Decreasing Minimization on M-convex Sets: Algorithms and Applications

András Frank[*] and Kazuo Murota[†]

July 2020 / July 2021

Abstract

This paper is concerned with algorithms and applications of decreasing minimization on an M-convex set, which is the set of integral elements of an integral base-polyhedron. Based on a recent characterization of decreasingly minimal (dec-min) elements, we develop a strongly polynomial algorithm for computing a dec-min element of an M-convex set. The matroidal feature of the set of dec-min elements makes it possible to compute a minimum cost dec-min element, as well. Our second goal is to exhibit various applications in matroid and network optimization, resource allocation, and (hyper)graph orientation. We extend earlier results on semi-matchings to a large degree by developing a structural description of dec-min in-degree bounded orientations of a graph. This characterization gives rise to a strongly polynomial algorithm for finding a minimum edge-cost dec-min orientation.

Keywords: Network flows, Resource allocation, Graph orientation, Decreasing minimization, M-convex set, Polynomial algorithm.
Mathematics Subject Classification (2010): 90C27, 05C, 68R10
Running head: Decreasing Minimization: Algorithms and Applications

[*]MTA-ELTE Egerváry Research Group, Department of Operations Research, Eötvös University, Pázmány P. s. 1/c, Budapest, Hungary, H-1117. e-mail: `frank@cs.elte.hu`. ORCID: 0000-0001-6161-4848. The research was partially supported by the National Research, Development and Innovation Fund of Hungary (FK_18) – No. NKFI-128673.

[†]Department of Economics and Business Administration, Tokyo Metropolitan University, Tokyo 192-0397, Japan, e-mail: `murota@tmu.ac.jp`. Currently at The Institute of Statistical Mathematics, Tokyo 190-8562, Japan. ORCID: 0000-0003-1518-9152. The research was supported by CREST, JST, Grant Number JPMJCR14D2, Japan, and JSPS KAKENHI Grant Numbers JP26280004, JP20K11697.

Contents

1 Introduction

This paper is concerned with algorithms and applications of decreasing minimization on an M-convex set, which is the set of integral elements of an integral base-polyhedron. An element of a set of vectors, in general, is called **decreasingly minimal** (dec-min) if its largest component is as small as possible, within this, its second largest component is as small as possible, and so on. Decreasing minimization means the problem of finding a dec-min element of a given set of vectors (or even a cheapest dec-min element with respect to a given linear cost-function). When the given set of vectors consists of integral vectors, this problem is also referred to as discrete decreasing minimization. In the literature, typically the term lexicographic optimization is used, but we prefer "decreasing minimization" because we also consider its natural counterpart "increasing maximization," and the use of these two symmetric terms seems more appropriate to distinguish the two related notions. An element of a set of vectors is called **increasingly maximal** (inc-max) if its smallest component is as large as possible, within this, its second smallest component is as large as possible, and so on.

In the companion paper [21], the present authors have investigated the structural aspects of the discrete decreasing minimization on an M-convex set. Among others, the dec-min elements are characterized as those admitting no local improvement, where the precise meaning of a local improvement is defined formally in term of "1-tightening step" in Section 2. It is also shown that an element of an M-convex set is decreasingly minimal precisely if it is a minimizer of the sum of the squared components.

As dual objects to dec-min elements, the notions of canonical chain, canonical partition of the ground-set, and essential-value sequence were defined, and the structure of the set of all dec-min elements is described in terms of these dual objects. (An equivalent definition of these notions is given in Section 2.2.) We emphasize that the role of these dual objects is not merely to help us fully understand the problem from its dual side. Beyond this, the dual characterization reveals the fundamental feature of the primal problem that the set of dec-min elements itself forms an M-convex set, and, in fact, a rather special one arising from a matroid by translation. In addition, these dual objects are inherent in computing a dec-min element in strongly polynomial time and indispensable for efficient computation of a minimum weight dec-min element, as well.

The first goal of this paper is to develop, on the basis of the above-mentioned structural characterizations, a strongly polynomial algorithm for computing a dec-min element as well as the canonical chain of a given M-convex set. The second goal is to exhibit several applications. For example, we prove a conjecture of Borradaile et al. [7] on dec-min strongly connected orientations of undirected graphs. Our general approach makes it possible to solve algorithmically even the minimum edge-cost dec-min orientation problem when upper and lower bounds are imposed on the in-degrees and the orientation is expected to be k-edge-connected (or even (k, ℓ)-edge-connected). These orientation results form the basis of a major generalization of the so-called semi-matching problem initiated by Harvey et al. [31], which had been motivated by a resource allocation problem in computer science. Our approach is the first one that provides a strongly polynomial algorithm for the capacitated case, as well.

An algorithmic solution to a discrete counterpart of Megiddo's lexicographic flow problem [43, 44] is also developed. Yet another application of the structural results of [21] gives rise to an extension of a result of Levin and Onn [42] on finding k bases of a matroid on a ground-set S with n elements such that the degree-vector of the hypergraph formed by these k bases is decreasingly minimal. Our approach generalizes this problem to the case when one has k distinct matroids on S.

The paper is organized as follows. Algorithms for computing a dec-min element and the canonical chain are given in Section 2. In Section 3, various kinds of applications are shown, including those to matroids, network flows, arborescences, and connectivity augmentations. Sections 4, 5, and 6 are devoted to detailed account of applications to graph orientation problems.

1.1 Notation and terminology

We continue to use notation and terminology introduced in [21], while some additional ones are given here. Let S be a finite ground-set S with n elements. Two subsets X and Y of S are **intersecting** if $X \cap Y \neq \emptyset$, **properly intersecting** if none of $X \cap Y$, $X - Y$, and $Y - X$ is empty, and **crossing** if none of $X - Y$, $Y - X$, $X \cap Y$, and $S - (X \cup Y)$ is empty. For distinct elements s, t of S, a subset of S containing t but not s is called a $t\overline{s}$-**set**.

For a vector $x \in \mathbf{R}^S$ or a function $x : S \to \mathbf{R}$, we define the set-function $\widetilde{x} : 2^S \to \mathbf{R}$ by $\widetilde{x}(Z) := \sum[x(s) : s \in Z]$ $(Z \subseteq S)$. This is a modular function in the sense that $\widetilde{x}(X) + \widetilde{x}(Y) = \widetilde{x}(X \cap Y) + \widetilde{x}(X \cup Y)$ holds for every $X, Y \subseteq S$.

For any integral polyhedron $R \subseteq \mathbf{R}^S$, we use the notation $\overset{....}{R}$ to denote the set of integral elements of R, that is,

$$\overset{....}{R} := R \cap \mathbf{Z}^S, \tag{1.1}$$

where $\overset{....}{R}$ may be pronounced "dotted R." The notation is intended to refer intuitively to the set of lattice points of R.

For a set-function h, we allow it to have value $+\infty$ or $-\infty$. Unless otherwise stated, $h(\emptyset) = 0$ is assumed throughout. When $h(S)$ is finite, the **complementary function** $\overline{h}$ is defined by $\overline{h}(X) := h(S) - h(S - X)$. Observe that the complementary function of $\overline{h}$ is h itself.

Let b be a set-function with $b(\emptyset) = 0$, for which $b(X) = +\infty$ is allowed but $b(X) = -\infty$ is not. The submodular inequality for subsets $X, Y \subseteq S$ is defined by

$$b(X) + b(Y) \geq b(X \cap Y) + b(X \cup Y). \tag{1.2}$$

We say that b is (fully) **submodular** if this inequality holds for every pair of subsets $X, Y \subseteq S$ with finite b-values. When the submodular inequality is required only for intersecting (crossing) pairs of subsets, we say that b is **intersecting** (**crossing**) **submodular**. A set-function p is called (fully, intersecting, crossing) **supermodular** if $-p$ is (fully, intersecting, crossing) submodular. It follows from the definitions that if b is fully (or crossing) submodular with finite $b(S)$, then its complementary function $\overline{b}$ is a fully (or crossing) supermodular function, while $\overline{b}$ is not necessarily intersecting supermodular when b is intersecting submodular.

For a (fully) submodular integer-valued set-function b on S with $b(\emptyset) = 0$ and finite $b(S)$, the **base-polyhedron** B is defined by

$$B = B(b) = \{x \in \mathbf{R}^S : \widetilde{x}(S) = b(S),\ \widetilde{x}(Z) \leq b(Z)\ \text{for every}\ Z \subset S\}, \tag{1.3}$$

which is a (possibly unbounded) integral polyhedron in $\mathbf{R}^S$. Section 14 of book [16] provides an overview of basic properties of base-polyhedra; see also the book of Schrijver [53]. For example, it is a basic property that $B = B(b)$ is a non-empty integral polyhedron, and B uniquely determines its defining (fully) submodular function b, namely,

$$b(Z) = \max\{\widetilde{x}(Z) : x \in B\} \quad (= \max\{\widetilde{x}(Z) : x \in \overset{....}{B}\}).$$

By convention, the empty set is also considered a base-polyhedron. A (fully) supermodular integer-valued set-function p with $p(\emptyset) = 0$ and $p(S)$ finite also defines an integral base-polyhedron by

$$B = B'(p) = \{x \in \mathbf{R}^S : \widetilde{x}(S) = p(S),\ \widetilde{x}(Z) \geq p(Z)\ \text{for every}\ Z \subset S\}, \tag{1.4}$$

since the complementary function $b := \overline{p}$ of p is fully submodular and $B'(p) = B(b)$. In discrete convex analysis [46, 47], the set $\overset{....}{B}$ of integral elements of an integral base-polyhedron B is called an **M-convex set**.

