How Hard Is the LP Relaxation of the Potts Min-Sum Labeling Problem?

Daniel Průša and Tomáš Werner

Czech Technical University, Faculty of Electrical Engineering
Karlovo náměstí 13, 121 35 Prague 2, Czech Republic
{prusapa1,werner}@cmp.felk.cvut.cz

Abstract. An important subclass of the min-sum labeling problem (also known as discrete energy minimization or valued constraint satisfaction) is the pairwise min-sum problem with arbitrary unary costs and attractive Potts pairwise costs (also known as the uniform metric labeling problem). In analogy with our recent result, we show that solving the LP relaxation of the Potts min-sum problem is not significantly easier than that of the general min-sum problem and thus, in turn, the general linear program. This suggests that trying to find an efficient algorithm to solve the LP relaxation of the Potts min-sum problem has a fundamental limitation. Our constructions apply also to integral solutions, yielding novel reductions of the (non-relaxed) general min-sum problem to the Potts min-sum problem.

Keywords: Markov random field, discrete energy minimization, valued constraint satisfaction, linear programming relaxation, uniform metric labeling problem, Potts model.

1 Introduction

The *min-sum (labeling) problem*, also known as discrete energy minimization [15,5] or valued constraint satisfaction [16], has numerous applications in machine learning and computer vision and other fields, in particular as MAP inference in graphical models [17]. The problem has a natural LP relaxation [13,18,7,3,17], which underlies many algorithms to approximately solve the problem (see [5] and references therein). It is therefore of great practical importance to have efficient algorithms to solve this LP. Unfortunately, the simplex and interior point methods solving general LP are prohibitively inefficient for large-scale vision instances.

It is known that the LP relaxation of the pairwise min-sum problem with 2 labels reduces in linear time to max-flow [1,11]. Therefore, this problem can be solved very efficiently because the worst-case complexity of best known algorithms for max-flow is much better than for the general LP (though both are in the P complexity class). Our recent paper [10] showed that solving the LP relaxation of the pairwise min-sum problem with 3 or more labels (with some costs possibly infinite) is as hard as solving the general LP, precisely, the latter reduces to the former in linear time. This suggests that trying to find a very efficient algorithm to solve the LP relaxation may be futile.

This negative result raises the question whether there are any other useful subclasses of the min-sum problem for which the LP relaxation is significantly easier than the general linear program and therefore there is hope for efficient algorithms. In this paper,

X.-C. Tai et al. (Eds.): EMMCVPR 2015, LNCS 8932, pp. 57–70, 2015.

we show that this is unlikely for the class of pairwise min-sum problems with attractive Potts costs, which is also known as the uniform metric labeling problem [2,6,3,4].

We present two efficient reductions of the general pairwise min-sum problem to the Potts min-sum problem that preserve the LP relaxation. The first one (§4, §5) reduces the general min-sum problem with some costs possibly infinite to the Potts min-sum problem with 3 labels (the complexity of this reduction is given by Theorems 5 and 8). Combined with [10], this implies that solving the general system of linear inequalities reduces in linear time to the LP relaxation of the Potts min-sum problem with 3 labels (Corollary 6, our most surprising result) and that the general linear program reduces in better than quadratic time to the LP relaxation of the Potts min-sum problem with 3 labels (Corollary 9). The second one (§6) reduces the general min-sum problem with k labels and finite costs to the Potts min-sum problem with k labels (Theorem 11). The output costs in this reduction are typically much smaller than in the first reduction.

Though these results are somewhat weaker than for the general min-sum problem [10], they are far from obvious. They show that finding an efficient algorithm to solve the LP relaxation of the Potts min-sum problem is unlikely because this might mean improving the complexity of the best known algorithms for the general LP. An example of an algorithm specialized to the LP relaxation of the Potts min-sum problem is [9].

Our reductions straightforwardly apply also to the original non-relaxed min-sum problems, thus we obtain as side-results novel reductions from the general min-sum problem to the Potts one (Theorems 4, 7, and 10). These results can be seen analogical to, e.g., the construction [12] which reduces the general pairwise min-sum problem with finite costs to the pairwise min-sum problem with 2 labels.

2 Min-sum Problem and Its LP Relaxation

We denote $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$ and $\overline{\mathbb{Z}} = \mathbb{Z} \cup \{\infty\}$. Let (V, E) be a connected undirected graph, with *objects* V and *object pairs* $E \subseteq \binom{V}{2}$. Let K be a finite set of *labels*. Let $g_u\colon K \to \overline{\mathbb{Q}}$ and $g_{uv}\colon K \times K \to \overline{\mathbb{Q}}$ be unary and pairwise *cost functions*, where we adopt that $g_{uv}(k, \ell) = g_{vu}(\ell, k)$. The pairwise *min-sum problem* is defined as

$$\min_{\mathbf{k} \in K^V} \Big(\sum_{v \in V} g_u(k_u) + \sum_{\{u,v\} \in E} g_{uv}(k_u, k_v) \Big). \tag{1}$$

All the costs $g_u(k)$ and $g_{uv}(k, \ell)$ together will be understood as a vector $\mathbf{g} \in \overline{\mathbb{Q}}^I$ with $I = (V \times K) \cup \{\, \{(u, k), (v, \ell)\} \mid \{u, v\} \in E;\ k, \ell \in K \,\}$.

The *local marginal polytope* [17] is the set Λ of vectors $\boldsymbol{\mu} \in \mathbb{R}^I_+$ satisfying

$$\sum_{k \in K} \mu_u(k) = 1, \qquad u \in V \tag{2a}$$

$$\sum_{\ell \in K} \mu_{uv}(k, \ell) = \mu_u(k), \qquad u \in V,\ v \in N_u,\ k \in K \tag{2b}$$

where $N_u = \{\, v \mid \{u, v\} \in E \,\}$ are the neighbors of object u and we again adopt that $\mu_{uv}(k, \ell) = \mu_{vu}(\ell, k)$. The numbers $\mu_u(k), \mu_{uv}(k, \ell)$ are known as *pseudomarginals* [17]. Figure 1 illustrates the meaning of constraints (2) for one object pair.

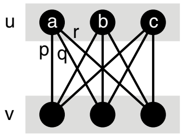

Fig. 1. Two objects forming an object pair $\{u, v\} \in E$. Objects $u \in V$ are depicted as boxes, labels (u, k) as nodes, and label pairs $\{(u, k), (v, \ell)\}$ as edges. Note the meaning of constraints (2): for unary pseudomarginals a, b, c and pairwise pseudomarginals p, q, r, (2b) enforces $a = p + q + r$ and (2a) enforces $a + b + c = 1$.

