Egerváry Research Group
on Combinatorial Optimization

Technical reportS

TR-2014-14. Published by the Egerváry Research Group, Pázmány P. sétány 1/C, H–1117, Budapest, Hungary. Web site: www.cs.elte.hu/egres . ISSN 1587–4451.

Minimizing Submodular Functions on Diamonds via Generalized Fractional Matroid Matchings

Satoru Fujishige, Tamás Király, Kazuhisa Makino,
Kenjiro Takazawa, and Shin-ichi Tanigawa

December 2014
Revised: January 2016, April 2017

Minimizing Submodular Functions on Diamonds via Generalized Fractional Matroid Matchings*

Satoru Fujishige**, Tamás Király***, Kazuhisa Makino**, Kenjiro Takazawa‡, and Shin-ichi Tanigawa**

Abstract

In this paper we show the first polynomial-time algorithm for the problem of minimizing submodular functions on the product of diamonds of finite size. 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 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

Let V be a finite set. 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

*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,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
‡Department of Industrial and Systems Engineering, Faculty of Science and Engineering, Hosei University, Tokyo 184-8584, Japan. takazawa@hosei.ac.jp

December 2014. Revised: January 2016, April 2017

obtained much later, independently by Schrijver [27] and by Iwata, Fleischer and Fujishige [12]. Since then, there have been several improvements in running time, e.g. [14, 25].

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 [15, 10], and it was posed as an open problem in [12] 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 [17, 18]. 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 [28, 29]. A result by Thapper and Živný [28] (see also [19]) 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] 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 [28]. 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 [12, 27], one can reduce the problem to the standard submodular function minimization if the underlying lattice is distributive. Krokhin and Larose [20] 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, one of the smallest non-distributive lattices, 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 [13]. A **diamond** is a lattice consisting of the minimal element, the maximal element, and an arbitrary finite number of pairwise incomparable elements called **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**. In this paper we will deal with only finite diamonds.

Previous work on the submodular function minimization problem on the product of diamonds is as follows. If the diamonds have at most two middle elements, then the lattice is distributive, and by the observation in [12, 27] we can use the standard submodular function minimization algorithm. However, if there exists a diamond with more than two middle elements, then this diamond is modular but not distributive, and hence the standard submodular function minimization algorithms cannot be directly applied. Kuivinen's pioneering research [21] gives a pseudo-polynomial algorithm for the minimization of submodular functions on diamonds of finite size.

Our main result is the first polynomial-time algorithm for the submodular function minimization problem on the product of diamonds of finite size. Let f be a submodular function on the direct product of n diamonds $U_1, \ldots, U_n$, and let $U = \bigcup_{i \in [n]} U_i$. Denote $|U|$ by m, and the maximal absolute value of f by M.

Theorem 1. *Let f be a submodular function on the direct product of a finite number of diamonds of finite size. A minimizer of f can be computed in a polynomial number of arithmetic steps and function evaluations in m and $\log M$.*

In preparation for detailed discussion, we introduce some notation and assumption. Denote the set of integers $\{1, \ldots, n\}$ by $[n]$. A subset $T \subseteq U$ is called a **transversal** if $|T \cap U_i| = 1$ for every $i \in [n]$. We denote by $\mathcal{T}$ the set of transversals, by T_{bottom} the transversal consisting of the minimal elements, and by T_{top} the transversal consisting of the maximal 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 our submodular function f 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 $T_1, T_2 \in \mathcal{T}$. Since adding a constant to $f(T)$ for each $T \in \mathcal{T}$ preserves submodularity, throughout the paper we assume $f(T_{\text{bottom}}) = 0$.

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 of $T \cap U_i$ in the lattice U_i. The following linear programming problem is closely related to the problem of minimizing f:

$$\begin{aligned} \text{maximize} \quad & cx \\ \text{subject to} \quad & a(T)x \leq f(T) \text{ for each } T \in \mathcal{T}, \\ & \sum_{i \in [n]} 2x_i = f(T_{\text{top}}), \\ & x \in \mathbb{R}^n. \end{aligned} \tag{1}$$

If this optimization problem can be solved in polynomial time, then, by the results of Grötschel, Lovász and Schrijver [8], the separation problem for (1) can also be solved in polynomial time. Based on this observation, we can solve the problem of minimizing f in polynomial time by binary search as follows.

We first remark that an upper bound of M (and thus the range of the binary search) can be computed in polynomial time. Note that the submodularity of f implies $f(T) \leq n \cdot \max\{f(T) \mid T \in \mathcal{T},\ |T \setminus T_{\text{bottom}}| \leq 1\}$ for each $T \in \mathcal{T}$, which provides a polynomial-time computable upper bound on values of f. On the other hand, the lower bound for values of f is obtained in the following manner. Let T_1 be a transversal consisting only of middle elements, and let $\mathcal{T}_1 = \{T \in \mathcal{T} \mid T \subseteq T_1 \cup T_{\text{bottom}} \cup T_{\text{top}}\}$. The family $\mathcal{T}_1$ induces a distributive lattice, so we can find $T_1^* \in \operatorname{argmin}\{f(T) \mid T \in \mathcal{T}_1\}$ using standard submodular minimization. Let T^* be an arbitrary minimizer of f. Then $T_1 \vee T^*$ and $T_1 \wedge T^*$ are in $\mathcal{T}_1$, so $f(T^*) \geq f(T_1 \vee T^*) + f(T_1 \wedge T^*) - f(T_1) \geq$

$2f(T_1^*) - f(T_1)$ by submodularity, where the last term can be computed in polynomial time.

Now, in the binary search, checking whether $\min_{T \in \mathcal{T}} f(T) \geq d$ for a given $d \leq 0$ can be done as follows. Add the same nonnegative value $-d$ to the function values of f except $f(T_{\rm bottom})$ to obtain a new function f'. It follows that f' is submodular, and hence checking whether $\min_{T \in \mathcal{T}} f(T) \geq d$ can be accomplished by checking whether a submodular function f' is nonnegative. If $f'(T_{\rm top}) < 0$, clearly f' is not nonnegative. Otherwise, construct a further new function f'' from f' by replacing $f'(T_{\rm top})$ with 0. Then f'' is again submodular, and hence the nonnegativity test for f' is done by checking whether $\mathbf{0}$ is feasible in (1) for f'', which is a special case of the separation problem.

Kuivinen's pseudo-polynomial time algorithm [21] also follows the same strategy. He reduced the problem of minimizing f to a linear programming problem over a distinct and larger polytope. He then showed that this linear programming can be solved in pseudo-polynomial time, again, with the aid of the ellipsoid method. The ellipsoid method is used in two phases: reducing the minimization of f to a linear maximization over a polytope; and solving the linear programming problem.

In contrast to this, in the present paper we give a combinatorial polynomial algorithm for solving the linear program (1).

Theorem 2. *Let f be a submodular function on the direct product of a finite number of diamonds of finite size. 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 [30], 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 [26], Fujishige and Iwata [5], and McCormick and Fujishige [23]. 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 [11], 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, which 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 $T \neq T'$. T and T' are said to be **comparable** if $T \preceq T'$ or $T' \preceq T$, and otherwise **incomparable**. Recall that T_{bottom} denotes the transversal formed by all minimal elements, i.e., $T_{\text{bottom}} = \{0_i \mid i \in V\}$. The transversal consisting of all maximal elements is denoted by T_{top}, i.e., $T_{\text{top}} = \{1_i \mid i \in V\}$. Given transversals T_1 and T_2 satisfying $T_1 \preceq T_2$, we define an interval $[T_1, T_2]$ of transversals by $[T_1, T_2] = \{T \in \mathcal{T} : T_1 \preceq T \preceq T_2\}$. 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 $T \cap U_i = \{0_i\}$, $a(T)_i = 2$ if $T \cap U_i = \{1_i\}$, and $a(T)_i = 1$ otherwise.

Let $f : \mathcal{T} \to \mathbb{Z}$ be a submodular function on the diamonds, that is, $f(T_1) + f(T_2) \geq f(T_1 \vee T_2) + f(T_1 \wedge T_2)$ for every $T_1, T_2 \in \mathcal{T}$. Recall that we assume $f(T_{\text{bottom}}) = 0$. 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 $x(X) = \sum_{i \in X} x_i$.

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

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

The dual program of ($\text{LP}^{=}$) is given by

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

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

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

Proof. We first show that there is an optimal dual solution y of (LP$^=$) whose support is a chain. This can be seen by the following standard argument. Consider the potential $\sum_{T\in\mathcal{T}} g(T)y_T$, where $g(T) := (\sum_{i\in V} a(T)_i)(\sum_{i\in V}(2-a(T)_i))$, and take a basic optimal dual solution y with the smallest potential. Suppose that two transversals $T, T' \in \text{supp}(y)$ are incomparable, i.e., neither $T \preceq T'$ nor $T' \preceq T$. Then set

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

where $\delta = \min\{y_T, y_{T'}\}$. The submodularity of f implies that the revised solution is still optimal. Moreover, the strict submodularity[1] of g implies that the potential value decreases, which contradicts the choice of y.

Now let y be an optimal dual solution such that its support is a chain $T_1 \prec T_2 \prec \cdots \prec T_k = T_{\text{top}}$. We define $\hat{y}$ by

$$\hat{y}_i = \sum_{j=i}^{k} y_{T_j} \qquad (i \in [k]).$$

Then the linear constraint in (D) can be written as

$$\hat{y}_{j_v} + \hat{y}_{j'_v} = c_v \qquad (v \in V),$$

where j_v (resp. j'_v) denote the first index j such that T_j contains a middle (resp. the top) element of U_v. Now this linear system is written as $A\hat{y} = c$ for some vertex-edge incidence matrix $A^\top$ of a graph with vertex set $[k]$. Since the determinant of any nonsingular submatrix of a vertex-edge incidence matrix of a connected graph belongs to $\{\pm 1, \pm 2\}$, $\hat{y}$ is half-integral. The half-integrality of y now immediately follows from that of $\hat{y}$ by definition. □

The following proposition enables us to focus on the case when all elements in c have distinct values.

[1] A function f is *strictly submodular* if $f(T) + f(T') > f(T \vee T') + f(T \wedge T')$ for any incomparable $T, T' \in \mathcal{T}$. The strict submodularity of g can be seen as follows.

In general, a continuous function $h : \mathbb{R} \to \mathbb{R}$ is said to be *strictly concave* if for any $\xi, \xi' \in \mathbb{R}$ with $\xi \leq \xi'$ and $\delta \in \mathbb{R}$ with $0 \leq \delta \leq \xi' - \xi$, $h(\xi + \delta) + h(\xi' - \delta) < h(\xi) + h(\xi')$. By setting $h(\eta) := \eta(2|V| - \eta)$ $(\eta \in \mathbb{R})$, h is strictly concave. Now for any $T, T' \in \mathcal{T}$, let $\xi = \sum_{i\in V} a(T)_i$ and $\xi' = \sum_{i\in V} a(T')_i$. Without loss of generality, assume that $\xi \leq \xi'$. Observe that, if T and T' are incomparable, then there is $\delta \in \mathbb{R}$ such that $0 \leq \delta \leq \xi' - \xi$, $\xi + \delta = \sum_{i\in V} a(T \vee T')_i$, and $\xi' - \delta = \sum_{i\in V} a(T \wedge T')_i$. We thus obtain $g(T) + g(T') = h(\xi) + h(\xi') > h(\xi + \delta) + h(\xi' - \delta) = g(T \vee T') + g(T \wedge T')$.

Proposition 4. *Let $c \in \mathbb{Z}^n$ and $\bar{M} \in \mathbb{R}$ with $\bar{M} \geq M(= \max\{|f(T)| : T \in \mathcal{T}\})$. Define $c' \in \mathbb{Z}^n$ by $c'_i = 6n^2\bar{M}c_i + i$ for $i \in V$. If $x \in \mathbb{R}^n$ is a basic optimal solution of* (LP$^=$) *with the objective function c', then x is also optimal for* (LP$^=$) *with the objective function c.*

Proof. We first remark that for any $x' \in P^=(f)$ and $i \in V$, $x'_i \leq f(\{1_i\})/2 \leq \bar{M}/2$ and $x'_i = x'(V) - x'(V - i) \geq -\bar{M}$ hold.

Suppose that x is not optimal for (LP$^=$) with the objective function c. Then for any basic optimal solution x' for (LP$^=$) with the objective function c we have $c'(x - x') = 6n^2\bar{M}(cx - cx') + \sum_{i\in V} i(x_i - x'_i) \leq -3n^2\bar{M} + \sum_{i\in V} \frac{3}{2}i\bar{M} < 0$, where for the second inequality we used the half-integrality of x, which follows from Proposition 3. This implies that x is not optimal for (LP$^=$) with the objective function c', a contradiction. □

Since an upper bound of M can be computed in polynomial time, 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 5. *Suppose that $c_i \neq c_{i'}$ for all distinct i and i' in V, and let y be a feasible solution of* (D) *whose support is a chain $T_{\text{bottom}} = T_0 \prec T_1 \prec \cdots \prec T_k = T_{\text{top}}$. 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

Given a feasible solution y for (D), we are interested in finding a new feasible solution with better objective value or certifying that y is optimal. We reduce this question to a combinatorial problem, called the **minimum 2-cover problem**. In Section 3.2, we show that by solving the minimum 2-cover problem one can determine if y is optimal, or one can find a direction that improves y if y is not optimal. Then in Section 3.3 we focus on solving the minimum 2-cover problem. We establish a min-max theorem (Theorem 11) for the minimum 2-cover problem, and we then show how to construct a minimum 2-cover by giving a constructive proof of Theorem 11. The canonical optimal solution given in the proof will be explicitly used in our algorithm for (LP$^=$) in Section 4.

3.1 Preliminaries to the minimum 2-cover problem

A rough sketch of the minimum 2-cover problem is as follows: given a 2-regular hypergraph $(V, \mathcal{E})$ and a function $f_Z : 2^Z \to \mathbb{Z}$ for each $Z \in \mathcal{E}$, we are asked to find sets $A_Z, B_Z \subseteq Z$ for each hyperedge $Z \in \mathcal{E}$ that minimize

$$\sum_{Z\in\mathcal{E}} (f_Z(A_Z) + f_Z(B_Z))$$

under the constraint that each element $v \in V$ is contained in exactly two sets among them. To be precise, the minimum 2-cover problem is defined on a 2-regular multiset family on V, which is a slightly more general set system than a hypergraph. Further, f_Z is not a usual set function: it is a multiset-functions defined on the family of the sub-multisets of $Z \in \mathcal{E}$ that satisfies a somewhat exotic version of submodularity inherited from f. We note that this submodularity is not based on lattice operations. Our proof of the min-max theorem for the minimum 2-cover problem—and our algorithm for finding a minimum 2-cover—will rely solely on this exotic submodularity of the functions f_Z, without explicitly using that these functions were derived from f. In this subsection, we show how f_Z is derived from f. A formal description of the minimum 2-cover problem will be given in Section 3.2.

