RIMS-1812

Minimizing Submodular Functions on Diamonds via Generalized Fractional Matroid Matchings

By

Satoru FUJISHIGE, Tamás KIRÁLY, Kazuhisa MAKINO, Kenjiro TAKAZAWA and Shin-ichi TANIGAWA

January 2015

京都大学　数理解析研究所

RESEARCH INSTITUTE FOR MATHEMATICAL SCIENCES

KYOTO UNIVERSITY, Kyoto, Japan

Minimizing Submodular Functions on Diamonds via Generalized Fractional Matroid Matchings*

Satoru Fujishige† Tamás Király‡ Kazuhisa Makino†
Kenjiro Takazawa† Shin-ichi Tanigawa†

January 5, 2015

Abstract

In this paper we show the first polynomial-time algorithm for the problem of minimizing submodular functions on the product of diamonds. This submodular function minimization problem is reduced to the membership problem for an associated polyhedron, which is equivalent to the optimization problem over the polyhedron, based on the ellipsoid method. The latter optimization problem is solved by polynomial number of solutions of subproblems, each being a generalization of the weighted fractional matroid matching problem. We give a combinatorial polynomial-time algorithm for this optimization problem by extending the result by Gijswijt and Pap [D. Gijswijt and G. Pap, An algorithm for weighted fractional matroid matching, J. Combin. Theory, Ser. B 103 (2013), 509–520].

1 Introduction

A set function $f : 2^V \to \mathbb{Z}$ is **submodular** if $f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y)$ for every $X, Y \subseteq V$. In the submodular function minimization problem, given an evaluation oracle for a submodular function f, we are asked to find a minimizer of f. For this problem, our goal is to find an algorithm with running time polynomial in $|V|$ and $\log \max_{X \subseteq V}\{|f(X)|\}$ that returns $X \in \mathrm{argmin}(f)$, assuming that the algorithm has access to an oracle that for any given X outputs $f(X)$.

It follows from the work of Grötschel, Lovász and Schrijver [8] on the equivalence of separation and optimization that such an algorithm can be obtained by using the ellipsoid method. Combinatorial strongly polynomial algorithms have only been obtained much later, independently by Schrijver [25] and by Iwata, Fleischer and

*Most of the research was done while the second author was a visiting researcher at Research Institute for Mathematical Sciences, Kyoto University.

†Research Institute for Mathematical Sciences, Kyoto University, Sakyo-ku, Kyoto 606-8502, Japan. {fujishig,makino,takazawa,tanigawa}@kurims.kyoto-u.ac.jp

‡Department of Operations Research, Eötvös University, Pázmány Péter sétány 1/C, 1117 Budapest, Hungary, and MTA-ELTE Egerváry Research Group on Combinatorial Optimization, Budapest, Hungary. tkiraly@cs.elte.hu

Fujishige [11]. Since then, there have been several improvements in running time, e.g. [23, 13].

The generalization that we consider in this paper concerns submodular functions on lattices. Given a finite lattice L, a function $f : L \to \mathbb{Z}$ is **submodular on** L if $f(x) + f(y) \geq f(x \vee y) + f(x \wedge y)$ for every $x, y \in L$. For modular lattices, such functions naturally arise when extending the Dulmage-Mendelsohn decompositions of generic matrices to generic partitioned matrices [14], and it was posed as an open problem in [11] to give an efficient algorithm for minimizing submodular functions on modular lattices.

Submodular functions on product lattices also got a lot of attention for the complexity classification of Max-CSP. The importance of submodular functions in this context was first pointed out by Cohen et al. [3], and then the connection was further investigated in [16, 17]. The systematic study of the complexity of Max-CSP was further extended to finite-valued CSP in [4], and a dichotomy theorem was finally obtained in [26, 27]. A result by Thapper and Živný [26] in turn implies the polynomial-time solvability of a special case of the submodular function minimization on the direct product of finite lattices, where the function is explicitly given as the sum of submodular functions of constant arity, i.e., the value of each function depends only on a constant number of lattices. Hirai [9] further introduced submodular functions on modular semi-lattices and discussed the solvability of the minimization problem in the constant arity case based on the result of [26]. However, as noted in most of the above literature, it is widely open whether the submodular function minimization problem on product lattices is tractable in the value oracle model.

As observed in [11, 25], one can reduce the problem to the standard submodular function minimization if the underlying lattice is distributive. Krokhin and Larose [18] showed that certain lattice operations preserve the tractability of the corresponding minimization problem in the value oracle model, and as a corollary they showed that the submodular function minimization on the product of the copies of the pentagon, a smallest non-distributive lattice, can be reduced to the standard submodular function minimization.

In this paper we shall consider the submodular function minimization problem on the product of diamonds, which is the remaining smallest non-distributive case and has an application to the Dulmage-Mendelsohn type decompositions of generic partitioned matrices consisting of two-by-two blocks [12]. A **diamond** is a lattice consisting of a minimal element, a maximal element, and an arbitrary finite number of pairwise incomparable **middle elements**: the meet (resp. join) of any two middle elements is the minimal (resp. maximal) element. A submodular function on the direct product of given diamonds $U_1, \ldots, U_n$ is simply called a **submodular function on diamonds**. If the diamonds have at most two middle elements, then the lattice is distributive, and by the observation in [11, 25] we can use the standard submodular function minimization algorithm in this case. However, a diamond with more than two middle elements is modular but not distributive, and hence we cannot directly apply the standard algorithms to the minimization of submodular functions on diamonds. A pseudo-polynomial algorithm for the minimization of submodular functions on diamonds was given by Kuivinen [19]. Our main result is the first polynomial-time algorithm.

Theorem 1. *Let f be a submodular function on the direct product of a finite number of diamonds $U_1, \ldots, U_n$. A minimizer of f can be computed in a polynomial number of arithmetic steps and function evaluations in m and $\log M$, where $m = \sum_{i=1}^n |U_i|$ and M is the maximum absolute function value.*

Let $U = \bigcup_{i=1}^n U_i$, and call $T \subseteq U$ a **transversal** if $|T \cap U_i| = 1$ for every $i \in [n]$, where $[n]$ denotes the set of integers $\{1, \ldots, n\}$. We denote by $\mathcal{T}$ the set of transversals and by $T_{\mathbf{0}}$ the transversal consisting of the minimal elements. There is a natural one-to-one correspondence between transversals and elements of the direct product lattice, which also defines operations $\wedge$ and $\vee$ on pairs of transversals. Thus a submodular function on diamonds can be considered as a function $f : \mathcal{T} \to \mathbb{Z}$ satisfying $f(T_1) + f(T_2) \geq f(T_1 \wedge T_2) + f(T_1 \vee T_2)$ for every $T_1, T_2 \in \mathcal{T}$. Throughout the paper we assume $f(T_{\mathbf{0}}) = 0$.

For a transversal $T \in \mathcal{T}$, let $a(T) \in \{0, 1, 2\}^n$ be a vector whose i-th element $a(T)_i$ is the rank of the unique element in $T \cap U_i$ in the lattice U_i. We consider the optimization problem:

$$\begin{array}{ll} \text{maximize} & cx \\ \text{subject to} & a(T)x \leq f(T) \text{ for each } T \in \mathcal{T}. \end{array} \tag{1}$$

If this can be solved in polynomial time, then by the results of Grötschel, Lovász and Schrijver [8] the minimization of submodular functions on diamonds can be solved in polynomial time using the ellipsoid method. Indeed, the problem of deciding whether f is nonnegative is a special case of the separation problem corresponding to (1), and submodularity is preserved by adding the same nonnegative integer to each function value except $f(T_{\mathbf{0}})$, and hence the minimization problem can be solved by applying binary search to the problem (1).

Kuivinen's pseudo-polynomial time algorithm [19] also follows the same strategy, where he considered a distinct and larger polytope and showed that the corresponding linear programming can be solved in pseudo-polynomial time, again, by the aid of the ellipsoid method. On the other hand in this paper we give a combinatorial algorithm for solving (1).

Theorem 2. *Let f be a submodular function on the direct product of a finite number of diamonds. Then there is a combinatorial algorithm for solving (1) that runs in a polynomial number of arithmetic steps and function evaluations in m and $\log M$.*

When f is derived from a matroid rank function, the polytope describing (1) coincides with the **fractional matroid matching polytope** introduced by Vande Vate [28], and the corresponding optimization problem (1) is known as the **weighted fractional matroid matching problem**, which was solved by Gijswijt and Pap [7]. The main restriction compared to our generalized problem is that the lattice function corresponding to fractional matroid matching is derived from a matroid rank function, and hence it is monotone nondecreasing and has maximum value at most $2n$. Also Gijswijt and Pap [7] used the unweighted algorithm of Chang, Llewellyn, and Vande Vate [1] as a subroutine, whereas we shall develop the corresponding theory for general submodular function on diamonds from scratch. Nevertheless, our algorithm makes use of several ideas from the Gijswijt-Pap paper.

A different extension of standard submodular minimization is the minimization of bisubmodular functions by Qi [24], Fujishige and Iwata [5], and Fujishige and McCormick [21]. Min-max theorems (without polynomial algorithms) were also given for the minimization of k-submodular functions, which is a common generalization of bisubmodular functions and multimatroid rank functions, by Huber and Kolmogorov [10], and for the more general class of transversal submodular functions by Fujishige and Tanigawa [6]. One of the exciting open problems is whether k-submodular functions can be minimized in polynomial time.

The rest of the paper is organized as follows. In Section 2 we describe the problem setting in detail. Section 3 introduces the minimum 2-cover problem that corresponds to the dual improvement of the optimization problem (1). Although a minimum 2-cover can be found in polynomial time using the ellipsoid method, we also present a combinatorial polynomial-time algorithm in Section 5, at the end of the paper, which leads to a combinatorial polynomial-time algorithm for the optimization problem (1). The algorithm for the optimization problem (1) is presented in Section 4, first in a pseudo-polynomial version, which is then transformed into a polynomial algorithm by a scaling technique.

2 Problem Setting

As described in Section 1, we consider submodular functions on the direct product of n diamonds $U_1, \ldots, U_n$. Let $V = [n]$. For $i \in V$, the minimal and maximal elements of U_i are denoted by 0_i and 1_i, respectively. Recall that $U = \bigcup_{i \in V} U_i$. A set $T \subseteq U$ is called a **sub-transversal** if $|T \cap U_i| \leq 1$ for every $i \in V$. Each sub-transversal can be identified with a transversal by extending it with 0_i for every U_i disjoint from the sub-transversal. Thus $\wedge$ and $\vee$ can be defined over the set of sub-transversals.

Recall that $\mathcal{T}$ denotes the set of all transversals. The partial order in the diamond induces a partial order on $\mathcal{T}$, denoted by $\preceq$. For two transversals T and T', we write $T \prec T'$ if $T \preceq T'$ and $T \neq T'$. Recall that $T_{\mathbf{0}}$ denotes the transversal formed by all minimal elements. The transversal consisting of all maximal elements is denoted by T_{top}. Given transversals T_1 and T_2 satisfying $T_1 \preceq T_2$, we use the notation $[T_1, T_2] = \{T \in \mathcal{T} : T_1 \preceq T \preceq T_2\}$. For a transversal $T \in \mathcal{T}$ and $i \in V$, recall that $a(T)_i \in \{0, 1, 2\}$ denotes the rank of $T \cap U_i$, i.e., $a(T)_i = 0$ if $T \cap U = \{0_i\}$, $a(T)_i = 2$ if $T \cap U = \{1_i\}$, and $a(T)_i = 1$ otherwise.

Let $f : \mathcal{T} \to \mathbb{Z}$ be a submodular function on the diamonds, that is, f satisfies $f(T_{\mathbf{0}}) = 0$ and $f(T_1) + f(T_2) \geq f(T_1 \vee T_2) + f(T_1 \wedge T_2)$ for every $T_1, T_2 \in \mathcal{T}$. We consider the following polyhedra defined by f:

$$P(f) = \{x \in \mathbb{R}^n : a(T)x \leq f(T) \ \forall T \in \mathcal{T}\},$$
$$P^{=}(f) = \{x \in \mathbb{R}^n : a(T)x \leq f(T) \ \forall T \in \mathcal{T},\ 2x(V) = f(T_{\text{top}})\},$$

where $x(V) = \sum_{i \in V} x_i$. In general, for $x \in \mathbb{R}^n$ and $X \subseteq V$, let $x(X) = \sum_{i \in X} x_i$.

Let $m = |U|$, $M = \max_{T \in \mathcal{T}} |f(T)|$, and $N = \max\{m, \lceil \log M \rceil\}$. Our goal is to give a combinatorial algorithm, with running time polynomial in N, that solves the following linear program for $c \in \mathbb{Z}^n$:

$$\begin{array}{lll} (\mathrm{LP}^{\leq}) & \text{maximize} & cx \\ & \text{subject to} & x \in P(f). \end{array}$$

To this end we focus on the following linear program and its dual:

$$\begin{array}{lll} (\mathrm{LP}^{=}) & \text{maximize} & cx \\ & \text{subject to} & x \in P^{=}(f), \end{array}$$

$$\begin{array}{llll} (\mathrm{D}) & \text{minimize} & \sum_{T \in \mathcal{T}} f(T) y_T & \\ & \text{subject to} & \sum_{T \in \mathcal{T}} a(T)_i y_T = c_i & \text{for each } i \in V, \\ & & y_T \geq 0 & \text{for each } T \in \mathcal{T} \setminus \{T_{\mathrm{top}}\}. \end{array}$$

For a dual feasible solution y, the **support** of y is defined as $\{T \in \mathcal{T} \mid y_T > 0\} \cup \{T_{\mathrm{top}}\}$. The following proposition implies that the linear systems describing $(\mathrm{LP}^{=})$ and $(\mathrm{LP}^{\leq})$ are half-TDI, and thus the basic solutions for $(\mathrm{LP}^{=})$ and $(\mathrm{LP}^{\leq})$ are half integral.

Proposition 3. *For a nonnegative integer vector $c \in \mathbb{Z}^n$, there is a half-integral dual optimal solution of* $(\mathrm{LP}^{=})$ *(and, respectively, of* $(\mathrm{LP}^{\leq})$*) whose support is a chain $T_1 \prec T_2 \prec \cdots \prec T_k = T_{\mathrm{top}}$.*

Proof. This is implicit in [7, Theorem 1]. □

Observe that once we have a polynomial-time algorithm for solving $(\mathrm{LP}^{=})$, then we can solve $(\mathrm{LP}^{\leq})$, e.g., by binary search. This can be seen as follows. For a real number t with $t \leq f(T_{\mathrm{top}})$, let f^t be the submodular function obtained from f such that $f^t(T) = t$ if $T = T_{\mathrm{top}}$, and $f^t(T) = f(T)$ otherwise. Define $g(t)$ by $g(t) = \max\{cx \mid x \in P^{=}(f^t)\}$ for $t \leq f(T_{\mathrm{top}})$. Then $g(t)$ is concave, and $\max_{t \leq f(T_{\mathrm{top}})} g(t)$ is attained by some integer since Proposition 3 implies that $2x(V)$ is integer. Thus one can compute $\max_{t \leq f(T_{\mathrm{top}})} g(t)$ in polynomial time, due to the following lower bound of $x(V)$ for an optimal solution x for $(\mathrm{LP}^{\leq})$.

Proposition 4. *Suppose that $c \in \mathbb{Z}^n$ is nonnegative. Then there is an optimal solution x for* $(\mathrm{LP}^{\leq})$ *satisfying $x(V) \geq -2M$.*

Proof. Let us take an optimal solution x for $(\mathrm{LP}^{\leq})$ with maximum $x(V)$. Note that for any transversals T and T' with $a(T)x = f(T)$ and $a(T')x = f(T')$ we have $a(T \vee T')x = f(T \vee T')$. Hence there is a transversal T such that $a(T)x = f(T)$ and $a(T)_i > 0$ for each $i \in V$, since otherwise x_i can be increased without decreasing the objective value. Let $V^- \subseteq V$ consist of the indices i for which $x_i < 0$, and let T^+ be the transversal obtained from T by replacing the element of $T \cap U_i$ by 0_i for every $i \in V^-$. Then $x(V) \geq x(V^-) \geq (a(T) - a(T^+))x = f(T) - a(T^+)x \geq f(T) - f(T^+) \geq -2M$. □

The following proposition implies that we may focus on the case when c has n distinct values.

Proposition 5. *For $c \in \mathbb{Z}^n$, define $c' \in \mathbb{Z}^n$ by $c'_i = 2n^2Mc_i + i$ for $i \in V$. If $x \in \mathbb{R}^n$ is an optimal solution of* (LP$^=$) *with the objective function c', then x is also optimal for* (LP$^=$) *with the objective function c.*

Proof. For any $x \in P^=(f)$ and $i \in V$, we have $x_i \le f(\{1_i\})/2 \le M/2$ and $x_i = x(V) - x(V - i) \ge -M$.

Suppose that x is not optimal for (LP$^=$) with the objective function c. Then for any optimal solution x' for (LP$^=$) with the objective function c we have $c'(x - x') = 2n^2M(cx - cx') + \sum_{i\in V} i(x_i - x'_i) \le -n^2M + \sum_{i\in V} \frac{3}{2}iM < 0$, which implies that x is not optimal for (LP$^=$) with the objective function c', a contradiction. □

Thus in the remainder of the paper we assume that c consists of n distinct values. By this assumption, we have the following property for dual feasible solutions.

Proposition 6. *Suppose that $c_i \ne c_{i'}$ for all distinct i and i' in V, and let y be a feasible solution of* (D) *whose support is a chain $T_{\mathbf{0}} = T_0 \prec T_1 \prec \cdots \prec T_k = T_{\mathrm{top}}$. For every $j \in [k]$, there is at most one index $i \in V$ such that $a(T_{j-1})_i = 0$ and $a(T_j)_i = 2$.*

Proof. Suppose that there are distinct such i and i' in V. Then we have $c_i = \sum_T a(T)_i y_T = \sum_T a(T)_{i'} y_T = c_{i'}$, a contradiction. □

3 The Minimum 2-cover Problem

In this section we shall show that finding a dual improvement direction reduces to a combinatorial problem, called the minimum 2-cover problem. We then show a min-max theorem for the minimum 2-cover problem and describe its relation to the fractional matroid matching problem. The canonical optimal solution given in the proof will be explicitly used in our algorithm for (LP$^=$) in Section 4.

3.1 A 2-cover and dual improvement

In this subsection, we introduce the minimum 2-cover problem, which is described by a 2-regular hypergraph on V and "almost submodular" set functions arising from a chain of transversals. We then show its relation to the dual improvement for (LP$^=$).

