algorithms. Following the pattern in Fig. A. 1 for dense graphs (a) and (b) MPLP ++ dominates all other algorithms by a considerable margin. For sparse graphs, (c) TRWS is the fastest. Both the dual $D(\phi)$ and time have been averaged over the entire dataset and normalized to 1.The curves have been normalized such that each instance of the dataset is weighed equally.


Fig. A.3. Degradation With Sparsity (Dual vs Iterations): (a)-(h) show graphs with decreasing average connectivity given as percentage of possible edges in figure subcaption. In (a)-(f) MP LP ++ outperforms TRWS. Handshake is resilient to graph sparsification even when $90 \%$ of the edges have been removed. For (g) and(h) TRWS outperforms handshake.

## B Formal Proofs

## B. 1 Additional Notation

To reduce clutter in the proofs we define relevant short-forms here. The set $[n]$ is the set of the first $n$ natural numbers, i.e. $[n]=\{1, \ldots, n\} .[m][n]$ denotes the Cartesian product $[m] \times[n]$. We denote the same reparametrized unaries of two different algorithms $\mathcal{A}$ and $\mathcal{B}$ as $\theta_{u}^{\mathcal{A}}$ and $\theta_{u}^{\mathcal{B}} . \theta \in \mathbb{R}^{\mathcal{I}}$ is the vector stacked up of unary $\theta_{u}$ and pairwise $\theta_{u v}$ potentials. Let $\mathcal{O}_{u}^{\varepsilon}(\theta)=\left\{s \mid \theta_{u}(s) \leq \min _{s^{\prime}} \theta_{u}\left(s^{\prime}\right)+\varepsilon\right\}$ be the set of labellings within $\varepsilon>0$ of the optimal $\theta_{u} . \mathcal{O}_{u v}^{\varepsilon}(\theta)$ is similarly defined for $\theta_{u v}$. Let also $\mathcal{O}^{\varepsilon}(\theta)=\left\{\mathcal{O}_{u}^{\varepsilon}(\theta) \mid \forall u \in \mathcal{V}\right\} \cup\left\{\mathcal{O}_{u v}^{\varepsilon}(\theta) \mid \forall u v \in \mathcal{E}\right\}$.

Definition 1. Tolerance factor $\varepsilon$ is the minimum value for which $\mathcal{O}^{\varepsilon}$ contains a consistent labelling, i.e. one with node-edge agreement.
$\varepsilon$ is a function of $\theta$ and will be written sometimes as $\varepsilon(\theta) . \mathcal{O}^{\varepsilon}(\theta)$ can also be thought of as the subset of unary and pairwise labels that are within $\varepsilon$ of the optimal labelling.
$\mathcal{O}_{u}^{0}(\theta)=\left\{s \mid \theta_{u}(s)=\min _{s^{\prime} \in \mathcal{Y}} \theta_{u}\left(s^{\prime}\right)\right\}$, then represents the set of optimal labellings. $\mathcal{O}_{u v}^{0}(\theta)$ can be similarly defined for pairwise potentials $\theta_{u v}$. Then, $\mathcal{O}^{0}(\theta)=$ $\left\{\mathcal{O}_{u}^{0}(\theta) \mid u \in \mathcal{V}\right\} \cup\left\{\mathcal{O}_{u v}^{0}(\theta) \mid u v \in \mathcal{E}\right\}$.

The MPLP ++ operator $\mathcal{H}$ can act both on $g_{u v}$ like pairwise costs and on the entire set of costs $\theta \in \mathbb{R}^{\mathcal{I}}$. In the former case, it is exactly as defined in Eqn. $(\mathcal{H})$. The latter case corresponds to an iteration of $\mathcal{H}$. The $i$-times composition operation of $\mathcal{H}$ on $\theta$ denotes $i$ iterations of $\mathcal{H}$ on $\theta$, i.e. $\underbrace{\mathcal{H}(\mathcal{H}(\ldots(\mathcal{H}(\theta))))}_{\text {i times }}=\mathcal{H}^{i}(\theta)$. Which is the case would be clear
from the argument the $\mathcal{H}$ operator takes. Also, to denote the resulting cost vector after $i$ iterations of $\mathcal{H}$ on $\theta$, we use notation $\theta^{i}=\mathcal{H}^{i}(\theta)=\mathcal{H}\left(\theta^{i-1}\right)$.

Let $\Theta_{u n}(\phi, \kappa)$ be a function that measures the max absolute difference between unary reparameterizations $\theta^{\phi}$ and $\theta^{\kappa}$, defined by $\Theta_{u n}(\phi, \kappa)=\max _{u \in \mathcal{V}, s \in \mathcal{Y}} \mid \theta_{u}^{\phi}(s)-$ $\theta_{v}^{\kappa}(s) \mid$. Likewise for pairwise costs, $\Theta_{p w}(\phi, \kappa)=\max _{u v \in \mathcal{E}, s t \in \mathcal{Y}^{2}}\left|\theta_{u v}^{\phi}(s, t)-\theta_{u v}^{\kappa}(s, t)\right|$.

We also assume that both unary and pairwise costs have been normalized, i.e. $\min _{s \in \mathcal{Y}} \theta_{u}(s)=0$ and $\min _{s t \in \mathcal{Y}^{2}} \theta_{u v}(s, t)=0$. A consequence of normalization is $\theta_{u} \geq 0$ and $\theta_{u v} \geq 0$. The subtracted cost does not affect the labelling and is added to a constant term.

Proposition 2. The following BCA-dominances hold: $\mathcal{H} \geq \mathcal{M} \geq \mathcal{U}$.
Proof. $\mathcal{H} \geq \boldsymbol{\mathcal { M }}:$ Let $g_{u v}=\theta_{u}+\theta_{v}+\theta_{u v}$. Consider the first line of $(\mathcal{H})$.

$$
\begin{equation*}
\theta_{u}^{\mathcal{H}}(s):=\theta_{u}^{\mathcal{M}}(s), \quad \theta_{v}^{\mathcal{H}}(s):=\theta_{v}^{\mathcal{M}}(s), \forall s \in \mathcal{Y} \tag{1}
\end{equation*}
$$

The first line assigns to the reparameterized MPLP ++ unaries, $\left(\theta_{u}^{\mathcal{H}}, \theta_{v}^{\mathcal{H}}\right)$ the unaries resulting from applying the MPLP update to $g_{u v}$. At this point we have $\mathcal{H}=\mathcal{M}$. Now, consider the subsequent equations of $\mathcal{H}$

$$
\begin{aligned}
\theta_{v}^{\mathcal{H}}(t) & :=\theta_{v}^{\mathcal{H}}(t)+\min _{s \in \mathcal{Y}}\left[g_{u v}(s, t)-\theta_{v}^{\mathcal{H}}(t)-\theta_{u}^{\mathcal{H}}(s)\right], \forall t \in \mathcal{Y} \\
\theta_{u}^{\mathcal{H}}(s) & :=\theta_{u}^{\mathcal{H}}(s)+\min _{t \in \mathcal{Y}}\left[g_{u v}(s, t)-\theta_{v}^{\mathcal{H}}(t)-\theta_{u}^{\mathcal{H}}(s)\right], \forall s \in \mathcal{Y} .
\end{aligned}
$$