We say that a submodular function b and a supermodular function p meet the **cross-inequality** for $X, Y \subseteq S$ if

$$b(X) - p(Y) \geq b(X - Y) - p(Y - X). \tag{1.5}$$

4

A pair (p, b) of set-functions is called **fully paramodular** (or we say that (p, b) is a **strong pair**) if b is fully submodular, p is fully supermodular, and the cross-inequality holds for every $X, Y \subseteq S$. For a strong pair (p, b), the polyhedron Q defined by

$$Q = Q(p, b) := \{x \in \mathbf{R}^S : p(Z) \leq \widetilde{x}(Z) \leq b(Z) \text{ for every } Z \subseteq S\} \tag{1.6}$$

is called a **generalized polymatroid** (**g-polymatroid**, for short), while (p, b) is called the **border pair** of Q. It is a basic fact [16] that Q is never empty. By convention, the empty set is also considered a g-polymatroid. Furthermore, a non-empty g-polymatroid Q uniquely determines its fully paramodular border pair, and Q is integral when p and b are integer-valued. Base-polyhedra are special g-polymatroids, namely, those for which $p(S) = b(S)$, and every g-polymatroid arises from a base-polyhedron by projection along a single coordinate axis.

To prove theorems on base-polyhedra, it is much easier to work with base-polyhedra defined by fully sub- or supermodular functions. For applications, however, it is fundamentally important that weaker set-functions (e.g. intersecting, crossing or even weaker submodular functions) may also define base-polyhedra (or M-convex sets), as well as g-polymatroids. (We shall use this fact frequently in Sections 5 and 6.) For example, if p is an integer-valued intersecting or crossing supermodular function, then $B'(p)$ can be proved to be an integral base-polyhedron (see, e.g. Theorem 15.3.4 in book [16]), which may, however, be empty. (This result, for example, underlies the fact that the in-degree vectors of k-edge-connected orientations of a $2k$-edge-connected graph form an M-convex set.)

When p is intersecting supermodular, b is intersecting submodular, and the cross-inequality holds for properly intersecting pairs of subsets X, Y, we say that (p, b) is a **weak pair** or **intersecting paramodular**. For such a border pair, $Q = Q(p, b)$ is known to be a (possibly empty) g-polymatroid. It should be noted, that the (unique) fully paramodular border pair defining Q can be concretely expressed with the help of the weak pair (p, b), but this formula is quite complicated (see, Corollary 15.3.4 in book [16]).

It is a non-trivial task to characterize the situation when $B(b)$ is non-empty for a crossing submodular function b but Fujishige [26] developed an elegant necessary and sufficient condition. A similar theorem can be formulated for g-polymatroids defined by a weak pair (p, b) (see, Theorem 15.3.13 in book [16]). It should, however, be emphasized that applications often need base-polyhedra or g-polymatroids defined by even weaker sub- or supermodular functions. For example, the in-degree vectors of k-edge-connected digraphs obtained by adding a given number of arcs to an input digraph $H = (V, A)$ form an M-convex set. (See, Theorem 17.2.9 and Section 15 of [16] for an overview. A recent paper of Bérczi and Frank [2] includes even more intricate constructions of M-convex sets appearing in graph connectivity augmentation problems.)

We assume that graphs or digraphs have no loops but parallel edges are allowed. Sometimes we refer to an edge of a digraph as an **arc**. For a digraph $D = (V, A)$, the **in-degree** of a node v is the number of arcs of D with head v. The in-degree $\varrho_D(Z) = \varrho(Z)$ of a subset $Z \subseteq V$ denotes the number of edges (= arcs) entering Z, where an arc uv is said to enter Z if its head v is in Z while its tail u is in $V - Z$. The **out-degree** $\delta_D(Z) = \delta(Z)$ is the number of arcs leaving Z, that is $\delta(Z) = \varrho(V - Z)$. The number of edges of a directed or undirected graph H induced by $Z \subseteq V$ is denoted by $i(Z) = i_H(Z)$. In an undirected graph $G = (V, E)$, the **degree** $d(Z) = d_G(Z)$ of a subset $Z \subseteq V$ denotes the number of edges connecting Z and $V - Z$, while $e(Z) = e_G(Z)$ denotes the number of edges with one or two end-nodes in Z. Clearly, $e(Z) = d(Z) + i(Z)$.

2 Algorithms

In this section, we consider algorithmic aspects of decreasing minimization over an M-convex set. In particular, we show how to compute efficiently a decreasingly minimal element along with its canonical chain and partition.

Let B be a non-empty integral base-polyhedron and let p denote the unique fully supermodular function for which $B = B'(p)$. Let m be an element of the M-convex set $\ddddot{B}$. We need some definitions introduced in [21]. A **1-tightening step** replaces m by $m' := m + \chi_s - \chi_t$, where s and t are elements of S for which $m(t) \geq m(s) + 2$ and m' belongs to $\ddddot{B}$. A subset $X \subseteq S$ is called m-**tight** (with respect to p) if $\widetilde{m}(X) = p(X)$. A subset $X \subseteq S$ is called an m-**top** set if $m(u) \geq m(v)$ holds whenever $u \in X$ and $v \in S - X$. We call an integral vector $x \in \mathbf{Z}^S$ **near-uniform** on a subset S' of S if its largest and smallest components on S' differ by at most 1, that is, if $x(s) \in \{\ell, \ell + 1\}$ for some integer ℓ for every $s \in S'$.

First we recall fundamental characterizations of a dec-min element of an M-convex set.

Theorem 2.1 ([21, Theorem 3.3]). *For an element m of an M-convex set $\ddddot{B} = \ddddot{B}'(p)$, the following four conditions are pairwise equivalent.*

(A) *There is no 1-tightening step for m.*

(B) *There is a chain $(\emptyset \subset) C_1 \subset C_2 \subset \cdots \subset C_\ell (= S)$ such that each C_i is an m-top and m-tight set (with respect to p) and m is near-uniform on each $S_i := C_i - C_{i-1}$ $(i = 1, 2, \ldots, \ell)$, where $C_0 := \emptyset$.*

(C1) *m is decreasingly minimal in $\ddddot{B}$.*

(C2) *m is increasingly maximal in $\ddddot{B}$.* ■

We mentioned already in Introduction that it is convenient to prove results for a base-polyhedron assuming that it is given by a fully sub- or supermodular set-function. In applications, however, a base-polyhedron is typically given by an intersecting or crossing (or even weaker) function. Therefore, in describing and analysing algorithms, we consider these weaker functions, as well.

Remark 2.1. One of the most fundamental algorithms of discrete optimization is for minimizing a submodular function, that is, for finding a subset Z of S for which $b(Z) = \min\{b(X) : X \subseteq S\}$. There are strongly polynomial algorithms for this problem; Schrijver [52] and Iwata et al. [36] are the first. We shall refer to such an algorithm as a **submod-minimizer** subroutine. The complexity of a significantly more efficient algorithm, due to Orlin [49], is $O(n^6)$ (where $n = |S|$) and this algorithm calls $O(n^5)$ times a routine which evaluates the submodular function in question. (An evaluation routine outputs the value $b(X)$ for any input subset $X \subseteq S$). A recent algorithm of Jiang [37] needs only $O(n^3)$ calls of an evaluation routine. Naturally, submodular function minimization and supermodular function maximization are equivalent. ■

2.1 The basic algorithm for computing a dec-min element

Our first goal is to describe a natural approach—the **basic algorithm**—for finding a decreasingly minimal element of an M-convex set $\ddddot{B}$. The basic algorithm is always finite, and if $B \subseteq \mathbf{R}_+^S$, it is pseudopolynomial in the sense that it is polynomial in $n + |p(S)|$. This means that the algorithm is polynomial in n when $B \subseteq \mathbf{R}_+^S$ and $|p(S)|$ can be bounded by a polynomial of n. This is the case, for example, in an application to strongly connected decreasingly minimal (=egalitarian) orientations. In the general case, where typical applications arise by defining p with a 'large' capacity function, a (more sophisticated) strongly polynomial-time algorithm will be described in Section 2.4. In order to find a dec-min element of an M-convex set $\ddddot{B}$, we assume that a subroutine is available to

$$\text{compute an integral element of } B. \tag{2.1}$$

We use this subroutine only once to obtain an initial point in the algorithm.

Suppose now that an integral member m of B is available. The algorithm needs a subroutine to

$$\text{decide for } m \in \ddddot{B} \text{ and for } s, t \in S \text{ if } m' := m + \chi_s - \chi_t \text{ belongs to } B. \tag{2.2}$$

By n^2 applications of subroutine (2.2), one can decide, for a given $m \in \overset{...}{B}$, whether there exists a 1-tightening step or not. Observe that the subroutine (2.2) is certainly available if we can

$$\text{decide for any } m' \in \mathbf{Z}^S \text{ whether or not } m' \text{ belongs to } B, \tag{2.3}$$

though applying this more general subroutine is clearly slower than a direct algorithm to realize (2.2). In this paper we mainly work with Subroutines (2.1) and (2.2), which are representation-free in the sense that they do not refer directly to the supermodular (or submodular) functions describing B. The realization of these subroutines depends on how the base-polyhedron B (or an M-convex set $\overset{...}{B}$) is given. Recall that a base-polyhedron can be described by a fully, intersecting, crossing, or even weaker functions. For example, when B is described by a fully supermodular function p, (2.1) can be carried out with n evaluations of p and (2.2) can be done with a single application of a submod-minimizer. See Remark 2.2 for more details.