Let $T_1 \prec T_2 \prec \cdots \prec T_k (= T_{\text{top}})$ be a chain of transversals, denoted by $\mathcal{C}$, that satisfies the property in Proposition 5. Let $T_0 = T_{\text{bottom}}$. For each $j \in [k]$, let Z_j be a multiset of elements in V which contains $i \in V$ with multiplicity $a(T_j)_i - a(T_{j-1})_i$. Let $\mathcal{E}(\mathcal{C}) = \{Z_1, \dots, Z_k\}$ be a family of those multisets. A multiset Z is associated with a characteristic function $\tilde{\chi}_Z$ on V, where $\tilde{\chi}_Z(i)$ is equal to the multiplicity of i in Z for each $i \in V$. As we only consider multisets in which each element has multiplicity 0, 1, or 2, it follows that $\tilde{\chi}_Z : V \to \{0, 1, 2\}$. Now it is not difficult to see that $(V, \mathcal{E}(\mathcal{C}))$ is 2-regular, i.e., $\sum_{Z \in \mathcal{E}(\mathcal{C})} \tilde{\chi}_Z(i) = 2$ for every $i \in V$.

For simplicity, we write $Y \subseteq Z$ if $\tilde{\chi}_Y \leq \tilde{\chi}_Z$ for multisets Y and Z. We also denote $\tilde{\chi}_Z x$ by $x(Z)$ for any $x \in \mathbb{R}^n$ and multiset Z. For $u, v \in V$, a multiset Y is called a $u\bar{v}$**-set** if Y contains u and avoids v.

By Proposition 5, each $Z \in \mathcal{E}(\mathcal{C})$ has at most one element with multiplicity two; denote this by v_Z° if it exists. For any multiset $Y \subseteq Z$, let $m_Z(Y) = \tilde{\chi}_Y(v_Z^\circ)$ if v_Z° exists, and $m_Z(Y) = 0$ otherwise.

If $Z_j \in \mathcal{E}(\mathcal{C})$ has no element of multiplicity two, then Z_j is a set and there is a natural one-to-one correspondence between 2^{Z_j} and the interval $[T_{j-1}, T_j]$. In contrast, if $v_{Z_j}^\circ$ exists, then such a natural correspondence does not exist. In this case, we associate a set $\mathcal{T}'_{Z_j}(Y)$ of transversals for each $Y \subseteq Z_j$. Formerly, for each $Y \subseteq Z_j$, let

$$\mathcal{T}'_{Z_j}(Y) = \{T \in [T_{j-1}, T_j] \mid a(T) = a(T_{j-1}) + \tilde{\chi}_Y\},$$
$$\mathcal{T}_{Z_j}(Y) = \operatorname{argmin}\{f(T) \mid T \in \mathcal{T}'_{Z_j}(Y)\}.$$

We simply write $m_Z(Y)$ and $\mathcal{T}_Z(Y)$ as $m(Y)$ and $\mathcal{T}(Y)$, respectively, if Z is clear from the context.

Consider $Z \in \mathcal{E}(\mathcal{C})$ for which v_Z° exists and a subset $Y \subseteq Z$ with $m(Y) = 1$. Each transversal in $\mathcal{T}(Y)$ contains one middle element in the diamond indexed by v_Z°. The set of middle elements that appear in at least one transversal in $\mathcal{T}(Y)$ is called the set of **shades** of Y, denoted by $\mathcal{S}(Y)$. If $m(Y) \neq 1$, then let $\mathcal{S}(Y) = \emptyset$. We say that two shade sets $\mathcal{S}(X)$ and $\mathcal{S}(Y)$ are *singly identical* if both of them consist of the same single element, i.e., $\mathcal{S}(X) = \mathcal{S}(Y)$ and $|\mathcal{S}(X)| = |\mathcal{S}(Y)| = 1$.

As we will see later in Proposition 6, the notions of shades and singly identical shade sets are essential in our algorithm.

Example. Consider the three diamonds indexed by $\{1, 2, 3\}$ given in Figure 1(a). We associate each middle element with a point and each top element with a line in a

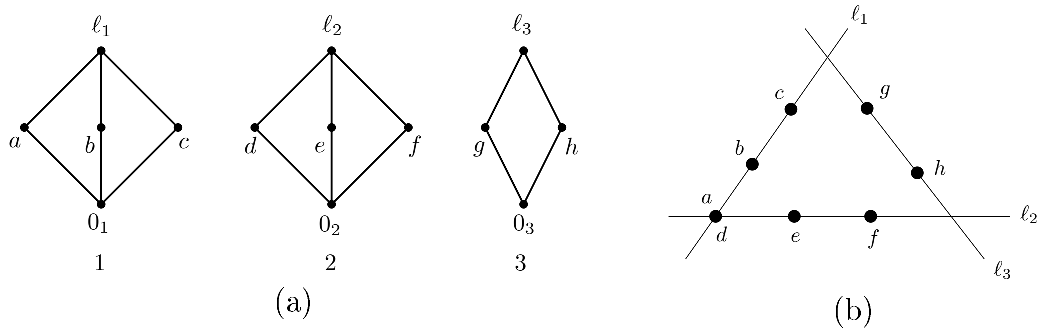

Figure 1: Example.

plane as shown in Figure 1(b). We define a function f' over the product of the three diamonds such that its value is equal to the affine dimension of the affine span of the corresponding elements in the plane, and let $f = f' + 1$ so that $f(T_{\text{bottom}}) = 0$. For example, $f((a, e, 0_3)) = 2, f((a, d, 0_3)) = 1$, $f((\ell_1, 0_2, 0_3)) = 2$, and $f((a, d, \ell_3)) = 3$. One can easily check that f is submodular.

Consider a chain $\mathcal{C} : T_0 \prec T_1 \prec T_2$ of transversals, where $T_0 = (0_1, 0_2, 0_3)$, $T_1 = (a, 0, g)$, and $T_2 = (\ell_1, \ell_2, \ell_3)$. Then it follows that $\mathcal{E}(\mathcal{C}) = \{Z_1, Z_2\}$, where $Z_1 = \{1, 3\}$ and $Z_2 = \{1, 2, 2, 3\}$, i.e.,

$$\tilde{\chi}_{Z_1} = \begin{pmatrix} 1 \\ 0 \\ 1 \end{pmatrix}, \quad \tilde{\chi}_{Z_2} = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}.$$

Hence $v^\circ_{Z_1}$ does not exist and $v^\circ_{Z_2} = 2$. Below we show examples of $m(Y_i)$ and $\mathcal{T}_{Z_2}(Y_i)$ for some $Y_i \subseteq Z_2$:

- $Y_1 = \{1\}$: $m(Y_1) = 0$ and $\mathcal{T}_{Z_2}(Y_1) = \{(\ell_1, 0, g)\}$;
- $Y_2 = \{2\}$: $m(Y_2) = 1$ and $\mathcal{T}_{Z_2}(Y_2) = \{(a, d, g)\}$;
- $Y_3 = \{1, 2\}$: $m(Y_3) = 1$ and $\mathcal{T}_{Z_2}(Y_3) = \{(\ell_1, d, g), (\ell_1, e, g), (\ell_1, f, g)\}$;
- $Y_4 = \{2, 2\}$: $m(Y_4) = 2$ and $\mathcal{T}_{Z_2}(Y_4) = \{(a, \ell_2, g)\}$.

The sets of shades of Y_2 and Y_3 are $\mathcal{S}(Y_2) = \{d\}$ and $\mathcal{S}(Y_3) = \{d, e, f\}$, respectively. □

For each $Z \in \mathcal{E}(\mathcal{C})$ define the multiset-function f_Z by

$$f_Z(Y) = f(T_Y) - f(T_{j-1}) \qquad (Y \subseteq Z),$$

where T_Y is an arbitrary member of $\mathcal{T}_Z(Y)$. Observe that in the case where $Z \in \mathcal{E}(\mathcal{C})$ has no element with multiplicity two, f_Z is a submodular set function on Z. In order to describe "submodularity" of f_Z for the case where v°_Z exists, we introduce operations $X \sqcap Y$ and $X \sqcup Y$ for multisets $X, Y \subseteq Z$. These operations are defined by using characteristic functions as follows:

$$\tilde{\chi}_{X\sqcup Y}(i) := \begin{cases} 2 & \text{if } i = v^\circ_Z,\ m(X) = m(Y) = 1, \text{ and} \\ & \mathcal{S}(X) \text{ and } \mathcal{S}(Y) \text{ are not singly identical,} \\ \max\{\tilde{\chi}_X(i), \tilde{\chi}_Y(i)\} & \text{otherwise,} \end{cases}$$

$$\tilde{\chi}_{X\sqcap Y}(i) := \begin{cases} 0 & \text{if } i = v^\circ_Z,\ m(X) = m(Y) = 1, \text{ and} \\ & \mathcal{S}(X) \text{ and } \mathcal{S}(Y) \text{ are not singly identical,} \\ \min\{\tilde{\chi}_X(i), \tilde{\chi}_Y(i)\} & \text{otherwise.} \end{cases}$$

These are the normal union and intersection operations unless $m(X) = m(Y) = 1$. If $m(X) = m(Y) = 1$, then the value of v°_Z depends on the shades of X and Y. Note also that $X \sqcap X \neq X$ and $X \sqcup X \neq X$ may hold.

Proposition 6. *For any $X, Y \subseteq Z$ and $x \in \mathbb{R}^n$, it holds that $f_Z(X) + f_Z(Y) \geq f_Z(X \sqcup Y) + f_Z(X \sqcap Y)$ and $x(X) + x(Y) = x(X \sqcup Y) + x(X \sqcap Y)$.*

Proof. For any $X, Y \subseteq Z$, one can choose $T_X \in \mathcal{T}(X)$ and $T_Y \in \mathcal{T}(Y)$ such that $T_X \vee T_Y \in \mathcal{T}(X \sqcup Y)$ and $T_X \wedge T_Y \in \mathcal{T}(X \sqcap Y)$.

Denote $Z = Z_j$. Then $f_Z(X) + f_Z(Y) = f(T_X) + f(T_Y) - 2f(T_{j-1}) \geq f(T_X \vee T_Y) + f(T_X \wedge T_Y) - 2f(T_{j-1}) = f_Z(X \sqcup Y) + f_Z(X \sqcap Y)$, where the inequality follows from the submodularity of f.

For the second equality, note that $a(T_X) + a(T_Y) = a(T_X \vee T_Y) + a(T_X \wedge T_Y)$. Then it follows that $x(X) + x(Y) = (\tilde{\chi}_X + \tilde{\chi}_Y) \cdot x = (a(T_X) + a(T_Y) - 2a(T_{j-1})) \cdot x = (a(T_X \vee T_Y) + a(T_X \wedge T_Y) - 2a(T_{j-1})) \cdot x = (\tilde{\chi}_{X\sqcup Y} + \tilde{\chi}_{X\sqcap Y}) \cdot x = x(X \sqcup Y) + x(X \sqcap Y)$. □

3.2 A 2-cover and dual improvement

In this subsection, we define the minimum 2-cover problem, an auxiliary problem obtained from a given chain of transversals satisfying the property in Proposition 5. We will show that an optimal solution of the problem corresponding to the chain formed by the support of a non-optimal solution y of (D) gives rise to an improving direction for y.

Let $\mathcal{C}$ be a chain of transversals that satisfies the property in Proposition 5. 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 $A_Z, B_Z \subseteq Z$ and $\sum_{Z\in\mathcal{E}(\mathcal{C})}(\tilde{\chi}_{A_Z}(i) + \tilde{\chi}_{B_Z}(i)) = 2$ 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

$$\sum_{Z\in\mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z)). \tag{2}$$

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 7. *Let y be a feasible solution for* (D) *whose support is a chain $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\mathrm{top}}$. Then y is optimal for* (D) *if and only if the trivial 2-cover is optimal for the minimum 2-cover problem for $(f, \mathcal{C})$.*

Lemma 7 directly follows from Theorem 11 and Lemma 20 in Section 4.1. Here we only prove the following easy part of Lemma 7.

Lemma 8. *Let y be a feasible solution for* (D) *whose support is a chain $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\mathrm{top}}$. If there is a 2-cover $\{(A_{Z_j}, B_{Z_j}) \mid j \in [k], Z_j \in \mathcal{E}(\mathcal{C})\}$ whose objective value is smaller than that of the trivial 2-cover, then a feasible solution y^ε of* (D) *whose objective value is smaller than that of y is constructed in the following manner: for each $j \in [k]$, choose arbitrary $T_j^A \in \mathcal{T}(A_{Z_j})$ and $T_j^B \in \mathcal{T}(B_{Z_j})$, and then define*

$$y^\varepsilon := y + \varepsilon \sum_{1 \le j \le k} (\chi_{T_j^A} + \chi_{T_j^B} - \chi_{T_{j-1}} - \chi_{T_j}), \tag{3}$$

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

Proof. Comparing the objective value of $\{(A_{Z_j}, B_{Z_j}) \mid j \in [k], Z_j \in \mathcal{E}(\mathcal{C})\}$ with that of the trivial 2-cover, we have that

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

Hence, if the objective value (2) of $\{(A_{Z_j}, B_{Z_j}) \mid j \in [k], Z_j \in \mathcal{E}(\mathcal{C})\}$ is smaller than that of the trivial 2-cover, then $\sum_T f(T) y_T^\varepsilon < \sum_T f(T) y_T$ holds.

To prove the feasibility of y^ε for (D), it suffices to show that

$$\sum_{1 \le j \le k} (a(T_j^A) + a(T_j^B) - a(T_{j-1}) - a(T_j))_v = 0 \tag{4}$$

for each $v \in V$. Take any $v \in V$. Without loss of generality we assume $a(T_j^A)_v \le a(T_j^B)_v$ for each $j \in [k]$. Then, since $\{(A_{Z_j}, B_{Z_j}) \mid j \in [k], Z_j \in \mathcal{E}(\mathcal{C})\}$ is a 2-cover, if v is in some A_{Z_j} then

$$\sum_{1 \le j \le k} (a(T_j^A) - a(T_{j-1}))_v = 1 = \sum_{1 \le j \le k} (a(T_j) - a(T_j^B))_v,$$

and otherwise

$$\sum_{1 \le j \le k} (a(T_j^A) - a(T_{j-1}))_v = 0 = \sum_{1 \le j \le k} (a(T_j) - a(T_j^B))_v.$$

Thus (4) holds. □

By Proposition 6, there always exists a minimum 2-cover that satisfies the following properties for every $Z \in \mathcal{E}(\mathcal{C})$:

(a) $A_Z \subseteq B_Z$, and

(b) if $m_Z(A_Z) = m_Z(B_Z) = 1$, then $\mathcal{S}(A_Z)$ and $\mathcal{S}(B_Z)$ are singly identical.