To $\theta_{v}^{\mathcal{H}}(t)$, a non-negative quantity $\min _{s \in \mathcal{Y}}\left[g_{u v}(s, t)-\theta_{v}^{\mathcal{H}}(t)-\theta_{u}^{\mathcal{H}}(s)\right]$ is added. Similarly, for $\theta_{u}^{\mathcal{H}}$. Thus, $\mathcal{H} \geq \mathcal{M}$.
$\mathcal{M} \geq \mathcal{U}: \theta_{u}^{\mathcal{M}}(s)=\min _{t^{\prime}} g_{u v}\left(s, t^{\prime}\right) \geq \min _{s^{\prime}, t^{\prime}} g_{u v}\left(s^{\prime}, t^{\prime}\right)=\theta_{u}^{\mathcal{U}}(s)$. Likewise, $\theta_{v}^{\mathcal{M}}(t)=\min _{s^{\prime}} g_{u v}\left(s^{\prime}, t\right) \geq \min _{s^{\prime}, t^{\prime}} g_{u v}\left(s^{\prime}, t^{\prime}\right)=\theta_{v}^{\mathcal{U}}(t)$. Thus, $\mathcal{M} \geq \mathcal{U}$.

Proposition 3. Updates $\mathcal{U}$ and $\mathcal{M}$ are monotonous. The update $\mathcal{H}$ is not monotonous.
Proof. Let $\theta_{u} \geq \theta_{u}^{\prime}$ and $\theta_{v} \geq \theta_{v}^{\prime}$. We define $g_{u v}=\theta_{u}+\theta_{v}+\theta_{u v}$ and $g_{u v}^{\prime}=\theta_{u}^{\prime}+\theta_{v}^{\prime}+$ $\theta_{u v}$. Thus we have $g_{u v} \geq g_{u v}^{\prime}$.
$\mathcal{U}$ is monotonous:
Performing update $\gamma_{\mathcal{U}}\left[g_{u v}\right] \rightarrow\left(\theta_{u}^{\mathcal{U}}, \theta_{v}^{\mathcal{U}}\right)$ (Eqn. $(\mathcal{U})$ ) yields

$$
\begin{equation*}
\theta_{u}^{\mathcal{U}}(s)=\theta_{v}^{\mathcal{U}}(t):=\min _{s^{\prime}, t^{\prime}} \frac{1}{2} g_{u v}\left(s^{\prime}, t^{\prime}\right) \tag{2}
\end{equation*}
$$

Likewise $\gamma_{\mathcal{U}}\left[g_{u v}^{\prime}\right] \rightarrow\left(\theta_{u}^{\mathcal{U}}, \theta_{v}^{\mathcal{U}}\right)$ yields

$$
\begin{equation*}
\theta_{u}^{\mathcal{U}}(s)=\theta_{v}^{\mathcal{U}}(t):=\min _{s^{\prime}, t^{\prime}} \frac{1}{2} g_{u v}^{\prime}\left(s^{\prime}, t^{\prime}\right) \tag{3}
\end{equation*}
$$

As $g_{u v} \geq g_{u v}^{\prime}$, we have $\min _{s^{\prime} t^{\prime}} g_{u v}\left(s^{\prime}, t^{\prime}\right) \geq \min _{s^{\prime} t^{\prime}} g_{u v}^{\prime}\left(s^{\prime}, t^{\prime}\right)$. This implies, $\theta_{u}^{\mathcal{U}} \geq \theta_{u}^{\mathcal{U}}, \quad \theta_{v}^{\mathcal{U}} \geq \theta_{v}^{\mathcal{U}}$. Hence proved.

## $\mathcal{M}$ is monotonous:

Performing update $\gamma_{\mathcal{M}}\left[g_{u v}\right] \rightarrow\left(\theta_{u}^{\mathcal{M}}, \theta_{v}^{\mathcal{M}}\right)($ Eqn. $\mathcal{M})$ yields

$$
\theta_{u}^{\mathcal{M}}(s):=\frac{1}{2} \min _{t^{\prime}} g_{u v}\left(s, t^{\prime}\right), \quad \theta_{v}^{\mathcal{M}}(t):=\frac{1}{2} \min _{s^{\prime}} g_{u v}\left(s^{\prime}, t\right)
$$

For $\gamma_{\mathcal{M}}\left[g_{u v}^{\prime}\right] \rightarrow\left(\theta_{u}^{\prime \mathcal{M}}, \theta_{v}^{\prime \mathcal{M}}\right)$ yields

$$
\theta_{u}^{\prime \mathcal{M}}(s):=\frac{1}{2} \min _{t^{\prime}} g_{u v}^{\prime}\left(s, t^{\prime}\right), \quad \theta_{v}^{\prime \mathcal{M}}(t):=\frac{1}{2} \min _{s^{\prime}} g_{u v}^{\prime}\left(s^{\prime}, t\right)
$$

As $g_{u v} \geq g_{u v}^{\prime}$, we have $\min _{t^{\prime}} g_{u v}\left(s, t^{\prime}\right) \geq \min _{t^{\prime}} g_{u v}^{\prime}\left(s, t^{\prime}\right) \& \min _{s^{\prime}} g_{u v}\left(s^{\prime}, t\right) \geq$ $\min _{s^{\prime}} g_{u v}^{\prime}\left(s^{\prime}, t\right)$. This implies $\theta_{u}^{\mathcal{M}} \geq \theta_{u}^{\prime \mathcal{M}}, \theta_{v}^{\mathcal{M}} \geq \theta_{v}^{\prime \mathcal{M}}$. Hence proved.

## $\mathcal{H}$ is not monotonous:

To prove that $\mathcal{H}$ is not monotonous we show a counter-example to the monotonous condition stated in theorem 3.