As long as possible, apply the 1-tightening step. By Theorem 2.1, when no more 1-tightening step is available, the current m is a decreasingly minimal member of $\overset{...}{B}$ and the algorithm terminates. In this connection we may recall the following characterization of a dec-min element of an M-convex set.

Theorem 2.2 ([21, Corollary 6.4]). *An element m of an M-convex set $\overset{...}{B}$ is a dec-min element of $\overset{...}{B}$ if and only if it is a minimizer of the square-sum $W(z) := \sum[z(s)^2 : s \in S]$ of z over the elements z of $\overset{...}{B}$.* ∎

Now observe that a single 1-tightening step strictly decreases the square-sum of the components. This implies that the number of 1-tightening steps is bounded by $W(m)$ when m is the initial member of $\overset{...}{B}$. In particular, the algorithm is finite.

This algorithm, however, may be quite inefficient as is demonstrated by the simple example where $|S| = 2$ and $B = \{(x_1, x_2) : x_1 + x_2 = 0\}$. Here if the initial member of $\overset{...}{B}$ is, for example, $m = (10^6, -10^6)$, then the algorithm needs 10^6 1-tightening steps. However, if B is in the non-negative orthant (which is often the case in applications), then the following reasonable bound can be given for the complexity. Since the square-sum of an arbitrary integral vector $z \geq 0$ with $\widetilde{z}(S) = p(S)$ is at most $p(S)^2$ and $\widetilde{z}(S) = p(S)$ holds for all members z of $\overset{...}{B}$, we conclude that the number of 1-tightening steps is at most $p(S)^2$. Therefore, if $B \subseteq \mathbf{R}_+^S$ and $|p(S)|$ is bounded by a polynomial of n, then the basic algorithm to compute a dec-min element of $\overset{...}{B}$ is strongly polynomial.

Although the basic algorithm is efficient when $B \subseteq \mathbf{R}_+^S$ and $|p(S)|$ is 'small' (that is, $|p(S)|$ is bounded by a power of n), it is not strongly polynomial when $|p(S)|$ is 'large'. We postpone, till Section 2.4, the description of a strongly polynomial algorithm for computing a dec-min element of an M-convex set $\overset{...}{B}$ defined by a general p. In the next section we show how the canonical chain as well as the essential value-sequence can be computed, once a dec-min element m is available. It is emphasized that these dual objects are indispensable and must be computed when we are interested in identifying the set of all dec-min elements of $\overset{...}{B}$ or in finding a minimum weight dec-min element (cf., [21, Section 5.3]).

Remark 2.2. In the algorithm above, we did not rely explicitly on the set-function defining the base-polyhedron B in question, apart from the single value $p(S)$. The only assumption was that B is non-empty and the subroutines (2.1) and (2.2) are available. When the base-polyhedron $B = B'(p)$ is given by a fully supermodular function p, then the subroutine (2.1) can be realized by a version of the polymatroid greedy algorithm of Edmonds [9] that needs n evaluations of p (but not a submod-minimizer subroutine). If p is intersecting supermodular, then an algorithm of Frank and Tardos [24] requires n calls of a submod-minimizer. The same paper includes an algorithm for the case when p is crossing supermodular, and this algorithm requires n^2 calls of a submod-minimizer. Concerning the other subroutine (2.2), we note that $m' = m + \chi_s - \chi_t$ is in $B = B'(p)$ precisely if there is no m-tight

$t\bar{s}$-set (with respect to p), and this is true even when B is defined by a crossing supermodular function p. Therefore, the subroutine (2.2) can be realized with a single call of a submod-minimizer if p is fully supermodular, n calls if p is intersecting supermodular, and n^2 calls if p is crossing supermodular. ∎

Remark 2.3. Theorem 2.2 shows that a dec-min element of an M-convex set is characterized as a minimizer of the square-sum $W(z)$. This naturally suggests the approach of minimizing $W(z)$ to find a dec-min element. Minimizing the square-sum $W(z)$ over an M-convex set is a special case of minimizing an M-convex function, for which a local improvement algorithm works (in finite steps) [47]. The basic algorithm described above corresponds to the special case of this local improvement algorithm for M-convex function minimization. It is also noted that Fujishige [25] (see also [27, Section 8.2]) solved the continuous case of the square-sum minimization in polynomial time, which problem is equivalent to finding the unique minimum norm point m^* of a base-polyhedron B. In [20] we proved a theorem formalizing the intuitive feeling that all the dec-min elements of $\overset{....}{B}$ are in the neighbourhood of m^*, and this result makes it possible to develop an alternative algorithm to compute a dec-min element of $\overset{....}{B}$. ∎

2.2 Computing the essential value-sequence and the canonical chain

In this section, we describe an algorithm that assigns a chain of subsets of S, a partition of S, and a strictly decreasing sequence of integers to a given dec-min element m of an M-convex set.

Let $B = B'(p)$ be again a (non-empty) integral base-polyhedron whose unique (fully) supermodular bounding function is p. In the algorithm, we assume that we can compute the smallest m-tight set $T_m(u) = T_m(u; p)$ containing a given element $u \in S$. Since we have $T_m(u) = \{s : m + \chi_s - \chi_u \in B\}$, $T_m(u)$ is indeed computable by at most n applications of Subroutine (2.2).

Algorithm 2.3. Given a dec-min element m of $\overset{....}{B}$, the following procedure computes a chain $C^* = \{C_1, C_2, \ldots, C_q\}$ with $C_1 \subset C_2 \subset \cdots \subset C_q (= S)$ and a partition $\mathcal{P}^* = \{S_1, S_2, \ldots, S_q\}$ of S along with a sequence $\beta_1 > \beta_2 > \cdots > \beta_q$ of integers.

1. Let β_1 denote the largest value of m. Let $C_1 := \bigcup\{T_m(u) : m(u) = \beta_1\}$, $S_1 := C_1$, and $i := 2$.

2. In the general step $i \geq 2$, the pairwise disjoint non-empty sets $S_1, S_2, \ldots, S_{i-1}$ and a chain $C_1 \subset C_2 \subset \cdots \subset C_{i-1}$ have already been computed along with the values $\beta_1 > \beta_2 > \cdots > \beta_{i-1}$. If $C_{i-1} = S$, set $q := i - 1$ and stop. Otherwise, let

$$\beta_i := \max\{m(s) : s \in S - C_{i-1}\},$$
$$C_i := \bigcup\{T_m(u) : m(u) \geq \beta_i\},$$
$$S_i := C_i - C_{i-1},$$

and go to the next step for $i := i + 1$. ∎

It was proved in [21] (Corollary 5.4) that these sequences do not depend on the choice of m. Therefore the chain C^* is called the **canonical chain** belonging to $\overset{....}{B}$, the partition $\mathcal{P}^*$ is the **canonical partition** of $\overset{....}{B}$, while the sequence $\{\beta_1, \beta_2, \ldots, \beta_q\}$ is called the **essential value-sequence** of $\overset{....}{B}$. We emphasize that Algorithm 2.3 is strongly polynomial for arbitrary p (independently of the magnitude of $|p(S)|$), provided that a dec-min element m of $\overset{....}{B}$ is already available as well as the subroutine (2.2).

It is in order here to emphasize the significance of this algorithm for computing these dual objects. By Theorem 2.4 below, Algorithm 2.3 enables us to computationally capture the set of all dec-min elements. Concisely, the matroid associated with dec-min elements, as in Theorem 2.5 below, can be identified by this algorithm. Here a **matroidal M-convex set** means [21] the translation of the incidence vectors of bases of a matroid by an integral vector.

Theorem 2.4 ([21, Corollary 5.2]). *An element m of an M-convex set $\overset{....}{B}$ is decreasingly minimal if and only if each $C_i \in C^*$ is m-tight (with respect to p) and $\beta_i - 1 \leq m(s) \leq \beta_i$ holds for each $s \in S_i$ ($i = 1, \ldots, q$).* ∎

Theorem 2.5 ([21, Theorem 5.7]). *The set of dec-min elements of an M-convex set $\overset{....}{B}$ is a matroidal M-convex set.* ∎

We shall we use Theorem 2.4 in Sections 4.3 and 5.1, and Theorem 2.5 in Sections 4.5 and 5.2.

Adaptation to the intersection with a box   Algorithm 2.3 can be adapted to the case when we have specific upper and lower bounds on the members of $\overset{....}{B} = \overset{...}{B}'(p)$. Let $f : S \rightarrow \mathbf{Z} \cup \{-\infty\}$ and $g : S \rightarrow \mathbf{Z} \cup \{+\infty\}$ be bounding functions with $f \leq g$ and let $T(f, g) := \{x \in \mathbf{R}^S : f \leq x \leq g\}$ denote the box defined by f and g. It is a basic fact that the intersection of an integral base-polyhedron with an integral box is another integral base-polyhedron (see, Theorem 14.3.9 in book [16] for this statement in a more general context). Therefore the intersection $B^\square := B \cap T(f, g)$ is also a (possibly empty) integral base-polyhedron. Assume that $B^\square$ is non-empty.

Let m be an element of $\overset{....}{B^\square}$ ($= B^\square \cap \mathbf{Z}^S$). Let $T_m(u)$ denote the smallest m-tight set containing u with respect to p, and let $T_m^\square(u)$ be the smallest m-tight set containing u with respect to $p^\square$.