Let $T_1 \prec T_2 \prec \cdots \prec T_k (= T_{\mathrm{top}})$ be a chain of transversals, denoted by $\mathcal{C}$, and let $T_0 = T_{\mathbf{0}}$. Define a family $\mathcal{E}(\mathcal{C}) = \{Z_1, \ldots, Z_k\}$ of multisets of elements in V such that Z_j contains $i \in V$ with multiplicity 1 if $a(T_j)_i = a(T_{j-1})_i + 1$, and with multiplicity 2 if $a(T_j)_i = a(T_{j-1})_i + 2$. As we only consider multisets where each element has multiplicity 0, 1, or 2, a multiset Z can be identified with its characteristic function $\chi_Z : V \to \{0, 1, 2\}$. Observe that $(V, \mathcal{E}(\mathcal{C}))$ is a 2-regular hypergraph, i.e., $\sum_{Z\in\mathcal{E}(\mathcal{C})} \chi_Z(i) = 2$ for every $i \in V$. For simplicity we write $Y \subseteq Z$ if $\chi_Y \le \chi_Z$, and $x(Z)$ denotes $\chi_Z x$ for any $x \in \mathbb{R}^n$. For $u, v \in V$, a multiset Y is called a **$u\bar{v}$-set** if Y contains u and avoids v.

By Proposition 6, each Z_j has at most one element with multiplicity two; denote this by $v^\circ_{Z_j}$ if it exists. For any multiset $Y \subseteq Z_j$, let $m_{Z_j}(Y) = \chi_Y(v^\circ_{Z_j})$ if $v^\circ_{Z_j}$

exists, and otherwise $m_{Z_j}(Y) = 0$. We simply write $m(Y)$ if Z_j is clear from the context. For each $Y \subseteq Z_j$, there is a unique transversal T_Y between T_{j-1} and T_j that corresponds to Y, i.e.,

$$T_Y = T_{j-1} \vee \bigcup_{v \in Y} (T_j \cap U_v),$$

unless $m(Z_j) = 2$ and $m(Y) = 1$. If $m(Z_j) = 2$ and $m(Y) = 1$, then several transversals of the form

$$T_{j-1} \vee \left(\{u\} \cup \bigcup_{v \in Y, v \neq v^\circ_{Z_j}} (T_j \cap U_v) \right)$$

for some middle element u in the diamond corresponding to $v^\circ_{Z_j}$ may correspond to Y; let T_Y be the one for which $f(T_Y)$ is the smallest. The middle element u of T_Y in the diamond corresponding to $v^\circ_{Z_j}$ is called a **shade** of Y. If more than one $f(T_Y)$ value is minimum, then Y has more than one shade.

For each $Z_j \in \mathcal{E}(\mathcal{C})$, define f_{Z_j} by

$$f_{Z_j}(Y) = f(T_Y) - f(T_{j-1}) \qquad (Y \subseteq Z_j).$$

Observe that if $Z \in \mathcal{E}(\mathcal{C})$ has no element with multiplicity two, then f_Z is a standard submodular set function on Z. In order to describe "submodularity" of f_Z, for multisets $X, Y \subseteq Z$, define $X \vee Y$ and $X \wedge Y$ as the multisets corresponding to $\chi_{X \vee Y}$ and $\chi_{X \wedge Y}$, where for each $i \in V$

$$\chi_{X \vee Y}(i) := \begin{cases} 2 & \text{if } i = v^\circ_Z,\ m(X) = m(Y) = 1, \text{ and} \\ & \quad \text{the sets of the shades of } X, Y \text{ are not identical,} \\ \max\{\chi_X(i), \chi_Y(i)\} & \text{otherwise,} \end{cases}$$

$$\chi_{X \wedge Y}(i) := \begin{cases} 0 & \text{if } i = v^\circ_Z,\ m(X) = m(Y) = 1, \text{ and} \\ & \quad \text{the sets of the shades of } X, Y \text{ are not identical,} \\ \min\{\chi_X(i), \chi_Y(i)\} & \text{otherwise.} \end{cases}$$

The following proposition establishing submodularity of f_Z is now straightforward.

Proposition 7. *For any $X, Y \subseteq Z$, $f_Z(X) + f_Z(Y) \geq f_Z(X \vee Y) + f_Z(X \wedge Y)$.*

We are now ready to define the minimum 2-cover problem. For a chain of transversals $\mathcal{C}$, a **2-cover** of $(V, \mathcal{E}(\mathcal{C}))$ is a family of multiset pairs $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ satisfying $A_Z \subseteq B_Z \subseteq Z$ and $\sum_{Z \in \mathcal{E}(\mathcal{C})} (\chi_{A_Z}(i) + \chi_{B_Z}(i)) = 2$ for all $i \in V$. An example of a 2-cover is $\{(\emptyset, Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$, which is called a **trivial 2-cover**. In **the minimum 2-cover problem**, given f and $\mathcal{C}$, we are asked to find a 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ that minimizes $\sum_{Z \in \mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z))$. Note that the objective value of the trivial 2-cover is $f(T_{\text{top}})$.

The following lemma gives an explicit link between the dual improvement for (LP$^=$) and the minimum 2-cover problem.

Lemma 8. *Let y be a feasible solution for* (D) *whose support is a chain $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\text{top}}$. The trivial 2-cover is optimal for the minimum 2-cover problem for $(f, \mathcal{C})$ if and only if y is optimal for* (D)*. If the objective value of 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ is smaller than that of the trivial one, then the following y^ε is a feasible solution of* (D) *whose objective value is better than that of y:*

$$y^\varepsilon := y + \varepsilon \sum_{1 \le j \le k} (\chi_{T_{A_{Z_j}}} + \chi_{T_{B_{Z_j}}} - \chi_{T_{j-1}} - \chi_{T_j}), \tag{2}$$

where $\varepsilon > 0$ is chosen so that $y^\varepsilon_T \ge 0$ holds for every $T \in \mathcal{T} \setminus \{T_{\text{top}}\}$.

Proof. Comparing the objective value of $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ with that of the trivial 2-cover, we have that

$$\begin{aligned}
&\sum_{Z \in \mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z)) - \sum_{Z \in \mathcal{E}(\mathcal{C})} (f_Z(\emptyset) + f_Z(Z)) \\
&= \sum_{j=1}^{k} (f(T_{A_{Z_j}}) - f(T_{j-1}) + f(T_{B_{Z_j}}) - f(T_{j-1})) - \sum_{j=1}^{k} (f(T_j) - f(T_{j-1})) \\
&= \sum_{j=1}^{k} (f(T_{A_{Z_j}}) + f(T_{B_{Z_j}}) - f(T_{j-1}) - f(T_j)).
\end{aligned}$$

Hence, if the objective value of $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ is smaller, then $\sum_T f(T) y^\varepsilon_T < \sum_T f(T) y_T$. To show that y^ε is a feasible solution of (D), it suffices to show that

$$\sum_{1 \le j \le k} (a(T_{A_{Z_j}}) - a(T_{j-1}))_v = \sum_{1 \le j \le k} (a(T_j) - a(T_{B_{Z_j}}))_v \quad \text{for each } v \in V.$$

This follows from the property that $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ is a 2-cover: if v is in some A_Z, then both sides are 1, and otherwise both sides are 0. Thus we can conclude that if y is optimal for (D), then the trivial 2-cover is optimal.

For the other direction, suppose that y is not optimal for (D). We show that the trivial 2-cover is not optimal. Since y is not optimal, the following system of inequalities has a solution z:

$$\sum_{T \in \mathcal{T}} f(T) z_T < 0, \tag{3}$$

$$\sum_{T \in \mathcal{T}} a(T)_v z_T = 0 \quad \text{for each } v \in V, \tag{4}$$

$$z_T \ge 0 \qquad \text{for each } T \in \mathcal{T} \setminus \mathcal{C}. \tag{5}$$

Let $\text{supp}(z) = \{T \in \mathcal{T} \setminus \mathcal{C} \mid z_T > 0\}$. We claim that there is a solution z such that $\text{supp}(z)$ forms a chain compatible with $\mathcal{C}$. This can be seen by the following standard uncrossing argument.

- First we modify z so that each transversal in $\text{supp}(z)$ is compatible with $\mathcal{C}$. This can be done inductively from T_1 through T_k as follows. Suppose that each transversal in $\mathrm{supp}(z)$ is compatible with $T_1, \dots, T_{i-1}$, and a transversal $T \in \mathrm{supp}(z)$ crosses with T_i. Then set

$$\begin{aligned} z_T &:= 0, & z_{T_i} &:= z_{T_i} - z_T, \\ z_{T \vee T_i} &:= z_{T \vee T_i} + z_T, & z_{T \wedge T_i} &:= z_{T \wedge T_i} + z_T. \end{aligned}$$

Note that the resulting z satisfies (3), (4) and (5). Since $T \vee T_i$ and $T \wedge T_i$ are compatible with $T_1, \dots, T_i$, after a finite number of steps, we get a desired z.

- Next we modify z so that $\mathrm{supp}(z)$ forms a chain. If two transversals $T, T' \in \mathrm{supp}(z)$ are crossing, then set

$$\begin{aligned} z_T &:= z_T - \delta, & z_{T'} &:= z_{T'} - \delta, \\ z_{T \vee T'} &:= z_{T \vee T'} + \delta, & z_{T \wedge T'} &:= z_{T \wedge T'} + \delta, \end{aligned}$$

where $\delta = \min\{z_T, z_{T'}\}$. Note that the resulting z satisfies (3), (4) and (5), and $T \vee T'$ and $T \wedge T'$ are compatible with $\mathcal{C}$. Consider the potential $\sum_{T \in \mathcal{T} \setminus \mathcal{C}} g(T) z_T$, where $g(T) := (\sum_{i \in V} a(T)_i)(\sum_{i \in V}(2 - a(T)_i))$. This is nonnegative, and the strict submodularity of g implies that $\mathrm{supp}(z)$ becomes a chain after modifying z finitely many times.

Let z be a solution for (3), (4) and (5) such that $\mathrm{supp}(z)$ is a chain $\mathcal{C}'$ compatible with $\mathcal{C}$. Since $\mathrm{supp}(z)$ is a chain $\mathcal{C}'$ compatible with $\mathcal{C}$, we may further assume that for each $T \in \mathrm{supp}(z)$ there is $Y \subseteq Z \in \mathcal{E}(\mathcal{C})$ such that $T_Y = T$.

Denote $\mathcal{C}^* = \mathcal{C} \cup \mathcal{C}' : T_1^* \prec \cdots \prec T_l^* (= T_{\mathrm{top}})$. Recall that $(V, \mathcal{E}(\mathcal{C}^*))$ is 2-regular, i.e., each vertex is contained in exactly two hyperedges, where a hyperedge is counted twice if it contains the vertex with multiplicity two. Hence if we define $\hat{z}$ by

$$\hat{z}_i = \sum_{j=i}^{l} z_{T_j^*} \qquad (i \in [l]),$$

then (4) can be written as

$$\hat{z}_{j_v} + \hat{z}_{j'_v} = 0 \qquad \text{for each } v \in V, \tag{6}$$

where j_v and j'_v denote the indices of the hyperedges in $\mathcal{E}(\mathcal{C}^*)$ that contain v.

In general, for an undirected graph (that may contain loops and parallel edges), the vertex-edge incidence matrix has nonzero kernel if and only if the graph has a bipartite connected component. In particular, the kernel has dimension one if and only if the graph contains exactly one bipartite connected component.

Now (6) is written as $A\hat{z} = 0$, where A is the vertex-edge incidence matrix of a graph G with vertex set $[l]$. We say that an edge of G is **nonzero** if $\hat{z}$-values of the endvertices are nonzero. Then we may assume that the subgraph of G induced by the nonzero edges is connected since otherwise (i.e., if more than one component exists) we can consider z inducing only one component among those. This implies that there is $\varepsilon > 0$ such that $\hat{z}_i \in \{-\varepsilon, 0, \varepsilon\}$ for all $i \in [l]$.

Take any two consecutive transversals $T_{i-1} \prec T_i$ in $\mathcal{C}$, and consider the interval between T_{i-1} and T_i in $\mathcal{C}^*$. Since $\mathcal{C}^*$ is a refinement of $\mathcal{C}$, denote $T_{i-1} = T_{j-1}^* \prec T_j^* \prec$

$\cdots \prec T^*_{j+s} = T_i$. Note that $Z_i \in \mathcal{E}(\mathcal{C})$ is decomposed into $s+1$ hyperedges in $\mathcal{E}(\mathcal{C}^*)$. Moreover, since $z_T > 0$ for any $T \in \mathcal{C}^* \setminus \mathcal{C}$, we have

$$\hat{z}_j > \cdots > \hat{z}_{j+s}, \tag{7}$$

in particular, s must be at most two, since $\hat{z}_j \in \{-\varepsilon, 0, \varepsilon\}$ for each $j \in [l]$. That is, Z_i is decomposed into at most three hyperedges $Z^*_{i,-\varepsilon}, Z^*_{i,0}, Z^*_{i,+\varepsilon}$, whose corresponding $\hat{z}$-values are $-\varepsilon, 0, \varepsilon$, respectively (the non-existing ones are considered to be empty). Note that $Z^*_{i,0}$ may contain $v^\circ_{Z_i}$ with multiplicity two. Then let

$$A^*_{Z_i} = Z^*_{i,+\varepsilon} \text{ and } B^*_{Z_i} = Z^*_{i,+\varepsilon} \vee Z^*_{i,0},$$

for each $Z_i \in \mathcal{E}(\mathcal{C})$. By (6) we have $\sum_{Z_i \in \mathcal{E}(\mathcal{C})} (\chi_{A^*_{Z_i}}(v) + \chi_{B^*_{Z_i}}(v)) = 2$ for $v \in V$, i.e., $\{(A^*_{Z_i}, B^*_{Z_i}) \mid Z_i \in \mathcal{E}(\mathcal{C})\}$ is a 2-cover.

Suppose that $A^*_{Z_i}, B^*_{Z_i}, T_{i-1}, T_i$ are all distinct for every i. Then by $\hat{z}_j \in \{-\varepsilon, 0, \varepsilon\}$ we have

$$z_{T_{A^*_i}} = z_{T_{B^*_i}} = \varepsilon, z_{T_k} = -\varepsilon, \text{ and } z_{T_i} = -2\varepsilon \text{ for } 1 \le i \le k-1. \tag{8}$$

Therefore

$$\begin{aligned}
&\left(\sum_{Z_i \in \mathcal{E}(\mathcal{C})} (f_{Z_i}(A^*_{Z_i}) + f_{Z_i}(B^*_{Z_i})) - \sum_{Z_i \in \mathcal{E}(\mathcal{C})} (f_{Z_i}(\emptyset) + f_{Z_i}(Z_i)) \right) \varepsilon \\
&= \sum_i (f(T_{A^*_{Z_i}}) + f(T_{B^*_{Z_i}}) - f(T_{i-1}) - f(T_i))\varepsilon \\
&= \sum_i (f(T_{A^*_{Z_i}}) z_{T_{A^*_i}} + f(T_{B^*_{Z_j}}) z_{T_{B^*_i}} + f(T_j) z_{T_i}) \\
&= \sum_{T \in \mathcal{T}} f(T) z_T < 0.
\end{aligned}$$

The same conclusion holds even if some of $A^*_{Z_i}, B^*_{Z_i}, T_{i-1}, T_i$ coincide by merging the corresponding z values given in (8). Hence the trivial 2-cover is not optimal. □

3.2 A min-max theorem of the minimum 2-cover problem

Since the minimum 2-cover problem corresponds to the dual improvement problem of (D), we now turn to solving the minimum 2-cover problem. In this subsection we shall give an optimality characterization.

Let f be a submodular function on diamonds and $\mathcal{C} : T_1 \prec T_2 \prec \cdots \prec T_k = T_{\text{top}}$ be a chain of transversals. Let $T_0 = T_{\mathbf{0}}$. The following polyhedron associated with f and $\mathcal{C}$ will play a key role in establishing a min-max theorem for the minimum 2-cover problem:

$$P(f, \mathcal{C}) = \{x \in \mathbb{R}^n \mid (a(T) - a(T_{j-1}))x \le f(T) - f(T_{j-1}) \ \forall j \in [k], \forall T \in [T_{j-1}, T_j]\}.$$

With the aid of the hypergraph $(V, \mathcal{E}(\mathcal{C}))$ and its associated functions f_Z's, this polyhedron $P(f, \mathcal{C})$ is restated as

$$P(f, \mathcal{C}) = \{x \in \mathbb{R}^n \mid x(Y) \le f_Z(Y) \ \forall Y \subseteq Z, \ \forall Z \in \mathcal{E}(\mathcal{C})\}.$$

We remark that the membership problem in $P(f, \mathcal{C})$ is solved in polynomial time. Since Z has at most one element with multiplicity two, f_Z can be minimized by calling a submodular function minimization algorithm as many times as the cardinality of the diamond corresponding to v_Z°, implying that, for a given $x \in \mathbb{R}^n$, one can decide whether $x \in P(f, \mathcal{C})$ in polynomial time.

Now the following min-max formula holds for the minimum 2-cover problem and a linear program over $P(f, \mathcal{C})$.

Theorem 9. *Let f be a submodular function on diamonds and $\mathcal{C} : T_1 \prec T_2 \prec \cdots \prec T_k = T_{\rm top}$ be a chain of transversals. Then*

$$\begin{aligned} &\max\{2x(V) \mid x \in P(f, \mathcal{C})\} \\ &= \min\left\{ \sum_{Z\in\mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z)) : \text{ a 2-cover } \{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\} \right\}. \end{aligned} \tag{9}$$

It is not difficult to see that, for an arbitrary $x \in P(f, \mathcal{C})$ and a 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$, we have that

$$2x(V) = \sum_{Z\in\mathcal{E}(\mathcal{C})} (x(A_Z) + x(B_Z)) \leq \sum_{Z\in\mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z)). \tag{10}$$

A full proof of Theorem 9 is given in Section 3.3. Before the proof, let us describe its relation to the fractional matroid matching problem.

Indeed, the left-hand side of (9) can be considered as a generalization of the fractional matchoid problem. In the **matchoid problem** introduced by Edmonds (cf. Jenkyns [15]), we are given an undirected graph $G = (X, E)$ and a matroid $\mathcal{M}_v$ on the set $\delta_G(v)$ of edges incident to v for each $v \in X$, and asked to find a set $F \subseteq E$ of maximum size such that $F \cap \delta_G(v)$ is independent in $\mathcal{M}_v$ for each $v \in X$. The fractional version of the matchoid problem is reduced to the maximization of $x(V)$ over $P(f, \mathcal{C})$ as follows. Let $V = E$ and $\mathcal{E} = \{\delta_G(v) \mid v \in X\}$. Note that $(V, \mathcal{E})$ forms a 2-regular hypergraph. The chain $\mathcal{C}$ is defined so that $\mathcal{E} = \mathcal{E}(\mathcal{C})$, and f is a submodular function on the product of diamonds, each of which corresponds to an edge in E and has two middle elements corresponding to its endvertices, satisfying that f_Z is the rank function of $\mathcal{M}_v$ for each $Z = \delta_G(v) \in \mathcal{E}(\mathcal{C})$.