If a 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ satisfies the above two properties, then the support of the new dual solution y^ϵ defined in the statement of Lemma 8 forms a chain again. We will update dual solutions by using the **canonical 2-cover** defined later, which always satisfies the two properties (a) and (b) (see Claim 18).

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

In Section 3.2, we have seen that the minimum 2-cover problem corresponds to the improvement of solutions for (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 $\mathcal{C} : T_1 \prec T_2 \prec \cdots \prec T_k = T_{\mathrm{top}}$ be a chain of transversals. Let $T_0 = T_{\mathrm{bottom}}$. 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 ($Z \in \mathcal{E}(\mathcal{C})$), this polyhedron $P(f, \mathcal{C})$ can be rewritten as

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

Remark 9. We remark that the membership problem in $P(f, \mathcal{C})$ is polynomial-time solvable. If $Z \in \mathcal{E}(\mathcal{C})$ has no element of multiplicity two, then Z is a set and f_Z is a submodular set function. If $Z \in \mathcal{E}(\mathcal{C})$ has an element of multiplicity two, then for each element s in $U_{v_Z^\circ}$, we define a standard submodular set function f_Z^s on $Z \setminus \{v_Z^\circ\}$ (considered as a set) by $f_Z^s(Y) = f(T_Y \vee \{s\}) - f(T_{j-1})$ for each $Y \subseteq Z \setminus \{v_Z^\circ\}$, where T_Y is the unique transversal in $\mathcal{T}(Y)$. For fixed $Z \in \mathcal{E}(\mathcal{C})$ and $s \in U_{v_Z^\circ}$, we can minimize f_Z^s by standard submodular minimization. This implies that we can minimize f_Z for each $Z \in \mathcal{E}(\mathcal{C})$; moreover, we can solve the separation and line search problems over $P(f, \mathcal{C})$ in polynomial time [24].

Remark 10. It should be noted that the submodularity described in Proposition 6 is the only property of the functions f_Z used to prove the results on the minimum 2-cover problem (Theorem 11 and the algorithm in Section 5). The relation between f and f_Z is not used in any other way. Therefore, when reading the proof of Theorem 11, the reader does not have to keep track of the correspondence between transversals and multisets.

Our first result is the following min-max formula for the minimum 2-cover problem.

Theorem 11. *Let f be a submodular function on diamonds and let $\mathcal{C} : T_1 \prec T_2 \prec \cdots \prec T_k = T_{\mathrm{top}}$ be a chain of transversals satisfying the property in Proposition 5. Then*

$$\begin{aligned} &\max\{2x(V) \mid x \in P(f, \mathcal{C})\} \\ &= \min\left\{ \sum_{Z \in \mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z)) \;\middle|\; \text{a 2-cover } \{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\} \right\}. \end{aligned} \tag{5}$$

It is not difficult to see that the left-hand side is at most the right-hand side in (5): 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

$$2x(V) = \sum_{Z \in \mathcal{E}(\mathcal{C})} (x(A_Z) + x(B_Z)) \leq \sum_{Z \in \mathcal{E}(\mathcal{C})} (f_Z(A_Z) + f_Z(B_Z)). \tag{6}$$

A full proof of Theorem 11 is given in Section 3.4.

Before the proof, let us describe the relation of the left-hand side of (5) to the fractional matchoid problem. In the **matchoid problem**, introduced by Edmonds (cf. Jenkyns [16]), we are given an undirected graph $G = (V_G, E_G)$ and a matroid $\mathcal{M}_u$ on the set $\delta_G(u)$ of edges incident to u for each $u \in V_G$. The aim is to find a set $F \subseteq E_G$ of maximum size such that $F \cap \delta_G(u)$ is independent in $\mathcal{M}_u$ for each $u \in V_G$.

Denote $V_G = \{u_1, \ldots, u_k\}$, and $Z_j = \delta_G(u_j)$ for $j \in [k]$. Let $V = E_G$ and let $\mathcal{E} = \{Z_j \mid j \in [k]\}$; note that $(V, \mathcal{E})$ forms a 2-regular hypergraph. Let r_{Z_j} denote the rank function of the matroid $\mathcal{M}_{u_j}$. The fractional version of the matchoid problem can be written as

$$\max\{x(V) \mid x \in \mathbb{R}^V_+,\ x(Y) \leq r_Z(Y)\ \forall Y \subseteq Z,\ \forall Z \in \mathcal{E}\}.$$

This is very similar to the left-hand side of (5). In fact, in the following manner, one can construct a submodular function f on the product of $|V|$ diamonds (each with two middle elements) and a chain of transversals $\mathcal{C}$ such that $\mathcal{E} = \mathcal{E}(\mathcal{C})$ and $f_Z = r_Z$ for every $Z \in \mathcal{E}$. For each $e = uv \in E_G$, we label the two middle elements of the diamond U_e by u and v. For a transversal T and an index $j \in [k]$, let

$$Y(T)_j = \{e \in Z_j \mid T \cap U_e \text{ is the top element or the middle element labeled by } u_j\}.$$

Define f by $f(T) = \sum_{j=1}^k r_{Z_j}(Y(T)_j)$; this function is submodular on the product of the diamonds. To define the chain $\mathcal{C} = \{T_1, \ldots, T_k\}$ of transversals, consider an edge $e = u_i u_l$ in E_G with $i < l$. Transversal T_j contains the top element of U_e if $j \geq l$, it contains the middle element corresponding to u_i if $i \leq j < l$, and contains the bottom element if $j < i$. It is easy to check that $\mathcal{E} = \mathcal{E}(\mathcal{C})$ and $f_Z = r_Z$ for every $Z \in \mathcal{E}$.

We remark here that the matchoid problem is known to be equivalent to the matroid matching problem (see, e.g., [22]), but this fractional version of the matchoid problem is not equivalent to the fractional matroid matching problem in the sense of Vande Vate [30].

3.4 A proof of Theorem 11 and the canonical 2-cover

Here we give a proof of Theorem 11 by using 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 (5). The obtained minimum 2-cover is called the canonical 2-cover, and plays an important role in our algorithm for ($\mathrm{LP}^=$). 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 the optimal x^* and the canonical 2-cover, which avoids the ellipsoid method.

3.4.1 Proof idea

Although our strategy follows the standard argument (e.g., for the matroid intersection problem), the presence of multisets makes the problem much more complicated. In order to overview the basic idea, in this subsection we shall give a sketch of the proof for the case when there is no vertex of multiplicity two in any $Z \in \mathcal{E}(\mathcal{C})$, i.e., each Z is a set. Since the chain $\mathcal{C}$ does not change in the proof, we denote $\mathcal{E}(\mathcal{C})$ by $\mathcal{E}$.

Let $x \in P(f, \mathcal{C})$ and $Z \in \mathcal{E}$. A set $Y \subseteq Z$ is called (x, Z)**-tight** (or, simply Z**-tight** if x is clear) if $x(Y) = f_Z(Y)$. Observe that if there is a 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}$ such that A_Z and B_Z are (x, Z)-tight for every $Z \in \mathcal{E}$, then both x and $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}$ are optimal in (5) since the inequality in (6) holds with equality.

For $Z \in \mathcal{E}$, define $S_Z \subseteq Z$ by

$$S_Z = \{v \in Z \mid \text{there is no } (x, Z)\text{-tight set containing } v\}. \tag{7}$$

Since each $v \in V$ belongs to exactly two distinct sets $Z, Z' \in \mathcal{E}$, if $v \in S_Z \cap S_{Z'}$, then we can increase $x(v)$ while maintaining $x \in P(f, \mathcal{C})$. Hence we assume that there exists no $v \in V$ contained in distinct S_Z and $S_{Z'}$.

By Proposition 6, if X and Y are Z-tight, then $X \sqcup Y$ and $X \sqcap Y$ (or equivalently $X \cup Y$ and $X \cap Y$ here) are Z-tight for any $X \subseteq Z$ and $Y \subseteq Z$. Hence, for each $v \in Z \setminus S_Z$, there is a unique minimal Z-tight set containing v, which is denoted by $D_Z(v)$. Note that f_Z is a submodular function over 2^Z, and hence $D_Z(v)$ can be computed in polynomial time for any v using a standard submodular function minimization algorithm.

Now construct a directed graph $(V, \bigcup_{Z \in \mathcal{E}} E_Z)$, where a directed arc set E_Z is defined by $E_Z = \{uv \mid u, v \in Z \setminus S_Z, u \in D_Z(v)\}$ for each $Z \in \mathcal{E}$. 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 it satisfies the following four conditions:

- the directions of the arcs alternate along the walk,
- consecutive arcs have different colors,
- the last arc is a forward arc and its head (the last vertex of the walk) is in $\bigcup_{Z \in \mathcal{E}} S_Z$, and

- the first arc is a backward arc and its head (the first vertex of the walk) is in $\bigcup_{Z\in\mathcal{E}} S_Z$.

See Figure 2(a) for an example. One can show that, if the value of x is alternately increased and decreased by a small amount through an augmenting walk, then $x(V)$ is increased while maintaining $x \in P(f, \mathcal{C})$. Therefore, if x attains $\max\{x(V) \mid x \in P(f, \mathcal{C})\}$, then there is no augmenting walk.

We say that a walk consisting of arcs in $\bigcup_{Z\in\mathcal{E}} E_Z$ is a **partial augmenting walk** (abbreviated as **PAW**) if it satisfies the first three conditions in the definition of augmenting walks. 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

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

Now suppose that x attains $\max\{x'(V) \mid x' \in P(f, \mathcal{C})\}$, and hence there is no augmenting walk. For each $Z \in \mathcal{E}$, define $A_Z, B_Z \subseteq Z$ by

$$A_Z = Q_Z \cup \bigcup_{Z' \neq Z} ((R_{Z'} \cup S_{Z'}) \cap Z), \quad \text{and} \quad B_Z = Z \setminus \bigcup_{Z' \neq Z} A_{Z'}.$$

One can show that, if there is no augmenting walk, then $A_Z \cap A_{Z'} = \emptyset$ for any $Z' \neq Z$, and hence $A_Z \subseteq B_Z$ for each Z. We claim that A_Z and B_Z are Z-tight. The Z-tightness of A_Z follows by showing that there is no arc in E_Z entering A_Z. To see this, assume E_Z has an edge uv with $u \notin A_Z$ and $v \in A_Z$. If $v \in R_{Z'} \cup S_{Z'}$ for some Z' with $Z' \neq Z$, then there is a forward PAW starting from u and hence $u \in Q_Z$, contradicting $u \notin A_Z$. If $v \in Q_Z$, then, by the definition of Q_Z, E_Z has an edge vw with $w \in R_{Z'} \cup S_{Z'}$ for some $Z' \neq Z$. Since $v \in D_Z(w)$ and $u \in D_Z(v)$, we get $u \in D_Z(w)$, implying $uw \in E_Z$ or $u = w$. By $w \in R_{Z'} \cup S_{Z'}$, we get $u \in Q_Z \cup R_{Z'} \cup S_{Z'}$, which is a contradiction. Thus we have shown that there is no arc in E_Z entering A_Z. Observe finally that $S_Z \cap A_Z = \emptyset$ since otherwise there would be an augmenting walk. Therefore, we have $\bigcup_{v\in A_Z} D_Z(v) = A_Z$, which is Z-tight as each $D_Z(v)$ is Z-tight. A similar argument works for showing the Z-tightness of B_Z.

Since $\{(A_Z, B_Z) : Z \in \mathcal{E}\}$ is a 2-cover, the inequality holds with equality in (6). Thus the proof is completed for the special case where every $Z \in \mathcal{E}$ has no vertex with multiplicity two.

3.5 A proof of Theorem 11

Let us begin proving Theorem 11. Let $x \in P(f, \mathcal{C})$ and $Z \in \mathcal{E}$. Recall that a multiset $Y \subseteq Z$ is called **(x, Z)-tight** (or, simply **Z-tight** if x is clear) if $x(Y) = f_Z(Y)$, and recall also that S_Z is defined by (7). 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 $v = v_Z^\circ \in S_Z$, then we can increase $x(v)$ while maintaining $x \in P(f, \mathcal{C})$, and hence we assume that this is not the case.

By Proposition 6, for any $X \subseteq Z$ and $Y \subseteq Z$, if X and Y are Z-tight then $X \sqcup Y$ and $X \sqcap Y$ are Z-tight. Hence, for each $v \in Z \setminus (S_Z \cup \{v_Z^\circ\})$, there is a unique minimal

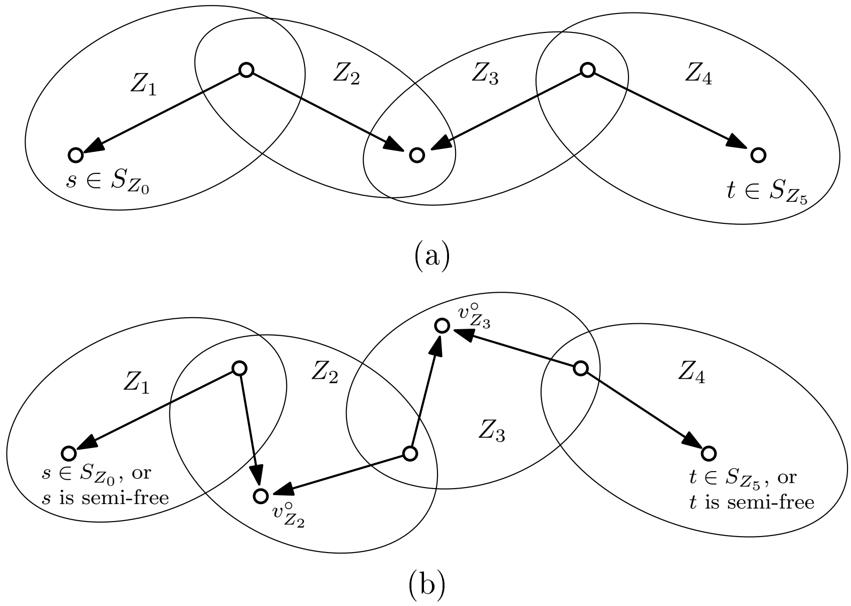

Figure 2: Examples of augmenting walks: (a) the case when there is no vertex of multiplicity two and (b) general case.

Z-tight multiset containing v, which is denoted by $D_Z(v)$. For v°_Z, if there is no Z-tight Y with $m(Y) = 2$, then a minimal Z-tight set containing v°_Z is also unique and has a unique shade, since otherwise there would be Z-tight sets X and Y whose shade sets are not singly identical and $X \sqcup Y$ would be Z-tight with $m(X \sqcup Y) = 2$. We say that v°_Z is **semi-free** if there is no Z-tight Y with $m(Y) = 2$, and the minimal Z-tight set is denoted by $D_Z(v^\circ_Z)$. Although f_Z is not an ordinary submodular set function, it can be minimized by a standard submodular function minimization algorithm, see Remark 9. This implies that $D_Z(v)$ can be computed in polynomial time for any v. If v°_Z is not semi-free, then the minimal Z-tight multiset Y with $m(Y) = 2$ is unique, which is denoted by D°_Z. This can also be computed in polynomial time by standard submodular function minimization.