Claim 2.6.

$$T_m^\square(u) = \begin{cases} \{u\} & \text{if} \quad m(u) = f(u), \\ T_m(u) - \{v : m(v) = g(v)\} & \text{if} \quad m(u) > f(u). \end{cases}$$

Proof. We have $T_m^\square(u) = \{s : m - \chi_u + \chi_s \in B^\square\}$. Since $B^\square = B \cap T(f, g)$, we have $m - \chi_u + \chi_s \in B^\square$ if and only if (i) $m - \chi_u + \chi_s \in B$ and (ii) $m - \chi_u + \chi_s \in T(f, g)$ hold. For $s \neq u$, (i) holds if and only if $s \in T_m(u)$, and (ii) holds if and only if $m(u) > f(u)$ and $m(s) < g(s)$. Hence follows the claim. ∎

The claim implies that Algorithm 2.3 can be adapted easily to compute the canonical chain and partition belonging to $\overset{....}{B^\square}$ along with its essential value-sequence.

Our next goal is to describe a strongly polynomial algorithm to compute a dec-min element of $\overset{....}{B}$ in the general case when no restriction is imposed on the magnitude of $|p(S)|$. To this end, we need an algorithm to maximize $\lceil p(X)/|X| \rceil$, which is given in Section 2.3. The strongly polynomial algorithm for computing a dec-min element is described in Section 2.4.

2.3 Maximizing $\lceil p(X)/|X| \rceil$ with the Newton–Dinkelbach algorithm

In this section we describe a variant of the Newton–Dinkelbach (ND) algorithm to compute the maximum of $\lceil p(X)/|X| \rceil$. We assume that p is an integer-valued set-function on a ground-set S with $n \geq 1$ elements, $p(\emptyset) = 0$, and $p(S)$ is finite ($p(X)$ may be $-\infty$ for some X but never $+\infty$).

An excellent overview by Radzik [50] presents fundamental properties of the ND-algorithm, describing (among others) a strongly polynomial algorithm for minimizing (or maximizing) the ratio of two modular set-functions. For the problem of maximizing $\lceil p(X)/|X| \rceil$ (or $p(X)/|X|$), the ND-algorithm terminates in at most n iterations, which follows from the observation of Topkis [55] that the function $h(\mu) := \max\{p(X) - \mu|X| : X \subseteq S\}$ is a convex, piecewise-linear function with at most n breakpoints A recent paper by Goemans et al. [28] establishes a quadratic $O(n^2)$ bound on the number of iterations of the ND-algorithm for maximizing $p(X)/a(X)$ over X with $a(X) > 0$, where a is an arbitrary modular set function. We present a variant of the ND-algorithm whose specific feature is that it works throughout with integers $\lceil p(X)/|X| \rceil$. This has the advantage that the proof is simpler than the original one working with the fractions $p(X)/|X|$.

The algorithm works if a subroutine is available to

$$\text{find a subset} \quad X \subseteq S \quad \text{maximizing} \quad p(X) - \mu|X| \quad \text{for any fixed integer } \mu. \tag{2.4}$$

This subroutine will actually be needed only for special values of μ when $\mu = \lceil p(X)/\ell \rceil$ (where $X \subseteq S$ and $1 \leq \ell \leq n$). We do not have to assume that p is supermodular and the only requirement for the ND-algorithm is that Subroutine (2.4) be available. Via a submod-minimizer this is certainly the case when p happens to be supermodular (cf., Remark 2.1).

In several applications, the requested general purpose submod-minimizer can be superseded by a direct and more efficient algorithm such as the one for network flows or for matroid partition. The subroutine (2.4) is also available in the more general case (needed in applications) when p is only crossing supermodular. Indeed, for a given ordered pair of elements $s, t \in S$, the restriction of p on the family of $\bar{st}$-sets is fully supermodular, and therefore we can apply a submod-minimizer to each of the $n(n-1)$ ordered pairs (s, t) to get the requested maximum of $p(X) - \mu|X|$.

We call a value μ **good** if $\mu|X| \geq p(X)$ for every $X \subseteq S$. A value that is not good is called **bad**. Clearly, a sufficiently large μ is good. Our goal is to compute the minimum $\mu_{\min}$ of the good integers. This number is nothing but the maximum of $\lceil p(X)/|X| \rceil$ over non-empty subsets of S.

Let $\mu_0 := \lceil p(S)/|S| \rceil - 1$. This (possibly negative) number is bad and the algorithm starts with μ_0. Let

$$X_0 \in \arg\max\{p(X) - \mu_0|X| : X \subseteq S\},$$

that is, X_0 is a set maximizing the function $p(X) - \mu_0|X|$. Note that the badness of μ_0 implies that $p(X_0) > \mu_0|X_0|$.

The procedure determines one by one a series of pairs (μ_j, X_j) for subscripts $j = 1, 2, \ldots$ where each integer μ_j is a tentative candidate for μ while X_j is a non-empty subset of S. Suppose that the pair (μ_{j-1}, X_{j-1}) has already been determined for a subscript $j \geq 1$. Let μ_j be the smallest integer for which $\mu_j|X_{j-1}| \geq p(X_{j-1})$, that is,

$$\mu_j := \left\lceil \frac{p(X_{j-1})}{|X_{j-1}|} \right\rceil.$$

If μ_j is bad, that is, if there is a set $X \subseteq S$ with $p(X) - \mu_j|X| > 0$, then let

$$X_j \in \arg\max\{p(X) - \mu_j|X| : X \subseteq S\},$$

that is, X_j is a set maximizing the function $p(X) - \mu_j|X|$. (If there are more than one maximizing set, we can take any). Since μ_j is bad, we have $X_j \neq \emptyset$ and $p(X_j) - \mu_j|X_j| > 0$.

Claim 2.7. *If μ_j is bad for some subscript $j \geq 0$, then $\mu_j < \mu_{j+1}$.*

Proof. The badness of μ_j means that $p(X_j) - \mu_j|X_j| > 0$, from which

$$\mu_{j+1} = \left\lceil \frac{p(X_j)}{|X_j|} \right\rceil = \left\lceil \frac{p(X_j) - \mu_j|X_j|}{|X_j|} \right\rceil + \mu_j > \mu_j. \qquad \blacksquare$$

Since there is a good μ and the sequence μ_j is strictly monotone increasing by Claim 2.7, there will be a first subscript $h \geq 1$ for which μ_h is good. The algorithm terminates by outputting this μ_h (and in this case X_h is not computed).

Theorem 2.8. *If h is the first subscript during the run of the algorithm for which μ_h is good, then $\mu_{\min} = \mu_h$ (that is, μ_h is the requested smallest good μ-value) and $h \leq n$.*

Proof. Since μ_h is good and μ_h is the smallest integer for which $\mu_h|X_{h-1}| \geq p(X_{h-1})$, the set X_{h-1} certifies that no good integer μ can exist which is smaller than μ_h, that is, $\mu_{\min} = \mu_h$.

Claim 2.9. *If μ_j is bad for some subscript $j \geq 1$, then $|X_{j-1}| > |X_j|$.*

Proof. As μ_j $(= \lceil p(X_{j-1})/|X_{j-1}| \rceil)$ is bad, we obtain that

$$p(X_j) - \mu_j|X_j| > 0 = p(X_{j-1}) - \frac{p(X_{j-1})}{|X_{j-1}|}|X_{j-1}|$$

$$\geq p(X_{j-1}) - \left\lceil \frac{p(X_{j-1})}{|X_{j-1}|} \right\rceil |X_{j-1}| = p(X_{j-1}) - \mu_j|X_{j-1}|,$$

from which we get

$$p(X_j) - \mu_j|X_j| > p(X_{j-1}) - \mu_j|X_{j-1}|. \tag{2.5}$$

Since X_{j-1} maximizes $p(X) - \mu_{j-1}|X|$, we have

$$p(X_{j-1}) - \mu_{j-1}|X_{j-1}| \geq p(X_j) - \mu_{j-1}|X_j|. \tag{2.6}$$

By adding up (2.5) and (2.6), we obtain

$$(\mu_j - \mu_{j-1})|X_{j-1}| > (\mu_j - \mu_{j-1})|X_j|.$$

As μ_j is bad, so is μ_{j-1}, and hence, by applying Claim 2.7 to $j-1$ in place of j, we obtain that $\mu_j > \mu_{j-1}$, from which we arrive at $|X_{j-1}| > |X_j|$, as required. ∎

Claim 2.9 implies that $n \geq |X_0| > |X_1| > \cdots > |X_{h-1}| \geq 1$, from which $h \leq n$ follows. ∎∎

2.4 Computing a dec-min element in strongly polynomial time

In order to compute a dec-min element of an M-convex set $\overset{....}{B} = \overset{....}{B}'(p)$, our first task is to compute the smallest integer β_1 for which $\overset{....}{B}$ has an element with largest component β_1. Theorem 4.1 of [21] asserts that $\beta_1 = \max\{\lceil p(X)/|X| \rceil : \emptyset \neq X \subseteq S\}$. By applying the ND-algorithm described in Section 2.3, we can compute β_1 in strongly polynomial time. Note that, by Theorem 2.8, the algorithm terminates after at most n applications of Subroutine (2.4).

For any number β, a vector is said to be β-**covered** if each of its components is at most β. An element m of $\overset{....}{B}$ is called a **max-minimizer** if its largest component is as small as possible. A max-minimizer element m is said to be **pre-dec-min** in $\overset{....}{B}$ if the number of its largest components is as small as possible. Obviously, a dec-min element is pre-dec-min, and a pre-dec-min element is a max-minimizer.