The LP relaxation of problem (1) reads

$$\min\{\, \mathbf{g}^\top \boldsymbol{\mu} \mid \boldsymbol{\mu} \in \Lambda \,\} \tag{3}$$

where, if some costs (components of $\mathbf{g}$) are infinite, we define $0 \cdot \infty = 0$ in the scalar product $\mathbf{g}^\top \boldsymbol{\mu}$. If $\boldsymbol{\mu} \in \{0, 1\}^I$ then (3) solves (1) exactly.

Reparameterizations of a vector $\mathbf{g} \in \overline{\mathbb{Q}}^I$ is a vector $\mathbf{g}' \in \overline{\mathbb{Q}}^I$ given by

$$g'_u(k) = g_u(k) - \sum_{v \in N_u} \varphi_{uv}(k) \tag{4a}$$

$$g'_{uv}(k, \ell) = g_{uv}(k, \ell) + \varphi_{uv}(k) + \varphi_{vu}(\ell) \tag{4b}$$

where $\boldsymbol{\varphi} = (\varphi_{uv}(k) \in \mathbb{R} : u \in V,\ v \in N_u,\ k \in K)$. We have $\mathbf{g}^\top \boldsymbol{\mu} = \mathbf{g}'^\top \boldsymbol{\mu}$ for every $\boldsymbol{\varphi}$ and every $\boldsymbol{\mu}$ satisfying (2), thus reparameterizations preserve the objective of (1) and its LP relaxation. Consider a *lower bound*

$$L(\mathbf{g}) = \sum_{u \in V} \min_{k \in K} g_u(k) + \sum_{\{u,v\} \in E} \min_{k, \ell \in K} g_{uv}(k, \ell) \tag{5}$$

on the true optimal value (1). The dual to the LP (3) can be expressed [18] as maximizing the lower bound over reparameterizations, i.e., maximizing $L(\mathbf{g}')$ over $\boldsymbol{\varphi}$.

If the pairwise cost functions g_{uv} in (1) are metric while the unary cost functions g_u remains arbitrary, the problem (1) has been called the *metric labeling problem* [2,6,3,4]. Its special case is the *uniform metric* or the *attractive Potts interaction*

$$g_{uv}(k, \ell) = h_{uv} [\![k \neq \ell]\!] \tag{6}$$

where $h_{uv} \geq 0$ and $[\![k \neq \ell]\!]$ equals 1 if $k \neq \ell$ and 0 otherwise. We will refer to problem (1) with pairwise costs (6) as the *Potts min-sum problem.*

3 Summary of Results

This section gives the overview of our contributions in this paper, after formulating previous closely related results that we obtained in [10].

As is usual in computational complexity, we will use the notions of problem (a set of instances), instance, and reduction. We start this section by defining the following problems, by specifying their instances (inputs) and solutions (outputs). Rather than more common decision problems, which map strings over an alphabet to the answers {yes, no}, we formulate our problems as *function problems*, which map strings over an alphabet to strings over an alphabet.

Problem: MinSum(Y) where $Y \subseteq \overline{\mathbb{Q}}$
Instance: $(V, E, K, \mathbf{g})$ where $\mathbf{g} \in Y^I$. (Thus, Y specifies the set of values the costs can take. E.g., in MinSum$(\mathbb{Z})$ the costs can take values from $\mathbb{Z}$ rather than from $\overline{\mathbb{Q}}$.)
Solution: If the optimal value of problem (1) is finite, it returns an optimal argument $\mathbf{k} \in K^V$. Otherwise, it answers 'infeasible'.

Problem: MinSum(Y)-LP
Instance: $(V, E, K, \mathbf{g})$ where $\mathbf{g} \in Y^I$ and $Y \subseteq \overline{\mathbb{Q}}$.
Solution: If the LP (3) is feasible, it returns an optimal argument $\boldsymbol{\mu} \in [0, 1]^I$. If (3) is infeasible, it answers 'infeasible'.

Problem: Potts
Instance: $(V, E, K, \mathbf{g})$ where $\mathbf{g} \in \mathbb{Q}^I$ and pairwise costs in $\mathbf{g}$ have the form (6).
Solution: An optimal argument $\mathbf{k} \in K^V$ of problem (1).

Problem: Potts-LP
Instance: $(V, E, K, \mathbf{g})$ where $\mathbf{g} \in \mathbb{Q}^I$ and pairwise costs in $\mathbf{g}$ have the form (6).
Solution: An optimal argument $\boldsymbol{\mu} \in [0, 1]^I$ of problem (3).

Problem: LinIneq
Instance: $(\mathbf{A}, \mathbf{b})$ where $\mathbf{A} \in \mathbb{Z}^{m \times n}$, $\mathbf{b} \in \mathbb{Z}^m$.
Solution: If the system $\{\, \mathbf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0} \,\}$ has a solution, it returns its arbitrary solution. Otherwise, it answers 'infeasible'.

Problem: LinProg
Instance: $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ where $\mathbf{A} \in \mathbb{Z}^{m \times n}$, $\mathbf{b} \in \mathbb{Z}^m$, $\mathbf{c} \in \mathbb{Z}^n$.
Solution: If the linear program $\min\{\, \mathbf{c}^\top \mathbf{x} \mid \mathbf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0} \,\}$ is feasible and bounded, it returns a solution $\mathbf{x} \in \mathbb{Q}^n$. If the program is infeasible, it answers 'infeasible'. If the program is unbounded, it answers 'unbounded'.

Instance Sizes. In general, the size of a problem instance is the length of the (binary) string needed to encode it. We will use $\langle x \rangle$ to denote the size of a number $x \in \mathbb{Z}$. Using one bit for the sign, storing x takes $\langle x \rangle = \lceil \log_2(|x| + 1) \rceil + 1$ bits. For a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}^n$, we define its size to be $\langle \mathbf{x} \rangle = \sum_{i=1}^{n} \langle x_i \rangle$. We will use this definition of size for vectors $\mathbf{g}$ and $\mathbf{c}$.