Consider the following unary and pairwise costs,

$$
\theta_{u}=\binom{4}{0} \quad \theta_{v}=\binom{2}{0} \quad \theta_{u}^{\prime}=\binom{0}{0} \quad \theta_{v}^{\prime}=\binom{0}{0} \quad \theta_{u v}=\left(\begin{array}{ll}
0 & 1 \\
7 & 5
\end{array}\right)
$$

These costs satisfy $\theta_{u} \geq \theta_{u}^{\prime}$ and $\theta_{v} \geq \theta_{v}^{\prime}$.
Now, applying the $\mathcal{H}$ operation to $g_{u v}$ and $g_{u v}^{\prime}$, we get the following potentials

$$
\theta_{u}^{\mathcal{H}}=\binom{2.5}{2.5} \quad \theta_{v}^{\mathcal{H}}=\binom{3.5}{2.5} \theta_{u}^{\prime \mathcal{H}}=\binom{0}{4} \quad \theta_{v}^{\prime \mathcal{H}}=\binom{0}{1}
$$

We thus have $\theta_{u}^{\mathcal{H}} \nsupseteq \theta_{u}^{\prime \mathcal{H}}$, providing the necessary counter-example.
Theorem 2. Let BCA-update $\gamma$ dominate BCA-update $\mu$ and let $\mu$ be monotonous. Then a BCA-iteration with $\gamma$ dominates a BCA-iteration with $\mu$.

Proof. From the definition of dominance we have, if $\gamma$ dominates $\mu(\gamma \geq \mu)$ then $\gamma\left[g_{u v}\right] \geq \mu\left[g_{u v}\right], \forall g_{u v}$.

By the definition of monotonous, if $\mu$ is monotonous, $g_{u v}^{1} \geq g_{u v}^{2} \Longrightarrow \mu\left[g_{u v}^{1}\right] \geq$ $\mu\left[g_{u v}^{2}\right]$.

Now during the $1^{\text {st }}$ iteration we have three cases,

- Case A: Nodes $u$ and $v$ of edge $u v$ have not been reparametrized before.
- Case B: Nodes $u$ and $v$ of edge $u v$ have been reparametrized before.
- Case C: Only node $u$ or $v$ of edge $u v$ have been reparametrized before.

If case $\mathbf{A}$ we have by dominance $\gamma\left[g_{u v}\right] \geq \mu\left[g_{u v}\right]$.
If case $\mathbf{B}$, the unaries $u$ and $v$ have been reparametrized. Let the reparametrized unaries for $\gamma$ be $\left(\theta_{u}^{\gamma}, \theta_{v}^{\gamma}\right)$ and for $\mu$ be $\left(\theta_{u}^{\mu}, \theta_{v}^{\mu}\right)$. From $\gamma \geq \mu$ we know $\theta_{u}^{\gamma} \geq \theta_{u}^{\mu}$ and
$\theta_{v}^{\gamma} \geq \theta_{v}^{\mu}$. Then, $g_{u v}^{\gamma}=\theta_{u}^{\gamma}+\theta_{v}^{\gamma}+\theta_{u v}$ and $g_{u v}^{\mu}=\theta_{u}^{\mu}+\theta_{v}^{\mu}+\theta_{u v}$. It follows $g_{u v}^{\gamma} \geq g_{u v}^{\mu}$. Consider the chain of inequalities

$$
\begin{equation*}
\gamma\left[g_{u v}^{\gamma}\right] \geq \mu\left[g_{u v}^{\gamma}\right] \geq \mu\left[g_{u v}^{\mu}\right] \tag{4}
\end{equation*}
$$

The first is true by $\gamma \geq \mu$ dominance. The second is true by $\mu$-monotonicity. Thus for case B also, $\gamma$ results in reparametrized unaries that are co-ordinate wise greater than $\mu$.

Case $\mathbf{C}$, can be proven in much the same way as case $B$.

Theorem 1. The MPLP ++ algorithm converges to node-edge agreement.
To prove node-edge agreement we have to show that

$$
\begin{equation*}
\lim _{i \rightarrow \infty} \varepsilon\left(\mathcal{H}^{i}(\theta)\right)=0 \tag{5}
\end{equation*}
$$

where $\varepsilon$ is the tolerance factor defined in § B.1.
By saying that MPLP++ converges to node-edge agreement, we mean that as the algorithm progresses $\varepsilon$ tends to 0 and the set of labels belonging to $\mathcal{O}^{\varepsilon}$ is sequentially pruned only to leave an optimal labelling satisfying the node-edge agreement condition, converting $\mathcal{O}^{\varepsilon}$ to $\mathcal{O}^{0}$.

The proof is dependent on several lemmas which are sequentially proved.

## Lemma B.1. $\mathcal{H}$ is a continuous function.

Proof. Stated differently, we show the proposed reparameterizations in this paper are continuous. To generalize the result across all proposed reparameterizations we recall the idea of an oracle call from the main paper, i.e. operations of the type $\min _{t \in \mathcal{Y}} g_{u v}(s, t)$, $\forall s \in \mathcal{Y}$. We show the continuity of the first equation of the MPLP ++ operation only ( $1^{s t}$ equation of Eqn. $(\mathcal{H})$ in the main paper). The other equations of MPLP ++ and other algorithms can be proved similarly.

We prove continuity for only one label of one unary $\operatorname{cost}$ (label $s$ and unary $u$ ). Let $\delta>0$. Similar proofs hold for all other unaries and pairwise costs. Consider two edges denoted by the triplet $\left(\theta_{u}, \theta_{v}, \theta_{u v}\right)$ and $\left(\theta_{u}^{\prime}, \theta_{v}, \theta_{u v}\right)$, such that

$$
\begin{equation*}
\left|\theta_{u}^{\prime}(s)-\theta_{u}(s)\right|<\delta \tag{6}
\end{equation*}
$$

Then, $g_{u v}:=\theta_{u}+\theta_{v}+\theta_{u v}$ and $g_{u v}^{\prime}:=\theta_{u}^{\prime}+\theta_{v}+\theta_{u v}$. By (6), we have

$$
\begin{equation*}
\forall t\left|g_{u v}^{\prime}(s, t)-g_{u v}(s, t)\right|<\delta \tag{7}
\end{equation*}
$$

Applying the first operation of MPLP ++ to $g_{u v}$ and $g_{u v}^{\prime}$, we get

$$
\begin{align*}
\theta_{u}^{\mathcal{H}}(s) & :=\frac{1}{2} \min _{t \in \mathcal{Y}} g_{u v}(s, t)  \tag{8}\\
\theta_{u}^{\prime \mathcal{H}}(s) & :=\frac{1}{2} \min _{t \in \mathcal{Y}} g_{u v}^{\prime}(s, t) \tag{9}
\end{align*}
$$

Let $\left|\theta_{u}^{\prime \mathcal{H}}(s)-\theta_{u}^{\mathcal{H}}(s)\right|<\nu$, then by choosing $\delta:=\nu / 2$, we have
$\forall \theta \in \mathbb{R}^{\mathcal{I}}, \forall \nu>0, \exists \delta>0, \forall \theta^{\prime} \in \mathbb{R}^{\mathcal{I}},\left|\theta_{u}^{\prime}(s)-\theta_{u}(s)\right|<\delta \Longrightarrow\left|\theta_{u}^{\prime \mathcal{H}}(s)-\theta_{u}^{\mathcal{H}}(s)\right|<\nu$