Given the value of β_1, a β_1-covered element m of $\overset{....}{B}$ can easily be computed with a greedy-type algorithm as follows. Since there is a β_1-covered member of B, the vector $(\beta_1, \beta_1, \ldots, \beta_1)$ belongs to the so-called supermodular polyhedron $S'(p) := \{x : \widetilde{x}(X) \geq p(X) \text{ for every } X \subseteq S\}$. Consider the elements of S in an arbitrary order $\{s_1, \ldots, s_n\}$. Let $m(s_1) := \min\{z : (z, \beta_1, \beta_1, \ldots, \beta_1) \in S'(p)\}$. In the general step, if the components $m(s_1), \ldots, m(s_{i-1})$ have already been determined, let

$$m(s_i) := \min\{z : (m(s_1), m(s_2), \ldots, m(s_{i-1}), z, \beta_1, \beta_1, \ldots, \beta_1) \in S'(p)\}. \tag{2.7}$$

This computation can be carried out by n applications of a subroutine for a submodular function minimization.

Given a β_1-covered integral element m of B, our next goal is to modify m by a series of 1-tightening steps to obtain a pre-dec-min element of $\overset{....}{B}$. To this end, we consider one by one those elements t of S for which $m(t) = \beta_1$, and check whether a 1-tightening step can be applied at t with some element $s \in S$. That is, for each element $s \in S$ with $m(s) \leq \beta_1 - 2$, we check whether there is no m-tight $t\overline{s}$-set. If we find such an s, then $m' := m - \chi_t + \chi_s \in \overset{....}{B}$ and m' has one less coordinates with value β_1. (Note that the largest component of m' is also β_1 as β_1 was chosen to be the smallest upper bound). If no such element s exists, then

$$\beta_1 \geq m(s) \geq \beta_1 - 1 \text{ holds for each } s \in T_m(t), \tag{2.8}$$

11

where $T_m(t)$ is the smallest m-tight set containing t. While updating $m := m'$ in the first case, we iterate the above procedure for each element t of S with $m(t) = \beta_1$.

The element of $\overset{....}{B}$ obtained from this series of modifications, which is denoted by m, has the property (2.8) for all $t \in S$ with $m(t) = \beta_1$. Theorem 4.2 of [21] states that a β_1-covered element m of $\overset{....}{B}$ is pre-dec-min precisely if $m(s) \geq \beta_1 - 1$ for each $s \in S_1(m)$, where $S_1(m) = \cup\{T_m(t) : m(t) = \beta_1\}$. This and (2.8) imply that the final vector obtained by the above procedure is a pre-dec-min element of $\overset{....}{B}$ indeed.

Note that for a given pair (t, s) of elements, deciding whether a 1-tightening step is applicable or not can be done by a single submod-minimization, and hence deciding whether a given t with $m(t) = \beta_1$ admits a 1-tightening step reducing $m(t)$ can be done by at most n calls of a submod-minimizer. Altogether, the procedure above needs at most n 1-tightening steps which can be carried out by n^2 calls of a submod-minimizer.

Recall that $T_m(t)$ denoted the unique smallest tight set containing t when p is (fully) supermodular. But $T_m(t)$ can be described without explicitly referring to p since an element $s \in S$ belongs to $T_m(t)$ precisely if $m' := m - \chi_t + \chi_s$ is in B, and this is computable by the subroutine (2.2). In particular, we can compute $S_1(m)$ by (2.2). We use a short-hand notation $S_1 := S_1(m)$.

Let B_1 denote the restriction of the base-polyhedron B to S_1 and B_1' the contraction of B by S_1. Theorem 4.6 of [21] states that, for $m_1 \in \mathbf{Z}^{S_1}$ and $m_1' \in \mathbf{Z}^{S-S_1}$, (m_1, m_1') is a dec-min element of $\overset{....}{B}$ precisely if m_1 is a dec-min element of $\overset{....}{B}_1$ and m_1' is a dec-min element of $\overset{....}{B}_1'$. Let $m_1 := m|S_1$ for the pre-dec-min element m constructed above. Since m_1 is near-uniform on S_1, it is a dec-min element of $\overset{....}{B}_1$. Hence, if m_1' is a dec-min element of $\overset{....}{B}_1'$, then (m_1, m_1') is a dec-min element of $\overset{....}{B}$. Such a dec-min element m_1' can be computed by applying iteratively the computation described above for computing m_1. In this way we can compute a dec-min element of $\overset{....}{B}$.

It is worth mentioning that the restriction of the pre-dec-min element m to $S_1 = S_1(m)$ is the same as the restriction of the dec-min element (found by the above algorithm) to S_1. Therefore this S_1 is the first member of the canonical partition belonging to $\overset{....}{B}$.

The above algorithm computes a dec-min element of $\overset{....}{B}$ in strongly polynomial time. The subroutine (2.1) is called only once at the beginning of the algorithm, and the running time of the algorithm is governed by the number of calls of (2.2) and (2.4). First assume that we are given a fully supermodular function p to describe B. To determine m_1 and S_1, we need (i) n calls of (2.4) to compute the value β_1 by the ND-algorithm, where one call of (2.4) requires a single application of a submod-minimizer, (ii) n applications of a submod-minimizer to compute a β_1-covered element by the greedy-type algorithm, and (iii) n^2 calls of (2.2) to compute a pre-dec-min element, where one call of (2.2) requires a single application of a submod-minimizer. Therefore, we can determine m_1 and S_1 with $n + n + n^2 = O(n^2)$ applications of a submod-minimizer. We repeat the above procedure on $S - S_1$, $S - (S_1 \cup S_2)$, and so on. Hence the above algorithm finds a dec-min element and the canonical partition with $O(n^3)$ applications of a submod-minimizer, provided that we are given a fully supermodular function p to describe the base-polyhedron B. Even when B is given in terms of an intersecting or crossing supermodular function, this algorithm is strongly polynomial. Namely, the total number of submod-minimizer calls is $O(n^4)$ if p is intersecting supermodular and $O(n^5)$ if p is crossing supermodular (see Remark 2.2). Finally we recall Remark 2.1 for the complexity of a submod-minimizer.

3 Applications

3.1 Background

There are two major sources of applicability of the structural results on decreasing minimization on an M-convex set. One of them relies on the fact that the class of integral base-polyhedra is closed under several operations. For example, a face of a base-polyhedron is also a base-polyhedron, and so is the intersection of an integral box with a base-polyhedron B. Also, the sum of integral base-polyhedra

$B_1, \ldots, B_k$ is a base-polyhedron B which has, in addition, the integer decomposition property meaning that any integral element of B can be obtained as the sum of k integral elements by taking one from each B_i. This latter property implies that the sum of M-convex sets is M-convex. We also mention the important operation of taking an aggregate of a base-polyhedron, to be introduced below in Section 3.2.

The other source of applicability is based on the fact that not only fully super- or submodular functions can define base-polyhedra but some weaker functions as well. For example, if p is an integer-valued crossing (in particular, intersecting) supermodular function with finite $p(S)$, then $B = B'(p)$ is a (possibly empty) integral base-polyhedron (and $\overset{...}{B}$ is an M-convex set). This fact will be exploited in solving dec-min orientation problems when both degree-constraints and edge-connectivity requirements must be fulfilled. In some cases even weaker set-functions can define base-polyhedra. This is why we can solve dec-min problems concerning edge- and node-connectivity augmentations of digraphs.

3.2 Applications to matroids

Levin and Onn [42] solved algorithmically the following problem: *Find k bases of a matroid M on a ground-set S such that the sum of their characteristic vectors be decreasingly minimal.* Their approach, however, does not seem to work in the following natural extension. Suppose we are given k matroids $M_1, \ldots, M_k$ on a common ground-set S, and our goal is to find a basis B_i of each matroid M_i in such a way that the vector $\sum [\chi_{B_i} : i = 1, \ldots, k]$ is decreasingly minimal. Let B_Σ denote the sum of the base-polyhedra of the k matroids. By a theorem of Edmonds [9], the integral elements of B_Σ are exactly the vectors of form $\sum [\chi_{B_i} : i = 1, \ldots, k]$ where B_i is a basis of M_i. Therefore the problem is to find a dec-min element of $\overset{...}{B}_\Sigma$. This can be done by the basic algorithm described in Section 2.1. Let us see how the requested subroutines are available in this special case. The algorithm starts with an arbitrary member m of $\overset{...}{B}_\Sigma$ which is obtained by taking a basis B_i from each matroid M_i, and these bases define $m := \sum_i \chi_{B_i}$.

To realize Subroutine (2.2), we mentioned that it suffices to realize Subroutine (2.3), which requires for a given integral vector m' with $\widetilde{m}'(S) = \sum_i r_i(S)$ to decide whether m' is in $\overset{...}{B}_\Sigma$ or not. But this can simply be done by Edmonds' matroid intersection algorithm [11] (see also Section 13.1.2 in [16]. Namely, let $S_1, \ldots, S_k$ be disjoint copies of S and M'_i an isomorphic copy of M_i on S_i. Let N_1 be the direct sum of matroids M'_i on ground-set $S' := S_1 \cup \cdots \cup S_k$. Let N_2 be a partition matroid on S' in which a subset Z is a basis if it contains exactly $m'(s)$ members of the k copies of s for each $s \in S$. Then m' is in $\overset{...}{B}_\Sigma$ precisely if N_1 and N_2 have a common basis.

In conclusion, with the help of Edmonds' matroid intersection algorithm, Subroutine (2.2) is available, and hence the basic algorithm can be applied.