We remark here that the matchoid problem is known to be equivalent to the matroid matching problem (see, e.g., [20]), but this fractional version of the matchoid problem is not equivalent to the fractional matching problem discussed in [28, 7].

3.3 A proof of Theorem 9 and a canonical 2-cover

Theorem 9 can be proved by showing the half-integrality of a maximizer by applying an argument in [7]. Here we shall give an alternative proof by an augmenting-walk approach, which implies a dynamic programming algorithm for computing a minimum 2-cover based on an optimal solution x^* of the left-hand side of (9). The obtained minimum 2-cover is called the canonical 2-cover, and plays an important role in our algorithm for (LP$^=$). Note that x^* can be found in polynomial time by the ellipsoid method, since the membership problem in $P(f,\mathcal{C})$ is solved in polynomial time. In Section 5, we shall give a combinatorial algorithm for computing x^* and the canonical 2-cover, which avoids the use of the ellipsoid method.

Let us begin proving Theorem 9. Let $x \in P(f,\mathcal{C})$ and $Z \in \mathcal{E} \equiv \mathcal{E}(\mathcal{C})$. A multiset $Y \subseteq Z$ is called an (x,Z)**-tight set** (or, simply a Z**-tight set** if x is clear) if $x(Y) = f_Z(Y)$. We remark here that, if there is a 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ such that A_Z and B_Z are (x,Z)-tight, then both x and $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ are optimal in (9).

Let

$$S_Z = \{v \in Z : \text{there is no } (x,Z)\text{-tight set containing } v\}.$$

The vertices in S_Z are called **free** in Z. If there exists $v \in S_Z \cap S_{Z'}$ with $Z \neq Z'$ or if $v_Z^\circ \in S_Z$, then we can increase $x(v)$ maintaining $x \in P(f,\mathcal{C})$, and hence we assume that this is not the case. For $v \in Z \setminus (S_Z \cup \{v_Z^\circ\})$, let $D_Z(v)$ be the smallest Z-tight set containing v. If there is no Z-tight set Y with $m(Y) = 2$, then the smallest Z-tight set containing v_Z° is also unique; v_Z° is called **semi-free** in this case, and the smallest Z-tight set is denoted by $D_Z(v_Z^\circ)$. Note that $D_Z(v)$ can be computed in polynomial time for any v using a standard submodular function minimization algorithm. If v_Z° is not semi-free, then let D_Z° be the smallest Z-tight set Y with $m(Y) = 2$. This can also be computed in polynomial time by the standard submodular function minimization.

To describe the possible modifications of x, we extend the idea of alternating paths and augmenting paths. For an intuitive description of the basic idea let us assume that there is no vertex of multiplicity two in each $Z \in \mathcal{E}$. For each $Z \in \mathcal{E}$, define a set E_Z of directed arcs by $E_Z = \{uv \mid u, v \in Z \setminus S_Z, u \in D_Z(v)\}$. An arc in E_Z is said to be **colored in** Z. Then a walk consisting of arcs in $\bigcup_{Z\in\mathcal{E}} E_Z$ is said to be augmenting if

- the directions of the arcs alternate along the walk,
- consecutive arcs have different colors,
- the first arc is a backward arc with the first vertex in S_Z, and
- the last arc is a forward arc with the last vertex in S_Z.

One can easily check that, if the value of x is alternately increased and decreased by a small amount though an augmenting walk, then $x(V)$ is increased while x remains feasible.

If there exists a vertex v_Z° of multiplicity two, we need a more careful definition for arcs incident to v_Z° and augmenting walks. The arc set E_Z is now defined to be the union of E_Z^0, E_Z^1 and E_Z^2 defined as follows. The arc set E_Z^0 consists of arcs not incident to v_Z° and is defined by

$$E_Z^0 = \{uv : u, v \in Z \setminus (S_Z \cup \{v_Z^\circ\}),\ u \in D_Z(v) \setminus \{v\}\}.$$

The arc sets E_Z^1 and E_Z^2 consist of arcs incindent to v_Z°, and their definitions depend on whether v_Z° is semi-free:

- If v°_Z is semi-free, let

$$E^1_Z = \{uv^\circ_Z : u \in D_Z(v^\circ_Z) \setminus \{v^\circ_Z\}\} \cup \{v^\circ_Z v : v^\circ_Z \in D_Z(v) \setminus \{v\}\},$$
$$E^2_Z = \emptyset.$$

- If v°_Z is not semi-free, let

$$\begin{aligned} E^1_Z =& \{uv^\circ_Z \mid u \in D^\circ_Z \setminus \{v^\circ_Z\} : \exists Z\text{-tight } v^\circ_Z\bar{u}\text{-set } Y'\} \\ & \cup \{v^\circ_Z v \mid v \in Z \setminus \{v^\circ_Z\} : m(D_Z(v)) = 1\}, \end{aligned}$$

$$\begin{aligned} E^2_Z =& \{uv^\circ_Z \mid u \in D^\circ_Z \setminus \{v^\circ_Z\} : \nexists Z\text{-tight } v^\circ_Z\bar{u}\text{-set } Y'\} \\ & \cup \{v^\circ_Z v \mid v \in Z \setminus \{v^\circ_Z\} : m(D_Z(v)) = 2\}. \end{aligned}$$

As above, arcs in E_Z will be referred to as arcs of color Z. Note that arcs with distinct colors are regarded as different arcs, and hence the resulting digraph on V may have parallel arcs. An arc in E^2_Z is called a **special arc**.

We also introduce a **label** with each arc in $E^1_Z \cup E^2_Z$ for the definition of augmenting walks. This labeling will be defined based on the following fact.

Claim 10. *If $v^\circ_Z u \in E^1_Z$, then $D_Z(u)$ has a unique shade.*

If v°_Z is not semi-free and $uv^\circ_Z \in E^1_Z$, then there is a unique smallest Z-tight $v^\circ_Z\bar{u}$-set Y with $m(Y) = 1$, and Y has a unique shade.

Proof. For the first claim, suppose that $D_Z(u)$ has more than one shade. Then there would be a Z-tight $u\bar{v^\circ_Z}$-set that corresponds to the intersection of the two transversals corresponding to $D_Z(u)$ and having distinct middle elements in the diamond of v°_Z. This contradicts $v^\circ_Z \in D_Z(u)$.

For the second claim, suppose there are two distinct minimal Z-tight $v^\circ_Z\bar{u}$-sets X and Y with $m(X) = m(Y) = 1$. Then by the minimality the shades of X and Y are different. Hence $(X \vee Y) \wedge D^\circ_Z$ is a Z-tight set smaller than D°_Z with $m((X \vee Y) \wedge D^\circ_Z) = 2$, contradicting the minimality of D°_Z. The same argument also implies that the minimal Z-tight $v^\circ_Z\bar{u}$-set has a unique shade. □

Based on this claim we shall assign a label $\ell(e)$ on each arc $e \in E^1_Z$ as follows:

$$\ell(e) = \begin{cases} \text{the shade of } D_Z(v) & \text{if } e = v^\circ_Z v, \\ \text{the shade of } D_Z(v^\circ_Z) & \text{if } v^\circ_Z \text{ is semi-free and } e = vv^\circ_Z, \\ \text{the shade of the smallest } Z\text{-tight } v^\circ_Z\bar{v}\text{-set} & \text{if } v^\circ_Z \text{ is not semi-free and } e = vv^\circ_Z. \end{cases}$$

Also, with each arc in E^2_Z, we assign a unique label, a label different from those on other arcs.

We now give a precise definition of augmenting walks. Here, in a walk, each arc may be traced more than once. A **partial augmenting walk** (PAW) is a walk that consists of arcs in $\bigcup_{Z\in\mathcal{E}} E_Z$ with the following properties:

- the last vertex is a semi-free or free vertex in some Z,
- the directions of the arcs alternate along the walk, with the last arc being a forward arc,

- consecutive arcs have different colors if the shared vertex belongs to two distinct hyperedges in $\mathcal{E}(\mathcal{C})$,
- consecutive arcs have different labels if they belong to E_Z^1 and the shared vertex is v_Z°.

Note that neither a free vertex nor a semi-free vertex can be an intermediate vertex of a PAW. Note also that a PAW may use an arc in E_Z^2 twice consecutively.

A **forward partial augmenting walk** is a PAW whose first arc is forward, while a **backward partial augmenting walk** is a PAW whose first arc is backward. For $Z \in \mathcal{E}$, let

$$Q_Z = \{v \in Z : \text{there is a forward PAW starting at } v \text{ by an arc in } E_Z\},$$
$$R_Z = \{v \in Z : \text{there is a backward PAW starting at } v \text{ by an arc in } E_Z\}.$$

Note that Q_Z and R_Z are sets, not multisets, and can be computed by dynamic programming. By definition, $(Q_Z \cup R_Z) \cap S_Z = \emptyset$.

With this definition for PAWs, augmenting walks are PAWs of the two types below:

- a backward PAW starting at a free vertex $v \in S_Z$ by an arc in $E_{Z'}$ with $Z' \neq Z$; and
- a backward PAW starting at a semi-free vertex.

An **augmenting walk** is defined to be a walk of the above types that does not have a shortcut.

The length of an augmenting walk is bounded as follows.

Claim 11. *Each vertex appears at most four times in an augmenting walk.*

Proof. Suppose that a vertex v appears more than four times. Then at least six incoming arcs or six outgoing arcs at v are used in the augmenting walk. Without loss of generality we assume that six incoming arcs at v are used, and let $u_i v$ for $1 \leq i \leq 6$ be those incoming arcs at v in the ordering of the walk such that $u_1 v$ and $u_2 v$, $u_3 v$ and $u_4 v$, and $u_5 v$ and $u_6 v$ are consecutive in the walk (where $u_i = u_j$ may hold).

If $u_1 v$ and $u_6 v$ have different colors or different labels, then there is a shortcut using $u_6 v$ next to $u_1 v$. Hence they should have the same color and the same label. Similarly the colors and the labels of $u_1 v$ and $u_4 v$ should be the same. Since $u_3 v$ and $u_4 v$ have different colors or different labels, $u_3 v$ and $u_6 v$ also have different colors or different labels. Hence there is shortcut using $u_6 v$ next to $u_3 v$, a contradiction. □

Let W be an augmenting walk with the vertex sequence $v_1, v_2, \ldots, v_l$. The **augmentation** of x through W by $\varepsilon > 0$ is to reset x by

$$x := x + \varepsilon \left(\sum_{1 \leq i \leq \lceil l/2 \rceil} \chi_{v_{2i-1}} - \sum_{1 \leq i \leq \lfloor l/2 \rfloor} \chi_{v_{2i}} \right).$$

By the definition of augmenting walks, l is always odd, and thus an augmentation increases $x(V)$. Moreover, the following claim implies that there does not exist an augmenting walk if $x(V)$ is maximized.

Claim 12. *If $\varepsilon > 0$ is sufficiently small, then x is still feasible for* (LP$^=$) *after augmentation.*

Proof. Suppose that an augmentation is performed through an augmenting walk W. It suffices to prove that $x(Y)$ does not increase for any (x, Z)-tight set Y. Let $W_Z^i = E_Z^i \cap W$ $(i = 0, 1, 2)$. For every $uv \in E_Z^0$ with $v \in Y$, the minimality of $D_Z(v)$ implies $u \in Y$. This means that the contribution of $uv \in W_Z^0$ to the increase of $x(Y)$ is nonpositive.

To prove that the total contribution of arcs in $W_Z^1 \cup W_Z^2$ to the increase of $x(V)$ is nonpositive, we shall show that the contribution of two consecutive arcs of W at v_Z° is nonpositive. Due to the definition of the augmentation, if the total contribution of the two consecutive arcs of W at v_Z° is positive, then one of the following cases occurs. Recall that $x(Y) = \chi_Y x$ for $Y \subseteq Z$.

(i) $m(Y) = 0$ and the two consecutive arcs are $v_Z^\circ u, v_Z^\circ w$ with $u \in Y$ or $w \in Y$;

(ii) $m(Y) = 1$ and the two consecutive arcs are $v_Z^\circ u, v_Z^\circ w$ with $u \in Y$ and $w \in Y$.

(iii) $m(Y) = 2$ and the two consecutive arcs are uv_Z°, wv_Z° with $u \in Y$ and $w \notin Y$;

(iv) $m(Y) = 1$ and the two consecutive arcs are uv_Z°, wv_Z° with $u \notin Y$ and $w \notin Y$;

We shall show that none of the above four cases can happen.

If (i) occurs with $u \in Y$, then $D_Z(u) \subseteq Y$. Since $v_Z^\circ \notin Y$ by $m(Y) = 0$, arc $v_Z^\circ u$ does not exist, a contradiction.

Suppose that (ii) occurs. Since $v_Z^\circ u$ exists, $v_Z^\circ \in D_Z(u)$. Moreover, the shade of $D_Z(u)$ is equal to the shade of Y since otherwise $D_Z(u) \wedge Y$ would be a smaller tight set containing u. By the same reason, the shade of $D_Z(w)$ is equal to the shade of Y. These in turn imply that $v_Z^\circ u$ and $v_Z^\circ w$ have the same label, a contradiction.

If (iii) occurs, then $D_Z^\circ \subseteq Y$. Since $w \notin Y$, arc wv_Z° does not exist, a contradiction.

Suppose that (iv) occurs. If $uv_Z^\circ \in W_Z^2$ or $wv_Z^\circ \in W_Z^2$, one can reach a contradiction as in case (iii). Hence $uv_Z^\circ \in W_Z^1$ and $wv_Z^\circ \in W_Z^1$. Therefore, since $u \notin Y$ and $w \notin Y$, the labels of uv_Z° and wv_Z° are equal to the shade of Y, contradicting that the two consecutive arcs of W have different labels. □

Now we show that if there is no augmenting walk then we can determine a 2-cover $\{(A_Z, B_Z) : Z \in \mathcal{E}(\mathcal{C})\}$ from the sets Q_Z, R_Z and S_Z such that A_Z and B_Z are Z-tight for each $Z \in \mathcal{E}(\mathcal{C})$, i.e., an optimal 2-cover $\{(A_Z, B_Z) : Z \in \mathcal{E}(\mathcal{C})\}$. To this end we need the following two claims.

Claim 13. *Suppose that Z contains three distinct elements u, v, w with $uv, vw \in E_Z$. Then the following statements hold.*

- *$uw \in E_Z$ unless $v = v_Z^\circ$, and uv_Z° and $v_Z^\circ w$ belong to E_Z^1 with the same label. Moreover, if $u = v_Z^\circ$ and $uv \in E_Z^2$ then $uw \in E_Z^2$, and if $w = v_Z^\circ$ and $vw \in E_Z^2$ then $uw \in E_Z^2$.*
- *If $u = v_Z^\circ$ and both uv and uw are in E_Z^1, then they have the same label. Similarly, if $w = v_Z^\circ$ and both vw and uw are in E_Z^1, then they have the same label.*

Proof. First we consider the case $v \neq v_Z^\circ$. If $w \neq v_Z^\circ$, then $v \in D_Z(w)$ and any Z-tight set containing v should contain u. Hence $u \in D_Z(w)$ and $uw \in E_Z$. If $u = v_Z^\circ$ and $uv \in E_Z^2$, then $u \in D_Z(v) \subseteq D_Z(w)$ and $m(D_Z(w)) = 2$, which implies $uw \in E_Z^2$. If both uv and uw are in E_Z^1, then $D_Z(v) \subseteq D_Z(w)$ implies that uv and uw have the same label.

If $w = v_Z^\circ$, then $v \in D_Z^\circ$, and as any Z-tight set containing v should contain u, u is also in D_Z°, and hence $uw \in E_Z$. If $vw \in E_Z^2$, then there is no Z-tight $v_Z^\circ\bar{v}$-set, and hence $u \in D_Z(v)$ implies that there is no Z-tight $v_Z^\circ\bar{u}$-set, i.e., $uw = uv_Z^\circ \in E_Z^2$. On the other hand, if both vw and uw are in E_Z^1, then a Z-tight set containing v_Z° that avoids u must also avoid v since $u \in D_Z(v)$. Therefore vw and uw have the same label.

Finally, suppose that $v = v_Z^\circ$ and $u \notin D_Z(w)$. Then $m(D_Z(w)) = 1$ and $uv \in E_Z^1$. Hence there is a unique smallest Z-tight $v_Z^\circ\bar{u}$-set $Y' \subseteq D_Z(w)$ with $m(Y') = 1$. Since Y' has the same shade as $D_Z(w)$, uv_Z° and $v_Z^\circ w$ have the same label. (Recall that the label of uv_Z° is the shade of Y' and the label of $v_Z^\circ w$ is the shade of $D_Z(w)$.) □

Claim 14. *If no augmenting walk exists for a feasible x, then the following statements hold for each $Z \in \mathcal{E}(\mathcal{C})$, where Z' and Z'' denote hyperedges in $\mathcal{E}(\mathcal{C})$ distinct from Z.*

(a) $Q_Z \cap Q_{Z'} = \emptyset$, $R_Z \cap R_{Z'} = \emptyset$, $S_Z \cap R_{Z'} = \emptyset$.

(b) $Q_Z \cap R_Z = \emptyset$ *if* v_Z° *is semi-free, and* $Q_Z \cap R_Z \subseteq \{v_Z^\circ\}$ *otherwise.*

(c) *If* $v_Z^\circ \in Q_Z \cup R_Z$, *then the first arcs of partial augmenting walks starting at* v_Z° *all have the same label.*

(d) *If* $vv_Z^\circ \in E_Z^2$, *then* $v \notin Q_{Z'} \cup R_Z \cup S_Z$. *Moreover,* $v_Z^\circ \in Q_Z \cup R_Z$ *implies* $v \in Q_Z \cup R_{Z'}$.

(e) *If* $vv_Z^\circ \in E_Z^1$, *then* $v \in Q_{Z'} \cup R_Z$ *implies* $v_Z^\circ \in Q_Z \cup R_Z$. *Moreover if* v_Z° *is semi-free, then* $v \notin Q_{Z'} \cup R_Z$

(f) *If* $v_Z^\circ v \in E_Z^2$, *then* $v \notin Q_Z \cup R_{Z'} \cup S_{Z'}$.

(g) *If* $v_Z^\circ v \in E_Z^1$, *then* $v \in Q_Z \cup R_{Z'} \cup S_{Z'}$ *implies* $v_Z^\circ \in Q_Z \cup R_Z$.