For $\mathbf{A}$ and $\mathbf{b}$ we use a slightly different definition of size. The pair $(\mathbf{A}, \mathbf{b})$ can be seen as the extended matrix $\bar{\mathbf{A}} = [\mathbf{A} \,|\, \mathbf{b}] \in \mathbb{Z}^{m \times (n+1)}$. Encoding $\bar{\mathbf{A}}$ by storing all its entries (including zeros) would take $L = \sum_{i=1}^{m} \sum_{j=1}^{n+1} \langle a_{ij} \rangle$ bits. This would describe the dense encoding of $\bar{\mathbf{A}}$. However, we define

$$\langle \bar{\mathbf{A}} \rangle = \sum_{i=1}^{m} \sum_{j=1}^{n+1} \lceil \log_2(|\bar{a}_{ij}| + 1) \rceil. \tag{7}$$

As zero entries $\bar{a}_{ij} = 0$ do not contribute to $\langle \bar{\mathbf{A}} \rangle$, this describes a *sparse* encoding of $\bar{\mathbf{A}}$. Note that $\langle \bar{\mathbf{A}} \rangle \leq L$, therefore (7) covers both sparse and dense encoding because $\bar{\mathbf{A}}$ will always describe input (never output) instances of our reductions. Indeed, for every $f \colon \mathbb{N} \to \mathbb{N}$, if the complexity of a reduction is $\mathcal{O}(f(\langle \bar{\mathbf{A}} \rangle))$ then it is also $\mathcal{O}(f(L))$.

For convenience, we defined the instance of POTTS and POTTS-LP the same way as for MINSUM(Y)-LP and MINSUM(Y), namely by the tuple $(V, E, K, \mathbf{g})$ with $\mathbf{g} \in \mathbb{Q}^I$. However, the components of $\mathbf{g}$ are not independent since they satisfy (6). This must be taken into account when computing $\langle \mathbf{g} \rangle$ for POTTS and POTTS-LP.

Existing Results. The results obtained in [10] can be formulated as follows.

Theorem 1. LINPROG *reduces in linear time to* MINSUM$(\overline{\mathbb{Z}})$-LP *with 3 labels.*

Theorem 2. LINPROG *reduces in quadratic time to* MINSUM$(\mathbb{Z})$-LP *with 3 labels.*

Theorem 3. LININEQ *reduces in linear time to* MINSUM$(\{0, \infty\})$-LP *with 3 labels.*

Theorem 3 is not explicitly stated in [10]. It holds because LININEQ is LINPROG with $\mathbf{c} = \mathbf{0}$, in which case the output min-sum problem has costs in $\{0, \infty\}$ [10, §5].

Contributions. Our contributions in this paper are two reductions of the general min-sum problem to the Potts min-sum problem that preserve both the optimum of (1) and the optimum of its LP relaxation (3). These reductions lead to the following results.

Theorem 4. MINSUM$(\{0, \infty\})$ *reduces in linear time to* POTTS *with 3 labels.*

Theorem 5. MINSUM$(\{0, \infty\})$-LP *reduces in linear time to* POTTS-LP *with 3 labels.*

Corollary 6. LININEQ *reduces in linear time to* POTTS-LP *with 3 labels.*

Proof. Combine Theorem 3 and Theorem 5. $\square$

Theorem 7. MINSUM$(\overline{\mathbb{Z}})$ *with p object pairs, k labels and size L reduces in time $\mathcal{O}(pk^2L)$ to* POTTS *with 3 labels.*

Theorem 8. MINSUM$(\overline{\mathbb{Z}})$-LP *with p object pairs, k labels and size L reduces in time $\mathcal{O}(pk^2L)$ to* POTTS-LP *with 3 labels.*

Corollary 9. LINPROG *reduces in quadratic time to* POTTS-LP *with 3 labels.*

Proof. By Theorem 8, MINSUM$(\overline{\mathbb{Z}})$-LP reduces in quadratic time to POTTS-LP with 3 labels, because $pk^2 = \mathcal{O}(L)$ and so $\mathcal{O}(pk^2L) \subseteq \mathcal{O}(L^2)$. This is combined with Theorem 1. $\square$

Theorem 10. MINSUM$(\mathbb{Z})$ *with k labels and size L reduces in time $\mathcal{O}(k^2L)$ to* POTTS *with k labels.*

Theorem 11. MINSUM$(\mathbb{Z})$-LP *with k labels and size L reduces in time $\mathcal{O}(k^2L)$ to* POTTS-LP *with k labels.*

In §4 we will describe our first reduction for input costs in $\{0, \infty\}$ and thereby prove Theorems 4 and 5. In §5 we generalize this to arbitrary costs, proving thus Theorems 7 and 8. In §6, we describe our second reduction and prove Theorems 10 and 11.

4 Encoding a Local Marginal Polytope

Consider the polyhedron

$$P = \{\, \mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} = \mathbf{b},\ \mathbf{x} \geq \mathbf{0} \,\} \tag{8}$$

where $\mathbf{A} \in \{-1, 0, 1\}^{m \times n}$ and $\mathbf{b} \in \{0, 1\}^m$ satisfy the following conditions:

(P1) $P \subseteq [0, 1]^n$.
(P2) Each row of the matrix $[-\mathbf{A} \mid \mathbf{b}\,]$ contains exactly one positive number.
(P3) Each row of $\mathbf{A}$ contains at most d non-zeros.
(P4) $\mathbf{A}$ has in total $\mathcal{O}(n)$ non-zeros.

Every local marginal polytope with k labels and p object pairs has this form, where $m = \mathcal{O}(kp)$, $n = \mathcal{O}(k^2 p)$, $d = k + 1$. Moreover, MINSUM($\{0, \infty\}$)-LP is equivalent to the problem that decides whether P is non-empty and if so, it finds an $\mathbf{x} \in P$.

In this section, we prove Theorems 4 and 5 by constructing, from the input polytope (8), the output Potts min-sum problem. More precisely, we construct a *reparameterized* Potts min-sum problem $(V, E, K, \mathbf{g})$, whose costs will satisfy

$$\begin{aligned} g_u(k) &= 0, && \forall u \in V,\ \forall k \in K && \text{(9a)}\\ g_{uv}(k, \ell) &= 2[\![k \neq \ell]\!] + \varphi_{uv}(k) + \varphi_{vu}(\ell), && \forall \{u, v\} \in E;\ \forall k, \ell \in K && \text{(9b)}\\ \min_{k, \ell \in K} g_{uv}(k, \ell) &= 0, && \forall \{u, v\} \in E && \text{(9c)} \end{aligned}$$

Note that (9) implies $L(\mathbf{g}) = 0$. By complementary slackness, any $\boldsymbol{\mu} \in \Lambda$ and any $\mathbf{g}$ of the form (9) are simultaneously optimal to (3) and its dual if and only if

$$g_{uv}(k, \ell)\, \mu_{uv}(k, \ell) = 0, \qquad \forall \{u, v\} \in E;\ \forall k, \ell \in K. \tag{10}$$

Moreover, the output min-sum problem will be designed such that if $P \neq \emptyset$ then $\mathbf{g}$ is dual-optimal, i.e., $\min\{\, \mathbf{g}^\top \boldsymbol{\mu} \mid \boldsymbol{\mu} \in \Lambda \,\} = L(\mathbf{g}) = 0$.