Lemma B.2. The dual function $D$ is a continuous function in it's input $\theta$.
Proof. Consider two different reparameterizations $\theta^{\phi}, \theta^{\kappa} \in \mathbb{R}^{\mathcal{I}}$. To prove $D($.$) is a$ continuous function, it suffices to show $\forall \nu>0, \exists \delta>0$, such that $\left|\theta^{\phi}-\theta^{\kappa}\right|<$ $\delta \Longrightarrow\left|D\left(\theta^{\phi}\right)-D\left(\theta^{\kappa}\right)\right|<\nu$. To do so, we need to recall from § B.1, $\Theta_{p w}(\phi, \kappa)=$ $\max _{u v \in \mathcal{E}, s t \in \mathcal{Y}^{2}}\left|\theta_{u v}^{\phi}(s, t)-\theta_{u v}^{\kappa}(s, t)\right|$ and $\Theta_{u n}(\phi, \kappa)=\max _{u \in \mathcal{V}, s \in \mathcal{Y}}\left|\theta_{u}^{\phi}(s)-\theta_{u}^{\kappa}(s)\right|$.

We then have

$$
\begin{aligned}
\mid D\left(\theta^{\phi}\right)- & D\left(\theta^{\kappa}\right)\left|=\left|\left(\sum_{u \in \mathcal{V}} \min _{s} \theta_{u}^{\phi}(s)+\sum_{u v \in \mathcal{E}} \min _{(s, t)} \theta_{u v}^{\phi}(s, t)\right)-\left(\sum_{u \in \mathcal{V}} \min _{s} \theta_{u}^{\kappa}(s)+\sum_{u v \in \mathcal{E}} \min _{s, t} \theta_{u v}^{\kappa}(s, t)\right)\right|\right. \\
& \left.=\mid \sum_{u \in \mathcal{V}}\left(\min _{s} \theta_{u}^{\phi}(s)-\min _{s} \theta_{u}^{\kappa}(s)\right)+\sum_{u v \in \mathcal{E}}\left(\min _{s, t}^{\phi} \theta_{u v}(s, t)\right)-\min _{s, t}^{\kappa} \theta_{u v}^{\kappa}(s, t)\right) \mid \\
\leq & \left.\left|\sum_{u \in \mathcal{V}}\left(\min _{s} \theta_{u}^{\phi}(s)-\min _{s} \theta_{u}^{\kappa}(s)\right)\right|+\mid \sum_{u v \in \mathcal{E}}\left(\min _{s, t} \theta_{u v}^{\phi}(s, t)\right)-\min _{s, t}^{\kappa} \theta_{u v}^{\kappa}(s, t)\right) \mid \\
\leq \mid & \sum_{u \in \mathcal{V}} \Theta_{u n}(\phi, \kappa)\left|+\left|\sum_{u v \in \mathcal{E}} \Theta_{p w}(\phi, \kappa)\right| \leq\left|\mathcal { V } \left\|\Theta _ { u n } ( \phi , \kappa ) \left|+\left|\mathcal{E} \| \Theta_{p w}(\phi, \kappa)\right|\right.\right.\right.\right.
\end{aligned}
$$

Thus for $\left|D\left(\theta^{\phi}\right)-D\left(\theta^{\kappa}\right)\right|<\nu$, we need to choose a $\delta$ as a function of $\left|\Theta_{p w}\right|<\frac{1}{|\mathcal{E}|}$ and $\left|\Theta_{u n}\right|<\frac{1}{|\mathcal{V}|}$. This can be done as $\mathcal{H}$ is continuous (by proposition B.1) and pointwisemaxima are also continuous [53].

Lemma B.3. Tolerance factor $\varepsilon$ is continuous on $\theta$.
Proof. Let $\theta^{\phi}$ and $\theta^{\kappa}$ be two reparameterizations. We prove $\varepsilon$ is continuous for unaries only. The proof for pairwise terms is similar. Let $\Theta_{u n}(\phi, \kappa)=\delta$. From the definition of $\Theta_{u n}$ we have $\max _{u \in \mathcal{V}, s \in \mathcal{Y}}\left|\theta_{u}^{\phi}(s)-\theta_{u}^{\kappa}(s)\right|=\delta$. Thus,

$$
\begin{equation*}
\left|\theta_{u}^{\phi}(s)-\theta_{u}^{\kappa}(s)\right| \leq \delta \tag{11}
\end{equation*}
$$

(11) can be rewritten as

$$
\begin{equation*}
\theta_{u}^{\kappa}(s)-\delta \leq \theta_{u}^{\phi}(s) \leq \theta_{u}^{\kappa}(s)+\delta \tag{12}
\end{equation*}
$$

Now consider the set $\mathcal{O}_{u}^{\varepsilon}(\theta)=\left\{s \mid \theta_{u}(s) \leq \min _{s^{\prime}}\left(\theta_{u}\left(s^{\prime}\right)+\varepsilon(\theta)\right)\right\}$. For a unary in $\mathcal{O}_{u}^{\varepsilon}(\theta)$

$$
\begin{equation*}
\theta_{u}^{\phi}(s) \leq \min _{s^{\prime}}\left(\theta_{u}^{\phi}\left(s^{\prime}\right)+\varepsilon(\theta)\right) \tag{13}
\end{equation*}
$$

Substituting (12) in (13) we get,

$$
\begin{equation*}
\theta_{u}^{\kappa}(s)-\delta \leq \min _{s^{\prime}}\left(\theta_{u}^{\kappa}\left(s^{\prime}\right)+\delta+\varepsilon(\theta)\right) \Longrightarrow \theta_{u}^{\kappa}(s) \leq \min _{s^{\prime}}\left(\theta_{u}^{\kappa}\left(s^{\prime}\right)+\varepsilon(\theta)+2 \delta\right) \tag{14}
\end{equation*}
$$

Thus if $s$ satisfies (13), it also satisfies (14), which has $\varepsilon^{\prime} \geq \varepsilon+2 \delta$.
Lemma B.4. $D\left(\mathcal{H}^{i+1}(\theta)\right) \geq D\left(\mathcal{H}^{i}(\theta)\right), \forall i$, i.e. the $\mathrm{MPLP}++$ reparametrization never decreases the dual $D$.

Proof. Let's consider $D$ to be fixed for all variables except the block $D_{u v}$. Also, we assume that the costs have been normalized, i.e. $\min _{s \in \mathcal{Y}} \theta_{u}(s)=0, \min _{t \in \mathcal{Y}} \theta_{v}(t)=0$ and $\min _{s t \in \mathcal{Y}^{2}} \theta_{u v}(s, t)=0$. So, we have $D_{u v}(\theta)=0$. We thus have to show that $D_{u v}(\mathcal{H}(\theta)) \geq D_{u v}(\theta)=0$.