(h) *If* $uv \in E_Z^0$, *then* $u \in R_Z \cup Q_{Z'}$ *implies* $v \in R_Z \cup Q_{Z''}$, *and* $v \in Q_Z \cup R_{Z'}$ *implies* $u \in Q_Z \cup R_{Z''}$.

Proof. If $v \in Q_Z \cap Q_{Z'}$, then there are forward PAWs from v with the initial arcs colored in Z and Z', respectively. Then their concatenation is an augmenting walk, contradicting that there is no augmenting walk. Similarly, if $R_Z \cap R_{Z'} \neq \emptyset$ or $S_Z \cap R_{Z'}$, one can find an augmenting walk. Thus (a) holds.

We next prove (b). If v_Z° is semi-free, then it is not in R_Z since otherwise a backward PAW starting from v_Z° would be an augmenting walk. Suppose that there is $v \in Q_Z \cap R_Z$ such that $v \neq v_Z^\circ$. Let W_1 and W_2 be forward/backward PAWs starting at v. Recall that $S_Z \cap (Q_Z \cup R_Z) = \emptyset$ for each Z, and hence a PAW does not pass through a free vertex. Therefore, when tracing W_1 and W_2 from v to the ends, there is a vertex v' such that the next vertices of v' in the two walks are distinct. Note that $v' \in Q_{Z'} \cap R_{Z'}$ for some Z'. Let v'' be the preceding vertex of v', and u_1 and u_2 be the vertices next to v' in the walks W_1 and W_2, respectively. Without loss of generality we assume that $v''v', u_1v' \in W_1$ and $v'v'', v'u_2 \in W_2$.

Suppose that $v' \neq v^\circ_{Z'}$. Then $u_1u_2 \in E_{Z'}$ by Claim 13. If neither $u_1 \neq v^\circ_{Z'}$ nor $u_2 \neq v^\circ_{Z'}$, one can find an augmenting walk using u_1u_2, a contradiction. If $u_1 = v^\circ_{Z'}$ (resp. $u_2 = v^\circ_{Z'}$), then Claim 13 implies that either $u_1u_2 \in E^2_{Z'}$ or u_1u_2 has the same label as u_1v' (resp. $v'u_2$). Hence in both cases there is an augmenting walk using u_1u_2, a contradiction.

Therefore assume that $v' = v^\circ_{Z'}$. If u_1v' and $v'u_2$ have different labels, then $u_1u_2 \in E^0_{Z'}$ by Claim 13 and one can find an augmenting walk, a contradiction. If they have the same label, then there is a unique Z'-tight set X containing v' and u_2 and avoiding u_1. Since u_1v' and $v''v'$ have different labels as they are consecutive in W_1, we have $v'' \in X$. This however implies that $v'v''$ has the same label as that of u_1v', which is equal to the label of $v'u_2$. This contradicts that $v'v''$ and $v'u_2$ are consecutive in W_2.

We next prove (c). Clearly the first arcs of two backward PAWs or two forward PAWs have the same label since otherwise there would be an augmenting walk by concatenating them. Suppose that the first arc uv°_Z of a backward PAW and the first arc $v^\circ_Z w$ of a forward PAW have different labels. If $u \neq w$, then $uw \in E^0_Z$ by Claim 13, and one can find an augmenting walk. If $u = w$, then $u \in Q_{Z'} \cap R_{Z'}$ with $u \neq v^\circ_{Z'}$, which contradicts (b).

For (d), if $v \in Q_{Z'}$, then there is an augmenting walk using special arc vv°_Z consecutively, and hence $v \notin Q_{Z'}$. It also holds that $v \notin S_Z$ by the existence of special arc vv°_Z. If $v \in R_Z$, then let $uv \in E_Z$ be the initial arc of a backward PAW starting at v. If $u \neq v^\circ_Z$, then by Claim 13 $uv^\circ_Z \in E^2_Z$, and thus an augmenting walk exists. On the other hand, if $u = v^\circ_Z$, then $v^\circ_Z \in Q_Z$ and there is another arc $v^\circ_Z w$ with $w \in R_{Z'}$ and $w \neq v$. By Claim 13 and $vv^\circ_Z \in E^2_Z$, $vw \in E_Z$ holds, and hence $v \in Q_Z$, which implies $v \in Q_Z \cap R_Z$. This contradicts (b).

Let us check the latter claim of (d). If $v^\circ_Z \in R_Z$, then $v \in Q_Z$ by $v \notin Q_{Z'}$. On the other hand, if $v^\circ_Z \in Q_Z$, then there is $v^\circ_Z w \in E_Z$ with $w \in R_{Z'}$. When $w \neq v$, $vw \in E_Z$ by Claim 13 and we have that $v \in Q_Z$. When $w = v$, we get $v \in R_{Z'}$.

For (e), if $v \in Q_{Z'}$, then $v^\circ_Z \in R_Z$. If $v \in R_Z$, then let $uv \in E_Z$ be the first arc of a backward PAW starting at v. If $u \neq v^\circ_Z$, then $u \in Q_{Z'}$ and $uv^\circ_Z \in E_Z$ by Claim 13. Hence $v^\circ_Z \in R_Z$ holds. On the other hand, if $u = v^\circ_Z$, then $v^\circ_Z \in Q_Z$ holds.

Now suppose that v°_Z is semi-free. If $v \in Q_{Z'}$, then there is an augmenting walk ending at v°_Z. If $v \in R_Z$, then as above let $uv \in E_Z$ be the first arc of a backward PAW starting at v. If $u \neq v^\circ_Z$, then there is an augmenting walk ending at v°_Z since uv°_Z exists. If $u = v^\circ_Z$, then v°_Z is an intermediate vertex of a PAW, which is not possible since v°_Z is semi-free.

For (f), if $v \in R_{Z'} \cup S_{Z'}$, then there is an augmenting walk that uses a special arc $v^\circ_Z v$. If $v \in Q_Z$, then let vu be the initial arc of a forward PAW starting at v. If $u \neq v^\circ_Z$, then $u \in R_{Z'}$ and $v^\circ_Z u \in E^2_Z$ by Claim 13. Hence there is an augmenting walk that uses $v^\circ_Z u$. If $u = v^\circ_Z$, then $v^\circ_Z \in R_Z$, and there is $wv^\circ_Z \in E_Z$ with $w \in Q_{Z'}$. Then, by Claim 13, $wv \in E_Z$ and we have that $v \in R_Z \cap Q_Z$. This contradicts (b).

Assertions (g) and (h) can be checked in the same manner as (e). □

We are now ready to show that (x, Z)-tight sets A_Z and B_Z are obtained when no augmenting walk exists.

Claim 15. *For $Z \in \mathcal{E}(\mathcal{C})$, let $V_Z^\circ = \{v_Z^\circ\}$ (with multiplicity one) if v_Z° is semi-free or $v_Z^\circ \in R_Z$, and $V_Z^\circ = \emptyset$ otherwise. If no augmenting walk exists, then the following multisets A_Z and B_Z are (x, Z)-tight:*

$$\begin{aligned} A_Z &= Q_Z \cup V_Z^\circ \cup \bigcup_{Z' \neq Z} ((R_{Z'} \cup S_{Z'}) \cap Z), \\ B_Z &= Z \setminus \bigcup_{Z' \neq Z} A_{Z'}, \end{aligned} \tag{11}$$

where B_Z contains v_Z° with multiplicity two if A_Z does not contain v_Z°.

Proof. We first prove that A_Z is Z-tight. Properties (a) and (b) of Claim 14 imply that $A_Z \cap A_{Z'} = \emptyset$. It is also easy to see that $d^{in}_{E_Z^2}(A_Z) = 0$ because of (d) and (f), where $d^{in}_{E_Z^i}(A_Z)$ denotes the number of arcs in E_Z^i incoming to A_Z. It follows from (h) that $d^{in}_{E_Z^0}(A_Z) = 0$.

If $v_Z^\circ \notin A_Z$, then $d^{in}_{E_Z^1}(A_Z) = 0$ also holds by (g), and hence $d^{in}_{E_Z}(A_Z) = 0$. This means that $D_Z(v) \subseteq A_Z$ for every $v \in A_Z$, and $A_Z = \bigcup_{v \in A_Z} D_Z(v)$ is Z-tight by submodularity. Similarly, if v_Z° is semi-free, $d^{in}_{E_Z^1}(A_Z) = 0$ holds and $A_Z = \bigcup_{v \in A_Z} D_Z(v)$ implies that A_Z is Z-tight.

The remaining case is when $v_Z^\circ \in Q_Z \cup R_Z$. Suppose that $v_Z^\circ \in Q_Z$. Then there is $v_Z^\circ v \in E_Z^1$ with $v \in R_{Z'} \cup S_{Z'}$. Since $v_Z^\circ v \in E_Z^1$, $m(D_Z(v)) = 1$. By (c) all partial augmenting walks starting from v_Z° start with the same label, that is, for each $u \in A_Z \setminus \{v_Z^\circ\}$ with $v_Z^\circ \in D_Z(u)$, $m(D_Z(u)) = 1$ holds and the shade of $D_Z(u)$ is equal to the shade of $D_Z(v)$. This means that $m(\bigcup_{u \in A_Z \setminus \{v_Z^\circ\}} D_Z(u)) = 1$ and $A_Z = \bigcup_{u \in A_Z \setminus \{v_Z^\circ\}} D_Z(u)$. Hence A_Z is Z-tight. On the other hand suppose that $v_Z^\circ \in R_Z \setminus Q_Z$. Then there is $v v_Z^\circ \in E_Z^1$ with $v \in Q_{Z'}$, and there is a unique smallest Z-tight $v_Z^\circ \bar{v}$-set Y with $m(Y) = 1$. Note that $u \in Y$ implies $u \in Q_Z \subseteq A_Z$, and hence $Y \subseteq A_Z$. Note also that if $v_Z^\circ \in D_Z(u)$ for $u \in A_Z$ then $m(D_Z(u)) = 1$ by (f), and the shade of $D_Z(u)$ is equal to the shade of Y since otherwise $vu \in E_Z$ by Claim 13 and thus $u \in R_Z$, which contradicts $u \in A_Z$ by (a) and (b). Hence $m(Y \cup \bigcup_{u \in A_Z \setminus \{v_Z^\circ\}} D_Z(u)) = 1$ and and $A_Z = Y \cup \bigcup_{u \in A_Z \setminus \{v_Z^\circ\}} D_Z(u)$. This completes the proof for the Z-tightness of A_Z.

The proof that B_Z is Z-tight is as follows. If $uv \in E_Z^0$ and $u \in A_{Z'}$, then $v \in R_Z$ by (h), which means that $v \in A_{Z''}$ for some $Z'' \neq Z$. Hence $d^{in}_{E_Z^0}(B_Z) = 0$. If $m(B_Z) = 2$, then v_Z° is not semi-free and $v_Z^\circ \notin Q_Z \cup R_Z$. Hence $uv_Z^\circ \in E_Z^1 \cup E_Z^2$ implies $u \notin A_{Z'}$ by (d) and (e), and thus $D_Z^\circ \subseteq B_Z$. This means that $B_Z = \bigcup_{u \in B_Z \setminus \{v_Z^\circ\}} D_Z(u) \cup D_Z^\circ$, and hence B_Z is Z-tight. If v_Z° is semi-free, then $D_Z(v_Z^\circ) \cap A_{Z'} = \emptyset$ for any Z' with $Z' \neq Z$ by (e). Therefore $B_Z = \bigcup_{u \in B_Z} D_Z(u)$.

Finally, consider the case when $v_Z^\circ \in Q_Z \cup R_Z$. Suppose that $v_Z^\circ \in Q_Z$. Note that, if $u \in B_Z \setminus A_Z$ satisfies $v_Z^\circ \in D_Z(u)$, then $v_Z^\circ u \in E_Z^1$ by (c), implying that $m(D_Z(u)) = 1$. Moreover the shade of $D_Z(u)$ is equal to the shade of A_Z, since otherwise $u \in R_Z$, contradicting $u \in B_Z$. Hence $B_Z = A_Z \cup \bigcup_{u \in B_Z \setminus A_Z} D_Z(u)$. On the other hand, suppose that $v_Z^\circ \in R_Z \setminus Q_Z$. Then there is $v v_Z^\circ \in E_Z^1$ with $v \in Q_{Z'}$. For any $u \in B_Z \setminus A_Z$ with $v_Z^\circ \in D_Z(u)$, the shade of $D_Z(u)$ is equal to the label of vv_Z°, since otherwise $vu \in E_Z$ by Claim 13 and $u \in R_Z$ follows, contradicting $u \in B_Z$. Hence the shade of $D_Z(u)$ is the shade of A_Z, which means that $B_Z = A_Z \cup \bigcup_{u \in B_Z \setminus A_Z} D_Z(u)$. □

Observe that $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ defined by (11) forms a 2-cover. We call it the **canonical 2-cover**. Now it is not difficult to see that the canonical 2-cover attains equality in (9).

Proof of Theorem 9. For a maximizer x in (9) and the canonical 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$, A_Z and B_Z are (x, Z)-tight for each $Z \in \mathcal{E}(\mathcal{C})$ by Claim 15. Thus the inequality in (10) holds with equality, and the theorem follows. □

4 Algorithms for (LP$^=$)

In this section we describe our algorithms for (LP$^=$). In Section 4.1 we give an algorithm whose running time is polynomial in m and M. We then apply a scaling technique to improve the complexity to be polynomial in m and $\log M$ in Section 4.2.

4.1 A pseudo-polynomial time algorithm

Let us begin with an important property of $P(f, \mathcal{C})$. A chain $\mathcal{C}^*$ of transversals is said to be a refinement of a chain $\mathcal{C}$ of transversals if each element in $\mathcal{C}$ appears in $\mathcal{C}^*$.

Lemma 16. *Let $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\text{top}}$ be a chain of transversals. If $\mathcal{C}^*$ is a refinement of $\mathcal{C}$, then $P(f, \mathcal{C}^*) \subseteq P(f, \mathcal{C})$. In particular, $P(f, \mathcal{C}) \subseteq P(f)$.*

Proof. It suffices to prove $P(f, \mathcal{C}^*) \subseteq P(f, \mathcal{C})$ for a chain $\mathcal{C}^* : T_1 \prec \dots T_{j-1} \prec T^* \prec T_j \prec \cdots \prec T_k = T_{\text{top}}$. Suppose that $x \in P(f, \mathcal{C}^*)$. Then $(a(T) - a(T_{j-1}))x \leq f(T) - f(T_{j-1})$ for every $T \in [T_{j-1}, T^*]$, and $(a(T) - a(T^*))x \leq f(T) - f(T^*)$ for every $T \in [T^*, T_j]$. We show that $(a(T) - a(T_{j-1}))x \leq f(T) - f(T_{j-1})$ for every $T \in [T_{j-1}, T_j]$.

Indeed, since $a(T^*) + a(T) = a(T^* \vee T) + a(T^* \wedge T)$ and $f(T^*) + f(T) \geq f(T^* \vee T) + f(T^* \wedge T)$, we have

$$\begin{aligned}(a(T) - a(T_{j-1}))x &= (a(T^* \vee T) - a(T^*))x + (a(T^* \wedge T) - a(T_{j-1}))x \\ &\leq f(T^* \vee T) - f(T^*) + f(T^* \wedge T) - f(T_{j-1}) \leq f(T) - f(T_{j-1}).\end{aligned}$$

□

The algorithm first constructs a dual feasible solution y and tries to improve y keeping the feasibility. The algorithm terminates if it finds a maximizer x of $\max\{x(V) \mid x \in P(f, \mathcal{C})\}$ satisfying $2x(V) = f(T_{\text{top}})$, where $\mathcal{C}$ is the support of the current y, or finds a direction along which the dual objective value can be made arbitrarily small. In the former case both x and y are optimal (see Lemma 17 below), while in the latter case we can conclude that $P^=(f) = \emptyset$.

Lemma 17. *Let y be a feasible solution of* (D) *whose support is a chain $\mathcal{C}$, and let $x \in \operatorname{argmax}\{x(V) \mid x \in P(f, \mathcal{C})\}$. Then, x and y are optimal solutions for* ($\mathrm{LP}^=$) *and* (D)*, respectively, if and only if $2x(V) = f(T_{\mathrm{top}})$.*

Proof. Clearly $2x(V) = f(T_{\mathrm{top}})$ should hold if x is optimal for ($\mathrm{LP}^=$). Suppose that $2x(V) = f(T_{\mathrm{top}})$ holds. By Lemma 16, we have that $x \in P(f, \mathcal{C}) \subseteq P(f)$, and hence $x \in P^=(f)$. Summing up $(a(T_j) - a(T_{j-1}))x \leq f(T_j) - f(T_{j-1})$ for $j = 1, \ldots, k$, we have $2x(V) = \sum_{j=1}^{k}(a(T_j) - a(T_{j-1}))x \leq \sum_{j=1}^{k}(f(T_j) - f(T_{j-1})) = f(T_{\mathrm{top}}) = 2x(V)$, and hence $(a(T_j) - a(T_{j-1}))x = f(T_j) - f(T_{j-1})$ for $j = 1, \ldots, k$. Thus, $a(T_j)x = f(T_j)$ for $j = 1, \ldots, k$, implying that x and y satisfy the complementary slackness condition. □

Now the algorithm is described as follows.

Algorithm

Initialization: Assume $c_1 > c_2 > \cdots > c_n$. Let $c_{n+1} = 0$. For each $j = 1, \ldots, n$, let T'_j be the transversal with $T'_j \cap U_i = \{1_i\}$ for $i = 1, \ldots, j$ and $T'_j \cap U_i = \{0_i\}$ for $i = j+1, \ldots, n$, and let y be a feasible dual solution defined by

$$y_T = \begin{cases} c_j - c_{j+1} & \text{if } T = T'_j, \\ 0 & \text{otherwise.} \end{cases}$$

Iteration:

Step 1. Let $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\mathrm{top}}$ be the support chain of y. Find $x \in \operatorname{argmax}\{x(V) \mid x \in P(f, \mathcal{C})\}$ and the canonical 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$. If $2x(V) = f(T_{\mathrm{top}})$, then output x as an optimal solution. Otherwise, go to Step 2.

Step 2. Let $\bar{\varepsilon} = \sup\{\varepsilon \in \mathbb{R}_+ \mid y^{\varepsilon}_T \geq 0\ \forall T \in \mathcal{T} \setminus \{T_{\mathrm{top}}\}\}$, where

$$y^{\varepsilon} := y + \varepsilon \sum_{1 \leq j \leq k} (\chi_{T_{A_{Z_j}}} + \chi_{T_{B_{Z_j}}} - \chi_{T_{j-1}} - \chi_{T_j}).$$

If $\bar{\varepsilon} = +\infty$, then output "$P^=(f) = \emptyset$ ". Otherwise, set $y \leftarrow y^{\bar{\varepsilon}}$ and go back to Step 1.