To describe possible modifications of x, we use the technique of alternating walks as shown in the last subsection. When 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 the union of E^0_Z, E^1_Z and E^2_Z, defined as follows. The arc set E^0_Z, consisting of arcs not incident to v°_Z, is defined by

$$E^0_Z = \{uv \mid u, v \in Z \setminus (S_Z \cup \{v^\circ_Z\}),\ u \in D_Z(v) \setminus \{v\}\}.$$

The arc sets E^1_Z and E^2_Z consist of arcs incident to v°_Z, and their definitions depend on whether v°_Z is semi-free:

- If v°_Z is semi-free, let

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

- If v_Z° is not semi-free, let

$$E_Z^1 = \{uv_Z^\circ \mid u \in D_Z^\circ \setminus \{v_Z^\circ\} : \exists Z\text{-tight } v_Z^\circ \bar{u}\text{-multiset } Y'\} \\ \cup \{v_Z^\circ v \mid v \in Z \setminus \{v_Z^\circ\} : m(D_Z(v)) = 1\},$$

$$E_Z^2 = \{uv_Z^\circ \mid u \in D_Z^\circ \setminus \{v_Z^\circ\} : \nexists Z\text{-tight } v_Z^\circ \bar{u}\text{-multiset } Y'\} \\ \cup \{v_Z^\circ v \mid v \in Z \setminus \{v_Z^\circ\} : m(D_Z(v)) = 2\}.$$

An arc in E_Z^2 is called a **special arc**.

As above, arcs in E_Z will be referred to as arcs of **color** Z. Note that there may exist arcs having the same initial and terminal vertices but distinct colors. Those arcs are regarded as distinct arcs, and hence the resulting digraph on V may have parallel arcs.

We also introduce a **label** with each arc in E_Z^1 for the definition of augmenting walks. This labeling will be defined, based on the following lemma.

Lemma 12. *If $v_Z^\circ u \in E_Z^1$, then $D_Z(u)$ has a unique shade; furthermore, if a Z-tight set Y with $m(Y) = 1$ contains u, then it has the same shade as $D_Z(u)$.*

If v_Z° is not semi-free and $uv_Z^\circ \in E_Z^1$, then there is a unique minimal Z-tight $v_Z^\circ \bar{u}$-set Y with $m(Y) = 1$ that has a unique shade; furthermore any Z-tight $v_Z^\circ \bar{u}$-set Y' with $m(Y') = 1$ has the same shade as Y.

Proof. For the first claim, suppose that $|\mathcal{S}(D_Z(u))| \geq 2$. Then $D_Z(u) \sqcap D_Z(u)$ is a Z-tight $u\bar{v_Z^\circ}$-set by the definition of $\sqcap$, which contradicts $v_Z^\circ u \in E_Z^1$. If Y is a Z-tight set with $m(Y) = 1$ that contains u and $\mathcal{S}(Y)$ is not singly identical to $\mathcal{S}(D_Z(u))$, then $Y \sqcap D_Z(u)$ is a Z-tight $u\bar{v_Z^\circ}$-set, again a contradiction.

For the second claim, suppose that there are Z-tight $v_Z^\circ \bar{u}$-sets X and Y with $m(X) = m(Y) = 1$ and $|\mathcal{S}(X) \cup \mathcal{S}(Y)| \geq 2$. Then as $X \sqcup Y$ avoids u, $(X \sqcup Y) \sqcap D_Z^\circ$ is a Z-tight multiset smaller than D_Z° with $m((X \sqcup Y) \sqcap D_Z^\circ) = 2$, contradicting the minimality of D_Z°. Hence all such sets have the same unique shade, and their intersection is the unique minimal Z-tight $v_Z^\circ \bar{u}$-set. □

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

$$\ell(e) = \begin{cases} \text{the shade of } D_Z(v) & \text{if } e = v_Z^\circ v, \\ \text{the shade of } D_Z(v_Z^\circ) & \text{if } v_Z^\circ \text{ is semi-free and } e = vv_Z^\circ, \\ \text{the shade of the minimal } Z\text{-tight } v_Z^\circ \bar{v}\text{-multiset} & \text{if } v_Z^\circ \text{ is not semi-free and } e = vv_Z^\circ. \end{cases}$$

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 of the walk 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^1_Z and the shared vertex is v°_Z.

Note also that a PAW may use an arc in E^2_Z twice consecutively. We also remark the following.

Claim 13. *Neither a free vertex nor a semi-free vertex can be an intermediate vertex of a PAW.*

Proof. By definition a free vertex in Z is incident to no arc in E_Z, and hence it cannot be an intermediate vertex of a PAW by the third property.

If v°_Z is a semi-free vertex, then every arc incident to v°_Z has the same label. Thus it cannot be an intermediate vertex of a PAW by the fourth property. □

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

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

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

With this definition for PAWs, an **augmenting walk** is defined to be a PAW of type (i) or (ii) below that does not have any shortcut:

(i) a backward PAW starting at a free vertex $v \in S_Z$ by an arc in $E_{Z'}$ with $Z' \neq Z$;

(ii) a backward PAW starting at a semi-free vertex.

See Figure 2 (b) for a simple example. The length of an augmenting walk is bounded as follows.

Claim 14. *Each vertex appears at most four times in any 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_iv for $1 \leq i \leq 6$ be those incoming arcs at v in the order of the walk. Note that each of the pairs of arcs u_1v and u_2v, u_3v and u_4v, and u_5v and u_6v is consecutive in the walk. Note also that $u_i = u_j$ may hold.

If u_1v and u_6v have different colors or different labels, then there is a shortcut using u_6v next to u_1v. Hence they should have the same color and the same label. Similarly the colors and the labels of u_1v and u_4v should be the same. Thus, u_4v and u_6v have the same color and label. This implies that there is a shortcut using u_6v next to u_3v, 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

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

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 when $x(V)$ is maximized.

Claim 15. *If there exists an augmenting walk W, then the augmentation of x through W by sufficiently small $\varepsilon > 0$ maintains that $x \in P(f, \mathcal{C})$.*

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 multiset 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 $x(Y)$ 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 four cases occurs.

(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 by Lemma 12. 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 $wv_Z^\circ \in W_Z^2$, then one can reach a contradiction as in case (iii). Hence it holds that $uv_Z^\circ \in W_Z^1$ and $wv_Z^\circ \in W_Z^1$. Since $u \notin Y$ and $w \notin Y$, the labels of uv_Z° and wv_Z° are equal to the unique shade of Y by Lemma 12, 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 $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}$ from the sets Q_Z, R_Z and S_Z such that A_Z and B_Z are Z-tight for each $Z \in \mathcal{E}$, i.e., an optimal 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}$. To this end we need the following two claims.

Claim 16. *Suppose that Z contains three distinct elements u, v, w with $uv, vw \in E_Z$. Then the following statements hold.*

(i) If $u = v_Z^\circ$ and $uv \in E_Z^2$, then $uw \in E_Z^2$.

(ii) If $u = v_Z^\circ$ and both uv and uw are in E_Z^1, then they have the same label.

(iii) If $w = v_Z^\circ$ and $vw \in E_Z^2$, then $uw \in E_Z^2$.

(iv) If $w = v_Z^\circ$ and both vw and uw are in E_Z^1, then they have the same label.

(v) If $uw \notin E_Z$, then it holds that $v = v_Z^\circ$, and both uv_Z° and $v_Z^\circ w$ belong to E_Z^1 with the same label.

Proof. First suppose that v_Z° is neither v nor w. Then $v \in D_Z(w)$ and any Z-tight multiset containing v should contain u. Hence $u \in D_Z(w)$ and $uw \in E_Z$ holds. Suppose further that $u = v_Z^\circ$. If $uv \in E_Z^2$, then $u \in D_Z(v) \subseteq D_Z(w)$ and $m(D_Z(w)) = 2$, which implies $uw \in E_Z^2$. Thus (i) holds. 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 by Lemma 12, implying (ii).

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

Finally, to complete the proof of (v), suppose that $v = v_Z^\circ$ and $u \notin D_Z(w)$. Since uv exists, it holds that $m(D_Z(w)) = 1$ and $uv \in E_Z^1$. Hence there is a unique minimal Z-tight $v_Z^\circ\bar{u}$-set $Y' \subseteq D_Z(w)$ with $m(Y') = 1$. Since Y' has the same shade as $D_Z(w)$ by Lemma 12, we conclude that 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 17. *If no augmenting walk exists for $x \in P(f, \mathcal{C})$, then the following statements hold for each $Z \in \mathcal{E}(\mathcal{C})$, where Z' and Z'' denote members of $\mathcal{E}(\mathcal{C})$ 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 all of the initial arcs of the partial augmenting walks starting from v_Z° 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* $v^\circ_Z v \in E^2_Z$*, then* $v \notin Q_Z \cup R_{Z'} \cup S_{Z'}$.

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

(h) *If* $uv \in E^0_Z$*, 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. We first prove (a). If $v \in Q_Z \cap Q_{Z'}$, then there are forward PAWs starting 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, one can find an augmenting walk if $R_Z \cap R_{Z'} \neq \emptyset$ or $S_Z \cap R_{Z'}$. 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 $v \neq v^\circ_Z$. Let W_1 and W_2 be forward/backward PAWs starting from v. Recall that a PAW does not pass through a (semi-)free vertex by Claim 13. Therefore, when tracing W_1 and W_2 from v to the ends, there is a vertex v' such that the vertices succeeding v' in the two walks are distinct. Note that $v' \in Q_{Z'} \cap R_{Z'}$ holds for some Z'. Denote the vertex preceding v' by v'' (if $v' \neq v$), and denote the vertices succeeding v' in the walks W_1 and W_2 by u_1 and u_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 16. If neither $u_1 \neq v^\circ_{Z'}$ nor $u_2 \neq v^\circ_{Z'}$, one can find an augmenting walk by concatenating a part of W_1 from u_1, edge u_1u_2, and a part of W_2 from u_2. This is a contradiction. If $u_1 = v^\circ_{Z'}$ (resp. $u_2 = v^\circ_{Z'}$), then Claim 16 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 one can easily find an augmenting walk using u_1u_2, a contradiction.

Therefore suppose that $v' = v^\circ_{Z'} \neq v$. If u_1v' and $v'u_2$ have different labels, then $u_1u_2 \in E^0_{Z'}$ by Claim 16 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 by the definition of labels. Since u_1v' and $v''v'$ have different labels as they are consecutive in W_1, we have $v'' \in X$. However, this 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. This completes the proof of (b).

We next prove (c). Clearly the initial 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 initial arc uv°_Z of a backward PAW and the initial arc $v^\circ_Z w$ of a forward PAW have different labels. If $u \neq w$, then $uw \in E^0_Z$ by Claim 16, 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 the former claim of (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 from v. If $u \neq v^\circ_Z$, then by Claim 16 $uv^\circ_Z \in E^2_Z$, and thus an augmenting walk exists by using uv°_Z consecutively. 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 16, since $vv_Z^\circ \in E_Z^2$, $vw \in E_Z$ 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 $v_Z^\circ \in R_Z$, then $v \in Q_Z$. On the other hand, if $v_Z^\circ \in Q_Z$, then there is $v_Z^\circ w \in E_Z$ with $w \in R_{Z'}$. When $w \neq v$, $vw \in E_Z$ follows from Claim 16 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 $v_Z^\circ \in R_Z$. If $v \in R_Z$, then let $uv \in E_Z$ be the initial arc of a backward PAW starting from v. If $u \neq v_Z^\circ$, then $u \in Q_{Z'}$ and $uv_Z^\circ \in E_Z$ by Claim 16. Hence $v_Z^\circ \in R_Z$ holds. On the other hand, if $u = v_Z^\circ$, then $v_Z^\circ \in Q_Z$ 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 initial arc of a backward PAW starting from v. If $u \neq v_Z^\circ$, then there is an augmenting walk ending at v_Z° since uv_Z° exists. If $u = v_Z^\circ$, then v_Z° is an intermediate vertex of a PAW, which is not possible by Claim 13 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 $v_Z^\circ v$. If $v \in Q_Z$, then let vu be the initial arc of a forward PAW starting from v. If $u \neq v_Z^\circ$, then $u \in R_{Z'}$ and $v_Z^\circ u \in E_Z^2$ by Claim 16. Hence there is an augmenting walk that uses $v_Z^\circ u$. If $u = v_Z^\circ$, then $v_Z^\circ \in R_Z$, and there is $wv_Z^\circ \in E_Z$ with $w \in Q_{Z'}$. Then, by Claim 16, $wv \in E_Z$ 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 define a 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ attaining the minimum in (5). Suppose that no augmenting walk exists. For $Z \in \mathcal{E}(\mathcal{C})$, let

$$V_Z^\circ = \begin{cases} \{v_Z^\circ\} \text{ (with multiplicity one)} & \text{if } v_Z^\circ \text{ is semi-free or } v_Z^\circ \in R_Z, \\ \emptyset & \text{otherwise.} \end{cases}$$

Now define multisets A_Z and B_Z by

$$\begin{aligned} A_Z &= Q_Z \cup V_Z^\circ \cup \bigcup_{Z' \neq Z} ((R_{Z'} \cup S_{Z'}) \cap Z), \\ B_Z &= Z \setminus \bigcup_{Z' \neq Z} A_{Z'}, \end{aligned} \tag{8}$$

where B_Z contains v_Z° with multiplicity two if A_Z does not contain v_Z°, otherwise both A_Z and B_Z contain it with multiplicity one.

Properties (a) and (b) of Claim 17 imply that $A_Z \cap A_{Z'} = \emptyset$ for any distinct Z and Z'. Thus $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ defined by (8) forms a 2-cover, which we call the **canonical 2-cover**.

The fact that the canonical 2-cover attains the minimum in (5) follows from the following claim.

Claim 18. *If no augmenting walk exists for $x \in P(f, \mathcal{C})$, then the following hold for the multisets A_Z and B_Z defined in (8) for each $Z \in \mathcal{C}$:*

- *$A_Z \subseteq B_Z$;*
- *A_Z and B_Z are (x, Z)-tight;*

- *If A_Z contains v_Z° then $\mathcal{S}(A_Z)$ and $\mathcal{S}(B_Z)$ are singly identical.*

Proof. The first claim follows from the definition (8) and the fact that $A_Z \cap A_{Z'} = \emptyset$ for any distinct Z and Z'.