Another natural problem concerns a single matroid M on a ground-set T. Suppose we are given a partition $\mathcal{P} = \{T_1, \ldots, T_n\}$ of T and we consider the intersection vector $(|Z \cap T_1|, \ldots, |Z \cap T_n|)$ assigned to a basis Z of M. The problem is to find a basis for which the intersection vector is decreasingly minimal.

To solve this problem, we recall an important construction of base-polyhedra, called the aggregate. Let T be a ground-set and B_T an integral base-polyhedron in $\mathbf{R}^T$. Let $\mathcal{P} = \{T_1, \ldots, T_n\}$ be a partition of T into non-empty subsets and let $S = \{s_1, \ldots, s_n\}$ be a set whose elements correspond to the members of $\mathcal{P}$. The aggregate B_S of B_T is defined as follows.

$$B_S := \{(y_1, \ldots, y_n) : \text{there is an } x \in B_T \text{ with } y_i = \widetilde{x}(T_i) \ (i = 1, \ldots, n)\}. \tag{3.1}$$

A basic theorem concerning base-polyhedra (see, for example, Theorems 14.2.12 and 14.2.13 in book [16]) states that B_S is a base-polyhedron, moreover, for each integral member $(y_1, \ldots, y_n)$ of B_S, the vector x in (3.1) can be chosen integer-valued. In other words,

$$\overset{...}{B}_S := \{(y_1, \ldots, y_n) : \text{there is an } x \in \overset{...}{B}_T \text{ with } y_i = \widetilde{x}(T_i) \ (i = 1, \ldots, n)\}. \tag{3.2}$$

We call $\ddddot{B}_S$ the **aggregate** of $\ddddot{B}_T$.

Returning to our matroid problem, let B_T denote the base-polyhedron of matroid M. Then the problem is nothing but finding a dec-min element of $\ddddot{B}_S$.

We can apply the basic algorithm (concerning M-convex sets) for this special case since the requested subroutines are available through standard matroid algorithms. Namely, Subroutine (2.1) is available since for any basis Z of M, the intersection vector assigned to Z is nothing but an element of $\ddddot{B}_S$.

To realize Subroutine (2.2), we mentioned that it suffices to realize Subroutine (2.3). Suppose we are given a vector $y \in \mathbf{Z}_+^S$ (Here y stands for m' in (2.3)). Suppose that $\widetilde{y}(S) = r(T)$ (where r is the rank-function of matroid M) and that $y(s_i) \leq |T_i|$ for $i = 1, \ldots, n$.

Let $G = (S, T; E)$ denote a bipartite graph where $E = \{ts_i : t \in T_i, i = 1, \ldots, n\}$. By this definition, the degree of every node in T is 1 and hence the elements of E correspond to the elements of M. Let M_1 be the matroid on E corresponding to M (on T). Let M_2 be a partition matroid on E in which a set $F \subseteq E$ is a basis if $d_F(s_i) = y(s_i)$, where $d_F(s_i)$ denotes the number of edges in F for which s_i is an end-node. By this construction, the vector y is in $\ddddot{B}_S$ precisely if the two matroids M_1 and M_2 have a common basis. This problem is again tractable by Edmonds' matroid intersection algorithm.

As a special case, we can find a spanning tree of a (connected) directed graph for which its in-degree-vector is decreasingly minimal. Since the family of unions of k disjoint bases of a matroid forms also a matroid, we can also compute k edge-disjoint spanning trees in a digraph whose union has a decreasingly minimal in-degree vector.

Another special case is when we want to find a spanning tree of a connected bipartite graph $G = (S, T; E)$ whose in-degree vector restricted to S is decreasingly minimal.

3.3 Applications to flows

3.3.1 A base-polyhedron associated with net-in-flows

Let $D = (V, A)$ be a digraph endowed with integer-valued bounding functions $f : A \to \mathbf{Z} \cup \{-\infty\}$ and $g : A \to \mathbf{Z} \cup \{+\infty\}$ for which $f \leq g$. We call a vector (or function) z on A **feasible** if $f \leq z \leq g$. The **net-in-flow** Ψ_z of z is a vector on V and defined by $\Psi_z(v) = \varrho_z(v) - \delta_z(v)$, where $\varrho_z(v) := \sum[z(uv) : uv \in A]$ and $\delta_z(v) := \sum[z(vu) : uv \in A]$. If m is the net-in-flow of a vector z, then we also say that z is an m-**flow**.

A variation of Hoffman's classic theorem on feasible circulations [34] is as follows.

Lemma 3.1. *An integral vector $m : V \to \mathbf{Z}$ is the net-in-flow of an integral feasible vector (or in other words, there is an integer-valued feasible m-flow) if and only if $\widetilde{m}(V) = 0$ and*

$$\varrho_f(Z) - \delta_g(Z) \leq \widetilde{m}(Z) \ \text{holds whenever} \ Z \subseteq V, \tag{3.3}$$

where $\varrho_f(Z) := \sum[f(a) : a \in A$ and a enters $Z]$ and $\delta_g(Z) := \sum[g(a) : a \in A$ and a leaves $Z]$. ∎

Define a set-function p_{fg} on V by

$$p_{fg}(Z) := \varrho_f(Z) - \delta_g(Z).$$

Then p_{fg} is (fully) supermodular (see, e.g. Proposition 1.2.3 in [16]). Consider the base-polyhedron $B_{fg} := B'(p_{fg})$ and the M-convex set $\ddddot{B}_{fg}$. By Lemma 3.1 the M-convex set $\ddddot{B}_{fg}$ consists exactly of the net-in-flow integral vectors m.

By the algorithm described in Section 2, we can compute a decreasingly minimal element of $\ddddot{B}_{fg}$ in strongly polynomial time. By relying on a variant of the strongly polynomial push-relabel algorithm described in book [16] (see, Section 6.1.3), one can check whether or not (3.3) holds. If it does not, then this variant can compute a set most violating (3.3) (that is, a maximizer of $\varrho_f(Z) - \delta_g(Z) -$

$\widetilde{m}(Z)$), while if (3.3) does hold, then the push-relabel algorithm computes an integral valued feasible m-flow. Therefore the requested oracles in the general algorithm for computing a dec-min element are available through a network flow algorithm, and we do not have to rely on a general-purpose submodular function minimizing oracle.

For the sake of an application of this algorithm to capacitated dec-min orientations in Section 4.2, we remark that the algorithm can also be used to compute a dec-min element of the M-convex set obtained from $\overset{....}{B}_{fg}$ by translating it with a given integral vector.

3.3.2 Discrete version of Megiddo's flow problem

Megiddo [43], [44] considered the following problem. Let $D = (V, A)$ be a digraph endowed with a non-negative capacity function $g : A \rightarrow \mathbf{R}_+$. Let S and T be two disjoint non-empty subsets of V. Megiddo described an algorithm to compute a feasible flow from S to T with maximum flow amount M for which the net-in-flow vector restricted on S is (in our terms) increasingly maximal. Here a feasible flow is a vector x on A for which $\Psi_x(v) \leq 0$ for $v \in S$, $\Psi_x(v) \geq 0$ for $v \in T$, and $\Psi_x(v) = 0$ for $v \in V - (S \cup T)$. The flow amount x is $\sum[\Psi_x(t) : t \in T]$.

We emphasize that Megiddo solved the continuous (fractional) case and did not consider the corresponding discrete (or integer-valued) flow problem. To our knowledge, this natural optimization problem has not been investigated so far.

To provide a solution, suppose that g is integer-valued. Let $f \equiv 0$ and consider the net-in-flow vectors belonging to feasible vectors. These form a base-polyhedron B_1 in $\mathbf{R}^V$. Let B_2 denote the base-polyhedron obtained from B_1 by intersecting it with the box defined by $z(v) \leq 0$ for $v \in S$, $z(v) \geq 0$ for $v \in T$ and $z(v) = 0$ for $v \in V - (S \cup T)$.

The restriction of B_2 to S is a g-polymatroid Q in $\mathbf{R}^S$. And finally, we can consider the face of Q defined by $\widetilde{z}(S) = -M$. This is a base-polyhedron B_3 in $\mathbf{R}^S$, and the discrete version of Megiddo's flow problem is equivalent to finding an inc-max element of $\overset{....}{B}_3$. (Recall that an element of an M-convex set is dec-min precisely if it is inc-max.)

It can be shown that in this case again the general submodular function minimizing subroutine used in the algorithm to find a dec-min element of an M-convex set can be replaced by a max-flow min-cut algorithm.

A recent paper [22] addresses a more general problem to find an integral feasible flow that is dec-min on an arbitrarily specified edge set.

3.4 Further applications

3.4.1 Root-vectors of arborescences

A graph-example comes from packing arborescences. Let $D = (V, A)$ be a digraph and $k > 0$ an integer. We say that a non-negative integral vector $m : V \rightarrow \mathbf{Z}_+$ is a **root-vector** if there are k edge-disjoint spanning arborescences such that each node $v \in V$ is the root of $m(v)$ arborescences. Edmonds [10] classic result on disjoint arborescences implies that m is a root-vector if and only if $\widetilde{m}(V) = k$ and $\widetilde{m}(X) \geq k - \varrho(X)$ holds for every subset X with $\emptyset \subset X \subset V$. Define set-function p by $p(X) := k - \varrho(X)$ if $\emptyset \subset X \subseteq V$ and $p(\emptyset) := 0$. Then p is intersecting supermodular, so $B'(p)$ is an integral base-polyhedron. The intersection B of $B'(p)$ with the non-negative orthant is also a base-polyhedron, and the theorem of Edmonds is equivalent to stating that a vector m is a root-vector if and only if m is in $\overset{....}{B}$.