We will depict min-sum problems by diagrams, as in Figure 1. In addition, we adopt the following conventions: non-zero values of $\varphi_{uv}(k)$ are written next to node (u, k) on the side of object $v \in N_u$, and an edge $\{(u, k), (v, \ell)\}$ is visible if $g_{uv}(k, \ell) = 0$ and invisible if $g_{uv}(k, \ell) > 0$. Assuming $P \neq \emptyset$, (10) thus says that $\boldsymbol{\mu} \in \Lambda$ is optimal to (3) if and only if pairwise pseudomarginals are zero on invisible edges.

4.1 Elementary Constructions

Similarly as in [10], we will construct the output min-sum problem by gluing certain smaller problems, each of them encoding a simple operation. We refer to these small problems as *elementary constructions*. Each elementary construction is a standalone min-sum problem whose costs $\mathbf{g}$ satisfy (9) and are optimal to the dual LP.

We will use the following elementary constructions (see Figure 2):

– SWAP encodes a swap of two unary pseudomarginals, one of them zero. More precisely, the LP relaxation (3) of this min-sum problem achieves its optimal value (zero) if and only if the unary pseudomarginals linked by visible edges are equal and the unary pseudomarginals in the crossed-out labels are zero.

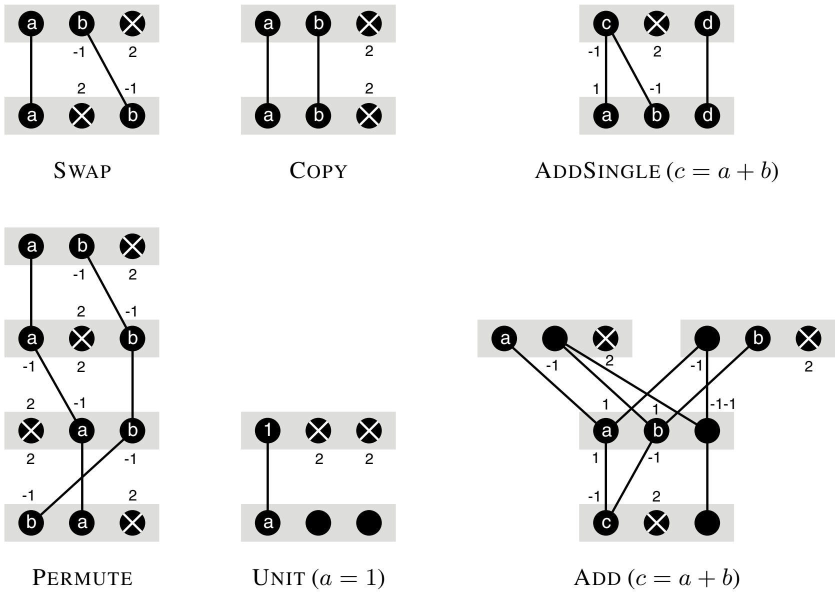

Fig. 2. Elementary constructions

- PERMUTE applies SWAP several times to arbitrarily permute all the three unary pseudomarginals, one of them zero. The figure shows one possible permutation.
- COPY copies all the three unary pseudomarginals, one of them zero, from one object to another object.
- UNIT enforces the value of a unary pseudomarginal to be 1. The other two unary pseudomarginals are necessarily zero.
- ADDSINGLE adds two unary pseudomarginals in a single object and copies the result in another object. The third (possibly nonzero) unary pseudomarginal is copied.
- ADD adds two unary pseudomarginals in two different objects. This is done by gluing three ADDSINGLE constructions.

For each elementary constructions (considered as a standalone min-sum problem), the LP relaxation is tight, i.e., the optimal values of (3) and (1) coincide.

4.2 The Encoding Algorithm

Using the elementary constructions, we now describe the algorithm to construct the output min-sum problem $(V, E, K, \mathbf{g})$ from the polytope P. First, we rewrite the system $\mathbf{Ax} = \mathbf{b}$ by moving negative terms to the right-hand side as

$$a_{i1}^{+}x_1 + \cdots + a_{in}^{+}x_n = a_{i1}^{-}x_1 + \cdots + a_{in}^{-}x_n + b_i, \qquad i = 1, \ldots, m \qquad (11)$$

where $a_{ij}^{+}, a_{ij}^{-} \in \{0, 1\}$ and $a_{ij} = a_{ij}^{+} - a_{ij}^{-}$. Note that condition (P2) says that the RHS of (11) has exactly one non-zero term. This in turn ensures that both sides of (11) are not greater than 1 for every $\mathbf{x} \in P$, thus all intermediate results are representable by pseudomarginals. We denote the labels as $K = \{1, 2, 3\}$.

The algorithm is initialized by setting $V = \{1, \ldots, n\}$ and $E = \emptyset$. Each variable x_j in (8) will be represented by unary pseudomarginal $\mu_j(1)$. Then each equation (11) is encoded one after another. The i-th equation is encoded as follows:

1. Construct a unary pseudomarginal with the value equal to the LHS of (11). This is done by repeatedly applying ADD, possibly permuting labels by PERMUTE.
2. Construct a unary pseudomarginal with value equal to the RHS of (11). Recall that the RHS of (11) has exactly one non-zero term. If $a_{ij}^- = 1$ for some j and $b_i = 0$, we already have the desired pseudomarginal, namely $\mu_j(1)$. If $a_{ij} = 0$ for all j and $b_i = 1$, we prepare a pseudomarginal with value $b_i = 1$ using UNIT.
3. Enforce equality of both sides of (11) using COPY, permuting labels when necessary by PERMUTE.