We prove the second claim. In the proof we denote by $d_Z^i(Y)$ the number of arcs in E_Z^i incoming to Y for each $Y \subseteq Z$, and denote by $d_Z(Y)$ the number of arcs in E_Z incoming to Y. We first prove that

$$A_Z \text{ is } Z\text{-tight, and it has a single shade if } m(A_Z) = 1. \tag{9}$$

To see this we prove

$$d_Z^2(A_Z) = 0. \tag{10}$$

Suppose to the contrary that $d_Z^2(A_Z) > 0$. Then v_Z° is not semi-free by definition of E_Z^2, and $v_Z^\circ \in A_Z$ implies $v_Z^\circ \in Q_Z \cup R_Z$ by definition of A_Z. However, Claim 17(d) implies $v \in Q_Z \cup R_{Z'} \subseteq A_Z$ for any $vv_Z^\circ \in E_Z^2$, contradicting $d_Z^2(A_Z) > 0$. Thus (10) holds. Similarly, it follows from Claim 17(h) that

$$d_Z^0(A_Z) = 0. \tag{11}$$

If $v_Z^\circ \notin A_Z$, then $d_Z^1(A_Z) = 0$ also holds by Claim 17(g), and hence $d_Z(A_Z) = 0$. This means that $D_Z(v) \subseteq A_Z$ for every $v \in A_Z$, so $A_Z = \bigcup_{v \in A_Z} D_Z(v)$ and it is Z-tight by submodularity, implying (9). Similarly, if v_Z° is semi-free, then $d_Z^1(A_Z) = 0$ holds and $A_Z = \bigcup_{v \in A_Z} D_Z(v)$ implies that A_Z is Z-tight, and it has a single shade if $m(A_Z) = 1$. Thus (9) holds for this case.

The remaining case for proving (9) is when $v_Z^\circ \in A_Z$ and v_Z° is not semi-free. Then by definition of A_Z we have $v_Z^\circ \in Q_Z \cup R_Z$. We prove that

$$\begin{aligned}&\text{there is a } Z\text{-tight set } Y \subseteq A_Z \text{ with } m(Y) = 1 \text{ and } |\mathcal{S}(Y)| = 1,\\ &\text{and for every } u \in A_Z \setminus \{v_Z^\circ\} \text{ with } v_Z^\circ \in D_Z(u),\ m(D_Z(u)) = 1 \text{ holds}\\ &\text{and } \mathcal{S}(D_Z(u)) \text{ and } \mathcal{S}(Y) \text{ are singly identical.}\end{aligned} \tag{12}$$

For (12) we have two subcases.

Suppose that $v_Z^\circ \in Q_Z$. Then there is $v_Z^\circ v \in E_Z$ with $v \in R_{Z'} \cup S_{Z'}$. By Claim 17(f) $v_Z^\circ v \in E_Z^1$ holds, and hence $m(D_Z(v)) = 1$ and $|\mathcal{S}(D_Z(v))| = 1$. By (11) we have $D_Z(v) \subseteq A_Z$. Hence let $Y = D_Z(v)$. To see (12), take $u \in A_Z$ with $v_Z^\circ \in D_Z(u)$. By $v_Z^\circ \in Q_Z$, $v_Z^\circ u$ is contained in some partial augmenting walk, and hence by Claim 17(c) $v_Z^\circ u$ has the same label as $v_Z^\circ v$, that is, $m(D_Z(u)) = 1$ holds and the shade of $D_Z(u)$ is the same as the shade of $D_Z(v) (= Y)$, implying (12).

Suppose next that $v_Z^\circ \in R_Z \setminus Q_Z$. Then there is $vv_Z^\circ \in E_Z$ with $v \in Q_{Z'}$. By Claim 17(d), $vv_Z^\circ \in E_Z^1$ holds, and hence there is a unique smallest Z-tight $v_Z^\circ\bar{v}$-set Y with $m(Y) = 1$, which has a single shade by Lemma 12. Note that $u \in Y$ implies $u \in Q_Z \subseteq A_Z$, and hence $Y \subseteq A_Z$. To see (12), take $u \in A_Z$ with $v_Z^\circ \in D_Z(u)$. By Claim 17(f) and $u \in A_Z$, $v_Z^\circ u \notin E_Z^2$ holds, and hence $m(D_Z(u)) = 1$ should holds by definition of E_Z^1. If $\mathcal{S}(D_Z(u))$ and $\mathcal{S}(Y)$ are not singly identical, then arc vu exists in E_Z by Claim 16 and thus $u \in R_Z$, which contradicts $u \in A_Z$ by Claim 17(a)(b). Thus we have shown (12).

By (11) and (12), $A_Z = Y \cup \bigcup_{u \in A_Z \setminus \{v_Z^\circ\}} D_Z(u)$, $m(A_Z) = 1$ and $|\mathcal{S}(A_Z)| = 1$. This proves that A_Z is Z-tight and it has a single shade, and the proof of (9) is completed.

A proof of the fact that B_Z is Z-tight is given as follows. First we claim

$$d_Z^0(B_Z) = 0. \tag{13}$$

Indeed, 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 $Z'' \neq Z$.

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 Claim 17(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 Claim 17(e). Therefore $B_Z = \bigcup_{u \in B_Z} D_Z(u)$, and its single shade is that of $D_Z(v_Z^\circ)$.

It remains to consider the case when $v_Z^\circ \in Q_Z \cup R_Z$. We prove that

$$\begin{aligned}&\text{for any } u \in B_Z \setminus A_Z \text{ with } v_Z^\circ \in D_Z(u),\ m(D_Z(u)) = 1,\\ &\text{and } \mathcal{S}(D_Z(u)) \text{ and } \mathcal{S}(A_Z) \text{ are singly identical.}\end{aligned} \tag{14}$$

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 $u \in B_Z$.

On the other hand, suppose that $v_Z^\circ \in R_Z \setminus Q_Z$. Then there is $vv_Z^\circ \in E_Z^1$ with $v \in Q_{Z'}$. Note that A_Z is a Z-tight $v_Z^\circ \bar{v}$-tight set and hence the shade of A_Z is equal to the label of $vv_Z^\circ \in E_Z^1$ by Lemma 12. Also, 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 16 and $u \in R_Z$ follows, contradicting $u \in B_Z$. Hence the shade of $D_Z(u)$ is the shade of A_Z, and we have shown (14). This completes the proof of the second claim.

By (13) and (14), $B_Z = A_Z \cup \bigcup_{u \in B_Z \setminus A_Z} D_Z(u)$, and $\mathcal{S}(B_Z)$ is singly identical to $\mathcal{S}(A_Z)$. This completes the proof of the third claim. □

Now it is not difficult to see that the canonical 2-cover attains equality in (5).

Proof of Theorem 11. For a maximizer x in (5) and the canonical 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}$, A_Z and B_Z are (x, Z)-tight for each $Z \in \mathcal{E}$ by Claim 18. Thus the inequality in (6) 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 $m = |U|$ and $M = \max_T |f(T)|$. 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 19. *Let $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\rm top}$ 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 $\mathcal{C}^* : T_1 \prec \ldots T_{j-1} \prec T^* \prec T_j \prec \cdots \prec T_k = T_{\rm top}$. 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 $T \in [T_{j-1}, T_j]$, which implies $x \in P(f,\mathcal{C})$. 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 while keeping the feasibility. The algorithm terminates if it finds $x \in \operatorname{argmax}\{x'(V) \mid x' \in P(f,\mathcal{C})\}$ satisfying $2x(V) = f(T_{\rm top})$, 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 20 below), while in the latter case we can conclude that $P^=(f) = \emptyset$.

Lemma 20. *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* (LP$^=$) *and* (D)*, respectively, if and only if $2x(V) = f(T_{\rm top})$.*

Proof. Clearly $2x(V) = f(T_{\rm top})$ should hold if x is optimal for (LP$^=$). Suppose that $2x(V) = f(T_{\rm top})$ holds. By Lemma 19, 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 $2x(V) = \sum_{j=1}^k (a(T_j) - a(T_{j-1}))x \leq \sum_{j=1}^k (f(T_j) - f(T_{j-1})) = f(T_{\rm top}) = 2x(V)$, 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. □

The details of the algorithm are described in Algorithm 1. The first two lines are the **initialization step**, and the main **iteration** starts from Line 3.

The correctness of Algorithm 1 follows from Theorem 11 combined with Lemmas 21 and 22 below. Proofs of Lemmas 21 and 22 will be given in Section 4.3.

Lemma 21. *Let x and $y^{\bar{\varepsilon}}$ be those at the end of an iteration (Line 14) for a given y, and let $\bar{\mathcal{C}}$ be the chain corresponding to the support of y. Then, it holds that $x \in P(f,\bar{\mathcal{C}})$.*

Algorithm 1 Pseudo-polynomial time algorithm

Input: A submodular function f over the direct product of diamonds $U_1, \ldots, U_n$ and $c \in \mathbb{Z}^n$ with $c_1 > c_2 > \cdots > c_n$.

Output: An optimal solution x of (LP$^=$) or conclude $P^=(f) = \emptyset$.

1: For each $j = 1, \ldots, n$, let T'_j be a transversal defined by $T'_j \cap U_i = \{1_i\}$ for $i = 1, \ldots, j$ and $T'_j \cap U_i = \{0_i\}$ for $i = j+1, \ldots, n$.

2: Let y be a dual feasible solution defined by

$$y_T = \begin{cases} \frac{c_j - c_{j+1}}{2} & \text{if } T = T'_j \text{ for some } j \in \{1, \ldots, n\}, \\ 0 & \text{otherwise,} \end{cases}$$

where $c_{n+1} = 0$.

3: **loop**

4: Let $\mathcal{C} : T_1 \prec \cdots \prec T_k = T_{\text{top}}$ be a chain corresponding to the support of y.

5: **if** $x(V) < \max\{x'(V) \mid x' \in P(f, \mathcal{C})\}$ *($x(V) = -\infty$ at the initial iteration)* **then**

6: Let x be a point in $\operatorname{argmax}\{x'(V) \mid x' \in P(f, \mathcal{C})\}$

7: **if** $2x(V) = f(V)$ **then**

8: **return** x

9: Compute the canonical 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ with respect to x.

10: Take $T_j^A \in \mathcal{T}(A_{Z_j})$ and $T_j^B \in \mathcal{T}(B_{Z_j})$ for each $Z_j \in \mathcal{E}(\mathcal{C})$.

11: Compute $\bar{\varepsilon} = \sup\{\varepsilon \in \mathbb{R}_+ \mid y_T^\varepsilon \geq 0 \ \forall T \in \mathcal{T} \setminus \{T_{\text{top}}\}\}$, where

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

12: **if** $\bar{\varepsilon} = +\infty$ **then**

13: **return** "$P^=(f) = \emptyset$ ".

14: $y \leftarrow y^{\bar{\varepsilon}}$.

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

Lemma 22. *Let $\mathcal{C} : T_1 \prec \cdots \prec T_k$, x, A_Z, B_Z, and $y^{\bar{\varepsilon}}$ are those at the end of an iteration (Line 14) for a given y, 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 $\mathcal{C}$ in the next iteration for $\bar{y} := y^{\bar{\varepsilon}}$. Suppose that $x = \bar{x}$ and the canonical 2-cover $\{(A_Z, B_Z) \mid Z \in \mathcal{E}(\mathcal{C})\}$ with respect to x is not trivial. Then it holds that*

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

If this inclusion relation holds with equality, then

$$\sum_{j=1}^{\bar{k}} \tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \sum_{j=1}^{\bar{k}} \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}| < \sum_{j=1}^{k} \tilde{a}(T_{j-1})|A_{Z_j}| - \sum_{j=1}^{k} \tilde{a}(T_j)|Z_j \setminus B_{Z_j}|. \quad (15)$$

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

Lemma 23. *Let f be an integer-valued submodular function on the product of n diamonds. Then Algorithm 1 terminates after $O(n^4M)$ iterations and outputs an optimal solution for* (LP$^=$) *or verifies that* (LP$^=$) *is infeasible, where $M = \max_T |f(T)|$.*

Proof. One can easily check that the vector y obtained at the initialization step is a feasible solution for the dual program (D). For a feasible solution y given in an iteration, let x be the vector defined at Line 6. Then $x \in \mathrm{argmax}\{x'(V) \mid x' \in P(f, \mathcal{C})\}$ for the support chain $\mathcal{C}$ of y. Note that $x(V) \geq -4nM$ by Theorem 11.

If $2x(V) = f(T_{\mathrm{top}})$, then x is an optimal solution for (LP$^=$) by Lemma 20.

If $2x(V) < f(T_{\mathrm{top}})$, then by Theorem 11 the canonical 2-cover is nontrivial. (Note that the value of the trivial 2-cover is $f(T_{\mathrm{top}})$.) Furthermore, y^ε is feasible for any $0 \leq \varepsilon \leq \bar{\varepsilon}$ by Lemma 8, and hence (D) is unbounded if $\bar{\varepsilon}$ is unbounded.

Assume that $\bar{\varepsilon}$ is bounded. Let $\bar{y} = y^{\bar{\varepsilon}}$, let $\bar{\mathcal{C}}$ be the support chain of $\bar{y}$, and let $\bar{x} \in \mathrm{argmax}\{x'(V) \mid x' \in P(f, \bar{\mathcal{C}})\}$. By Lemma 21, we have $x(V) = \max\{x(V) \mid x \in P(f, \mathcal{C})\} \leq \max\{x(V) \mid x \in P(f, \bar{\mathcal{C}})\} = \bar{x}(V)$. Theorem 11 implies that $x(V)$ and $\bar{x}(V)$ are half-integral. Thus $\bar{x}(V) > x(V)$ implies $\bar{x}(V) \geq x(V) + 1/2$. By Lemma 22, $\bar{x}(V) > x(V)$ occurs after $O(n^3)$ iterations. Therefore, after $O(n^4M)$ iterations, $2x(V)$ attains $f(T_{\mathrm{top}})$, which implies that x is optimal by Lemma 20. □

As we remarked in Section 3, $x \in \mathrm{argmax}\{x'(V) \mid x' \in P(f, \mathcal{C})\}$ 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 function $f^{(t)} : \mathcal{T} \to \mathbb{Z}$. For a function $f : \mathcal{T} \to \mathbb{Z}$ and a nonnegative integer t, define $f^{(t)} : \mathcal{T} \to \mathbb{Z}$ by

$$f^{(t)}(T) = \left\lceil \frac{f(T)}{2^t} \right\rceil - \tilde{a}(T)^2 \qquad (T \in \mathcal{T}). \tag{16}$$

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 idea to work.

Lemma 24. *For a submodular function $f : \mathcal{T} \to \mathbb{Z}$ and a nonnegative integer t, it holds that*

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

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

(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 that T and T' have 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 $\tilde{a}(\cdot)^2$ and $\tilde{a}$, we get the desired relation.