The correctness of the algorithm follows from Theorem 9 and Lemma 8 combined with Lemmas 18 and 19 below. Proofs of Lemmas 18 and 19 will be given in Section 4.3.

Lemma 18. *Let $\bar{y} = y^{\bar{\varepsilon}}$ be the new dual solution in Step* 2 *and $\bar{\mathcal{C}}$ be the support chain of $\bar{y}$. Then $x \in P(f, \bar{\mathcal{C}})$, where x is the one obtained in Step* 1*.*

For a transversal T, let $\tilde{a}(T) = \sum_{i=1}^{n} a(T)_i$.

Lemma 19. *Let $\bar{y} = y^{\bar{\varepsilon}}$ be the new dual solution in Step* 2 *and let $A_{\bar{Z}}$, $B_{\bar{Z}}$, $\bar{x}$, and $\bar{\mathcal{C}} : \bar{T}_1 \prec \cdots \prec \bar{T}_{\bar{k}}$ be the counterparts of A_Z, B_Z, x, and $\mathcal{C}$ for $\bar{y}$, respectively.*

Suppose that $x(V) = \bar{x}(V)$ *and the canonical* 2*-cover* $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ *with respect to* x *is not trivial. Then*

$$\bigcup_{\bar{Z}\in\mathcal{E}(\bar{\mathcal{C}})} A_{\bar{Z}} \supseteq \bigcup_{Z\in\mathcal{E}(\mathcal{C})} A_Z.$$

If this holds with equality, then

$$\sum_{j=1}^{\bar{k}} \tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \sum_{j=1}^{\bar{k}} \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}| < \sum_{j=1}^{k} \tilde{a}(T_{j-1})|A_{Z_j}| - \sum_{j=1}^{k} \tilde{a}(T_j)|Z_j \setminus B_{Z_j}|.$$

Assuming those two lemmas, we now show the correctness of our algorithm.

Lemma 20. *Let* f *be an integer-valued submodular function on the product of* n *diamonds. Then the algorithm terminates after* $O(n^4M)$ *iterations and outputs an optimal solution for* ($\mathrm{LP}^=$) *or verifies that* ($\mathrm{LP}^=$) *is infeasible, where* $M = \max_T |f(T)|$.

Proof. Let y be a feasible solution for (D) obtained in the middle of iterations, and let $x \in \operatorname{argmax}\{x(V) \mid x \in P(f, \mathcal{C})\}$ for the support chain $\mathcal{C}$ of y. Note that $x(V) \geq -4nM$ by Theorem 9.

If $2x(V) = f(T_{\mathrm{top}})$, then x is an optimal solution for ($\mathrm{LP}^=$) by Lemma 17. If $2x(V) < f(T_{\mathrm{top}})$, then y is not optimal, and by Lemma 8 the canonical 2-cover is nontrivial; furthermore, y^ε is feasible for any $0 \leq \varepsilon \leq \bar{\varepsilon}$, and (D) is unbounded if $\bar{\varepsilon}$ is unbounded.

Assume that $\bar{\varepsilon}$ is bounded. Let $\bar{y} = y^{\bar{\varepsilon}}$, $\bar{\mathcal{C}}$ be the support chain of $\bar{y}$, and $\bar{x} \in \operatorname{argmax}\{x(V) \mid x \in P(f, \bar{\mathcal{C}})\}$. By Lemma 18, we have $x(V) = \max\{x(V) \mid x \in P(f, \mathcal{C})\} \leq \max\{x(V) \mid x \in P(f, \bar{\mathcal{C}})\} = \bar{x}(V)$. Due to the half-integrality of the linear system describing $P(f, \mathcal{C})$ (that follows from Theorem 9), $\bar{x}(V) > x(V)$ implies $\bar{x}(V) \geq x(V) + 1/2$. By Lemma 19, $\bar{x}(V) > x(V)$ occurs after $O(n^3)$ iterations. Therefore after $O(n^4M)$ iterations $2x(V)$ attains $f(T_{\mathrm{top}})$. □

As we remarked in Section 3, an optimal solution of the LP over $P(f, \mathcal{C})$ and the canonical 2-cover with respect to x can be found in polynomial time. Hence the algorithm is a pseudo-polynomial time algorithm.

4.2 A polynomial-time algorithm: a scaling technique

The number of iterations of the algorithm described in Section 4.1 is $O(n^4M)$, and thus it may be exponential in $N = \lceil \max\{m, \log M\} \rceil$. In order to obtain an algorithm with running time polynomial in N, we use a technique based on scaling the values of f.

Instead of the original problem, we may consider the problem determined by the modified set function $f^{(t)} : \mathcal{T} \to \mathbb{Z}$ defined by

$$f^{(t)}(T) = \left\lceil \frac{f(T)}{2^t} \right\rceil - \tilde{a}(T)^2 \qquad (T \in \mathcal{T}). \tag{12}$$

The idea is that if we know an optimal solution in $P^=(f^{(t)})$, then we can compute an optimal solution in $P^=(f^{(t-1)})$ using a smaller number of iterations. The following lemma establishes the properties of $f^{(t)}$ that are needed for this to work.

Lemma 21. *For a submodular function $f : \mathcal{T} \to \mathbb{Z}$,*

(i) *$f^{(t)}$ is submodular;*

(ii) *$2f_Z^{(t)} \leq f_Z^{(t-1)}$ for any chain $\mathcal{C}$ of transversals and any $Z \in \mathcal{E}(\mathcal{C})$;*

(iii) *if $P^{=}(f^{(t)})$ is empty, then so is $P^{=}(f)$.*

Proof. Proof of (i): It suffices to show that $\tilde{a}(T)^2+\tilde{a}(T')^2+2 \leq \tilde{a}(T\vee T')^2+\tilde{a}(T\wedge T')^2$ for any incomparable T and T'. Suppose there are distinct middle elements $v \in T$ and $v' \in T'$ in a diamond U_i. Then

$$\begin{aligned}
\tilde{a}(T)^2 &= \tilde{a}(T-v)^2 + 2\tilde{a}(T-v) + 1,\\
\tilde{a}(T')^2 &= \tilde{a}(T'-v')^2 + 2\tilde{a}(T'-v') + 1,\\
\tilde{a}(T \vee T')^2 &= \tilde{a}((T-v) \vee (T'-v'))^2 + 4\tilde{a}((T-v) \vee (T'-v')) + 4,\\
\tilde{a}(T \wedge T')^2 &= \tilde{a}((T-v) \wedge (T'-v'))^2.
\end{aligned}$$

Due to the supermodularity and the monotonicity of $\tilde{a}(\cdot)^2$, we get the desired relation. If there is no such a diamond, then there are two distinct diamonds U_i and U_j on which T and T' are incomparable. Then by applying the same argument one can obtain the desired relation.

Proof of (ii): Let $g^{(t)}(T) = \lceil f(T)/2^t \rceil$. Note that for any T

$$-1 \leq g^{(t-1)}(T) - 2g^{(t)}(T) \leq 0. \tag{13}$$

Let $\mathcal{C} : T_1 \prec T_2 \prec \cdots \prec T_k = T_{\text{top}}$, and let us consider $f_{Z_j}^{(t)}(Y) = f^{(t)}(T_Y) - f^{(t)}(T_{j-1})$ for each $Y \subseteq Z_j$. Since $2f_{Z_j}^{(t)}(\emptyset) = f_{Z_j}^{(t-1)}(\emptyset) = 0$, we assume $Y \neq \emptyset$. Then $\tilde{a}(T_Y)^2 \geq \tilde{a}(T_{j-1})^2 + 1$, and hence by (13) we have

$$\begin{aligned}
&f_{Z_j}^{(t-1)}(Y) - 2f_{Z_j}^{(t)}(Y)\\
&= (g^{(t-1)}(T_Y) - 2g^{(t)}(T_Y)) - (g^{(t-1)}(T_{j-1}) - 2g^{(t)}(T_{j-1})) + (\tilde{a}(T_Y)^2 - \tilde{a}(T_{j-1})^2) \geq 0.
\end{aligned}$$

Proof of (iii): Suppose that $x \in P^{=}(f)$. Let $r = \lceil f(T_{\text{top}})/2^t \rceil - f(T_{\text{top}})/2^t$ and let $x' = x/2^t - (2n - r/2n)\mathbf{1}$. Then

$$2x'(V) = \frac{2x(V)}{2^t} - (4n^2 - r) = \frac{f(T_{\text{top}})}{2^t} - 4n^2 + r = f^{(t)}(T_{\text{top}}),$$

and for any $T \in \mathcal{T} \setminus \{T_{\text{top}}\}$ we have

$$a(T)x' = \frac{a(T)x}{2^t} - \left(2n - \frac{r}{2n}\right)\tilde{a}(T) \leq \frac{a(T)x}{2^t} - \tilde{a}(T)^2 \leq f^{(t)}(T).$$

Thus $x' \in P^{=}(f)$. □

Theorem 22. *Let f be an integer-valued submodular function on the product of n diamonds, m be the sum of the sizes of all diamonds, and $M = \max_T |f(T)|$. Then there is an algorithm that solves* (LP$^{=}$) *with running time $O(\text{poly}(m) \log M)$.*

Proof. The following scaling algorithm has polynomial running time:

- Start with $t = \lceil \log M \rceil$. We can find $x \in P^{=}(f^{(t)})$ maximizing cx in time polynomial in m using the algorithm of Section 4.1.
- Suppose that for some t we have found $x \in P^{=}(f^{(t)})$ maximizing cx by the algorithm of Section 4.1. If $P^{=}(f^{(t)}) = \emptyset$, then we are done, because $P^{=}(f) = \emptyset$ by Lemma 21 (iii). Otherwise let $\bar{x}$ and $\bar{y}$ be the optimal primal and dual solutions obtained by the algorithm. Let $\bar{\mathcal{C}}$ be the support chain of $\bar{y}$. Then note that $\bar{x} \in P(f^{(t)}, \bar{\mathcal{C}})$.
- Note that $\bar{y}$ is a feasible dual solution of (D) for $f^{(t-1)}$ since in (D) we only change the objective function. Thus, to find $x \in P(f^{(t-1)})$ maximizing cx, we can start the algorithm of Section 4.1 from $\bar{y}$. Since $\bar{x} \in P(f^{(t)}, \bar{\mathcal{C}})$, it holds that $2\bar{x} \in P(f^{(t-1)}, \bar{\mathcal{C}})$ by Lemma 21 (ii). Therefore we have that

$$\begin{aligned} & f^{(t-1)}(T_{\text{top}}) - \max\{2x(V) \mid x \in P(f^{(t-1)}, \bar{\mathcal{C}})\} \leq f^{(t-1)}(T_{\text{top}}) - 4\bar{x}(V) \\ & = f^{(t-1)}(T_{\text{top}}) - 2f^{(t)}(T_{\text{top}}) \leq \tilde{a}(T_{\text{top}})^2 = 4n^2, \end{aligned}$$

  where the last inequality follows from (13). This implies that with $O(n^5)$ iterations the algorithm of Section 4.1 can solve (LP$^{=}$) for $f^{(t-1)}$. (Recall that $x(V)$ increases by 1/2 after $O(n^3)$ iterations.)
- Continuing this process, we get the primal and dual optimal $\bar{x}$ and $\bar{y}$ for $f^{(0)}$, with support $\bar{\mathcal{C}}$ such that $\bar{x} \in P(f^{(0)}, \bar{\mathcal{C}})$. Since $P(f^{(0)}, \bar{\mathcal{C}}) \subseteq P(f, \bar{\mathcal{C}})$, the algorithm of Section 4.1 can solve (LP$^{=}$) for f from $\bar{y}$ in $O(n^5)$ iterations.

Since each iteration can be done in $O(\text{poly}(m))$ time, the algorithm solves (LP$^{=}$) in $O(\text{poly}(m) \log M)$ time. □

We remark that it is not necessary to know N in advance: the algorithm in 4.1 is polynomial if the difference between $f^{(t)}(T_{\text{top}})$ and the objective value of the initial dual solution is polynomial. Thus instead of starting with $t = N$, we can start with the smallest t for which this property holds.

4.3 Proofs of Lemmas 18 and 19

In this subsection we shall give the proofs of Lemmas 18 and 19. This completes the proof of the correctness of our algorithm.

4.3.1 Proof of Lemma 18

We first prove a lemma below, from which Lemma 18 straightforwardly follows. For a chain $\mathcal{C} : T_0 \prec T_1 \prec \dots T_k = T_{\text{top}}$ and $x \in \mathbb{R}^n$, a transversal $T \in [T_{j-1}, T_j]$ is said to be $(x, \mathcal{C})$**-tight** if $(a(T) - a(T_{j-1}))x = f(T) - f(T_{j-1})$.

Lemma 23. *Let $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\text{top}}$ be a chain of transversals. If $\mathcal{C}^*$ is a refinement of $\mathcal{C}$, $x \in P(f, \mathcal{C})$, and every $T^* \in \mathcal{C}^* \setminus \mathcal{C}$ is $(x, \mathcal{C})$-tight, then $x \in P(f, \mathcal{C}^*)$, and every $(x, \mathcal{C}^*)$-tight transversal T is also $(x, \mathcal{C})$-tight.*

Proof. Let $T^* \in \mathcal{C}^* \cap [T_{j-1}, T_j]$. We have that x satisfies $(a(T) - a(T_{j-1}))x \leq f(T) - f(T_{j-1})$ for every $T \in [T_{j-1}, T_j]$, and $(a(T^*) - a(T_{j-1}))x = f(T^*) - f(T_{j-1})$. Thus

$$(a(T) - a(T^*))x = (a(T) - a(T_{j-1}))x - f(T^*) + f(T_{j-1}) \leq f(T) - f(T^*)$$

holds for any $T \in [T^*, T_j]$, and equality holds if and only if $(a(T) - a(T_{j-1}))x = f(T) - f(T_{j-1})$. □

Proof of Lemma 18. Let $\mathcal{C}^* = \mathcal{C} \cup \bar{\mathcal{C}}$. Note that $\mathcal{C}^*$ is a chain. The (x, Z)-tightness of A_Z and B_Z implies that $T_{A_{Z_j}}$ and $T_{B_{Z_j}}$ are $(x, \mathcal{C})$-tight, and hence $x \in P(f, \mathcal{C}^*)$ by Lemma 23. As $\mathcal{C}^*$ is a refinement of $\bar{\mathcal{C}}$, $x \in P(f, \bar{\mathcal{C}})$ follows from Lemma 16. □

4.3.2 Proof of Lemma 19

In the proof of Lemma 19 we shall use the following two lemmas on properties of hypergraphs obtained by refining the underlying chain.

Lemma 24. *Let $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\mathrm{top}}$ and $\mathcal{C}^* : T_1 \prec \cdots \prec T_{j-1} \prec T^* \prec T_j \prec \cdots \prec T_k$ be chains. If $x \in P(f, \mathcal{C}^*)$, $T \in [T_{j-1}, T_j]$, and T is $(x, \mathcal{C})$-tight, then both $T^* \wedge T$ and $T^* \vee T$ are $(x, \mathcal{C}^*)$-tight.*

Proof. By submodularity, $f(T \wedge T^*) + f(T \vee T^*) - f(T^*) \leq f(T)$, and thus

$$\begin{aligned} &a(T)x - a(T_{j-1})x = (a(T \wedge T^*) - a(T_{j-1}))x + (a(T \vee T^*) - a(T^*))x \\ &\leq f(T \wedge T^*) - f(T_{j-1}) + f(T \vee T^*) - f(T^*) \leq f(T) - f(T_{j-1}) = a(T)x - a(T_{j-1})x. \end{aligned}$$

As equality holds throughout, we have $(a(T \wedge T^*) - a(T_{j-1}))x = f(T \wedge T^*) - f(T_{j-1})$ and $(a(T \vee T^*) - a(T^*))x = f(T \vee T^*) - f(T^*)$. □

Lemma 25. *Let $\mathcal{C}^*$ be a refinement of a chain $\mathcal{C}$ of transversals, x be a common maximizer of $\max\{x(V) \mid x \in P(f, \mathcal{C})\}$ and $\max\{x(V) \mid x \in P(f, \mathcal{C}^*)\}$, and E and E^* be the corresponding auxiliary arcs with respect to x on $(V, \mathcal{E}(\mathcal{C}))$ and on $(V, \mathcal{E}(\mathcal{C}^*))$, respectively. Then a partial augmenting walk in E^* is a partial augmenting walk in E.*

Proof. It suffices to deal with a refinement $\mathcal{C}^* : T_1 \prec \cdots \prec T_{j-1} \prec T^* \prec T_j \prec \cdots \prec T_k$. By Lemma 24, a free vertex in $(V, \mathcal{E}(\mathcal{C}^*))$ is also free in $(V, \mathcal{E}(\mathcal{C}))$, a semi-free vertex in $(V, \mathcal{E}(\mathcal{C}^*))$ is also semi-free in $(V, \mathcal{E}(\mathcal{C}))$, and if $uv \in E^*$ then $uv \in E$. This implies that the statement holds if there is no vertex of multiplicity two.

Let $Z = Z_j$, and let Z_1^* and Z_2^* be the two hyperedges in $\mathcal{E}(\mathcal{C}^*)$ that decompose Z, in this order. If v_Z° exists, then it may belong to both Z_1^* and Z_2^* with multiplicity one. Hence some edges in $E_Z^1 \cup E_Z^2$ may move to $E_{Z_i^*}^0$, and we have to prove that this does not lead to a PAW in E^* that is not a PAW in E.

Suppose that there is a PAW W in E^* that is not a PAW in E. According to the rule of augmentation walks, there are two possibilities: (i) W passes through two consecutive arcs which are of the form $uv_Z^\circ, vv_Z^\circ \in E_Z^1$ for some $Z \in \mathcal{E}(\mathcal{C})$ with the same label; (ii) W passes through two consecutive arcs which are of the form $v_Z^\circ v, v_Z^\circ w \in E_Z^1$ for some $Z \in \mathcal{E}(\mathcal{C})$ with the same label. We shall show that none of these can happen.