The aggregated potential $g_{u v}=\theta_{u}+\theta_{v}+\theta_{u v}$ can be written in the form of a $\mathcal{Y} \times \mathcal{Y}$ matrix

$$
\left[\begin{array}{ccc}
r_{1}+\Delta_{1,1} & \ldots & r_{1}+\Delta_{1,|\mathcal{Y}|} \\
\vdots & \ddots & \vdots \\
r_{|\mathcal{Y}|}+\Delta_{|\mathcal{Y}|, 1} & \ldots & r_{|\mathcal{Y}|}+\Delta_{|\mathcal{Y}|,|\mathcal{Y}|}
\end{array}\right]
$$

where $r_{s}$ is the row minimum and $\Delta_{s, t} \geq 0$. So, $\Delta_{s, t}$ is 0 for all elements of the row that are equal to $r_{s}$. As each element of row $s$ of $g_{u v}$ contains $\theta_{u}(s)$, we know that $r_{s} \geq \theta_{u}(s), \forall s \in \mathcal{Y}$.

Now, consider the $1^{\text {st }}$ equation of the MPLP $++\operatorname{operation}(\mathcal{H})$

$$
\begin{aligned}
\theta_{u}^{\mathcal{H}}(s) & :=\frac{1}{2} \min _{t \in \mathcal{Y}}\left[g_{u v}(s, t)\right]=\frac{1}{2} \min _{t \in \mathcal{Y}} r_{s}+\Delta_{s, t}=\frac{r_{s}}{2}, \forall s \in \mathcal{Y} \\
\theta_{v}^{\mathcal{H}}(t) & :=\frac{1}{2} \min _{s \in \mathcal{Y}}\left[g_{u v}(s, t)\right]=\frac{1}{2} \min _{s \in \mathcal{Y}} r_{s}+\Delta_{s, t}, \forall t \in \mathcal{Y}
\end{aligned}
$$

It is $r_{s} \geq \theta_{u}(s) \geq 0$ and therefore, $\theta_{u}^{\mathcal{H}}(s)=r_{s} / 2 \geq 0$. Likewise, $\theta_{v}^{\mathcal{H}}(t)=$ $\frac{1}{2} \min _{s \in \mathcal{Y}}\left[r_{s}+\Delta_{s, t}\right]$, where $r_{s} \geq 0$ and $\Delta_{s, t} \geq 0 \forall t$, thus $\theta_{v}^{\mathcal{H}}(t) \geq 0$. Hence, $D(\mathcal{H}(\theta)) \geq 0$.

Lemma B.5. $\mathcal{H}$ converges to a fixed point.
Proof. Initially, the dual bound $D\left(\theta^{\phi}\right)$ is computed using (3) which takes the min over all unaries $\theta_{u}(s)$ and pairwise terms $\theta_{u v}(s, t)$ individually. Thus, $D\left(\theta^{\phi}\right)$ is bounded unless for any $\theta_{u}(s)$ or $\theta_{u v}(s, t)$ all the elements are $\infty$. If the dual is unbounded, then by strong duality [53] the primal is also unbounded, and the only energy attainable is $\infty$. On the other hand, if $D\left(\theta^{\phi}\right)$ is bounded, by the LP duality theorem, the primal energy $E(y \mid \theta)$ serves as an upper bound for the $D\left(\theta^{\phi}\right)$.

We have proven that the BCA-update $\mathcal{H}$ brings about a monotonic improvement of $D\left(\theta^{\phi}\right)$. As each algorithm involves performing these updates over a sequence of edges $e^{i}$ covering the entire graph, we end up at the end of every iteration with a non-decreasing dual $D\left(\theta^{\phi}\right)$. Thus, each algorithm generates an increasing sequence of $D\left(\theta^{\phi}\right)$ which is bounded from above and by the Monotone Convergence Theorem for real numbers $\mathbb{R}$ [54] (Section 3.3) the algorithm converges.

Lemma B.6. $\mathcal{H}$ is bounded w.r.t. $D$ : For any $\theta$ there exists an $M \in \mathbb{R}$ such that $D\left(\mathcal{H}^{i}(\theta)\right)<M$ for any $i$.

Proof. As $D$ is a dual LP, we know from LP-Duality Theorem [53] that $D\left(\mathcal{H}^{i}(\theta)\right) \leq$ $E(y \mid \theta)$, where $E$ is as in (1) and $y$ is the labelling. Also, an alternative upper bound can be constructed as follows. Let $M_{u}=\max _{s} \theta_{u}^{\mathcal{H}}(s), M_{u v}=\max _{s t} \theta_{u v}^{\mathcal{H}}(s, t)$. We then have $\left|\theta_{u}^{\mathcal{H}}(s)\right| \leq M_{u}, \forall s$ and $\left|\theta_{u v}^{\mathcal{H}}(s, t)\right| \leq M_{u v} \forall s, t$. This also implies $\min _{s \in \mathcal{Y}} \theta_{u}^{\mathcal{H}}(s) \leq$ $M_{u}$ and $\min _{s t \in \mathcal{Y}^{2}} \theta_{u v}^{\mathcal{H}}(s, t) \leq M_{u v}$.

Thus we have

$$
D(\mathcal{H}(\theta))=\sum_{u \in \mathcal{V}} \min _{s} \theta_{u}^{\mathcal{H}}(s)+\sum_{u v \in \mathcal{E}} \min _{s, t} \theta_{u v}^{\mathcal{H}}(s, t) \leq \sum_{u \in \mathcal{V}} M_{u}+\sum_{u v \in \mathcal{E}} M_{u v}
$$

Therefore $D(\mathcal{H}(\theta))$ is always bounded by $M_{\mathcal{H}(\theta)}=\sum_{u \in \mathcal{V}} M_{u}+\sum_{u v \in \mathcal{E}} M_{u v}$. We know from Lemma B. 5 that $\mathcal{H}$ converges to a fixed point, thus after a certain number of iterations no changes occur in $\mathcal{H}^{i}(\theta)$. To compute the required bound, one must simply take $\max _{i} M_{\mathcal{H}^{i}(\theta)}$ over all reparameterizations that have occurred.