If there is no such 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

$$-1 \leq g^{(t)}(T) - 2g^{(t+1)}(T) \leq 0. \tag{17}$$

Let $\mathcal{C} : T_1 \prec T_2 \prec \cdots \prec T_k = T_{\text{top}}$. We show $2f_{Z_j}^{(t+1)}(Y) \leq f_{Z_j}^{(t)}(Y)$ for each $Y \subseteq Z_j$. Take $T_Y \in \mathcal{T}$ such that $f_{Z_j}^{(t)}(Y) = f^{(t)}(T_Y) - f^{(t)}(T_{j-1})$. If $Y \neq \emptyset$, then $2f_{Z_j}^{(t+1)}(\emptyset) = f_{Z_j}^{(t)}(\emptyset) = 0$ holds by definition. Otherwise, $\tilde{a}(T_Y)^2 \geq \tilde{a}(T_{j-1})^2 + 1$ holds, and hence by (17) we have

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

Proof of (iii): We prove that $P^{=}(f^{(t)})$ is not empty if $P^{=}(f)$ is not empty. Suppose that $x \in P^{=}(f)$. Let $r = \lceil f(T_{\text{top}})/2^t \rceil - f(T_{\text{top}})/2^t$ and let $x' = x/2^t - (2n - r/2n)\mathbf{1}$, where $\mathbf{1}$ denotes the all-one vector. 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 $x' \in P^{=}(f^{(t)})$. □

On the basis of Lemma 24, we establish our scaling algorithm with running time polynomial in m and $\log M$, which appears in the proof of the following theorem.

Theorem 25. *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 $t = \lceil \log M \rceil$.
- Find $x \in P^{=}(f^{(t)})$ maximizing cx by Algorithm 1, which can be done in time polynomial in m. If $P^{=}(f^{(t)}) = \emptyset$, then conclude that $P^{=}(f) = \emptyset$, which follows from Lemma 24 (iii). Otherwise, denote the optimal primal and dual solutions for $\max\{cx \mid x \in P^{=}(f^{(t)})\}$ by $\bar{x}$ and $\bar{y}$, respectively. Let $\bar{\mathcal{C}}$ be the support chain of $\bar{y}$; note that $\bar{x} \in P(f^{(t)}, \bar{\mathcal{C}})$.
- Solve (LP$^{=}$) for $f^{(t-1)}$ as follows. First 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 main loop of Algorithm 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 24 (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 (17). This implies that after $O(n^5)$ iterations the algorithm solves (LP$^{=}$) for $f^{(t-1)}$. (Recall that $x(V)$ increases by $1/2$ after $O(n^3)$ iterations.)
- Iterating the above process, we obtain 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 for this algorithm it is not necessary to know M in advance: the algorithm given in Section 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 $t = \lceil \log M \rceil$, we can start with the smallest t for which this property holds.

4.3 Proofs of Lemmas 21 and 22

In this subsection we shall give the proofs of Lemmas 21 and 22. This completes the proof of the correctness of our algorithm.

4.3.1 Proof of Lemma 21

Lemma 21 straightforwardly follows from the lemma below. For a chain $\mathcal{C} : T_0 \prec T_1 \prec \ldots T_k = T_{\text{top}}$ 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 26. *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 it holds that $x \in P(f, \mathcal{C}^*)$, 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 21. 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_j^A and T_j^B are $(x, \mathcal{C})$-tight for every j, and hence $x \in P(f, \mathcal{C}^*)$ by Lemma 26. As $\mathcal{C}^*$ is a refinement of $\bar{\mathcal{C}}$, $x \in P(f, \bar{\mathcal{C}})$ follows from Lemma 19. □

4.3.2 Proof of Lemma 22

In the proof of Lemma 22, we shall use the following two lemmas on properties of hypergraphs obtained by refining the underlying chain.

Lemma 27. *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 of transversals. 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 28. *Let* $\mathcal{C}^*$ *be a refinement of a chain* $\mathcal{C}$ *of transversals,* $x \in \operatorname{argmax}\{x'(V) \mid x' \in P(f, \mathcal{C})\} \cap \operatorname{argmax}\{x'(V) \mid x' \in P(f, \mathcal{C}^*)\}$*, 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 $\mathcal{C}^* : T_1 \prec \cdots \prec T_{j-1} \prec T^* \prec T_j \prec \cdots \prec T_k$. Let $Z = Z_j \in \mathcal{E}(\mathcal{C})$. By inserting T^* into $\mathcal{C}$, $Z \in \mathcal{E}(\mathcal{C})$ is decomposed into two hyperedges in $\mathcal{E}(\mathcal{C}^*)$. Let Z_1^* and Z_2^* be those hyperedges in $\mathcal{E}(\mathcal{C}^*)$ that decompose Z, in the order specified by $\mathcal{C}^*$, i.e., Z_1^* and Z_2^* correspond to intervals $T_{j-1} \preceq T^*$ and $T^* \preceq T_j$, respectively.

By Lemma 27, 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 any PAW W in E^* still forms a walk in E, and W satisfies the first three conditions for PAW in E. Therefore, if W does not pass through v_Z°, then it is a PAW in E.

We thus focus on the case when W passes through v_Z°. Suppose that W is not a PAW in E, which implies that W violates the fourth condition for PAW. We have two possibilities: (i) W passes through two consecutive arcs which are of the form

$uv_Z^\circ, vv_Z^\circ \in E_Z^1$ with the same label; and (ii) W passes through two consecutive arcs which are of the form $v_Z^\circ v, v_Z^\circ w \in E_Z^1$ with the same label. We shall show that none of these can happen.

Case (i): Since $uv_Z^\circ, vv_Z^\circ \in E_Z^1$ with the same label, there is a Z-tight set Y containing v_Z° but neither u nor v with $m(Y) = 1$.

If v_Z° exists in Z_i^* as a vertex of multiplicity two for some $i \in \{1, 2\}$ (i.e., $v_{Z_i^*}^\circ = v_Z^\circ$), then both u, v exist in Z_i^* as W is a PAW in E^*, but $Y \sqcap Z_i^*$ is Z-tight and hence $Y \sqcap Z_i^*$ would be a Z_i^*-tight set that contains v_Z° but neither u nor v. Thus $uv_{Z_i^*}^\circ, vv_{Z_i^*}^\circ \in E_Z^1$ are two consecutive arcs in W having the same label in E^*, contradicting that W is a PAW in E^*.

Therefore we may assume that v_Z° exists in both Z_1^* and Z_2^* as a vertex of multiplicity one. Then we have $uv_Z^\circ \in E_{Z_1^*}^0$ and $vv_Z^\circ \in E_{Z_2^*}^0$ as W is a PAW in E^*. Take $T_Y \in \mathcal{T}_Z(Y)$. By Lemma 27, $T_Y \wedge T^*$ and $T_Y \vee T^*$ are $(x, \mathcal{C}^*)$-tight. If T_Y and T^* have the same middle element at the diamond of v_Z°, then we have $T_Y \wedge T^* \in \mathcal{T}_{Z_1^*}(Y \cap Z_1^*)$ and hence $Y \cap Z_1^*$ is a Z_1^*-tight $\bar{u}v_Z^\circ$-set, which contradicts $uv_Z^\circ \in E_{Z_1^*}^0$. Otherwise $T_Y \vee T^* \in \mathcal{T}_{Z_2^*}(Y \cap Z_2^*)$ and hence $Y \cap Z_2^*$ is a Z_2^*-tight $\bar{v}v_Z^\circ$-set, which contradicts $vv_Z^\circ \in E_{Z_2^*}^0$. This completes the proof for Case (i).

Case (ii). The proof is similar to case (i). Since $v_Z^\circ v, v_Z^\circ w \in E_Z^1$ with the same label, $D_Z(v)$ and $D_Z(w)$ have the same shade.

If v_Z° exists in Z_i^* as a vertex of multiplicity two for some $i \in \{1, 2\}$, then $D_{Z_i^*}(v)$ and $D_{Z_i^*}(w)$ have the same shade by Lemma 27, contradicting that W is a PAW in E^*.

Suppose that v_Z° exists in both Z_1^* and Z_2^* as a vertex of multiplicity one. We have $v_Z^\circ v \in E_{Z_1^*}^0$ and $v_Z^\circ w \in E_{Z_2^*}^0$ as W is a PAW in E^*. Take $T_v \in \mathcal{T}_{Z^*}(D_Z(v))$ and $T_w \in \mathcal{T}_{Z^*}(D_Z(w))$. By Lemma 27, $T_v \wedge T^*$ and $T_w \vee T^*$ are $(x, \mathcal{C}^*)$-tight. If T_v and T^* have the same middle element at the diamond of v_Z°, then $(D_Z(w) \setminus \{v_Z^\circ\}) \cap Z_2^*$ would be a Z_2^*-tight $\overline{v_Z^\circ}w$-set since $T_w \vee T^* \in \mathcal{T}_{Z_2^*}((D_Z(w) \setminus \{v_Z^\circ\}) \cap Z_2^*)$. This contradicts $v_Z^\circ w \in E_{Z_2^*}^0$. Otherwise $(D_Z(v) \setminus \{v_Z^\circ\}) \cap Z_1^*$ would be a Z_1^*-tight $\overline{v_Z^\circ}v$-set by $T_v \wedge T^* \in \mathcal{T}_{Z_1^*}((D_Z(v) \setminus \{v_Z^\circ\}) \cap Z_1^*)$, contradicting $v_Z^\circ v \in E_{Z_1^*}^0$. □

In the following we shall use notation as given in Lemma 22. Let $\bar{\mathcal{C}} : \bar{T}_1 \prec \cdots \prec \bar{T}_{\bar{k}} = T_{\rm top}$ and $\mathcal{C}^* = \mathcal{C} \cup \bar{\mathcal{C}} : T_1^* \prec \cdots \prec T_{k^*}^* = T_{\rm top}$. 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 $\{(A_Z, B_Z) \mid Z \in \mathcal{E}\}, \{(A_{\bar{Z}}, B_{\bar{Z}}) \mid \bar{Z} \in \bar{\mathcal{E}}\}$, and $\{(A_{Z^*}, B_{Z^*}) \mid Z^* \in \mathcal{E}^*\}$ are all defined with respect to $x (= \bar{x})$.

The chain $\mathcal{C}^*$ is obtained from $\mathcal{C}$ by inserting T_j^A and T_j^B into the interval $[T_{j-1}, T_j]$ for each j, and the insertion decomposes $Z_j \in \mathcal{E}$ into three hyperedges. Those three hyperedges are corresponding to the three intervals $[T_{j-1}, T_j^A]$, $[T_j^A, T_j^B]$, and $[T_j^B, T_j]$ in $\mathcal{C}^*$, and they are denoted by $A_Z, B_Z \setminus A_Z$, and $Z \setminus B_Z$ in $\mathcal{E}^*$, respectively. (Some of the three may be empty.) Hence $\mathcal{E}^*$ can be partitioned into three subsets $\mathcal{E}_a^*, \mathcal{E}_b^*, \mathcal{E}_c^*$ such that $A_Z \in \mathcal{E}_a^*$, $B_Z \setminus A_Z \in \mathcal{E}_b^*$, and $Z \setminus B_Z \in \mathcal{E}_c^*$ for each $Z \in \mathcal{E}$. 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 according to the definition of A_Z), 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 from any $(Z \setminus B_Z)$-tight subset containing v_Z° one can construct a Z-tight subset Y with $m(Y) = 2$ by Lemma 26.

Lemma 29. *For any $Z^* \in \mathcal{C}^*$, the sets A_{Z^*} and B_{Z^*} have the following properties:*

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

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

(c) *if $Z^* \in \mathcal{E}_c^*$, 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. It follows from Lemma 26 that for each $Z = Z_j \in \mathcal{E}$,

$$\begin{aligned}&\text{a free vertex in } Z \text{ remains free in } A_Z, B_Z \setminus A_Z, \text{ or } B_Z \setminus Z \text{ in } (V, \mathcal{E}^*),\\ &\text{a semi-free vertex } v_Z^\circ \text{ in } Z \text{ remains semi-free in } B_Z \text{ in } (V, \mathcal{E}^*) \text{ if } m(A_Z) = 0,\\ &\text{a semi-free vertex } v_Z^\circ \text{ in } Z \text{ becomes free in } Z \setminus B_Z \text{ in } (V, \mathcal{E}^*) \text{ if } m(A_Z) = 1.\end{aligned} \tag{18}$$

Indeed, if v is free in Z in $(V, \mathcal{E})$ but not in $(V, \mathcal{E}^*)$ (say in A_Z), then $(V, \mathcal{E}^*)$ has an A_Z-tight set $Y \subseteq A_Z$ with $v \in Y$. Take $T_Y \in \mathcal{T}_{A_Z}(Y)$. Then T_Y is $(x, \mathcal{C}^*)$-tight, and by Lemma 26 it is also $(x, \mathcal{C})$-tight. Thus Y is Z-tight, and v is not free in Z, a contradiction. An identical argument works for the other cases, leading to (18).

We next claim

$$\text{the set of PAWs (w.r.t. } x(= \bar{x})\text{) does not change in } (V, \mathcal{E}) \text{ and } (V, \mathcal{E}^*). \tag{19}$$

If this is correct, then (18) and (19) immediately imply (a)–(c) by the definition of A_Z and B_Z. Statement (d) directly follows from (a)–(c). Thus our goal is to prove (19).

By Lemma 28, every PAW in $(V, \mathcal{E}^*)$ is also a PAW in $(V, \mathcal{E})$. It remains to show that every PAW in $(V, \mathcal{E})$ is also a PAW in $(V, \mathcal{E}^*)$.

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 any A_Z-tight set separating u from v would also be Z-tight by Lemma 26. 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 26. Thus, by (18), W remains a PAW in $(V, \mathcal{E}^*)$ if it does not pass through v_Z° in $(V, \mathcal{E})$.

If $m(A_Z) = 0$ (i.e., $m(B_Z) = 2$), then the definition of A_Z implies that v_Z° is not semi-free and v_Z° is not in $Q_Z \cup R_Z$. Thus W does not pass through v_Z°, and we are done.

We can therefore focus on the case when $m(A_Z) = 1$, i.e., $v_Z^\circ \in A_Z$. We have three cases for W to pass through 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), v_Z° is semi-free and W uses $uv_Z^\circ \in E_Z^1$. In this case we have that $u \in Q_Z$, and hence $u \in A_Z$. Furthermore, uv_Z° implies $uv_Z^\circ \in E_{A_Z}^0$ by Lemma 26. (If $uv_Z^\circ \notin E_{A_Z}^0$, then there would be an A_Z-tight set Y with $v_Z^\circ \in Y$ and $u \notin Y$, and Y would be a Z-tight set even in Z by Lemma 26, which is a contradiction.) Since v_Z° is free in $(V, \mathcal{E}^*)$ by (18), 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 $v \in Q_{Z'}$ with $Z \neq Z'$. 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 17 (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 26. Lemma 26 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 $w \in R_{Z'} \cup S_{Z'}$ with $Z \neq Z'$. 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 17 (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 26 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 are now ready to prove Lemma 22.