- In case (i) we have $uv_Z^\circ \in E^0_{Z_1^*}$ and $vv_Z^\circ \in E^0_{Z_2^*}$. Since $uv_Z^\circ, vv_Z^\circ \in E^1_Z$ with the same label, there is a Z-tight set Y containing v_Z° but neither u nor v. If the shade of Y is the same as the shade of Z_1^* in Z, then $Y \cap Z_1^*$ is a Z_1^*-tight $\bar{u}v_Z^\circ$-set by Lemma 24, which contradicts $uv_Z^\circ \in E^0_{Z_1^*}$. Otherwise $Y \cap Z_2^*$ is a Z_2^*-tight $\bar{v}v_Z^\circ$-set by Lemma 24, which contradicts $vv_Z^\circ \in E^0_{Z_2^*}$.
- For case (ii), the proof is similar to case (i). We have $v_Z^\circ v \in E^0_{Z_1^*}$ and $v_Z^\circ w \in E^0_{Z_2^*}$. Since $v_Z^\circ v, v_Z^\circ w \in E^1_Z$ with the same label, $D_Z(v)$ and $D_Z(w)$ have the same shade. If this shade is the same as the shade of Z_1^*, then $(D_Z(w) \setminus \{v_Z^\circ\}) \cap Z_2^*$ is a Z_2^*-tight $\overline{v_Z^\circ}w$-set by Lemma 24, contradicting $v_Z^\circ w \in E^0_{Z_2^*}$. Otherwise $(D_Z(v) \setminus \{v_Z^\circ\}) \cap Z_1^*$ is a Z_1^*-tight $\overline{v_Z^\circ}v$-set by Lemma 24, contradicting $v_Z^\circ v \in E^0_{Z_1^*}$.

□

In the following we shall use notation as given in Lemma 19, and assume that $x(V) = \bar{x}(V)$ (i.e., $\max\{x(V) \mid x \in P(f, \mathcal{C})\} = \max\{x(V) \mid x \in P(f, \bar{\mathcal{C}})\}$) and $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ is nontrivial. By Lemma 18, we can assume that $x = \bar{x}$. Let $\bar{\mathcal{C}} : \bar{T}_1 \prec \cdots \prec \bar{T}_{\bar{k}} = T_{\text{top}}$ and $\mathcal{C}^* = \mathcal{C} \cup \bar{\mathcal{C}} : T_1^* \prec \cdots \prec T_{k^*}^* = T_{\text{top}}$.

We shall consider three hypergraphs $(V, \mathcal{E})$, $(V, \bar{\mathcal{E}})$, and $(V, \mathcal{E}^*)$, where $\mathcal{E} = \mathcal{E}(\mathcal{C})$, $\bar{\mathcal{E}} = \mathcal{E}(\bar{\mathcal{C}})$, and $\mathcal{E}^* = \mathcal{E}(\mathcal{C}^*)$. The corresponding canonical 2-covers $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}, \{(A_{Z'}, B_{Z'}) \mid Z' \in \mathcal{E}'\}$, and $\{(A_{Z^*}, B_{Z^*}) \mid Z^* \in \mathcal{E}^*\}$ are all defined with respect to x. For the hyperedges corresponding to $\mathcal{C}^*$ we use the following additional notation: $\mathcal{E}_1^*$ is the set of hyperedges corresponding to $T_{Z_{A_j}} - T_{j-1}$ for some j, $\mathcal{E}_2^*$ is the set of those corresponding to $T_{B_{Z_j}} - T_{A_{Z_j}}$, and $\mathcal{E}_3^*$ is the set of those corresponding to $T_j - T_{Z_{B_j}}$ for some j. A hyperedge $Z \in \mathcal{E}$ is therefore decomposed into three hyperedges: $A_Z \in \mathcal{E}_1^*$, $B_Z \setminus A_Z \in \mathcal{E}_2^*$, and $Z \setminus B_Z \in \mathcal{E}_3^*$, where some of these may be empty. It should be noted that these hyperedges can be multisets: if $m(B_Z) = 2$, then v_Z° has multiplicity two in $B_Z \setminus A_Z$. However, if $v_Z^\circ \in A_Z$ (i.e., if v_Z° appears in a PAW or it is semi-free), then this vertex is both in A_Z and in $Z \setminus B_Z$ with multiplicity one. If v_Z° is semi-free, then it is free in $Z \setminus B_Z$, because any $(Z \setminus B_Z)$-tight subset containing v_Z° corresponds to a Z-tight subset Y with $m(Y) = 2$ by Lemma 23.

Lemma 26. *For $Z^* \in \mathcal{C}^*$, the sets A_{Z^*} and B_{Z^*} have the following properties:*

(a) *if $Z^* \in \mathcal{E}_1^*$, then $A_{Z^*} = B_{Z^*} = Z^*$;*

(b) *if $Z^* \in \mathcal{E}_2^*$, then $A_{Z^*} = \emptyset$, $B_{Z^*} = Z^*$;*

(c) *if $Z^* \in \mathcal{E}_3^*$, then $A_{Z^*} = B_{Z^*} = \emptyset$;*

(d) *$\bigcup_{Z^* \in \mathcal{E}^*} A_{Z^*} = \bigcup_{Z \in \mathcal{E}} A_Z$, and*

$$\sum_{j=1}^{k^*} \tilde{a}(T_{j-1}^*)|A_{Z_j^*}| - \sum_{j=1}^{k^*} \tilde{a}(T_j^*)|Z_j^* \setminus B_{Z_j^*}| = \sum_{j=1}^{k} \tilde{a}(T_{j-1})|A_{Z_j}| - \sum_{j=1}^{k} \tilde{a}(T_j)|Z_j \setminus B_{Z_j}|.$$

Proof. We prove that the set of PAWs (w.r.t. x) does not change in $(V,\mathcal{E})$ and $(V,\mathcal{E}^*)$. By Lemma 25 every PAW in $(V,\mathcal{E}^*)$ is also a PAW in $(V,\mathcal{E})$. We now show that every PAW in $(V,\mathcal{E})$ is also a PAW in $(V,\mathcal{E}^*)$. Let $Z = Z_j \in \mathcal{C}$, and let us consider the hyperedges $A_Z \in \mathcal{E}_1^*$, $B_Z \setminus A_Z \in \mathcal{E}_2^*$, and $Z \setminus B_Z \in \mathcal{E}_3^*$. It follows from Lemma 23 that a free vertex in $(V,\mathcal{E})$ remains free in $(V,\mathcal{E}^*)$ and a semi-free vertex v_Z° in Z in $(V,\mathcal{E})$ becomes free in $Z \setminus B_Z$ in $(V,\mathcal{E}^*)$ if $m(A_Z) = 1$.

Take any PAW W in $(V,\mathcal{E})$. Suppose that $uv \in E_Z^0$ is a forward arc in W, i.e., $u \in Q_Z$ and $v \in R_{Z'} \cup S_{Z'}$ (here, as before, Z' denotes the hyperedge of $\mathcal{E}$ that contains v and is not Z). Then both u and v are in A_Z, and $uv \in E_{A_Z}^0$ because an A_Z-tight set separating u from v would also be Z-tight by Lemma 23. Similarly, if $uv \in E_Z^0$ is a backward arc in W, then both u and v are in $Z \setminus B_Z$, and $uv \in E_{Z\setminus B_Z}^0$ by Lemma 23. The existence of these arcs implies that all forward arcs of an original PAW are in $E_{A_Z}^0$ for some $Z \in \mathcal{E}$, while all backward arcs are in $E_{Z\setminus B_Z}^0$ for some $Z \in \mathcal{E}$. As free vertices remain free, W remains a PAW in $(V,\mathcal{E}^*)$ if it does not pass through v_Z° in $(V,\mathcal{E})$.

If $m(B_Z) = 2$, then W does not pass through v_Z°, since v_Z° is not semi-free and it is not in $Q_Z \cup R_Z$. If $v_Z^\circ \in A_Z$, there are three possibilities for W to pass through v_Z° according to the direction of the arcs incident to v_Z°: (i) W ends at v_Z°; (ii) W passes through v_Z° by arcs uv_Z° and vv_Z°; (iii) W passes through v_Z° by arcs $v_Z^\circ v$ and $v_Z^\circ w$.

- In case (i), W uses $uv_Z^\circ \in E_Z^1$ and v_Z° is semi-free. In this case we have that $u \in Q_Z$, and hence $u \in A_Z$. Furthermore, $uv_Z^\circ \in E_{A_Z}^0$ holds because there is no Z-tight set containing v_Z° but not u by Lemma 23. Since v_Z° is free in $(V,\mathcal{E}^*)$, W is a PAW in $(V,\mathcal{E}^*)$.

- In case (ii), v_Z° is not semi-free, and $uv_Z^\circ, vv_Z^\circ \in E_Z^1 \cup E_Z^2$ with $u \in Q_Z$ and $v \in Q_{Z'}$. We have $u \in A_Z$ and $v \in Z \setminus B_Z$ by definition of the canonical 2-cover. Also, $vv_Z^\circ \in E_Z^1$ holds by Claim 14 (d). As these two arcs have different labels, there is no Z-tight set containing v_Z° but neither u nor v. This means that $uv_Z^\circ \in E_{A_Z}^0$ by Lemma 23. Lemma 23 also implies $vv_Z^\circ \in E_{Z\setminus B_Z}^0$, since otherwise there would be a Z-tight set Y with $m_Z(Y) = 2$ that avoids v. Therefore, W remains a PAW in $(V,\mathcal{E}^*)$.

- In case (iii), v_Z° is not semi-free, and $v_Z^\circ v, v_Z^\circ w \in E_Z^1 \cup E_Z^2$ with $v \in R_Z$ and $w \in R_{Z'} \cup S_{Z'}$. In this case, $v \in Z \setminus B_Z$ and $w \in A_Z$. It also holds that $v_Z^\circ w \in E_Z^1$ by Claim 14 (f). We have $v_Z^\circ w \in E_{A_Z}^0$ because $v_Z^\circ w$ does not enter a Z-tight set. On the other hand, $v_Z^\circ v \in E_{Z\setminus B_Z}^0$ follows from Lemma 23 because any Z-tight set Y with $A_Z \cup \{v\} \subseteq Y$ must have $m_Z(Y) = 2$. Indeed, $m_Z(Y) = 1$ would imply that in $(V,\mathcal{E})$ the label of $v_Z^\circ v$ is the same as the shade of A_Z which in turn is the same as the label of $v_Z^\circ w$, contradicting the assumption that these two arcs appear in a PAW. Therefore, W remains a PAW in $(V,\mathcal{E}^*)$.

We obtained that the set of PAWs does not change, which proves (a)–(c) by the definition of A_Z and B_Z. Statement (d) easily follows from (a)–(c). □

We are now ready to prove Lemma 19.

Proof of Lemma 19. By Lemma 25 every PAW in $(V, \mathcal{E}^*)$ is also a PAW in $(V, \bar{\mathcal{E}})$. Combined with Lemma 26, this implies that $\bigcup_{\bar{Z}\in\bar{\mathcal{E}}} A_{\bar{Z}} \supseteq \bigcup_{Z^*\in\mathcal{E}^*} A_{Z^*} = \bigcup_{Z\in\mathcal{E}} A_Z$.

From now on we assume that $\bigcup_{\bar{Z}\in\bar{\mathcal{E}}} A_{\bar{Z}} = \bigcup_{Z^*\in\mathcal{E}^*} A_{Z^*}$. Observe that each hyperedge of $\bar{\mathcal{E}}$ is obtained by taking the union (with multiplicities) of at most three hyperedges in $\mathcal{E}^*$, at most one from each of $\mathcal{E}_1^*, \mathcal{E}_2^*, \mathcal{E}_3^*$. Let $\bar{Z}_j \in \bar{\mathcal{E}}$ be obtained as the union of Z_j^1, Z_j^2, Z_j^3 with $Z_j^i \in \mathcal{E}_i^*$ $(i = 1, 2, 3)$, where Z_j^i may be empty. Denote the transversal corresponding to Z_j^i in the chain $\mathcal{C}^*$ by $T_{j_i}^*$. Since Lemma 26 implies $A_{Z^*} = Z^*$ for $Z^* \in \mathcal{E}_1^*$, we have that $Z_j^1 \subseteq A_{\bar{Z}_j}$. Moreover, since $\bigcup_{\bar{Z}\in\bar{\mathcal{E}}} A_{\bar{Z}} = \bigcup_{Z^*\in\mathcal{E}^*} A_{Z^*}$, we actually have

$$A_{\bar{Z}_j} = Z_j^1 \quad \text{and} \quad \bar{Z}_j \setminus B_{\bar{Z}_j} = Z_j^3. \tag{14}$$

Let $\bar{\mathcal{E}}'$ be the set of hyperedges of $\bar{\mathcal{E}}$ that are not in $\mathcal{E}^*$. Then for each $\bar{Z}_j \in \bar{\mathcal{E}}'$ we have

$$\begin{aligned} \bar{T}_{j-1} \prec T_{j_1-1}^* &\quad \text{if } Z_j^1 \neq \emptyset, \\ T_{j_3}^* \prec \bar{T}_j &\quad \text{if } Z_j^3 \neq \emptyset. \end{aligned} \tag{15}$$

Moreover at least one of $Z_j^1 \neq \emptyset$ and $Z_j^3 \neq \emptyset$ holds since each hyperedge $\bar{Z}_j$ in $\bar{\mathcal{E}}'$ is obtained as a union of at least two hyperedges in $\mathcal{E}^*$. Therefore, by (14) and (15), it holds that

$$\tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \tilde{a}(\bar{T}_j)|\bar{Z} \setminus B_{\bar{Z}_j}| < \tilde{a}(T_{j_1-1}^*)|Z_j^1| - \tilde{a}(T_{j_3}^*)|Z_j^3| \tag{16}$$

for each $\bar{Z}_j \in \bar{\mathcal{E}}'$. Lemma 26 also implies

$$\tilde{a}(T_{j_1-1}^*)|Z_j^1| - \tilde{a}(T_{j_3}^*)|Z_j^3| = \sum_i \left\{\tilde{a}(T_{i-1}^*)|A_{Z_i^*}| - \tilde{a}(T_i^*)|Z_i^* \setminus B_{Z_i^*}|\right\},$$

where the sum is taken over all $Z_i^* \in \mathcal{E}^*$ composing $\bar{Z}_j$. On the other hand, for each $\bar{Z}_j \in \bar{\mathcal{E}} \setminus \bar{\mathcal{E}}'$, we have $\bar{Z}_j = Z_j^i$ for some $i \in \{1, 2, 3\}$, and $\bar{T}_{j-1} = T_{j_i-1}^*$ and $\bar{T}_j = T_{j_i}^*$. Therefore

$$\begin{aligned} &\left(\sum_{j=1}^{\bar{k}} \tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \sum_{j=1}^{\bar{k}} \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}|\right) - \left(\sum_{j=1}^{k^*} \tilde{a}(T_{j-1}^*)|A_{Z_j^*}| - \sum_{j=1}^{k^*} \tilde{a}(T_j^*)|Z_j^* \setminus B_{Z_j^*}|\right) \\ &= \sum_{\bar{Z}_j\in\bar{\mathcal{E}}'} \left\{\left(\tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}|\right) - \left(\tilde{a}(T_{j_1-1}^*)|Z_j^1| - \tilde{a}(T_{j_3}^*)|Z_j^3|\right)\right\} \\ &< 0, \end{aligned}$$

where the last relation follows from (16) and the fact that the 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}$ is nontrivial and hence $\bar{\mathcal{E}}' \neq \emptyset$. By Lemma 26 (d), this proves the lemma. □

5 Solving the Minimum 2-cover Problem Combinatorially

As stated in Section 3.2, $x \in \operatorname{argmax}\{x(V) \mid x \in P(f, \mathcal{C})\}$ is found by the ellipsoid method, while in Section 3.3 we have discussed how to increase $x(V)$ for a given $x \in$

$P(f, \mathcal{C})$ by using an augmenting walk. In this section we shall show that the number of augmentations becomes $O(n^3)$ by choosing the lexicographically shortest walk in each augmentation, which implies a combinatorial polynomial-time algorithm for the minimum 2-cover problem.

Suppose that W is an augmenting walk with the vertex sequence $v_1, v_2, \ldots, v_l$, and let $d = \sum_{1 \leq i \leq \lceil l/2 \rceil} \chi_{v_{2i-1}} - \sum_{1 \leq i \leq \lfloor l/2 \rfloor} \chi_{v_{2i}}$. As defined in Section 3.3, the augmentation of x through W by ε is to reset by $x := x + \varepsilon d$. In the remainder of this section, the augmentation through W means the augmentation of x through W by ε^*, where $\varepsilon^* = \max\{\varepsilon \in \mathbb{R} \mid x + \varepsilon d \in P(f, \mathcal{C})\}$. We remark that ε^* can be computed by the standard submodular function minimization. Indeed, for $Z_j \in \mathcal{E}(\mathcal{C})$ and for each nonzero element s in $U_{v^\circ_{Z_j}}$, a set function $f^s_{Z_j} : 2^{Z_j} \to \mathbb{Z}$ (by regarding Z_j as a set) defined by

$$f^s_{Z_j}(Y) = \begin{cases} f(T_Y) - f(T_{j-1}) & \text{if } v^\circ_{Z_j} \notin Y, \\ f(T_{Y \setminus \{v^\circ_{Z_j}\}} \vee \{s\}) - f(T_{j-1}) & \text{otherwise} \end{cases} \qquad (Y \subseteq Z_j)$$

is submodular. Hence, when fixing $Z_j \in \mathcal{E}(\mathcal{C})$ and $s \in U_{v^\circ_{Z_j}}$, the computation of the maximum ε is reduced to the line search problem in a submodular polyhedron, which can be solved in strongly polynomial time [22]. Thus the desired ε^* can be obtained by picking the smallest value among the maximum ε values of the restricted line search problems over all $Z \in \mathcal{E}(\mathcal{C})$ and $s \in U_{v^\circ_Z}$.

We now define the lexicographical order of augmenting walks. Assume that a total order on V is given. For a (partial) augmenting walk W, the length of the walk is denoted by $|W|$. For two partial augmenting walks W_1 and W_2 starting from a common vertex v, W_1 is said to be lexicographically shorter than W_2, denoted $W_1 \prec W_2$, if $|W_1| < |W_2|$ or $|W_1| = |W_2|$ and the list of vertices from v to the end along W_1 is lexicographically smaller than that of W_2.

Every augmenting walk W has even length, and hence has a center vertex v_W. If W_1 and W_2 are the two walks from v_W to the two endpoints of W, then the **vertex list** of W is defined to be (W_1, W_2) if $W_1 \prec W_2$, and (W_2, W_1) otherwise.

For two augmenting walks W and W', W is said to be **lexicographically smaller** than W', denoted $W \prec W'$, if $|W| < |W'|$ or $|W| = |W'|$ and (W_1, W_2) is lexicographically smaller than (W'_1, W'_2), where (W_1, W_2) and (W'_1, W'_2) are the vertex lists of W and W', respectively.

Recall that v°_Z denotes the vertex of multiplicity two in $Z \in \mathcal{E}$. For each $Z \in \mathcal{E}$, $v \in Z \setminus \{v^\circ_Z\}$, and $x \in \mathbb{R}^V$, let $W_{\mathrm{f}}(v, Z, x)$ and $W_{\mathrm{b}}(v, Z, x)$ be the lexicographically shortest forward/backward PAW among those starting from v with the initial arc in E_Z in the auxiliary digraph with respect to x. On the other hand, for v°_Z, let $W_{\mathrm{f}}(v^\circ_Z, s, x)$ and $W_{\mathrm{b}}(v^\circ_Z, s, x)$ be the lexicographically shortest forward/backward PAW among those starting from v°_Z whose initial arc is special or has label s.