Lemma B.7. For any $\theta$ there exists $C>0$ such that $\left\|\mathcal{H}^{i}(\theta)\right\| \leq C\|\theta\|$ for any $i$.
Proof. We start off by showing that the lemma is true for one label of a unary. Let $\theta^{\phi}(s)$ be the reparameterization of $\theta_{u}(s)$. Consider the edge triplet $\left(\theta_{u}, \theta_{v}, \theta_{u v}\right)$. Let $m_{u}:=\min _{s \in \mathcal{Y}} \theta_{u}(s), M_{u}:=\max _{s \in \mathcal{Y}} \theta_{u}(s), M_{v}:=\max _{t \in \mathcal{Y}} \theta_{v}(t)$ and $M_{u v}:=$ $\max _{s t \in \mathcal{Y}^{2}} \theta_{u v}(s, t)$.

$$
\begin{aligned}
\theta_{u}^{\mathcal{H}}(s) & :=\frac{1}{2} \min _{t} g_{u v}(s, t) \\
\theta_{u}^{\mathcal{H}}(s) & :=\frac{1}{2} \theta_{u}(s)+\frac{1}{2} \min _{t}\left\{\theta_{v}(t)+\theta_{u v}(s, t)\right\} \\
\theta_{u}^{\mathcal{H}}(s) & \leq \frac{1}{2} \theta_{u}(s)+\frac{1}{2}\left(M_{v}+M_{u v}\right) \\
\theta_{u}^{\mathcal{H}}(s) & \leq\left(1+\frac{M_{v}+M_{u v}}{\theta_{u}(s)}\right) \theta_{u}(s) \leq\left(1+\frac{M_{v}+M_{u v}}{m_{u}}\right) \theta_{u}(s)
\end{aligned}
$$

Thus, we can choose $C_{u}^{i}=\left(1+\frac{M_{u}+M_{u v}}{m_{u}}\right)$ and get a bound $\left\|\theta_{u}^{\mathcal{H}}\right\| \leq C_{u}^{i}\left\|\theta_{u}\right\|$. As we have to find an upper bound $\forall \theta_{u}$ and $\forall \theta_{u v}$, we repeat the process and then take the max over all $C_{u}^{i}$ and $C_{u v}^{i}$ obtaining $C^{i}$. This gives us a $C^{i}$ that satisfies $\left\|\theta^{i+1}=\mathcal{H}\left(\theta^{i}\right)\right\| \leq$ $C^{i}\left\|\theta^{i}\right\|$. By Lemma B. 5 after a certain number of iterations $n, \theta^{n}$ converges, so to get the required result we simply simply take $C=\prod_{i=1}^{n} C^{i}$, giving $\left\|\mathcal{H}^{i}(\theta)\right\| \leq C\|\theta\|$ for any $i$.

Lemma B.8. $D(\mathcal{H}(\theta))=D(\theta)$ implies $\mathcal{O}^{0}(\mathcal{H}(\theta)) \subseteq \mathcal{O}^{0}(\theta)$, i.e. at a fixed point no new optimal labellings are achieved. If additionally $\varepsilon\left(\theta^{\mathcal{H}}\right)>0$, then $\mathcal{O}^{0}(\mathcal{H}(\theta)) \subset$ $\mathcal{O}^{0}(\theta)$, i.e. the inequality is strict.

Proof. Let $\theta^{\mathcal{H}}=\mathcal{H}(\theta)$, i.e. $\theta^{\mathcal{H}}$ is the reparameterized cost vector obtained after reparameterizing the original cost vector $\theta$. We prove the first assertion for unary cost $u$ only, the remaining costs can be proved similarly. To do so, we assume that only after reparameterization only the optimal label of unary cost $u$ has changed.

Let this label be $s^{*}=\arg \min _{s} \theta_{u}^{\mathcal{H}}(s)$. This implies $s^{*} \in \mathcal{O}\left(\theta^{\mathcal{H}}\right)$. Now, to prove the first assertion, we have to show the inclusion $s^{*} \in \mathcal{O}_{u}^{0}\left(\theta^{\mathcal{H}}\right) \Longrightarrow s^{*} \in \mathcal{O}_{u}^{0}(\theta)$.

Let us additionally define a function $\eta$ that takes as input a node index $p$ and outputs its optimal labelling, i.e. $k=\eta(v)$ implies $k=\arg \min _{s \in \mathcal{Y}} \theta_{p}(s)$. We define this function over all unary costs except $u$. As $\theta^{\mathcal{H}}$ and $\theta$ differ in only the label $u$, we can use $\eta$ for choosing the optimal labelling for both of them.

Since, $D\left(\theta^{\mathcal{H}}\right)=D(\theta)$, we have

$$
\begin{equation*}
\sum_{v \in \mathcal{V}} \min _{s \in \mathcal{Y}} \theta_{v}^{\mathcal{H}}(s)+\sum_{v w \in \mathcal{E}} \min _{s t \in \mathcal{Y}^{2}} \theta_{v w}^{\mathcal{H}}(s, t)=\sum_{v \in \mathcal{V}} \min _{s \in \mathcal{Y}} \theta_{v}(s)+\sum_{v w \in \mathcal{E}} \min _{s t \in \mathcal{Y}^{2}} \theta_{v w}(s, t) \tag{15}
\end{equation*}
$$

Since all the labels but the ones for $u$ are the same, the terms on both sides of (15) cancel out leaving

$$
\begin{equation*}
\min _{s \in \mathcal{Y}} \theta_{u}^{\mathcal{H}}(s)+\sum_{u v \mid v \in N b(u)} \min _{s t \in \mathcal{Y}^{2}} \theta_{u v}^{\mathcal{H}}(s, t)=\min _{s \in \mathcal{Y}} \theta_{u}(s)+\sum_{u v \mid v \in N b(u)} \min _{s t \in \mathcal{Y}^{2}} \theta_{u v}(s, t) \tag{16}
\end{equation*}
$$

Substituting $s^{*}=\arg \min _{s} \theta_{u}^{\mathcal{H}}(s)$ and $\eta$ into (16) we get

$$
\begin{equation*}
\theta_{u}^{\mathcal{H}}\left(s^{*}\right)+\sum_{u v \mid v \in N b(u)} \theta_{u v}^{\mathcal{H}}\left(s^{*}, \eta(v)\right)=\min _{s \in \mathcal{Y}} \theta_{u}(s)+\sum_{u v \mid v \in N b(u)} \theta_{u v}(s, \eta(v)) \tag{17}
\end{equation*}
$$

Now, by the definition of reparameterization we have, $\forall s \in \mathcal{Y}$

$$
\begin{equation*}
\theta_{u}^{\mathcal{H}}(s)+\sum_{u v \mid v \in N b(u)} \theta_{u v}^{\mathcal{H}}(s, \eta(v))=\theta_{u}(s)+\sum_{u v \mid v \in N b(u)} \theta_{u v}(s, \eta(v)) \tag{18}
\end{equation*}
$$

This also holds true for label $s^{*}$, thus

$$
\begin{equation*}
\theta_{u}^{\mathcal{H}}\left(s^{*}\right)+\sum_{u v \mid v \in N b(u)} \theta_{u v}^{\mathcal{H}}\left(s^{*}, \eta(v)\right)=\theta_{u}\left(s^{*}\right)+\sum_{u v \mid v \in N b(u)} \theta_{u v}\left(s^{*}, \eta(v)\right) \tag{19}
\end{equation*}
$$