Proof of Lemma 22. By Lemma 28, every PAW in $(V, \mathcal{E}^*)$ is also a PAW in $(V, \bar{\mathcal{E}})$. Combining this with Lemma 29(d), we have 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 $\bigcup_{\bar{Z}\in\bar{\mathcal{E}}} A_{\bar{Z}} = \bigcup_{Z^*\in\mathcal{E}^*} A_{Z^*}$, and we prove (15). 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 $\mathcal{E}_a^*, \mathcal{E}_b^*, \mathcal{E}_c^*$. Hence suppose that $\bar{Z}_j \in \bar{\mathcal{E}}$ is obtained as the union of $Z_{j_a}^*, Z_{j_b}^*, Z_{j_c}^*$ with $Z_{j_i}^* \in \mathcal{E}_i^*$ $(i \in \{a, b, c\})$, some of which may be empty. Each $Z_{j_i}^*$ corresponds to interval $[T_{j_i-1}^*, T_{j_i}^*]$ in the chain $\mathcal{C}^*$. (See Figure 3 for an example.) By Lemmas 28 and 29 we have

$$Z_{j_a}^* = A_{Z_{j_a}^*} \subseteq A_{\bar{Z}_j}. \tag{20}$$

Moreover, since $\bigcup_{\bar{Z}\in\bar{\mathcal{E}}} A_{\bar{Z}} = \bigcup_{Z^*\in\mathcal{E}^*} A_{Z^*}$ and each $A_{\bar{Z}}$ is disjoint from the other, (20) implies

$$A_{\bar{Z}_j} = Z_{j_a}^* \quad \text{and} \quad \bar{Z}_j \setminus B_{\bar{Z}_j} = Z_{j_c}^*. \tag{21}$$

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

$$\begin{aligned} \bar{T}_{j-1} \prec T_{j_a-1}^* \quad &\text{if } Z_{j_a}^* \neq \emptyset, \\ T_{j_b}^* \prec \bar{T}_j \quad &\text{if } Z_{j_c}^* \neq \emptyset. \end{aligned} \tag{22}$$

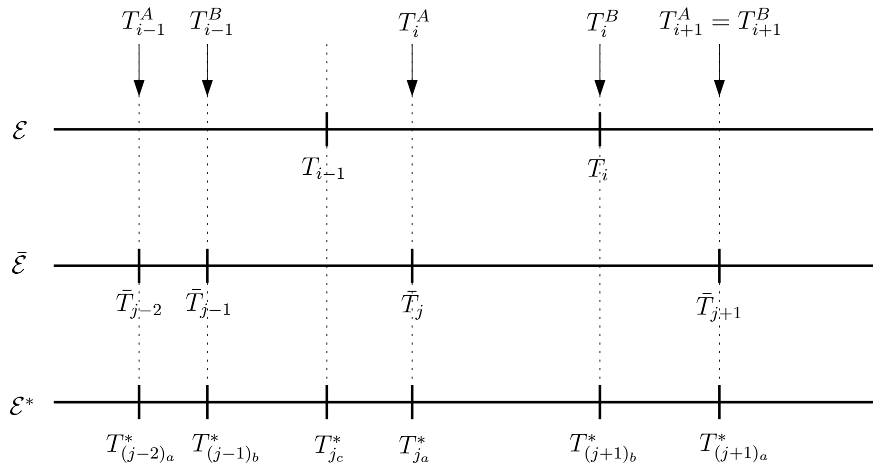

Figure 3: Proof of Lemma 22.

Moreover, at least one of $Z_{j_a} \neq \emptyset$ and $Z_{j_c} \neq \emptyset$ 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 (21) and (22), it holds that

$$\tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}| < \tilde{a}(T^*_{j_a-1})|Z^*_{j_a}| - \tilde{a}(T^*_{j_c})|Z^*_{j_c}| \tag{23}$$

for each $\bar{Z}_j \in \bar{\mathcal{E}}'$. By Lemma 29(a)–(c), we also have

$$\tilde{a}(T^*_{j_a-1})|Z^*_{j_a}| - \tilde{a}(T^*_{j_c})|Z^*_{j_c}| = \sum_{i \in \{a,b,c\}} \left\{ \tilde{a}(T^*_{j_i-1})|A_{Z^*_{j_i}}| - \tilde{a}(T^*_{j_i})|Z^*_{j_i} \setminus B_{Z^*_{j_i}}| \right\}. \tag{24}$$

Combining it with (23), we obtain, for each $\bar{Z}_j \in \bar{\mathcal{E}}'$

$$\tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}| < \sum_{i \in \{a,b,c\}} \left\{ \tilde{a}(T^*_{j_i-1})|A_{Z^*_{j_i}}| - \tilde{a}(T^*_{j_i})|Z^*_{j_i} \setminus B_{Z^*_{j_i}}| \right\}. \tag{25}$$

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 $i \in \{a, b, c\}$, and $\bar{T}_{j-1} = T^*_{j_i-1}$ and $\bar{T}_j = T^*_{j_i}$. Thus by (21) we get

$$\tilde{a}(\bar{T}_{j-1})|A_{\bar{Z}_j}| - \tilde{a}(\bar{T}_j)|\bar{Z}_j \setminus B_{\bar{Z}_j}| = \tilde{a}(T^*_{j_i-1})|A_{Z^*_{j_i}}| - \tilde{a}(T^*_{j_i})|Z^*_{j_i} \setminus B_{Z^*_{j_i}}| \tag{26}$$

for each $\bar{Z}_j \in \bar{\mathcal{E}} \setminus \bar{\mathcal{E}}'$.

Finally, note that $\bar{\mathcal{E}}' \neq \emptyset$ due to the construction of $\bar{y}$ from y in the algorithm. Therefore, summing up (25) and (26) yields

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

By Lemma 29 (d), this proves the lemma. □

5 Solving the Minimum 2-cover Problem Combinatorially

As stated in Section 3.3, $x \in \text{argmax}\{x'(V) \mid x' \in P(f, \mathcal{C})\}$ is found by the ellipsoid method, while in Section 3.4 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.

Recall that an augmentation is performed though an augmenting walk in auxiliary directed graph $(V, E = \bigcup_{Z \in \mathcal{E}} E_Z)$, where the definition of each arc set E_Z depends on x. Since x will be changed during augmentations, throughout this section, to avoid possible confusion we use $E_Z(x)$ and $E(x)$ to denote E_Z and E, respectively, with respect to x. Similarly, for $D_Z(v)$ and D_Z° we use $D_Z(v, x)$ and $D_Z^\circ(x)$ if those are defined with respect to x.

Suppose that W is an augmenting walk with the vertex sequence $v_1, v_2, \dots, v_l$ in $(V, E(x))$, and let $d = \sum_{1 \le i \le \lceil l/2 \rceil} \chi_{v_{2i-1}} - \sum_{1 \le i \le \lfloor l/2 \rfloor} \chi_{v_{2i}}$. As defined in Section 3.4, the **augmentation of** x **through** W **by** ε is to reset by $x := x + \varepsilon d$ and then the auxiliary digraph is updated. In the remainder of this section, the **augmentation through** W means the augmentation of x through W by ε^*, where $\varepsilon^* = \max\{\varepsilon \in \mathbb{R} \mid x + \varepsilon d \in P(f, \mathcal{C})\}$. By Remark 9, ε^* can be computed by line searches in submodular polyhedra, which can be solved in strongly polynomial time [24].

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 shorter** 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 $Z \in \mathcal{E}$, $v \in Z \setminus \{v_Z^\circ\}$, 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 $E_Z(x)$.

Recall that for each arc $e \in E_Z^1$, a label $\ell(e)$ was defined in Section 3.4. In this section, we also assign a 'special' label $\star$ to each special arc in E_Z^2 for simplicity of description. The 'special' label will be treated as a label different from any other label. For v_Z°, let $W_\mathrm{f}(v_Z^\circ, s, x)$ and $W_\mathrm{b}(v_Z^\circ, s, x)$ be the lexicographically shortest forward/backward PAW among those starting from v_Z° whose initial arc has label s or

$\star$.

Our main theorem (Theorem 38) is a direct consequence of Lemma 33 below. Before showing Lemma 33, we first establish several technical lemmas.

Lemma 30. *Suppose that $(V, E(x))$ has 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 $(V, E(x))$ has an augmenting walk W satisfying $|W| < |W_1| + |W_2|$, unless $v = v_Z^\circ$ and the initial arcs of W_1 and W_2 are both in $E_Z^1(x)$ and have the same label.*

Proof. This is implicit in the proof of Claim 17 (b). □

A pair of walks in Lemma 30 can be used as a certificate for the existence of a shorter augmenting walk. Lemma 30 also implies the following lemma.

Lemma 31. *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 30 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, x)-**PAW** is a forward PAW in $(V, E(x))$ 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, x)-**PAW** is a forward or backward PAW in $(V, E(x))$ starting at v_Z° with the initial arc labeled s or $\star$. A $(v_Z^\circ, \bar{s}, x)$-**PAW** is a forward or backward PAW starting at v_Z° with the initial arc not labeled s or with the initial arc labeled $\star$. (Thus, if $s = \star$, then every forward or backward PAW starting at v_Z° is a $(v_Z^\circ, \bar{s}, x)$-PAW). Lemma 30 further implies the following lemma.

Lemma 32. *If W_1 is a (v, Z, x)-PAW and W_2 is a (v, Z', x)-PAW with $Z \neq Z'$, then there is an augmenting walk W satisfying $|W| \leq |W_1| + |W_2|$.*

If W_1 is a (v_Z°, s, x)-PAW and W_2 is a $(v_Z^\circ, \bar{s}, x)$-PAW different from W_1, 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, x)-PAW. 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. Let $W(v_Z^\circ, \bar{s}, x)$ denote the lexicographically shortest $(v_Z^\circ, \bar{s}, x)$-PAW.

The following is the main technical lemma of this section.

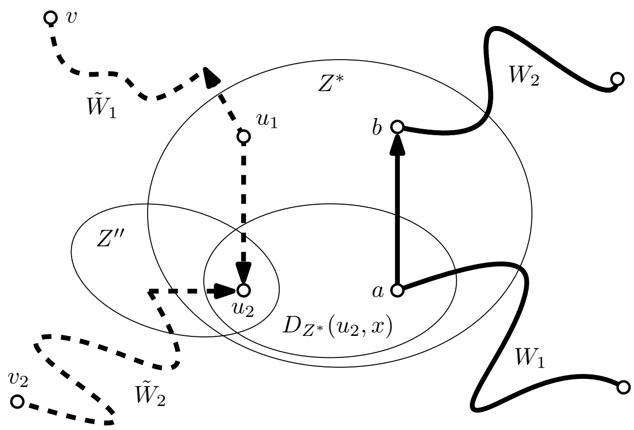

Figure 4: Proof of Lemma 33. The left dashed walk represents $W(v, Z, x')$ and the right bold walk represents W.

Lemma 33. *Let $x \in P(\mathcal{C}, f)$ and let x' be obtained from x by the augmentation through the lexicographically shortest augmenting walk W in $(V, E(x))$. Then, for each $Z \in \mathcal{E}$ and 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 33 when there is no vertex of multiplicity two. Let w_1 and w_2 be the endvertices of W. For each vertex $v \in Z$, let $c(v, Z)$ be the number of times $W(v, Z, x')$ passes through arcs in $E(x') \setminus E(x)$.

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 $E(x') \setminus E(x)$, 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., when u_1u_2 is passed in the backward direction). Let $Z^* \in \mathcal{E}$ be the color of u_1u_2 in $E(x')$, i.e., $u_1u_2 \in E_{Z^*}(x')$. 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 $\tilde{W}_2 = W_{\rm b}(u_2, Z'', x')$, where Z'' is the member of $\mathcal{E}$ containing u_2 and distinct from Z^*. See Figure 4. Then it holds that

$$|\tilde{W}_2| = |W_{\rm b}(u_2, Z'', x')| \geq |W(u_2, Z^*, x')| \geq |W(u_2, Z^*, x)|, \tag{27}$$

where the second inequality is from the definition of $W(u_2, Z^*, x')$ and the third inequality follows by induction. We say that an arc $ab \in W \cap E_{Z^*}(x)$ is **short** if

(s1) $|W_{\rm f}(a, Z^*, x)| \leq |\tilde{W}_2| + 1$, and (s2) 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 34. *Every arc $ab \in W \cap E_{Z^*}(x)$ leaving $D_{Z^*}(u_2, x)$ is short.*

Proof. Let W_1 be the first part of W from the initial vertex of W to 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 $a \in D_{Z^*}(u_2, x)$, we have that $au_2 \in E_{Z^*}(x)$ if $a \neq u_2$. Hence the concatenation of au_2 and W_1 is a (u_2, Z'', x)-PAW if $a \neq u_2$, while W_1 itself is a (u_2, Z'', x)-PAW if $a = u_2$. Thus, by Lemma 32 there is an augmenting walk W' in $(V, E(x))$ satisfying

$$|W'| \leq \begin{cases} |W(u_2, Z^*, x)| + |W_1| + 1 & \text{if } a \neq u_2, \\ |W(u_2, Z^*, x)| + |W_1| & \text{if } a = u_2. \end{cases} \tag{28}$$

Since W is the shortest, W is the concatenation of W_1 and $W_\mathrm{f}(a, Z^*, x)$. This we have $|W_\mathrm{f}(a, Z^*, x)| \leq |W(u_2, Z^*, x)| + 1 \leq |\tilde{W}_2| + 1$ by (27) and (28), where the equality may hold only if $a \neq u_2$. Thus ab satisfies (s1).

To see (s2) 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 $u_2 \in D_{Z^*}(u_2, x)$ while $b \notin D_{Z^*}(u_2, x)$), $W \preceq W'$ implies that $b \prec u_2$. Thus ab satisfies (s2). □

Claim 35. *There is a short arc ab in $W \cap E_{Z^*}(x)$ satisfying $u_1 \in D_{Z^*}(b, x)$.*

Proof. Since $u_1u_2 \in E(x')$ but $u_1u_2 \notin E(x)$, we have $u_1 \notin D_{Z^*}(u_2, x)$. We shall take a maximal (x, Z^*)-tight $\overline{u_1}u_2$-set Y with the property that every arc $a'b' \in W \cap E_{Z^*}(x)$ leaving Y is short. As $D_{Z^*}(u_2, x)$ satisfies this condition by Claim 34, such a maximal tight set Y exists. Since $u_1u_2 \in E(x')$ but $u_1u_2 \notin E(x)$, the arc set $W \cap E_{Z^*}(x)$ 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 $u_1 \in D_{Z^*}(b, x)$, which proves the claim.

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

Due to the choice of Y, this is trivial if $a' \in Y$. Assume $a' \in D_{Z^*}(b, x)$. Lemma 31 implies that $a' \neq b$, and hence we have $a'b \in E_{Z^*}(x)$, which in turn implies that

$$|W_\mathrm{f}(a', Z^*, x)| \leq |W_2| + 1 = |W_\mathrm{f}(a, Z^*, x)| \leq |\tilde{W}_2| + 1, \tag{29}$$

where the first inequality follows from the fact that the concatenation of $a'b$ and W_2 is a forward PAW starting from a', the second equation follows because the concatenation of ab and W_2 is $W_\mathrm{f}(a, Z^*, x)$, and the thrid inequality follows since ab is short. Thus condition (s1) holds for $a'b'$.