Therefore the general results on base-polyhedra can be specialized to obtain k disjoint spanning arborescences whose root-vector is decreasingly minimal.

3.4.2 Connectivity augmentations

Let $D = (V, A)$ be a directed graph and $k > 0$ an integer. We are interested in finding a so-called augmenting digraph $H = (V, F)$ of γ arcs for which $D + H$ is k-edge-connected or k-node-connected. In both cases, the in-degree vectors of the augmenting digraphs are the integral elements of an integral base-polyhedron [15], [17]. Obviously, the in-degree vectors of the augmented digraphs are the integral elements of an integral base-polyhedron.

Again, our results on general base-polyhedra can be specialized to find an augmenting digraph whose in-degree vector is decreasingly minimal.

4 Orientations of graphs

The literature is quite rich in graph orientation problems where the task is to orient the edges of an undirected graph such that the resulting digraph meet some expected properties. For an overview of graph orientation problems, see, for example, Chapter 61 of the book of Schrijver [53] or Chapter 9 of the book of Frank [16]. This latter one illuminates the deep connection between orientation problems and submodular optimization. Although quite general and efficient tools are exhibited in these books to manage graph orientation problems, they do not say anything about the orientation problem investigated first in a recent paper of Borradaile et al. [7], which is the decreasingly minimal orientation problem in our terms. It was actually their paper that triggered the whole research behind our present work. In this section, we exhibit how the theoretical background developed in [21] and the algorithms of Section 2 can be used to obtain major extensions of results in [7].

Let $G = (V, E)$ be an undirected graph. For $X \subseteq V$, let $i_G(X)$ denote the number of edges induced by X while $e_G(X)$ is the number of edges with at least one end-node in X. Then i_G is supermodular, e_G is submodular, and they are complementary functions, that is, $i_G(X) = e_G(V) - e_G(V - X)$. Let $B_G := B(e_G) = B'(i_G)$ denote the base-polyhedron defined by e_G or i_G.

We say that a function $m : V \to \mathbf{Z}$ is the **in-degree vector** of an orientation D of G if $\varrho_D(v) = m(v)$ for each node $v \in V$. An in-degree vector m obviously meets the equality $\widetilde{m}(V) = |E|$. The following basic result, sometimes called the Orientation lemma, is due to Hakimi [29].

Lemma 4.1 (Orientation lemma). *Let $G = (V, E)$ be an undirected graph and $m : V \to \mathbf{Z}$ an integral vector for which $\widetilde{m}(V) = |E|$. Then G has an orientation with in-degree vector m if and only if*

$$\widetilde{m}(X) \le e_G(X) \ \text{ for every subset } \ X \subseteq V, \tag{4.1}$$

which is equivalent to

$$\widetilde{m}(X) \ge i_G(X) \ \text{ for every subset } \ X \subseteq V. \tag{4.2}$$

∎

This immediately implies the following claim.

Claim 4.2. *The in-degree vectors of orientations of G are precisely the integral elements of base-polyhedron B_G ($= B(e_G) = B'(i_G)$), that is, the set of in-degree vectors of orientations of G is the M-convex set B_G.* ∎

The proof of Lemma 4.1 is algorithmic (see, e.g., Theorem 2.3.2 of [16]) and the orientation corresponding to a given m can be constructed easily.

4.1 Decreasingly minimal orientations

Due to Claim 4.2, we can apply the results on dec-min elements to the special base-polyhedron B_G. Borradaile et al. [7] called an orientation of G egalitarian if its in-degree vector is decreasingly minimal but we prefer the term **dec-min orientation** since an orientation with an increasingly maximal

in-degree vector also has an intuitive egalitarian feeling. Such an orientation is called **inc-max**. Theorem 2.1 immediately implies the following.

Corollary 4.3. *An orientation of G is dec-min if and only if it is inc-max.* ∎

Note that the term dec-min orientation is asymmetric in the sense that it refers to in-degree vectors. One could also aspire for finding an orientation whose out-degree vector is decreasingly minimal. But this problem is clearly equivalent to the in-degree version and hence in the present work we do not consider out-degree vectors, apart from a single exception in Section 4.5.

By Theorem 2.1, an element m of $\ddddot{B}_G$ is decreasingly minimal if and only if there is no 1-tightening step for m. What is the meaning of a 1-tightening step in terms of orientations?

Claim 4.4. *Let D be an orientation of G with in-degree vector m. Let t and s be nodes of G. The vector $m' := m + \chi_s - \chi_t$ is in B_G if and only if D admits a dipath from s to t.*

Proof. $m' \in B_G$ holds precisely if there is no $\bar{t}s$-set X which is tight with respect to i_G, that is, $\widetilde{m}(X) = i_G(X)$. Since $\varrho(Y) + i_G(Y) = \sum[\varrho(v) : v \in Y] = \widetilde{m}(Y)$ holds for any set $Y \subseteq V$, the tightness of X is equivalent to requiring that $\varrho(X) = 0$. Therefore $m' \in B_G$ if and only if $\varrho(Y) > 0$ holds for every $\bar{t}s$-set Y, which is equivalent to the existence of a dipath of D from s to t. ∎

Recall that a 1-tightening step at a member m of B_G consists of replacing m by m' provided that $m(s) \geq m(t) + 2$ and $m' \in B_G$. By Claim 4.4, a 1-tightening step at a given orientation of G corresponds to reorienting an arbitrary dipath from a node s to node t for which $\varrho(s) \geq \varrho(t) + 2$. Therefore, Theorem 2.1 immediately implies the following basic theorem of Borradaile et al. [7].

Theorem 4.5 (Borradaile et al. [7]). *An orientation D of a graph $G = (V, E)$ is decreasingly minimal if and only if no dipath exists from a node s to a node t for which $\varrho(t) \geq \varrho(s) + 2$.* ∎

Note that this theorem also implies Corollary 4.3. It immediately gives rise to an algorithm for finding a dec-min orientation. Namely, we start with an arbitrary orientation of G. We call a dipath **feasible** if $\varrho(t) \geq \varrho(s) + 2$ holds for its starting node s and end-node t. The algorithm consists of reversing feasible dipaths as long as possible. Since the sum of the squares of in-degrees always drops when a feasible dipath is reversed, and originally this sum is at most $|E|^2$, the dipath-reversing procedure terminates after at most $|E|^2$ reversals. By Theorem 4.5, when no more feasible dipath exists, the current orientation is dec-min. The basic algorithm concerning general base-polyhedra in Section 2.1 is nothing but an extension of the algorithm of Borradaile et al.

It should be noted that they suggested to choose at every step the current feasible dipath in such a way that the in-degree of its end-node t is as high as possible, and they proved that the algorithm in this case terminates after at most $|E||V|$ dipath reversals.

Note that we obtained Corollary 4.3 as a special case of a result on M-convex sets but it is also a direct consequence of Theorem 4.5.

4.2 Capacitated orientation

Consider the following capacitated version of the basic dec-min orientation problem of Borradaile et al. [7]. Suppose that a positive integer $\ell(e)$ is assigned to each edge e of G. Denote by G^+ the graph arising from G by replacing each edge e of G with $\ell(e)$ parallel edges. Our goal is to find a dec-min orientation of G^+. In this case, an orientation of G^+ is described by telling that, among the $\ell(e)$ parallel edges connecting the end-nodes u and v of e how many are oriented toward v (implying that the rest of the $\ell(e)$ edges are oriented toward u). In principle, this problem can be solved by applying the algorithm described above to G^+, and this algorithm is satisfactory when ℓ is small in the sense that its largest value can be bounded by a power of $|E|$. The difficulty in the general case is that the algorithm will be polynomial only in the number of edges of G^+, that is, in $\widetilde{\ell}(E)$, and hence this algorithm is not polynomial in $|E|$.

We show how the algorithm in Section 3.3.1 can be used to solve the decreasingly minimal orientation problem in the capacitated case in strongly polynomial time. To this end, let $D = (V, A)$ be an arbitrary orientation of G serving as a reference orientation. Define a capacity function g on A by $g(\vec{e}) := \ell(e)$, where $\vec{e}$ denotes the arc of D obtained by orienting e.

We associate an orientation of G^+ with an integral vector $z : A \rightarrow \mathbf{Z}_+$ with $z \leq g$ as follows. For an arc uv of D, orient $z(uv)$ parallel copies of $e = uv \in E$ toward v and $g(uv) - z(uv)$ parallel copies toward u. Then the in-degree of a node v is $m_z(v) := \varrho_z(v) + \delta_{g-z}(v) = \varrho_z(v) - \delta_z(v) + \delta_g(v)$. Therefore our goal is to find an integral vector z on A for which $0 \leq z \leq g$ and the vector m_z on V is dec-min. Consider the set of net-in-flow vectors $\{(\Psi_z(v) : v \in V) : 0 \leq z \leq g\}$. In Section 3.3.1, we proved that this is a base-polyhedron B_1. Therefore the set of vectors $(m_z(v) : v \in V)$ is also a base-polyhedron B arising from B_1 by translating B_1 with the vector $(\delta_g(v) : v \in V)$.

As remarked at the end of Section 3.3.1, a dec-min element of $\ddot{B}$ can be computed in strongly polynomial time by a variant (Section 6.1.3 in [16]) of the push-relabel subroutine for network flows (and not using a general-purpose submodular function minimizer).