Our main theorem (Theorem 35) is a direct consequence of Lemma 30 below. Before showing Lemma 30, we first establish several technical lemmas.

Lemma 27. *Suppose that there are forward and backward PAWs W_1 and W_2 starting at a vertex v with the initial arcs both colored in $Z \in \mathcal{E}$. Then there is an augmenting walk W satisfying $|W| < |W_1| + |W_2|$, unless $v = v^\circ_Z$ and the initial arcs of W_1 and W_2 are in E^1_Z and have the same label.*

Proof. This is implicit in the proof of Claim 14 (b). □

A pair of walks in Lemma 27 can be used as a certificate for the existence of a shorter augmenting walk. Lemma 27 also implies the following lemma.

Lemma 28. *Let W be the lexicographically shortest augmenting walk. Then arcs of W incident to each vertex are all incoming or all outgoing.*

Proof. Suppose to the contrary that a vertex v is incident to both an incoming arc and an outgoing arc of W. Assume for simplicity that $v \neq v_Z^\circ$. Then, by splitting W at the consecutive incoming pair at v, W can be considered as the concatenation of $W_{\mathrm{f}}(v, Z, x)$ and $W_{\mathrm{f}}(v, Z', x)$. Similarly W can be considered as the concatenation of $W_{\mathrm{b}}(v, Z, x)$ and $W_{\mathrm{b}}(v, Z', x)$. Hence $|W| = |W_{\mathrm{f}}(v, Z, x)| + |W_{\mathrm{f}}(v, Z', x)| = |W_{\mathrm{b}}(v, Z, x)| + |W_{\mathrm{b}}(v, Z', x)|$. On the other hand by Lemma 27 there are augmenting walks W_1 and W_2 such that $|W_1| < |W_{\mathrm{f}}(v, Z, x)| + |W_{\mathrm{b}}(v, Z, x)|$ and $|W_2| < |W_{\mathrm{f}}(v, Z', x)| + |W_{\mathrm{b}}(v, Z', x)|$. Thus $|W_1| + |W_2| < |W_{\mathrm{f}}(v, Z, x)| + |W_{\mathrm{b}}(v, Z, x)| + |W_{\mathrm{f}}(v, Z', x)| + |W_{\mathrm{b}}(v, Z', x)| = 2|W|$, contradicting that W is the lexicographically shortest augmenting walk.

The same argument clearly works when $v = v_Z^\circ$. □

Let v be a vertex that belongs to distinct $Z, Z' \in \mathcal{E}$. A (v, Z)**-PAW** is a forward PAW starting at v with the initial arc colored in Z or a backward PAW starting at v with the initial arc colored in Z'.

For v_Z° and a label s, a (v_Z°, s)**-PAW** is a forward or backward PAW starting at v_Z° with the initial arc labeled in s. A $(v_Z^\circ, \bar{s})$**-PAW** is a forward or backward PAW starting at v_Z° with the initial arc not labeled in s if s is a shade of some tight set, and otherwise a $(v_Z^\circ, \bar{s})$-PAW just means any forward or backward PAW starting at v_Z°.

Lemma 27 further implies the following lemma.

Lemma 29. *If there are a (v, Z)-PAW W_1 and a (v, Z')-PAW W_2 with $Z \neq Z'$, then there is an augmenting walk W satisfying $|W| \leq |W_1| + |W_2|$.*

If there are a (v_Z°, s)-PAW W_1 and a $(v_Z^\circ, \bar{s})$-PAW W_2, then there is an augmenting walk W satisfying $|W| \leq |W_1| + |W_2|$.

For $v \in Z$ with $v \neq v_Z^\circ$, let $W(v, Z, x)$ be the lexicographically shortest (v, Z)-PAW with respect to $x \in \mathbb{R}^n$. Note that $W(v, Z, x) = \min\{W_{\mathrm{f}}(v, Z, x), W_{\mathrm{b}}(v, Z', x)\}$, where the minimum is taken with respect to the lexicographical order. On the other hand, let $W(v_Z^\circ, \bar{s}, x)$ be the lexicographically shortest $(v_Z^\circ, \bar{s})$-PAW.

Lemma 30. *Let $x \in P(\mathcal{C}, f)$ and x' be obtained from x by the augmentation through the lexicographically shortest augmenting walk W. Then, for each vertex $v \in Z \setminus \{v_Z^\circ\}$,*

- *$|W(v, Z, x)| \leq |W(v, Z, x')|$, and*
- *if $|W(v, Z, x')| \leq |W|/2$, then $W(v, Z, x) \preceq W(v, Z, x')$.*

The corresponding relation also holds for $W(v_Z^\circ, \bar{s}, x')$ for any label s.

In order to describe our proof idea we first prove the case when there is no vertex of multiplicity two.

Proof of Lemma 30 when there is no vertex of multiplicity two. Let w_1 and w_2 be the endvertices of W, and E (resp. E') be the sets of arcs before (resp. after) the augmentation. Throughout the proof $D_Z(u)$ will denote the smallest (x, Z)-tight set containing u for each $u \in Z$. For each vertex $v \in Z$, let $c(v, Z)$ be the number of times that $W(v, Z, x')$ passes through arcs in $E' \setminus E$.

We shall prove the statement by induction on $c(v, Z)$, i.e., we assume that the statement holds for any v' and Z' with $c(v', Z') < c(v, Z)$. If $c(v, Z) = 0$, then the statement is trivial. Hence we may assume that there is an arc in $W(v, Z, x')$ which is in $E' \setminus E$, and let u_1u_2 be the first such arc when tracing $W(v, Z, x')$ from v. We assume that u_1u_2 is passed in the forward direction from v (and we omit the identical proof for the other case, i.e., u_1u_2 is passed in the backward direction). Let $Z^* \in \mathcal{E}$ be the color of u_1u_2 in E', i.e., $u_1u_2 \in E'_{Z^*}$. Let v_2 be the last vertex of $W(v, Z, x')$, $\tilde{W}_1$ be the part of $W(v, Z, x')$ from v to u_1, and $\tilde{W}_2$ be the part of $W(v, Z, x')$ from u_2 to v_2. Observe that $\tilde{W}_2 = W_{\mathrm{b}}(u_2, Z'', x')$, where Z'' is the element of $\mathcal{E}$ containing u_2 and distinct from Z^*. By induction it holds that

$$|\tilde{W}_2| \geq |W(u_2, Z^*, x)|. \tag{17}$$

(Note that $W_{\mathrm{b}}(u_2, Z'', x')$ is a (u_2, Z^*)-PAW.)

We say that an arc $ab \in W \cap E_{Z^*}$ is **short** if

(i) $|W_{\mathrm{f}}(a, Z^*, x)| \leq |\tilde{W}_2| + 1$, and

(ii) in the case $|\tilde{W}_2| \leq |W|/2 - 1$, $|W_{\mathrm{f}}(a, Z^*, x)| = |\tilde{W}_2| + 1$ implies $b \prec u_2$.

This definition is motivated by the following fact.

Claim 31. *Every arc $ab \in W \cap E_{Z^*}$ leaving $D_{Z^*}(u_2)$ is short.*

Proof. Let W_1 be the first part of W from the initial vertex of W until a when tracing W so that ab is in forward direction, and let W_2 be the latter part of W from b to the end vertex of W. Since $a \in D_{Z^*}(u_2)$, we have that $au_2 \in E_Z$ if $a \neq u_2$. Hence the concatenation of au_2 and W_1 is a (u_2, Z'')-PAW if $a \neq u_2$, while W_1 itself is a (u_2, Z'')-PAW if $a = u_2$. Thus, by Lemma 29 there is an augmenting walk W' satisfying

$$|W'| \leq \begin{cases} |W(u_2, Z^*, x)| + |W_1| + 1 & \text{if } a \neq u_2, \\ |W(u_2, Z^*, x)| + |W_1| & \text{if } a = u_2. \end{cases} \tag{18}$$

Since W is the shortest and W is the concatenation of W_1 and $W_{\mathrm{f}}(a, Z^*, x)$, it holds that $|W_{\mathrm{f}}(a, Z^*, x)| \leq |W(u_2, Z^*, x)| + 1 \leq |\tilde{W}_2| + 1$ by (17) and (18), where the equality may hold only if $a \neq u_2$. Thus ab satisfies (i).

To see (ii) suppose that $|\tilde{W}_2| \leq |W|/2 - 1$ and $|W_{\mathrm{f}}(a, Z^*, x)| = |\tilde{W}_2| + 1$. Then $a \neq u_2$, $|W| = |W'|$, and $|W_1| \geq |W|/2$. Hence the two walks W and W' have a common center on W_1. As $b \neq u_2$ (because $u_2 \in D_{Z^*}(u_2)$ while $b \notin D_{Z^*}(u_2)$), $W \preceq W'$ implies that $b \prec u_2$. Thus ab satisfies (ii). □

Claim 32. *The arc set $W \cap E_{Z^*}$ contains a short arc ab with $u_1 \in D_{Z^*}(b)$.*

Proof. Since $u_1u_2 \in E'$ but $u_1u_2 \notin E$, we have $u_1 \notin D_{Z^*}(u_2)$. We shall take a maximal (x, Z^*)-tight $\overline{u_1}u_2$-set Y with the property that every arc $a'b' \in W \cap E_{Z^*}$ leaving Y is short. As $D_{Z^*}(u_2)$ satisfies this condition by Claim 31, such a maximal tight set Y exists. Since $u_1u_2 \in E'$ but $u_1u_2 \notin E$, the arc set $W \cap E_{Z^*}$ contains an arc ab leaving Y. Let W_1 and W_2 be the parts of W before and after ab, respectively, when tracing W in the direction of ab. We shall show that $u_1 \in D_{Z^*}(b)$, which proves the claim.

Suppose to the contrary that $u_1 \notin D_{Z^*}(b)$. Then $Y' = Y \cup D_{Z^*}(b)$ is a Z^*-tight $\overline{u_1}u_2$-set that is larger than Y. We shall show that any arc $a'b' \in W \cap E_{Z^*}$ leaving Y' is short, contradicting the maximality of Y.

Due to the choice of Y, this is trivial if $a' \in Y$. Assume $a' \in D_{Z^*}(b)$. Lemma 28 implies that $a' \neq b$, and hence we have $a'b \in E_{Z^*}$, which in turn implies that

$$|W_{\mathrm{f}}(a', Z^*, x)| \leq |W_2| + 1 = |W_{\mathrm{f}}(a, Z^*, x)| \leq |\tilde{W}_2| + 1, \tag{19}$$

where the first inequality follows from the fact that the concatenation of $a'b$ and W_2 is a forward PAW starting from a' and the second inequality follows since ab is short. Thus condition (i) holds for $a'b'$.

To check condition (ii), suppose $|\tilde{W}_2| \leq |W|/2 - 1$. Let W_1' and W_2' be the parts of W before and after $a'b'$, respectively, when traversing W in the direction of $a'b'$. Note that $|W_2'| + 1 = |W_{\mathrm{f}}(a', Z^*, x)|$. Therefore, if $|W_{\mathrm{f}}(a', Z^*, x)| = |\tilde{W}_2| + 1$, then (19) implies $|W_2'| = |W_2| = |\tilde{W}_2| \leq |W|/2 - 1$. This in turn implies that $a \in W_1'$, $a' \in W_1$, and the center of W is between a and a'. As $a'b \in E_{Z^*}$, the concatenation of W_1', $a'b$, and W_2 is an augmenting walk. Note that, by $|W_2| = |W_2'|$, this walk has the same length and the same center as W. This cannot be lexicographically smaller than W, and thus $b' \preceq b$. On the other hand, we also have $b \prec u_2$ since ab is short, and $|W_{\mathrm{f}}(a, Z^*, x)| = |\tilde{W}_2| + 1$ holds. Thus we get $b' \prec u_2$, and (ii) holds for $a'b'$. □

Let ab be the arc guaranteed by Claim 32. We have two cases depending on whether $u_1 = b$ or not.

If $u_1 \neq b$, then, by $u_1 \in D_{Z^*}(b)$, the concatenation of $\tilde{W}_1, u_1b, W_2$ is a (v, Z)-PAW in E, denoted W'. We have $|W_2| = |W_{\mathrm{f}}(a, Z^*, x)| - 1 \leq |\tilde{W}_2|$ by the shortness, and thus $|W(v, Z, x)| \leq |W'| = |\tilde{W}_1| + |W_2| + 1 \leq |\tilde{W}_1| + |\tilde{W}_2| + 1 = |W(v, Z, x')|$.

Now suppose that $|W(v, Z, x')| \leq |W|/2$, which implies that $|\tilde{W}_2| \leq |W|/2 - 1$. If $|W_{\mathrm{f}}(a, Z^*, x)| \leq |\tilde{W}_2|$, then the above argument gives $|W(v, Z, x)| < |W(v, Z, x')|$, and hence $W(v, Z, x) \prec W(v, Z, x')$ Suppose that $|W_{\mathrm{f}}(a, Z^*, x)| = |\tilde{W}_2| + 1$. Then $b \prec u_2$ since ab is short and hence $W' \prec W(v, Z, x')$. Thus $W(v, Z, x) \prec W(v, Z, x')$ and the statement follows.

If $u_1 = b$, then let $\tilde{W}_1'$ be the walk tracing $\tilde{W}_1$ in the reversed order from u_1 to v. If $\tilde{W}_1'$ and W_2 never split when tracing them from u_1 (i.e., the vertex sequence of $\tilde{W}_1'$ coincides with an initial part of the vertex sequence of W_2, because a free or semi-free vertex cannot be an internal vertex of a PAW), then the remaining part of W_2 from v to the end is a (v, Z)-PAW. Therefore $|W(v, Z, x)| \leq |W_2| = |W_{\mathrm{f}}(a, Z^*, x)| - 1 \leq |\tilde{W}_2| < |W(v, Z, x')|$. On the other hand, if $\tilde{W}_1'$ and W_2 split at some vertex v', then the concatenation of $\tilde{W}_1'$ and W_2 followed by a shortcut at v' results in a (v, Z)-PAW (cf. the proof of Claim 14(b)). Thus $|W(v, Z, x)| \leq$

$|\tilde{W}_1|+|W_2| = |\tilde{W}_1|+|W_{\mathrm{f}}(a, Z^*, x)|-1 \leq |\tilde{W}_1|+|\tilde{W}_2| < |W(v, Z, x')|$, which completes the proof. □

Now we shall describe how to adapt the above proof to the general case.

Proof of Lemma 30. We shall check what happens when $v^{\circ}_{Z^*}$ appears in the above proof. Let E and E' be the arc sets with respect to x and x', respectively. The proof is done by induction on $c(v, Z)$ (and $c(v^{\circ}_{Z}, \bar{s})$ by extending the definition, where we regard two arcs $e \in E$ and $e' \in E'$ as different arcs if their labels are different even if the corresponding endvertices coincide). Let $\tilde{W}_1$ be the initial part of $W(v, Z, x')$ from v to u_1 and $\tilde{W}_2$ be the latter part from u_2 to the end. (When we prove the statement for $W(v^{\circ}_{Z}, \bar{s}, x)$, replace v with v°_{Z} and Z with $\bar{s}$ in the subsequent discussion.)

Take $u_1u_2 \in E'_{Z^*} \setminus E$, as in the case when there is no vertex of multiplicity two. If v_{Z^*} is semi-free, then the above proof works. Assume that v_{Z^*} is not semi-free.

For each arc in W incident to a vertex v, its **partner** (in W at v) is the arc in W adjacent at v. The label of an arc $e \in E^1_{Z^*} \cup E^2_{Z^*}$ is denoted by $\ell(e)$. For a (x, Z^*)-tight set Y with $m(Y) = 1$, the shade of Y is denoted by $sh(Y)$. When $u_1 = v^{\circ}_{Z^*}$, let $\tilde{s}$ be the label of the last arc of $\tilde{W}_1$ if $|\tilde{W}_1| \geq 1$, and let $\tilde{s} = s$ if $|\tilde{W}_1| = 0$ (i.e., $v = u_1 = v^{\circ}_{Z^*}$).

Claim 33. *Suppose that the lemma does not hold for $W(v, Z, x)$. Then there is a (x, Z^*)-tight set Y satisfying the following properties:*

(a) *$u_2 \in Y$;*

(b) *$u_1 \notin Y$ if $u_1 \neq v^{\circ}_{Z^*}$; otherwise $u_1 \notin Y$ or the shade of Y is $\tilde{s}$;*

(c) *every $ab \in E_{Z^*} \cap W$ leaving Y with $a \neq v^{\circ}_{Z^*}$ is short, i.e., it satisfies* (i) *and* (ii) *given above.*

(d) *every $v^{\circ}_{Z^*}b \in E_{Z^*} \cap W$ leaving Y is short, that is, the following properties hold unless $m(Y) = 1$ and the partner of $v^{\circ}_{Z^*}b$ at $v^{\circ}_{Z^*}$ has label equal to the shade of Y;*

- *$|W_{\mathrm{f}}(v^{\circ}_{Z^*}, \ell(v^{\circ}_{Z^*}b), x)| \leq |\tilde{W}_2| + 1$, and*
- *in the case $|\tilde{W}_2| \leq |W|/2 - 1$, $|W_{\mathrm{f}}(v^{\circ}_{Z^*}, \ell(v^{\circ}_{Z^*}b), x)| \leq |\tilde{W}_2| + 1$ implies $b \prec u_2$.*

Proof. We split the proof into two cases.

Case 1: $u_2 \neq v^{\circ}_{Z^*}$. We claim that $D_{Z^*}(u_2)$ satisfies the properties (a)–(d). Clearly (a) is satisfied. If u_1u_2 does not exist in E, then $u_1 \notin D_{Z^*}(u_2)$. If u_1u_2 exists in E, then $u_1 = v^{\circ}_{Z^*}$ should hold with $\ell(u_1u_2) = \tilde{s}$ since otherwise the lemma follows by induction. Then the shade of $D_{Z^*}(u_2)$ is $\tilde{s}$, i.e., (b) is satisfied. Property (c) can be checked by directly applying the proof of Claim 31 since $u_2 \neq v^{\circ}_{Z^*}$. To see (d), note that, if $v^{\circ}_{Z^*} \in D_{Z^*}(u_2)$, $v^{\circ}_{Z^*}u_2$ exists in E_{Z^*} with $\ell(v^{\circ}_{Z^*}u_2) = sh(D_{Z^*}(u_2))$. Observe also that one can apply the proof of Claim 31 to $v^{\circ}_{Z^*}b$ if the label of the partner of $v^{\circ}_{Z^*}b$ is not equal to $\ell(v^{\circ}_{Z^*}u_2$, which is equal to $D_{Z^*}(u_2)$. Thus (d) holds.