To check condition (s2), 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 (29) implies $|W_2'| = |W_2| = |\tilde{W}_2| \le |W|/2 - 1$. 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 (s2) holds for $a'b'$. □

Let ab be the arc guaranteed by Claim 35. We have two cases depending on whether $u_1 = b$ or not.

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

Now suppose that $|W(v, Z, x')| \le |W|/2$, which implies that $|\tilde{W}_2| \le |W|/2 - 1$. If $|W_{\mathrm{f}}(a, Z^*, x)| \le |\tilde{W}_2|$, 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 $W(v, Z, x) \preceq W' \prec W(v, Z, x')$ holds, and the statement follows.

If $u_1 = b$, then let $\tilde{W}_1^{-1}$ be the walk tracing $\tilde{W}_1$ in the reversed order from u_1 to v. If $\tilde{W}_1^{-1}$ and W_2 never split when tracing them from u_1 (that is, the vertex sequence of $\tilde{W}_1^{-1}$ 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 by Claim 13), then the remaining part of W_2 from v to the end is a (v, Z, x)-PAW. Therefore $|W(v, Z, x)| \le |W_2| = |W_{\mathrm{f}}(a, Z^*, x)| - 1 \le |\tilde{W}_2| < |W(v, Z, x')|$. On the other hand, if $\tilde{W}_1^{-1}$ and W_2 split at some vertex v', then the concatenation of $\tilde{W}_1^{-1}$ and W_2 followed by a shortcut at v' results in a (v, Z, x)-PAW (cf. the proof of Claim 17(b)). Thus $|W(v, Z, x)| \le |\tilde{W}_1| + |W_2| = |\tilde{W}_1| + |W_{\mathrm{f}}(a, Z^*, x)| - 1 \le |\tilde{W}_1| + |\tilde{W}_2| < |W(v, Z, x')|$, which completes the proof. □

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

Proof of Lemma 33. We shall check what happens when $v^\circ_{Z^*}$ appears in the above proof. Recall that $c(v, Z)$ denotes the number of times $W(v, Z, x')$ passes through arcs in $E(x') \setminus E(x)$ (where parallel arcs with distinct colors or labels are regarded as distinct arcs). Similarly we define $c(v^\circ_Z, \bar{s})$ to be the number of times $W(v^\circ_Z, \bar{s}, x')$ passes through arcs in $E(x') \setminus E(x)$. The proof is done by induction on $c(v, Z)$ and $c(v^\circ_Z, \bar{s})$.

As in the previous case, we take the first arc $u_1 u_2$ in $E(x') \setminus E$ when tracing $W(v, Z, x')$ from v. (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.) 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.

Let $Z^* \in \mathcal{E}$ be the element with $u_1 u_2 \in E_{Z^*}(x')$. If v_{Z^*} is semi-free, then one can apply the exactly same proof as that for the special case of no $v^\circ_{Z^*}$. Thus we focus on the case when v_{Z^*} is not semi-free.

For a (x, Z^*)-tight set Y with $m(Y) = 1$, the set of shades of Y is denoted by $\mathcal{S}(Y)$. Recall that, by Lemma 12, Y has a single shade if $v^\circ_{Z^*}u \in E^1_{Z^*}(x)$ for some $u \in Y$ or if $uv^\circ_{Z^*} \in E^1_{Z^*}(x)$ for some $u \notin Y$. If $m(Y) \neq 1$, then $\mathcal{S}(Y) = \emptyset$.

For each arc in W incident to a vertex v, its **partner** (in W at v) is the arc in W adjacent at v. 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., $v = u_1 = v^\circ_{Z^*}$).

Recall that in the last proof an arc $ab \in W \cap E_{Z^*}(x)$ was said to be short if it satisfies (s1) and (s2) given above. We can reuse the same definition for each arc ab if $a \neq v^\circ_{Z^*}$. If $a = v^\circ_{Z^*}$, we say that $v^\circ_{Z^*}b$ is **short** if

(s1) $|W_{\mathrm{f}}(v^\circ_{Z^*}, \ell(v^\circ_{Z^*}b), x)| \leq |\tilde{W}_2| + 1$, and

(s2) 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 $b \prec u_2$.

Claim 36. *Suppose that the lemma does not hold for $W(v, Z, x)$. Then there is a (x, Z^*)-tight multiset 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 unique shade of Y is $\tilde{s}$;*

(c) *every $ab \in W \cap E_{Z^*}(x)$ leaving Y with $a \neq v^\circ_{Z^*}$ is short;*

(d) *every $v^\circ_{Z^*}b \in W \cap E_{Z^*}(x)$ leaving Y is short unless $m(Y) = 1$ and the partner of $v^\circ_{Z^*}b$ at $v^\circ_{Z^*}$ has label equal to the shade of Y.*

Proof. We split the proof into two cases.

Case 1: $u_2 \neq v^\circ_{Z^*}$. We claim that $D_{Z^*}(u_2, x)$ satisfies the properties (a)–(d). Clearly (a) is satisfied. If u_1u_2 does not exist in $E(x)$, then $u_1 \notin D_{Z^*}(u_2, x)$. If u_1u_2 exists in $E(x)$, then $u_1 = v^\circ_{Z^*}$ should hold with $\ell(u_1u_2) = \tilde{s}$ since otherwise the lemma easily follows by induction. Then the shade of $D_{Z^*}(u_2, x)$ is $\tilde{s}$, i.e., (b) is satisfied. Property (c) can be checked by directly applying the proof of Claim 34 since $u_2 \neq v^\circ_{Z^*}$. To see (d), note that if $v^\circ_{Z^*} \in D_{Z^*}(u_2, x)$, then $v^\circ_{Z^*}u_2$ exists in $E_{Z^*}(x)$, and $\mathcal{S}(D_{Z^*}(u_2, x)) = \{\ell(v^\circ_{Z^*}u_2)\}$ if $m(D_{Z^*}(u_2, x)) = 1$. Observe also that one can apply the proof of Claim 34 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 the shade of $D_{Z^*}(u_2, x)$. 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 $W(v^\circ_{Z^*}, \bar{s}', x)$. If $e \in E^2_{Z^*}(x)$, then we claim that $D^\circ_{Z^*}(x)$, the smallest (x, Z^*)-tight set with multiplicity two, satisfies the desired properties. Note that $D^\circ_{Z^*}(x)$ exists since $v^\circ_{Z^*}$ is not semi-free. To see (b), suppose $u_1 \in D^\circ_{Z^*}(x)$. Then the concatenation of $\tilde{W}_1$, $u_1v^\circ_{Z^*}$, and $W(v^\circ_{Z^*}, \bar{s}', x)$ (and then applying Lemma 32 if necessary) will lead to a (v, Z, x)-PAW W'. Thus the lemma follows by applying the induction hypothesis to $W(v^\circ_{Z^*}, \bar{s}', x)$, and hence $u_1 \notin D^\circ_{Z^*}(x)$ holds by the assumption of the claim. We get (b). Since $e \in E^2_{Z^*}(x)$, (c) and (d) can be checked by applying the proof of Claim 34.

If $e \in E^1_{Z^*}(x)$, then let Y be the smallest (x, Z^*)-tight set whose shade is $\ell(e)$. To see (b) for Y, suppose $u_1 \in Y$. Then $\{\ell(u_1v^\circ_{Z^*})\} \neq \mathcal{S}(Y) = \{\ell(e)\}$. Hence the concatenation of $\tilde{W}_1$, $u_1v_{Z^*}$, and $W(v^\circ_{Z^*}, \bar{s}', x)$ (and then applying Lemma 32 if necessary) again leads to a (v, Z, x)-PAW W' shorter than $W(v, z, x')$, certifying the lemma. Hence we must have $u_1 \notin Y$ by the assumption of the claim, and we get (b). For (c) and (d), note that for any $a \in Y$ with $a \neq v^\circ_{Z^*}$ we have $av^\circ_{Z^*} \in E_{Z^*}(x)$ with $\{\ell(av^\circ_{Z^*})\} \neq \mathcal{S}(Y) = \{\ell(e)\}$. Hence one can apply the proof of Claim 34. □

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 36.

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

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

Claim 37. *The arc set $W \cap E_{Z^*}(x)$ 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 $\mathcal{S}(X_{ab}) \neq \{\tilde{s}\}$.*

Proof. Due to properties (a) and (b) of Y, there exists an arc ab in $W \cap E_{Z^*}(x)$ 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 36. To see this, take any $a'b' \in W \cap E_{Z^*}(x)$ 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 35 to see that $a'b'$ is short. Note that $a' \neq b$ by Lemma 31.

Case 1: Suppose $b = v^\circ_{Z^*}$. Let $\tilde{a}v^\circ_{Z^*}$ be the partner of $av^\circ_{Z^*}$ in W. If $\tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}(x)$, 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 $\tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}(x)$). Thus one can apply the proof of Claim 35 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 $|W_{\mathrm{f}}(v^\circ_{Z^*}, \ell(v^\circ_{Z^*}b'), x)| = |\tilde{W}_2| + 1$ imply $b' \prec u_2$ by applying the argument of the proof of Claim 35, i.e., $v^\circ_{Z^*}b'$ is short.

With this in mind, we now consider the following three cases to prove that $v^\circ_{Z^*}b'$ satisfies (d):

Case 2-1: If $v^\circ_{Z^*} \notin X_{ab}$, then clearly $\mathcal{S}(Y) = \mathcal{S}(Y \cup X_{ab})$, 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 $\mathcal{S}(X_{ab}) = \mathcal{S}(X_{ab} \cup Y)$. When $\mathcal{S}(X_{ab}) = \{\ell(v^\circ_{Z^*}\tilde{b'})\}$, there is nothing to prove for $v^\circ_{Z^*}b'$. When $\mathcal{S}(X_{ab}) \neq \{\ell(v^\circ_{Z^*}\tilde{b'})\}$, we have $\ell(v^\circ_{Z^*}b) \neq \ell(v^\circ_{Z^*}\tilde{b'})$, 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 $\mathcal{S}(Y) = \mathcal{S}(X_{ab})$, then $m(Y \cup X_{ab}) = m(Y) = 2$ or $\mathcal{S}(Y) = \mathcal{S}(X_{ab} \cup Y)$, and (d) holds for $v_{Z^*}b'$ as (d) holds for Y. Otherwise $m(Y) = 1$ and $\mathcal{S}(Y) \neq \mathcal{S}(X_{ab})$. If $\mathcal{S}(Y) \neq \{\ell(v^\circ_{Z^*}\tilde{b'})\}$, then the shortness of $v^\circ_{Z^*}b'$ follows since Y satisfies (d). Otherwise $\mathcal{S}(Y) \neq \mathcal{S}(X_{ab})$ implies $\ell(v^\circ_{Z^*}\tilde{b'}) \neq \ell(v^\circ_{Z^*}b)$, which implies that $v^\circ_{Z^*}b'$ is short, as it was shown earlier.

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 $\mathcal{S}(X_{ab}) \neq \{\tilde{s}\}$. Thus the claim holds. □

Now we are ready to complete the proof of the lemma. Let ab be the arc guaranteed in Claim 37. 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 Claim 36(b). In this case, $u_1 \in X_{ab} = D_{Z^*}(b, x)$, and $\mathcal{S}(X_{ab}) \neq \{\tilde{s}\}$ by Claim 37. Consequently, $v^\circ_{Z^*}b$ exists in E with $\ell(v^\circ_{Z^*}b) \neq \tilde{s}$. Hence the concatenation of $\tilde{W}_1, v^\circ_{Z^*}b$, and W_2 leads to a (v, Z, x)-PAW. Note that by the shortness of ab the existence of such a (v, Z, x)-PAW implies the lemma.

We may therefore 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, x)-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, x)-PAW if $\tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}(x)$. Otherwise, i.e., if $\tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}(x)$, we have $\{\ell(\tilde{a}v^\circ_{Z^*})\} = \mathcal{S}(X_{ab}) \neq \{\tilde{s}\}$ by Claim 37. Hence again the concatenation of $\tilde{W}_1$ and W_2 can be shortcut to a (v, Z, x)-PAW.

Suppose $u_1 = v^\circ_{Z^*} \neq b$. If $m(D_{Z^*}(b, x)) = 2$, then $v^\circ_{Z^*}b \in E^2_{Z^*}(x)$. If $m(D_{Z^*}(b, x)) = 1$, then $\{\ell(v^\circ_{Z^*}b)\} = \mathcal{S}(D_{Z^*}(b, x)) = \mathcal{S}(X_{ab}) \neq \{\tilde{s}\}$. In both cases, the concatenation of $\tilde{W}_1$, $v^\circ_{Z^*}b$ and W_2 leads to a (v, Z, x)-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, x)-PAW if $\tilde{a}v^\circ_{Z^*} \in E^2_{Z^*}(x)$. Otherwise, i.e., if $\tilde{a}v^\circ_{Z^*} \in E^1_{Z^*}(x)$, 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, x)-PAW.

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

We now state and prove the main theorem of this section.

Theorem 38. *Given $x \in P(f, \mathcal{C})$, $x(V)$ is maximized after $O(n^3)$ augmentations using lexicographically shortest augmenting walks,*

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 33, $|W(v_{W'}, Z, x)| \leq |W(v_{W'}, Z, x')|$ and $|W(v_{W'}, Z', x)| \leq |W(v_{W'}, Z', x')|$ hold, which means that the auxiliary digraph with respect to x contains an augmenting walk shorter than or equal to $|W'|$ by Lemma 32. This implies $|W| \leq |W'|$.

Suppose $|W| = |W'|$. Let $N_Z(x) = \{v \in Z \mid |W(v, Z, x)| \leq |W|/2\}$ and let $w(v, Z, x)$ be the second vertex of $W(v, Z, x)$. We define $N_Z(x')$ and $w(v, Z, x')$ similarly for x'. By Lemma 33, 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 also $w(v, Z, x) \preceq w(v, Z, x')$ for each $v \in N_Z(x')$. Moreover, since $W \neq W'$, at least one of the following three properties 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 neither (i) nor (ii) holds, 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 25 and 38, 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, Math. Program. 155 (2016), 1–55

[10] H. Hirai, Computing DM-decomposition of a partitioned matrix with rank-1 blocks, arXiv:1609.01934 (2016).

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

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

[13] 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.

[14] 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.

[15] 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.

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

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

[18] 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.

[19] V. Kolmogorov, J. Thapper and S. Živný, The power of linear programming for general-valued CSPs. SIAM J. Comput., 44 (2015), 1–36.

[20] 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.

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

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

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

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

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

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

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

[28] J. Thapper and S. Živný, The power of linear programming for valued CSPs, in: Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science (2012), 668–678.

[29] J. Thapper and S. Živný, The complexity of finite-valued CSPs, J. ACM 63 (2016), 37:1–37:33.

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