Figure 3 shows the constructed min-sum problem for an example polytope P. By construction, the output min-sum problem has the following properties:

- If $P \neq \emptyset$ then $\min\{\, \mathbf{g}^\top \boldsymbol{\mu} \mid \boldsymbol{\mu} \in \Lambda \,\} = 0$. For every $\boldsymbol{\mu}$ optimal to this problem, we have $\mathbf{x} = (\mu_1(1), \ldots, \mu_n(1)) \in P$.
- If $P \cap \{0,1\}^n \neq \emptyset$ then $\min\{\, \mathbf{g}^\top \boldsymbol{\mu} \mid \boldsymbol{\mu} \in \Lambda \cap \{0,1\}^I \,\} = 0$. For every $\boldsymbol{\mu}$ optimal to this problem, we have $\mathbf{x} = (\mu_1(1), \ldots, \mu_n(1)) \in P \cap \{0,1\}^n$.
- If $P = \emptyset$ then $\min\{\, \mathbf{g}^\top \boldsymbol{\mu} \mid \boldsymbol{\mu} \in \Lambda \,\} > 0$.

This proves Theorems 4 and 5, up to complexity.

4.3 Complexity of Encoding

Let us count the number of objects and object pairs in the output min-sum problem. Since for each elementary construction we have $|E| = \mathcal{O}(|V|)$ and the output problem is constructed by gluing elementary constructions, we have $|E| = \mathcal{O}(|V|)$. The variables $x_1, \ldots, x_n$ are represented by n objects. Each equation (11) is represented by $\mathcal{O}(d)$ objects. It follows from conditions (P3) and (P4) that $n = \mathcal{O}(dm)$. Thus, the total number of objects is $\mathcal{O}(n + dm) = \mathcal{O}(n)$. The time complexity of the algorithm is proportional to $|V|$, thus it is also $\mathcal{O}(n)$.

5 Encoding a Min-sum Problem

In this section, we show that any (integer) linear optimization over polyhedron (8),

$$\min\{\, \mathbf{c}^\top \mathbf{x} \mid \mathbf{x} \in P \cap \{0,1\}^n \,\}, \tag{12a}$$

$$\min\{\, \mathbf{c}^\top \mathbf{x} \mid \mathbf{x} \in P \,\}, \tag{12b}$$

can be efficiently reduced to the Potts min-sum problem with 3 labels. Since every local marginal polytope has the form (8), this will prove Theorems 7 and 8.

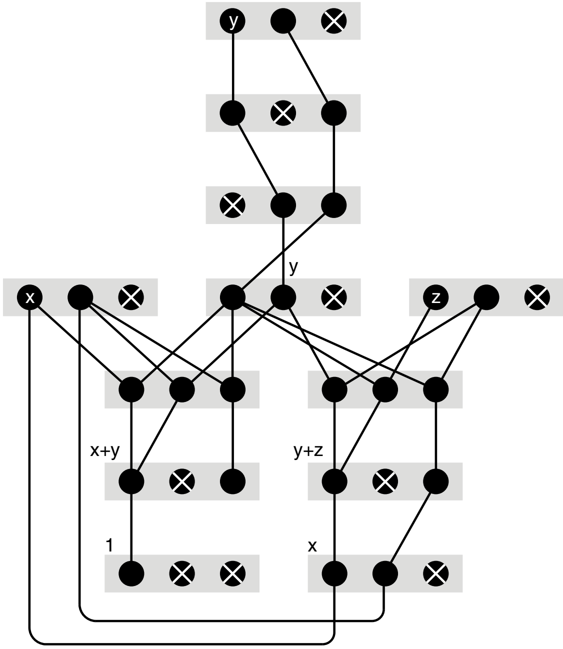

Fig. 3. The constructed reparameterized Potts min-sum problem that encodes the polytope $P = \{\, (x, y, z) \in [0, 1]^3 \mid x + y = 1,\ y + z = x \,\}$. The labels representing variables x, y, z have the variables written in them in white. The messages $\varphi_{uv}(k)$ are not shown, they are like in Figure 2.

The input of the reduction is a triplet $(\mathbf{A}, \mathbf{b}, \mathbf{c})$, where $(\mathbf{A}, \mathbf{b}) = \bar{\mathbf{A}}$ describes P. The output is a min-sum problem $(V, E, K, \mathbf{g})$, constructed as follows. First we construct min-sum problem $(V, E, K, \mathbf{g}')$ according to §4.2. Then we set $\mathbf{g} \in \mathbb{Z}^I$ as

$$g_j(k) = \begin{cases} c_j & \text{if } k = 1 \text{ and } j \le n \\ 0 & \text{otherwise} \end{cases} \tag{13a}$$

$$g_{ij}(k, \ell) = M g'_{ij}(k, \ell) \tag{13b}$$

where $M \in \mathbb{N}$ is a big enough number (derived below) to ensure that every optimal $\boldsymbol{\mu} \in [0, 1]^I$ and every integer optimal $\boldsymbol{\mu} \in \{0, 1\}^I$ to the output problem satisfies (10).

We first derive M for the simpler case, the ILP (12a). It suffices to set

$$M = C_\mathrm{u} - C_\ell + 1 \tag{14}$$

where

$$C_\ell = \sum_{j=1}^{n} \min\{0, c_j\}, \qquad C_\mathrm{u} = \sum_{j=1}^{n} \max\{0, c_j\} \tag{15}$$

is a lower and upper bound, respectively, on the optimal value of (12a).

Let us prove that every optimal solution $\boldsymbol{\mu}$ of (12a) satisfies (10). The smallest non-zero pairwise cost $g'_{uv}(k, \ell)$ is 1, thus the smallest non-zero $g_{uv}(k, \ell)$ is M. Assume that for some $\{u, v\} \in E$ and $k, \ell \in K$ we have $g_{uv}(k, \ell) > 0$ and $\mu_{uv}(k, \ell) = 1$. Then $\mathbf{g}^\top \boldsymbol{\mu} \geq M + C_\ell > C_\mathrm{u}$, which is a contradiction.

Let us derive the complexity of the reduction. We have $\langle M \rangle = \mathcal{O}(\langle \mathbf{c} \rangle)$, because we must consider the worst case when the sizes of $c_1, \ldots, c_n$ are very unequally distributed, e.g., $\langle c_1 \rangle = \mathcal{O}(\langle \mathbf{c} \rangle)$. Each unary cost $g_u(k)$ is a sum of at most $|V|$ values not greater than $2M$, hence $\langle g_u(k) \rangle = \mathcal{O}(\langle \mathbf{c} \rangle + \log |V|) = \mathcal{O}(\langle \mathbf{c} \rangle)$. Thus the description length of the output problem is[1] $\mathcal{O}(n \langle \mathbf{c} \rangle)$. This concludes the proof of Theorem 7.