Case 2: $u_2 = v^{\circ}_{Z^*}$. Let s' be the label of u_1u_2 in E', and let e be the initial arc of $W(v^{\circ}_{Z^*}, \overline{s'}, x)$. If $e \in E^2_{Z^*}$, we claim that $D^{\circ}_{Z^*}$, the smallest Z^*-tight set with multiplicity two, satisfies the desired properties. Note that $D^\circ_{Z^*}$ exists since $v^\circ_{Z^*}$ is not semi-free. To see (b), suppose $u_1 \in D^\circ_{Z^*}$. Then the concatenation of $\tilde{W}_1$, $u_1 v_{Z^*}$, and $W(v^\circ_{Z^*}, \bar{s}', x)$ (and then applying a shortcut if necessary) will lead to a (v, Z)-PAW W'. Thus the lemma follows by applying the induction hypothesis to $W(v^\circ_{Z^*}, \bar{s}', x)$. Hence assume $u_1 \notin D^\circ_{Z^*}$. Then (b) follows, and since $e \in E^2_{Z^*}$, (c) and (d) can be checked by applying the proof of Claim 31.

If $e \in E^1_{Z^*}$, then let Y be the smallest Z^*-tight set whose shade is $\ell(e)$. To see (b) for Y, suppose $u_1 \in Y$. Then $\ell(u_1 v^\circ_{Z^*}) \neq sh(Y) = \ell(e)$. Hence the concatenation of $\tilde{W}_1$, $u_1 v_{Z^*}$, and $W(v^\circ_{Z^*}, \bar{s}', x)$ (and then applying a shortcut if necessary) again leads to a (v, Z)-PAW W' shorter than $W(v, z, x')$, certifying the lemma. Hence assume $u_1 \notin Y$. Then (b) holds. For (c) and (d), note that for any $a \in Y$ with $a \neq v^\circ_{Z^*}$ we have $av^\circ_{Z^*} \in E_{Z^*}$ with $\ell(av^\circ_{Z^*}) \neq sh(Y) = \ell(e)$. Hence one can apply the proof of Claim 31. □

In what follows we assume that the lemma does not hold for $W(v, Z, x)$, and we shall take a maximal Y satisfying the properties (a)–(d) in Claim 33.

For an arc $av^\circ_{Z^*} \in W$, let $\tilde{a}v^\circ_{Z^*}$ be the partner of $av^\circ_{Z^*}$ in W at $v^\circ_{Z^*}$. For each arc ab, define X_{ab} by

$$X_{ab} = \begin{cases} D_{Z^*}(b) & \text{if } b \neq v^\circ_{Z^*}, \\ \text{the smallest } Z^*\text{-tight set containing } b \text{ and avoiding } \tilde{a} & \text{if } b = v^\circ_{Z^*},\ \tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}, \\ D^\circ_{Z^*} & \text{if } b = v^\circ_{Z^*},\ \tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}. \end{cases}$$

Claim 34. *The arc set $W \cap E_{Z^*}$ contains a short arc ab leaving Y such that*

- *$u_1 \in X_{ab}$, and*
- *if $u_1 = v^\circ_{Z^*}$ and $m(X_{ab}) = 1$, then $sh(X_{ab}) \neq \tilde{s}$.*

Proof. Due to properties (a) and (b) for Y, there exists an arc ab in $W \cap E_{Z^*}$ leaving Y. By (c) and (d), such ab can be chosen to be short. (Note that, if $m(Y) = 1$ and all arcs in $W \cap E_{Z^*}$ leaving Y are outgoing from $v^\circ_{Z^*}$, then there are at least two arcs $v^\circ_{Z^*}b$ and $v^\circ_{Z^*}b'$ leaving Y, at least one of which must be short.)

We now show that $Y \cup X_{ab}$ satisfies (c) and (d) in Claim 33. To see this, take any $a'b' \in W \cap E_{Z^*}$ leaving $Y \cup X_{ab}$. It suffices to consider the case when $b = v^\circ_{Z^*}$ or $a' = v^\circ_{Z^*}$, since otherwise one can directly apply the proof of Claim 32 to see that $a'b'$ is short. Note that $a' \neq b$ by Lemma 28.

Case 1: Suppose $b = v^\circ_{Z^*}$. Let $\tilde{a}v^\circ_{Z^*}$ be the partner of $av^\circ_{Z^*}$ in W. If $\tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}$, then $\tilde{a} \notin X_{ab}$ by the definition of X_{ab}. Since $a' \in X_{ab}$, we have that $\ell(a'v^\circ_{Z^*}) \neq \ell(\tilde{a}v^\circ_{Z^*})$ (which also holds if $\tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}$). Thus one can apply the proof of Claim 32 to see that $a'b'$ is short.

Case 2: Suppose $a' = v^\circ_{Z^*}$. Let $v^\circ_{Z^*}\tilde{b'}$ be the partner of $v^\circ_{Z^*}b'$ in W. Note that, if $v^\circ_{Z^*}b$ exists with $\ell(v^\circ_{Z^*}b) \neq \ell(v^\circ_{Z^*}\tilde{b'})$, then we have

$$\begin{aligned} |W_{\mathrm{f}}(v^\circ_{Z^*}, \ell(v^\circ_{Z^*}b'), x)| &= |W(v^\circ_{Z^*}, \overline{\ell(v^\circ_{Z^*}\tilde{b'})}, x)| \\ &\leq |W_{\mathrm{f}}(v^\circ_{Z^*}, \ell(v^\circ_{Z^*}b), x)| \leq |W_{\mathrm{f}}(a, Z^*, x)| \leq |\tilde{W}_2| + 1 \end{aligned}$$

where the last inequality follows from the shortness of ab. Moreover $|W(v, Z, x')| \leq |W|/2$ and $|W_{\rm f}(v^\circ_{Z^*}, \ell(v^\circ_{Z^*}b'), x)| = |\tilde{W}_2| + 1$ imply $b' \prec u_2$ by applying the argument of the proof of Claim 32, i.e., $v^\circ_{Z^*}b'$ is short.

With this in mind, we now consider three cases to prove that $v^\circ_{Z^*}b'$ satisfies (d): (2-1) $v^\circ_{Z^*} \notin X_{ab}$; (2-2) $v^\circ_{Z^*} \in X_{ab} \setminus Y$; (2-3) $v^\circ_{Z^*} \in X_{ab} \cap Y$. Case 2-1: If $v^\circ_{Z^*} \notin X_{ab}$, then clearly $sh(Y) = sh(Y \cup X_{ab})$, and hence (d) holds for $v^\circ_{Z^*}b'$ as (d) holds for Y. Case 2-2: If $v^\circ_{Z^*} \in X_{ab} \setminus Y$, then $sh(X_{ab}) = sh(X_{ab} \cup Y)$. When $sh(X_{ab}) = \ell(v^\circ_{Z^*}\tilde{b'})$, there is nothing to prove for $v^\circ_{Z^*}b'$. When $sh(X_{ab}) \neq \ell(v^\circ_{Z^*}\tilde{b'})$, we have $\ell(v^\circ_{Z^*}b) = sh(X_{ab}) \neq \ell(v^\circ_{Z^*}\tilde{b'})$, which implies that $v^\circ_{Z^*}b'$ is short as shown above. Case 2-3: Suppose $v^\circ_{Z^*} \in X_{ab} \cap Y$. If $m(Y) = 2$ or $sh(Y) = sh(X_{ab})$, then $m(Y \cup X_{ab}) = m(Y) = 2$ or $sh(Y) = sh(X_{ab} \cup Y)$, and (d) holds for $v_{Z^*}b'$ as (d) holds for Y. Otherwise $m(Y) = 1$ and $sh(Y) \neq sh(X_{ab})$. If $\ell(v^\circ_{Z^*}\tilde{b'}) \neq sh(Y)$, then the shortness of $v^\circ_{Z^*}b'$ follows since Y satisfies (d). Otherwise $\ell(v^\circ_{Z^*}\tilde{b'}) = sh(Y) \neq sh(X_{ab}) = \ell(v^\circ_{Z^*}b)$, implying that $v^\circ_{Z^*}b'$ is short.

This completes the proof for Case 2, and thus $Y \cup X_{ab}$ satisfies (c) and (d). Since $Y \cup X_{ab}$ clearly satisfies (a), the maximality of Y implies that $Y \cup X_{ab}$ violates (b). This means that if $u_1 \neq v^\circ_{Z^*}$ then $u_1 \in X_{ab}$ while if $u_1 = v^\circ_{Z^*}$ then $u_1 \in X_{ab}$ and $sh(X_{ab}) \neq sh(Y) = \tilde{s}$. Thus the claim holds. □

Now we are ready to complete the proof. Let ab be the arc guaranteed in Claim 34. Let W_2 be the latter part of W from b to the end when tracing ab in the forward direction.

If $u_1 \in Y$, then $b \neq u_1$ and $u_1 = v^\circ_{Z^*}$ by (b). In this case, $u_1 \in X_{ab} = D_{Z^*}(b)$, and $sh(X_{ab}) \neq \tilde{s}$ by Claim 34. In other words $v^\circ_{Z^*}b$ exists in E with $\ell(v^\circ_{Z^*}b) = sh(X_{ab}) \neq \tilde{s}$. Hence the concatenation of $\tilde{W}_1, v^\circ_{Z^*}b, W_2$ leads to a (v, Z)-PAW, implying the lemma by the shortness of ab.

Hence assume $u_1 \notin Y$. If $u_1 \neq v^\circ_{Z^*}$ and $b \neq v^\circ_{Z^*}$, then the concatenation of $\tilde{W}_1$, u_1b and W_2 leads to a (v, Z)-PAW (by taking a shortcut if necessary).

Suppose $u_1 = b = v^\circ_{Z^*}$. Then the concatenation of $\tilde{W}_1$ and W_2 can be shortcut to a (v, Z)-PAW if $\tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}$. Otherwise, i.e., if $\tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}$, we have $\ell(\tilde{a}v^\circ_{Z^*}) = sh(X_{ab}) \neq \tilde{s}$ by Claim 34. Hence again the concatenation of $\tilde{W}_1$ and W_2 can be shortcut to a (v, Z)-PAW.

Suppose $u_1 = v^\circ_{Z^*} \neq b$. Then $\ell(v^\circ_{Z^*}b) = sh(D_{Z^*}(b)) = sh(X_{ab}) \neq \tilde{s}$. Hence the concatenation of $\tilde{W}_1$, $v^\circ_{Z^*}b$ and W_2 leads to a (v, Z)-PAW.

Suppose $u_1 \neq b = v^\circ_{Z^*}$. Then the concatenation of $\tilde{W}_1$, $u_1v^\circ_{Z^*}$ and W_2 leads to a (v, Z)-PAW if $\tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}$. Otherwise, i.e., if $\tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}$, we have $u_1 \in X_{ab}$ and $\tilde{a} \notin X_{ab}$, implying $\ell(u_1v^\circ_{Z^*}) \neq \ell(\tilde{a}v^\circ_{Z^*})$. Thus the concatenation of $\tilde{W}_1$, $u_1v^\circ_{Z^*}$ and W_2 leads to a (v, Z)-PAW.

In every case we have found a (v, Z)-PAW which certifies the lemma by the shortness of ab. This completes the proof. □

The main theorem now follows.

Theorem 35. *Given $x \in P(f, \mathcal{C})$, after $O(n^3)$ augmentations through the lexicographically shortest augmenting walks, $x(V)$ is maximized.*

Proof. For simplicity we give a proof for the case when there is no vertex of multiplicity two, and we omit the straightforward extension of the proof to the general case.

Let $x' \in P(f, \mathcal{C})$ be obtained from x by an augmentation through the lexicographically shortest augmenting walk W. Let W' be the lexicographically shortest augmenting walk after the augmentation, and denote the center of W' by $v_{W'}$. Note that W' is the concatenation of $W(v_{W'}, Z, x')$ and $W(v_{W'}, Z', x')$. By Lemma 30, $|W(v_{W'}, Z, x)| \leq |W(v_{W'}, Z, x')|$ and $|W(v_{W'}, Z', x)| \leq |W(v_{W'}, Z', x')|$ hold, which means that the auxiliary digraph with resect to x contains an augmenting walk shorter than or equal to $|W'|$ by Lemma 29. This implies $|W| \leq |W'|$.

Suppose $|W| = |W'|$. Let $N_Z(x) = \{v \in Z \mid |W(v, Z, x)| \leq |W|/2\}$ and $w(v, Z, x)$ be the second vertex of $W(v, Z, x)$. Similarly, define $N_Z(x')$ and $w(v, Z, x')$ for the feasible solution x' after the augmentation. By Lemma 30, we have that $|W(v, Z, x)| \leq |W(v, Z, x')|$ for each $v \in V$, $N_Z(x') \subseteq N_Z(x)$, and $w(v, Z, x) \preceq w(v, Z, x')$ for each $v \in N_Z(x')$. Moreover, since $W \neq W'$, at least one of the following three holds: (i) $N_Z(x') \subsetneq N_Z(x)$; (ii) $w(v, Z, x) \prec w(v, Z, x')$ for some $v \in N_Z(x')$; (iii) $|W(v, Z, x)| < |W(v, Z, x')|$ for some $v \in N_Z(x')$. Indeed, if (i) and (ii) do not hold, then either $W(v, Z, x) = W_{\mathrm{f}}(v, Z, x)$ and $W(v, Z, x') = W_{\mathrm{b}}(v, Z', x')$ or $W(v, Z, x) = W_{\mathrm{b}}(v, Z', x)$ and $W(v, Z, x') = W_{\mathrm{f}}(v, Z, x')$ hold for some v, Z with $v \in N_Z(x')$. However a forward PAW and a backward PAW cannot have the same length, which implies (iii). Therefore the number of augmentations is $O(n^3)$. □

Combining Theorems 22 and 35, the proof of Theorem 2 is completed.

Acknowledgement

This work is partially supported by OTKA grant K109240, the János Bolyai Research Scholarship of the Hungarian Academy of Sciences, and JSPS Grant-in-Aid for Scientific Research (B) 25280004.

References

[1] S. Chang, D. Llewellyn and J. Vande Vate, Matching 2-lattice polyhedra: finding a maximum vector, Discrete Math. 237 (2001), 29–61.

[2] S. Chang, D. Llewellyn and J. Vande Vate, Matching 2-lattice polyhedra: duality and extreme points, Discrete Math. 237 (2001), 63–95.

[3] D. Cohen, M. Cooper, P. Jeavons and A. Krokhin, Supermodular functions and the complexity of Max CSP, Discrete Appl. Math. 149 (2005), 53–72.

[4] D. Cohen, M. Cooper, P. Jeavons and A. Krokhin, The complexity of soft constraint satisfaction, Artificial Intelligence 170 (2006), 983–1016.

[5] S. Fujishige and S. Iwata, Bisubmodular function minimization, SIAM J. Discrete Math. 19 (2006), 1065–1073.

[6] S. Fujishige and S. Tanigawa, A Min-max theorem for transversal submodular functions and its implications, SIAM J. Discrete Math. 28 (2013), 1855–1875.

[7] D. Gijswijt and G. Pap, An algorithm for weighted fractional matroid matching, J. Combin. Theory, Ser. B 103 (2013), 509–520.

[8] M. Grötschel, L. Lovász and A. Schrijver, The ellipsoid method and its consequences in combinatorial optimization, Combinatorica 1 (1981), 169–197.

[9] H. Hirai, Discrete convexity and polynomial solvability in minimum 0-extension problems, in: Proc. 24th Annual ACM-SIAM Symposium on Discrete Algorithms (2013), 1770–1788.

[10] A. Huber and V. Kolmogorov, Towards minimizing k-Submodular functions, in: Proc. 2nd International Symposium on Combinatorial Optimization (2012), 451–462.

[11] S. Iwata, L. Fleischer and S. Fujishige, A combinatorial strongly polynomial algorithm for minimizing submodular functions, J. ACM 48 (2001), 761–777.

[12] S. Iwata and K. Murota, A minimax theorem and a Dulmage–Mendelsohn type decomposition for a class of generic partitioned matrices, SIAM J. Matrix Anal. Appl. 16 (1995), 719–734.

[13] S. Iwata and J.B. Orlin, A simple combinatorial algorithm for submodular function minimization, in: Proc. 20th Annual ACM-SIAM Symposium on Discrete Algorithms (2009), 1230–1237.

[14] H. Ito, S. Iwata and K. Murota, Block-triangularizations of partitioned matrices under similarity/equivalence transformations, SIAM J. Matrix Anal. Appl. 15 (1994), 1226–1255.

[15] T.A. Jenkyns, Matchoids: A Generalization of Matchings and Matroids, Ph.D. Thesis, University of Waterloo, 1974.

[16] P. Jonsson, M. Klasson and A. Krokhin, The approximability of three-valued MAX-CSP, SIAM J. Comput. 35 (2006) 1329–1349.

[17] P. Jonsson, F. Kuivinen and J. Thapper, Min CSP on four elements: moving beyond submodularity, in: Proc. 20th International Conference on Principles and Practice of Constraint Programming, Lecture Notes in Computer Science 6876 (2011), 438–453.

[18] A. Krokhin and B. Larose, Maximizing supermodular functions on product lattices, with application to maximum constraint satisfaction, SIAM J. Discrete Math. 22 (2008), 312–328.

[19] F. Kuivinen, On the complexity of submodular function minimisation on diamonds, Discrete Optimization 8 (2011), 459–477.

[20] L. Lovász and M.D. Plummer, Matching Theory, Akadémiai Kiadó, 1986.

[21] S.T. McCormick and S. Fujishige, Strongly polynomial and fully combinatorial algorithms for bisubmodular function minimization, Math. Program. 122 (2010), 87–120.

[22] K. Nagano, A strongly polynomial algorithm for line search in submodular polyhedra, Discrete Optimization 4 (2007), 349–359.

[23] J.B. Orlin, A faster strongly polynomial time algorithm for submodular function minimization, Math. Program. 118 (2009), 237–251.

[24] L. Qi, Directed submodularity, ditroids and directed submodular flows, Math. Program. 42 (1988), 579–599.

[25] A. Schrijver, A combinatorial algorithm minimizing submodular functions in strongly polynomial time, J. Combin. Theory, Ser. B 80 (2000), 346–355.

[26] J. Thapper and S. Živný, The power of linear programming for valued CSPs, in: Proc. 44th ACM Symposium on the Theory of Computing (2012), 668–678.

[27] J. Thapper and S. Živný, The complexity of finite-valued CSPs, in: Proc. 45th ACM Symposium on the Theory of Computing (2013), 695–704.

[28] J. Vande Vate, Fractional matroid matchings, J. Combin. Theory, Ser. B 55 (1992), 133–145.