Now, equating the RHS of (19) and the RHS of (17), we get

$$
\begin{equation*}
\theta_{u}\left(s^{*}\right)+\sum_{u v \mid v \in N b(u)} \theta_{u v}\left(s^{*}, \eta(v)\right)=\min _{s \in \mathcal{Y}} \theta_{u}(s)+\sum_{u v \mid v \in N b(u)} \theta_{u v}(s, \eta(v)) \tag{20}
\end{equation*}
$$

Thus $s^{*}=\arg \min _{s \in \mathcal{Y}} \theta_{u}(s) \Longrightarrow s^{*} \in \mathcal{O}_{u}^{0}(\theta)$, proving the first assertion.
To prove the second assertion, we have to show that if $\epsilon\left(\theta^{\mathcal{H}}\right)>0$ and $D(\mathcal{H}(\theta))=$ $D(\theta)$, there exists $s^{\prime} \in \mathcal{O}_{u}^{0}(\theta)$ such that $s^{\prime} \notin \mathcal{O}_{u}^{0}\left(\theta^{\mathcal{H}}\right)$.

Since, $s^{\prime} \in \mathcal{O}_{u}^{0}(\theta)$, we have

$$
\begin{equation*}
\theta_{u}\left(s^{\prime}\right)+\sum_{v \in N b(u)} \theta_{u v}\left(s^{\prime}, \eta(v)\right)=\min _{s \in \mathcal{Y}} \theta_{u}(s)+\sum_{v \in N b(u)} \min _{s \in \mathcal{Y}} \theta_{u v}(s, \eta(v)) \tag{21}
\end{equation*}
$$

From the reparameterization relation we have

$$
\begin{equation*}
\theta_{u}\left(s^{\prime}\right)+\sum_{v \in N b(u)} \theta_{u v}\left(s^{\prime}, \eta(v)\right)=\theta_{u}^{\mathcal{H}}\left(s^{\prime}\right)+\sum_{v \in N b(u)} \theta_{u v}^{\mathcal{H}}\left(s^{\prime}, \eta(v)\right) \tag{22}
\end{equation*}
$$

Since $\varepsilon\left(\theta^{\mathcal{H}}\right)>0$, it has to be the case that

$$
\begin{equation*}
\theta_{u}^{\mathcal{H}}\left(s^{\prime}\right)+\sum_{v \in N b(u)} \theta_{u v}^{\mathcal{H}}\left(s^{\prime}, \eta(v)\right)>\min _{s \in \mathcal{Y}} \theta_{u}^{\mathcal{H}}(s)+\sum_{v \in N b(u)} \min _{s \in \mathcal{Y}} \theta_{u v}^{\mathcal{H}}(s, \eta(v)) \tag{23}
\end{equation*}
$$

otherwise $\varepsilon\left(\theta^{\mathcal{H}}\right)=0$. Thus, $s^{\prime} \notin \mathcal{O}_{u}^{0}\left(\theta^{\mathcal{H}}\right)$.
Lemma B.9. There exists an $n$, such that $D\left(\mathcal{H}^{n+m}(\theta)\right)=D\left(\mathcal{H}^{n}(\theta)\right)$ implies $\mathcal{O}^{0}\left(\mathcal{H}^{n+m}(\theta)\right)=\mathcal{O}^{0}\left(\mathcal{H}^{n}(\theta)\right), \forall m \geq 0$.

Proof. Let $\theta^{0}, \theta^{1}, \ldots$ be the sequence of vectors generated from $\mathcal{H}$ via $\theta^{i+1}=\mathcal{H}\left(\theta^{i}\right)$. After a certain number of iterations, we have to show that $\mathcal{O}^{0}\left(\theta^{n+m}\right)=\mathcal{O}^{0}\left(\theta^{n}\right)$, for all $m \geq 0$. From Lemma B. 8 we have after a fixed point has been reached at iteration $i, \mathcal{O}^{0}\left(\theta^{i}\right) \supset \mathcal{O}^{0}\left(\theta^{i+1}\right) \supset \mathcal{O}^{0}\left(\theta^{i+2}\right) \supset \ldots$. Since $\mathcal{O}^{0}(\theta)$ is finite, it cannot shrink indefinitely and after a certain number $(n)$ of iterations $\mathcal{O}^{0}\left(\theta^{n}\right)$ will become consistent. Yielding, $\mathcal{O}^{0}\left(\theta^{m+n}\right)=\mathcal{O}^{0}\left(\theta^{n}\right), \forall m \geq 0$.

Now, we are ready to prove Theorem 1
Combining the above lemmas, we introduce the notion of consistency-enforcing algorithms. An algorithm is consistency-enforcing if it satisfies Lemmas B.1, B.2, B.3, В.4, В.5, В.6, В.7, В.8, В.9, .

The MPLP ++ operator $\mathcal{H}$ is consistency enforcing then $\lim _{i \rightarrow \infty} \varepsilon\left(\mathcal{H}^{i}(\theta)\right)=0$.

$$
\begin{equation*}
\lim _{i \rightarrow \infty} \varepsilon\left(\mathcal{H}^{i}(\theta)\right)=0 \tag{24}
\end{equation*}
$$

Theorem 1. The $\mathrm{MPLP}++$ algorithm converges to node-edge agreement.
Proof. By virtue of Lemma B. 7 and Lemma B.4, the sequence $\theta^{i}=\mathcal{H}^{i}(\theta)$ is bounded. Therefore, by the Bolzano-Weierstrass Theorem [54], there exists a converging subsequence $\theta^{i(j)}, j=1,2, \ldots$, where $j>j^{\prime}$ implies $i(j)>i\left(j^{\prime}\right)$, i.e. the limit $\theta^{*}:=$ $\lim _{j \rightarrow \infty} \theta^{i(j)}$ exists. Let us show that it holds

$$
\begin{equation*}
\varepsilon\left(\theta^{*}\right)=0 \tag{25}
\end{equation*}
$$

for any converging subsequence of $\theta^{i}$.
Since due to convergence of $\mathcal{H}$ (shown in Lemma B.4) and Lemma B. 6 the sequence $D\left(\theta^{i}\right)$ is non-decreasing and bounded from above, and therefore converges to a limit point $D^{*}:=\lim _{i \rightarrow \infty} D\left(\theta^{i}\right)$. Therefore, it also holds

$$
\begin{equation*}
D^{*}=\lim _{j \rightarrow \infty} D\left(\theta^{i(j)}\right)=\lim _{j \rightarrow \infty} D\left(\theta^{i(j)+n}\right), \quad \forall n \geq 0 \tag{26}
\end{equation*}
$$