We now derive M for the more difficult case, the LP (12b). We first need a lemma.

Lemma 12. *Let $(x_1, \ldots, x_n)$ be a vertex of the polytope P defined by (8). For every $j = 1, \ldots, n$ we have $x_j = 0$ or $x_j \geq (d+1)^{-m/2}$.*

Proof. The proof is analogical to that of [10, Lemma 7].

At least one optimal solution to (12b) is attained at a vertex of P. The coordinates of a vertex are fractions, however, by Lemma 12, each non-zero coordinate is not less than $(d+1)^{-m/2}$. This means it suffices to choose

$$M = (C_\mathrm{u} - C_\ell)(d+1)^{m/2} + 1. \tag{16}$$

In the worst case, $\langle M \rangle = \mathcal{O}(\langle \mathbf{c} \rangle + m \log d) = \mathcal{O}(\langle \mathbf{c} \rangle)$. This proves the claimed complexity $\mathcal{O}(n \langle \mathbf{c} \rangle)$ and thus concludes the proof of Theorem 8. Note that while the number (16) is much larger then (14), asymptotically they have the same bit size.

6 Direct Encoding of a Min-sum Problem

The reduction described in §5 involves large output costs (14) and (16), which makes it impractical and affects its theoretical complexity. Here we present a more direct reduction, which does not produce large output costs but applies only to finite input costs. By that, we prove Theorems 10 and 11.

We construct a reparameterized Potts min-sum problem $(V', E', K, \mathbf{g}')$ that encodes an input min-sum problem $(V, E, K, \mathbf{g})$. Note that both problems have the same label set. Each object $u \in V$ of the input problem is represented by one object of the output problem, so that $V \subseteq V'$. Precisely, the unary pseudomarginals of the input problem are represented by unary pseudomarginals in objects V in the output problem, which automatically enforces normalization constraints (2a). Similarly, the unary costs of the input problem are copied to unary costs in objects V of the output problem.

Each object pair $\{u, v\} \in E$ of the input problem is replaced by the following construction (see Figure 4). For each input label pair $\{(u, k), (v, \ell)\}$ we introduce a new object $\{(u, k), (v, \ell)\}$ into V'. One selected label in object $\{(u, k), (v, \ell)\} \in V'$ of the output problem represents the label pair $\{(u, k), (v, \ell)\}$ of the input problem, such that

[1] The derived complexity can be improved if some additional knowledge is available. First, we may obtain better bounds on the optimal value of (12a) than (15). E.g., if a feasible solution $\mathbf{x}$ to (12a) can be obtained cheaply, it yields an upper bound $\mathbf{c}^\top \mathbf{x} \leq C_\mathrm{u}$. Second, $\langle M \rangle = \mathcal{O}(\langle \mathbf{c} \rangle)$ holds in the unfavorable case when the distribution of the sizes $\langle c_i \rangle$ is very non-uniform. Under some additional assumptions on $\mathbf{c}$, this worst-case bound can be made much smaller. Assume, e.g., that $\langle c_i \rangle \leq 2\langle \mathbf{c} \rangle / n$ for every i. Then $M \leq n 2^{2\langle \mathbf{c} \rangle / n}$ and $\langle M \rangle = \mathcal{O}(\langle \mathbf{c} \rangle / n + \log n)$. Thus the description length of the output problem would be only $\mathcal{O}(\langle \mathbf{c} \rangle + n \log n)$.

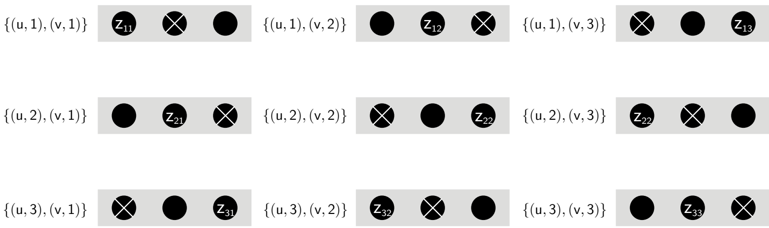

Fig. 4. Objects added to V' for one input object pair $\{u, v\} \in E$ and $|K| = 3$. For brevity, $\mu_{uv}(k, \ell)$ is denoted by $z_{k\ell}$.

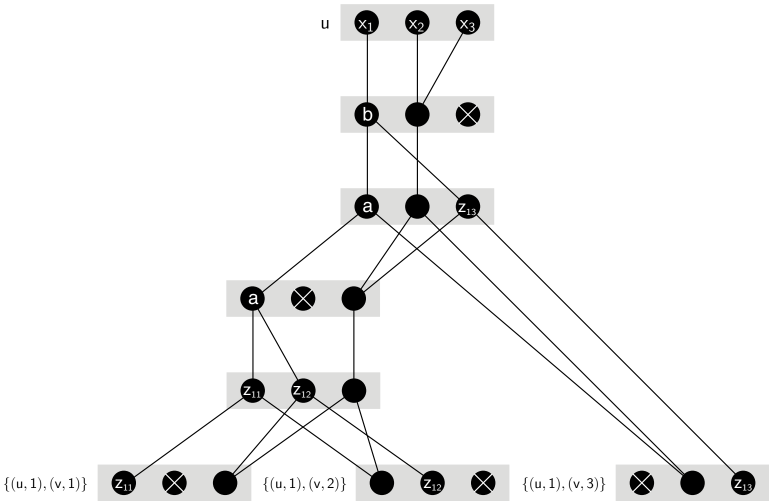

Fig. 5. The ADDK elementary construction, enforcing $z_{11} + z_{12} + z_{13} = x_1$. Note that $a = z_{11} + z_{12}$ and $b = a + z_{13}$. For brevity, $\mu_u(k)$ is denoted by x_k.

the unary pseudomarginal of this label represents the pseudomarginal $\mu_{uv}(k, \ell)$ of the input problem and the unary cost of this label equals the input cost $g_{uv}(k, \ell)$.