This implies

$$
\begin{equation*}
0=\lim _{j \rightarrow \infty} D\left(\theta^{i(j)}\right)-\lim _{j \rightarrow \infty} D\left(\theta^{i(j)+n}\right)=\lim _{j \rightarrow \infty} D\left(\theta^{i(j)}\right)-D\left(\mathcal{H}^{n}\left(\theta^{i(j)}\right)\right) \tag{27}
\end{equation*}
$$

Since $D$ is continuous it holds

$$
\begin{equation*}
0=\lim _{j \rightarrow \infty}\left(D\left(\theta^{i(j)}\right)-D\left(\mathcal{H}^{n}\left(\theta^{i(j)}\right)\right)\right) \tag{28}
\end{equation*}
$$

and therefore, (25) holds by virtue of Lemma B.9.
Since, $\varepsilon$ is a continuous function, (25) implies

$$
\begin{equation*}
\lim _{j \rightarrow \infty} \varepsilon\left(\theta^{i(j)}\right)=0 \tag{29}
\end{equation*}
$$

for any converging sub-sequence $\theta^{i(j)}$.
Now, considering the sequence $\varepsilon\left(\theta^{i}\right)$, we know by virtue of Lemma B. $1 s^{i}:=$ $\sup _{j \geq i} \varepsilon\left(\theta^{j}\right)$. Sequence $s^{i}$ is a monotonically non-increasing sequence of non-negative numbers and therefore it has a limit $s^{*}=\lim _{i \rightarrow \infty} s^{i}$.

According to the "Theorem of Superior and Inferior Limits" [54] there exists a subsequence $\varepsilon\left(\theta^{i^{\prime}(j)}\right)$ such that

$$
\begin{equation*}
\lim _{j \rightarrow \infty} \varepsilon\left(\theta^{i^{\prime}(j(k))}\right)=s^{*} \tag{30}
\end{equation*}
$$

The sequence $\theta^{i^{\prime}(j)}$ is bounded virtue of Lemma B. 7 and therefore contains a converging subsequence $\theta^{i^{\prime}(j(k))}$. For this subsequence it also holds,

$$
\begin{equation*}
\lim _{k \rightarrow \infty} \varepsilon\left(\theta^{i^{\prime}(j(k))}\right)=s^{*} \tag{31}
\end{equation*}
$$

At the same time, as proved in (29), for any converging subsequence it holds

$$
\begin{equation*}
0=\lim _{k \rightarrow \infty} \varepsilon\left(\theta^{i^{\prime}(j(k))}\right)=s^{*}=\lim _{i \rightarrow \infty} \sup _{k \geq i} \varepsilon\left(\theta^{k}\right) \tag{32}
\end{equation*}
$$

Finally, $0 \leq \varepsilon\left(\theta^{i}\right) \leq \sup _{k \geq i} \varepsilon\left(\theta^{k}\right)$ implies $\lim _{i \rightarrow \infty} \varepsilon\left(\theta^{i}\right)=0$.
Proposition 1. Reparametrization $\phi_{u \leftrightarrow v}$ maximizes $D_{u v}(\cdot)$ iff there exist $(s, t) \in \mathcal{Y}^{2}$ such that s minimizes $\theta_{u}^{\phi}(\cdot)$, t minimizes $\theta_{v}^{\phi}(\cdot)$ and $(s, t)$ minimizes $\theta_{u v}^{\phi}(\cdot, \cdot)$.

Proof. The if clause has been proven in ([20], Theorem 2). The only if clause can be proven as follows:

As mentioned in the main paper and shown in $[15,21]$ the Dual LP $D(\phi)$ is a concave, piecewise linear function.

For proving optimality of block $D_{u v}\left(\phi_{u \leftrightarrow v}\right)$ we need to show $\mathbf{0} \in \partial D_{u v}\left(\phi_{u \leftrightarrow v}\right)$, where $\partial D_{u v}\left(\phi_{u \leftrightarrow v}\right)$ is the super-differential of $D_{u v}\left(\phi_{u \leftrightarrow v}\right)$. Reconsidering the dual $D_{u v}\left(\phi_{u \leftrightarrow v}\right)$

$$
\begin{equation*}
D_{u v}\left(\phi_{u \leftrightarrow v}\right):=\min _{s t \in \mathcal{Y}^{2}} \theta_{u v}^{\phi}(s, t)+\min _{s \in \mathcal{Y}} \theta_{u}^{\phi}(s)+\min _{t \in \mathcal{Y}} \theta_{v}^{\phi}(t) \tag{33}
\end{equation*}
$$

Let $x_{u}^{\prime}=\arg \min _{s \in \mathcal{Y}} \theta_{u}^{\phi}(s), x_{v}^{\prime}=\arg \min _{t \in \mathcal{Y}} \theta_{v}^{\phi}(t)$ and $\left(x_{u}^{\prime \prime}, x_{v}^{\prime \prime}\right)=\arg \min _{s t \in \mathcal{Y}^{2}} \theta_{u v}^{\phi}(s, t)$. Taking the super-gradient of $D_{u v}\left(\phi_{u \leftrightarrow v}\right)$ w.r.t. $\phi_{u \rightarrow v}$, we have

$$
\frac{\partial D_{u v}\left(\phi_{u \leftrightarrow v}\right)}{\partial \phi_{u \rightarrow v}(s)}:= \begin{cases}0, & s \neq x_{u}^{\prime}, s \neq x_{u}^{\prime \prime}  \tag{34}\\ 0, & s=x_{u}^{\prime}, s=x_{u}^{\prime \prime} \\ 1, & s=x_{u}^{\prime}, s \neq x_{u}^{\prime \prime} \\ -1, & s \neq x_{u}^{\prime}, s=x_{u}^{\prime \prime}\end{cases}
$$

Likewise, taking the super-gradient of $D_{u v}\left(\phi_{u \leftrightarrow v}\right)$ w.r.t. $\phi_{v \rightarrow u}$, we have

$$
\frac{\partial D_{u v}\left(\phi_{u \leftrightarrow v}\right)}{\partial \phi_{v \rightarrow u}(t)}:= \begin{cases}0, & t \neq x_{v}^{\prime}, t \neq x_{v}^{\prime \prime}  \tag{35}\\ 0, & t=x_{v}^{\prime}, t=x_{v}^{\prime \prime} \\ 1, & t=x_{v}^{\prime}, t \neq x_{v}^{\prime \prime} \\ -1, & t \neq x_{v}^{\prime}, t=x_{v}^{\prime \prime}\end{cases}
$$

Thus, if $s=x_{u}^{\prime}=x_{u}^{\prime \prime}$ and $t=x_{v}^{\prime}=x_{v}^{\prime \prime}$, we have $0 \in \frac{\partial D_{u v}\left(\phi_{v \leftrightarrow u}\right)}{\partial \phi_{v \leftrightarrow u}(s, t)}$, proving the only if clause.