Each marginalization constraint (2b) is encoded by the ADDK construction, shown in Figure 5 for $|K| = 3$ labels. It is built from several constructions ADDSINGLE and ADD. For brevity, we denote $\mu_u(k)$ and $\mu_{uv}(k, \ell)$ by x_k and $z_{k\ell}$, respectively. The LP relaxation of ADDK attains zero optimal value if and only if $z_{k1} + z_{k2} + z_{k3} = x_k$, i.e., if and only if the marginalization constraint is satisfied.

Let $f(z_{k1}, z_{k2}, z_{k3}, x_k)$ denote the optimal value of the LP relaxation of ADDK subject to the constraint that the unary pseudomarginals $z_{k1}, z_{k2}, z_{k3}, x_k$ are fixed. As we said, if $z_{k1} + z_{k2} + z_{k3} = x_k$ then $f(z_{k1}, z_{k2}, z_{k3}, x_k) = 0$. Otherwise, one can show[2] that there is a small constant $C \in \mathbb{N}$ such that

$$Cf(z_{k1}, z_{k2}, z_{k3}, x_k) \geq |z_{k1} + z_{k2} + z_{k3} - x_k|. \tag{17}$$

[2] We omit the proof, which is long. For illustration, we state the similar claim for the ADDSINGLE construction (see Figure 2). Denoting by $f(a, b, c)$ the optimal value of the LP relaxation of ADDSINGLE subject to fixed a, b, c, it is easy to show that $f(a, b, c) \geq |a + b - c|$.

It is straightforward to generalize the ADDK construction to $|K| \geq 3$, e.g., by using more objects and adding $|K| - 3$ dummy labels to each object. Then we can write (17) as $R_{uv}(k) \geq |r_{uv}(k)|$ where

$$r_{uv}(k) = \sum_{\ell \in K} \mu_{uv}(k, \ell) - \mu_u(k). \tag{18}$$

Let us multiply all pairwise costs in each ADDK construction by CM_{uv}, where $M_{uv} \in \mathbb{N}$. Then the LP relaxation of the output min-sum problem can be written as

$$\min\Big\{\, \mathbf{g}^\top \boldsymbol{\mu} + \sum_{u \in V} \sum_{v \in N_u} \sum_{k \in K} M_{uv} R_{uv}(k) \;\Big|\; \boldsymbol{\mu} \in \mathbb{R}^I_+,\ \boldsymbol{\mu} \text{ satisfies (2a)} \Big\}. \tag{19}$$

The numbers M_{uv} ($u \in V$, $v \in N_u$) must be big enough to ensure that for every $\boldsymbol{\mu}$ optimal to (19) all the residuals $r_{uv}(k)$ vanish. It suffices to set

$$M_{uv} = M_{vu} = \left\lceil \tfrac{1}{2} \max_{k,\ell \in K} g_{uv}(k, \ell) \right\rceil + 1. \tag{20}$$

To prove this, observe that if unary pseudomarginals μ_u are fixed, one can optimize over pairwise pseudomarginals μ_{uv} separately for each $\{u, v\} \in E$. The rest follows from Proposition 13.

Proposition 13. *Consider a single pair $\{u, v\} \in E$. Let functions $\mu_u, \mu_v \colon K \to \mathbb{R}_+$ satisfy (2a). Let $g_{uv} \colon K \times K \to \mathbb{R}_+$. Every optimal μ_{uv} in the problem*

$$\min_{\mu_{uv} \colon K \times K \to [0,1]} \Big(\sum_{k,\ell \in K} g_{uv}(k, \ell) \mu_{uv}(k, \ell) + \sum_{k \in K} M_{uv}(|r_{uv}(k)| + |r_{vu}(k)|) \Big) \tag{21}$$

satisfies $r_{uv}(k) = r_{vu}(k) = 0$ for all $k \in K$.

Proof. Suppose that some of the numbers $r_{uv}(\cdot), r_{vu}(\cdot)$ are non-zero. We will show that then $\boldsymbol{\mu}$ cannot be optimal to (21). Since $\sum_k r_{uv}(k) = \sum_\ell r_{vu}(\ell)$, at least one of the following cases must occur. For each case, we show that by changing μ_{uv} (but keeping them feasible) the objective of (21) can be decreased:

1. $r_{uv}(k) > 0$ for some k, $r_{uv}(k') < 0$ for some k', $r_{uv}(\ell) = 0$ for all ℓ:
   Pick any ℓ such that $\mu_{uv}(k, \ell) > 0$. Because $r_{uv}(k) > 0$ and $r_{uv}(\ell) = 0$, we have $\mu_{uv}(k', \ell) < 1$. Decrease $\mu_{uv}(k, \ell)$ by a small $\delta > 0$ and increase $\mu_{uv}(k', \ell)$ by the same δ. This changes the objective by $g_{uv}(k', \ell)\delta - g_{uv}(k, \ell)\delta - 2M_{uv}\delta < 0$.
2. $r_{uv}(k) < 0$ for some k, $r_{vu}(\ell) < 0$ for some ℓ:
   Because $r_{uv}(k) < 0$, we have $\mu_{uv}(k, \ell) < 1$. Increase $\mu_{uv}(k, \ell)$ by a small $\delta > 0$. This changes the objective by $g_{uv}(k, \ell)\delta - 2M_{uv}\delta < 0$.
3. $r_{uv}(k) > 0$ for some k, $r_{vu}(\ell) > 0$ for some ℓ, $\mu_{uv}(k, \ell) > 0$:
   Decrease $\mu_{uv}(k, \ell)$ by a small $\delta > 0$. This decreases the objective by $2M_{uv}\delta + g_{uv}(k, \ell)\delta$.
4. $r_{uv}(k) > 0$ for some k, $r_{vu}(\ell) > 0$ for some ℓ, $\mu_{uv}(k, \ell) = 0$:
   Pick any k' and ℓ' such that $\mu_{uv}(k, \ell') > 0$ and $\mu_{uv}(k', \ell) > 0$. Such k' and ℓ' exist because $r_{uv}(k) > 0$ and $r_{vu}(\ell) > 0$. Then proceed as follows:

– If $\mu_{uv}(k', \ell') = 1$ then $r_{uv}(k') > 0$ and $r_{vu}(l') > 0$. Proceed as in case 3.
– If $\mu_{uv}(k', \ell') < 1$, decrease $\mu_{uv}(k, \ell')$ by a small $\delta > 0$, decrease $\mu_{uv}(k', \ell)$ by δ, and increase $\mu_{uv}(k', \ell')$ by δ. This changes the objective by $-g_{uv}(k, \ell')\delta - g_{uv}(k', \ell)\delta + g_{uv}(k', \ell')\delta - 2M_{uv}\delta < 0$. □

6.1 Complexity of the Reduction

Let us derive the complexity of the reduction. Clearly, $|V'| = \mathcal{O}(|V| + |K|^2|E|)$ and $|E'| = \mathcal{O}(|K|^2|E|)$. The cumulative size of all numbers M_{uv} ($\{u, v\} \in E$) is $\mathcal{O}(\langle \mathbf{g} \rangle)$. Each value M_{uv} appears as the Potts pairwise cost in $\mathcal{O}(|K|^2)$ object pairs, thus all the Potts pairwise costs are described by a vector of size $\mathcal{O}(|K|^2\langle \mathbf{g} \rangle)$. The cumulative size of the unary costs in $\mathbf{g}'$ is bounded by the sum of sizes of all messages. Every M_{uv} induces $\mathcal{O}(|K|^2)$ messages, each of them having the absolute value at most $2M_{uv}$. It means all the messages are described by a vector of size $\mathcal{O}(|K|^2\langle \mathbf{g} \rangle)$, which proves the output has the size $\mathcal{O}(|K|^2\langle \mathbf{g} \rangle)$. Note that the numbers (20) are (possibly much) smaller than (14) and (16). If $|K|$ is fixed, the complexity of the reduction is linear.

7 Conclusion

Our results (Corollaries 6 and 9, Theorem 11) suggest that solving the LP relaxation of the pairwise min-sum problem with attractive Potts costs cannot be expected much easier than solving the LP relaxation of the general min-sum problem.

This statement may sound misleading in case of reduction with higher than linear complexity, because in that case efficiency of solving the LP relaxation of the Potts min-sum problem does not fully translate to efficiency of solving the general LP. However, our argument is more subtle: if a new principle were invented to solve the LP relaxation of Potts min-sum problems (e.g., similar to network flow algorithms), it would mean this principle is applicable to an arbitrary LP. Since there are only few principles to solve general LPs in polynomial time, this is unlikely.

In particular, message passing algorithms do not solve the LP relaxation of a general min-sum problem exactly, but find only a local (with respect to block-coordinate updates) dual optimum. It would be desirable to modify these algorithms to alleviate this drawback. One might hope this might be easier for Potts min-sum than for general min-sum. However, inventing a message passing algorithm that avoids local optima for Potts min-sum problems would mean it can solve general LPs.

Besides the results for the LP relaxation, we obtained similar reductions for the non-relaxed problems (Theorems 4, 7, 10). These may have practical impact in the case of exact (e.g., branch-and-bound) solvers, which can be tuned only for Potts problems. Unfortunately, they may not be useful for approximate solvers (such as primal move-making algorithms [2]) or solvers obtaining persistency [8,14], because the reductions may not preserve approximation ratio or persistency.

Acknowledgment. This work has been supported by the Czech Science Foundation project P202/12/2071 (both authors) and the EC project FP7-ICT-270138 DARWIN (the second author only).

References

1. Boros, E., Hammer, P.L.: Pseudo-Boolean optimization. Discrete Applied Mathematics 123(1-3), 155–225 (2002)
2. Boykov, Y., Veksler, O., Zabih, R.: Fast approximate energy minimization via graph cuts. IEEE Trans. Pattern Analysis and Machine Intelligence 23(11), 1222–1239 (2001)
3. Chekuri, C., Khanna, S., Naor, J., Zosin, L.: A linear programming formulation and approximation algorithms for the metric labeling problem. SIAM Journal on Discrete Mathematics 18(3), 608–625 (2005)
4. Chuzhoy, J., Naor, J.: The hardness of metric labeling. SIAM J. Computation 36(5), 1376–1386 (2007)
5. Kappes, J.H., Andres, B., Hamprecht, F.A., Schnörr, C., Nowozin, S., Batra, D., Kim, S., Kausler, B.X., Lellmann, J., Komodakis, N., Rother, C.: A comparative study of modern inference techniques for discrete energy minimization problem. In: Conf. on Computer Vision and Pattern Recognition (2013)
6. Kleinberg, J., Tardos, E.: Approximation algorithms for classification problems with pairwise relationships: Metric labeling and markov random fields. J. ACM 49(5), 616–639 (2002)
7. Koster, A., van Hoesel, S.P.M., Kolen, A.W.J.: The partial constraint satisfaction problem: Facets and lifting theorems. Operations Research Letters 23(3-5), 89–97 (1998)
8. Kovtun, I.: Partial optimal labelling search for a NP-hard subclass of (max,+) problems. In: Conf. German Assoc. for Pattern Recognition, pp. 402–409 (2003)
9. Osokin, A., Vetrov, D., Kolmogorov, V.: Submodular decomposition framework for inference in associative Markov networks with global constraints. In: IEEE Conf. on Computer Vision and Pattern Recognition, pp. 1889–1896 (2011)
10. Průša, D., Werner, T.: Universality of the local marginal polytope. In: Conf. on Computer Vision and Pattern Recognition, pp. 1738–1743. IEEE Computer Society (2013)
11. Rother, C., Kolmogorov, V., Lempitsky, V.S., Szummer, M.: Optimizing binary MRFs via extended roof duality. In: Conf. on Computer Vision and Pattern Recognition (2007)
12. Schlesinger, D., Flach, B.: Transforming an arbitrary MinSum problem into a binary one. Tech. Rep. TUD-FI06-01, Dresden University of Technology, Germany (April 2006)
13. Shlezinger, M.I.: Syntactic analysis of two-dimensional visual signals in noisy conditions. Cybernetics and Systems Analysis 12(4), 612–628 (1976)
14. Swoboda, P., Savchynskyy, B., Kappes, J.H., Schnörr, C.: Partial optimality via iterative pruning for the Potts model. In: Conf. on Scale Space and Variational Methods in Computer Vision (2013)
15. 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 Trans. on Pattern Analysis and Machine Intelligence 30(6), 1068–1080 (2008)
16. Živný, S.: The Complexity of Valued Constraint Satisfaction Problems. Cognitive Technologies. Springer (2012)
17. Wainwright, M.J., Jordan, M.I.: Graphical models, exponential families, and variational inference. Foundations and Trends in Machine Learning 1(1-2), 1–305 (2008)
18. Werner, T.: A linear programming approach to max-sum problem: A review. IEEE Trans. Pattern Analysis and Machine Intelligence 29(7), 1165–1179 (2007)