4.3   Canonical chain and essential value-sequence for orientations

In Section 2.2, we described Algorithm 2.3 for an arbitrary M-convex set $\ddot{B}$ that computes, from a given dec-min element m of $\ddot{B}$, the canonical chain and essential value-sequence belonging to $\ddot{B}$. That algorithm needed an oracle for computing the smallest m-tight set $T_m(u)$ containing u. Here we show how this general algorithm can be turned into a pure graph-algorithm in the special case of dec-min orientations.

To this end, consider the special M-convex set, denoted by $\ddot{B}_G$, consisting of the in-degree vectors of the orientations of an undirected graph $G = (V, E)$. By the Orientation lemma, $B_G = B'(i_G)$ where $i_G(X)$ denotes the number of edges induced by X. Recall that i_G is a fully supermodular function. For an orientation D of G with in-degree vector m, the smallest m-tight set $T_m(t)$ (with respect to i_G) containing a node t will be denoted by $T_D(t)$.

Claim 4.6. *Let D be an arbitrary orientation of G with in-degree vector m. (A) A set $X \subseteq V$ is m-tight (with respect to i_G) if and only if $\varrho_D(X) = 0$. (B) The smallest m-tight set $T_D(t)$ containing a node t is the set of nodes from which t is reachable in D.*

Proof. We have
$$\varrho_D(X) + i_G(X) = \sum [\varrho_D(v) : v \in X] = \widetilde{m}(X) \geq i_G(X),$$
from which X is m-tight (that is, $\widetilde{m}(X) = i_G(X)$) precisely if $\varrho_D(X) = 0$, and Part (A) follows. Therefore the smallest m-tight set $T_D(t)$ containing t is the smallest set containing t with in-degree 0, and hence $T_D(t)$ is indeed the set of nodes from which t is reachable in D, as stated in Part (B). ■

By Claim 4.6, $T_D(t)$ is easily computable, and hence Algorithm 2.3 for general M-convex sets can easily be specialized to graph orientations. By applying Theorem 2.4 to $p := i_G$ and recalling from Claim 4.6 that C_i is m-tight, in the present case, precisely if $\varrho_D(C_i) = 0$, we obtain the following.

Theorem 4.7. *An orientation of D of G is dec-min if and only of $\varrho_D(C_i) = 0$ for each member C_i of the canonical chain and $\beta_i - 1 \leq \varrho_D(v) \leq \beta_i$ holds for every node $v \in S_i$ $(i = 1, \ldots, q)$.* ■

We remark that the members of the canonical partition computed by our algorithm for $\ddot{B}_G$ is exactly the non-empty members of the so-called density decomposition of G introduced by Borradaile et al. [8]. We also remark that by combining the approach of the present section with Section 4.2, the canonical partition and the essential value-sequence can be computed in strongly polynomial time even in the capacitated orientation problem.

4.4 Cheapest dec-min orientations

It is indicated in [21] that, in decreasing minimization on an M-convex set in general, we can construct an algorithm to compute a cheapest dec-min element with respect to a given (linear) cost-function on the ground-set. In the special case of dec-min orientations, this means that if c is a cost-function on the node-set of $G = (V, E)$, then we have an algorithm to compute a dec-min orientation of G for which $\sum[c(v)\varrho(v) : v \in V]$ is minimum.

But the question remains: what happens if, instead of a cost-function on the node-set, we have a cost-function c on $\vec{E}_2$, where $\vec{E}_2$ arises from E by replacing each element $e = uv (= vu)$ of E by two oppositely oriented arcs uv and vu, and we are interested in finding a cheapest orientation with specified properties? (As an orientation of e consists of replacing e by one of the two arcs uv and vu and the cost of its orientation is, accordingly, $c(uv)$ or $c(vu)$. Therefore we can actually assume that $\min\{c(uv), c(vu)\} = 0$.)

It is important to remark that the standard minimum cost in-degree specified orientation problem can be easily reduced, with a straightforward technique, to a minimum cost flow problem in a digraph with small integral capacities (see, e.g. Section 3.6.1 in book [16]). The same reduction works for min-cost in-degree constrained orientations, as well. Note that already the min-cost flow algorithm of Ford and Fulkerson [13] is strongly polynomial when the capacities are small integers (that is, we do not need here the significantly more sophisticated min-cost flow algorithm of Tardos [54] which is strongly polynomial for arbitrary capacities.) Actually, we shall need a version of this minimum cost orientation problem when some of the edges are already oriented, and this slight extension is also tractable by network flows.

Theorem 4.7 implies that the problem of finding a cheapest dec-min orientation is equivalent to finding a cheapest in-degree constrained orientation by orienting edges connecting C_i and $V - C_i$ toward $V - C_i$ ($i = 1, \ldots, q$). Here the in-degree constraints are given by $\beta_i - 1 \le \varrho_D(v) \le \beta_i$ for $v \in S_i$ ($i = 1, \ldots, q$).

Note that Harada et al. [30] provided a direct algorithm for the minimum cost version of the so-called semi-matching problem, which problem includes the minimum cost dec-min orientation problem. For this link, see Section 5.4.

We remark that by combining the approach of the present section with Section 4.2, a cheapest dec-min orientation can be computed in strongly polynomial time even in the capacitated orientation problem. In this case, however, one needs a strongly polynomial subroutine to compute a minimum cost feasible circulation. The first such algorithm is due to Tardos [54].

4.5 Orientation with dec-min in-degree vector and dec-min out-degree vector

We mentioned that dec-min and inc-max orientations always concern in-degree vectors. As an example to demonstrate the advantage of the general base-polyhedral view, we outline here one exception when in-degree vectors and out-degree vectors play a symmetric role. The problem is to characterize undirected graphs admitting an orientation which is both dec-min with respect to its in-degree vector and dec-min with respect to its out-degree vector.

For the present purposes, we let d_G denote the degree vector of G, that is, $d_G(v)$ is the number of edges incident to $v \in V$. (This notation differs from the standard set-function meaning of d_G.)

Let B_{in} denote the convex hull of the in-degree vectors of orientations of G, and B_{out} the convex hull of out-degree vectors of orientations of G. (Earlier B_{in} was denoted by B_G but now we have to deal with both out-degrees and in-degrees.) As before, $\overset{....}{B}_{\mathrm{in}}$ is the set of in-degree vectors of orientations of G, and $\overset{....}{B}_{\mathrm{out}}$ is the set of out-degree vectors of orientations of G. Let $\overset{....}{B}^{\bullet}_{\mathrm{in}}$ denote the set of dec-min in-degree vectors of orientations of G, and $\overset{....}{B}^{\bullet}_{\mathrm{out}}$ the set of dec-min out-degree vectors of orientations of G. By Theorem 2.5, both $\overset{....}{B}^{\bullet}_{\mathrm{in}}$ and $\overset{....}{B}^{\bullet}_{\mathrm{out}}$ are matroidal M-convex sets.

Note that the negative of a (matroidal) M-convex set is also a (matroidal) M-convex set, and the translation of a (matroidal) M-convex set by an integral vector is also a (matroidal) M-convex set.

Therefore $d_G - \ddddot{B}_{\text{out}}^\bullet$ is a matroidal M-convex set. Clearly, a vector m_{in} is the in-degree vector of an orientation D of G precisely if $d_G - m_{\text{in}}$ is the out-degree vector of D.

We are interested in finding an orientation whose in-degree vector is dec-min and whose out-degree vector is dec-min. This is equivalent to finding a member m_{in} of $\ddddot{B}_{\text{in}}^\bullet$ for which the vector $m_{\text{out}} := d_G - m_{\text{in}}$ is in the matroidal M-convex set $\ddddot{B}_{\text{out}}^\bullet$. But this latter is equivalent to requiring that m_{in} is in the M-convex set $d_G - \ddddot{B}_{\text{out}}^\bullet$. That is, the problem is equivalent to finding an element of the intersection of the matroidal M-convex sets $\ddddot{B}_{\text{in}}^\bullet$ and $d_G - \ddddot{B}_{\text{out}}^\bullet$. This latter problem can be solved by Edmonds matroid intersection algorithm [11].

5 In-degree constrained orientations of graphs

In this section we first describe an algorithm to find a dec-min in-degree constrained orientation. Second, we develop a complete description of the set of dec-min in-degree constrained orientations, which gives rise to an algorithm to compute a cheapest dec-min in-degree constrained orientation.

5.1 Computing a dec-min in-degree constrained orientation

Let $f : V \rightarrow \mathbf{Z} \cup \{-\infty\}$ be a lower bound function and $g : V \rightarrow \mathbf{Z} \cup \{+\infty\}$ an upper bound function for which $f \leq g$. We are interested in in-degree constrained orientations D of G, by which we mean that $f(v) \leq \varrho_D(v) \leq g(v)$ for every $v \in V$. Such an orientation is called (f, g)-**bounded**, and we assume that G has such an orientation. By a well-known theorem (see, Theorem 2.3.5 in [16]), such an orientation exists if and only if $i_G \leq \widetilde{g}$ and $\widetilde{f} \leq e_G$.

As before, let $\dddot{B}_G$ denote the M-convex set of the in-degree vectors of orientations of G, and let $\ddddot{B}_G^\square$ denote the intersection of $\dddot{B}_G$ with the integral box $T(f, g)$. That is, $\ddddot{B}_G^\square$ is the set of in-degree vectors of (f, g)-bounded orientations of G. Let D be an (f, g)-bounded orientation of G with in-degree vector m. We denote the smallest tight set containing a node t by $T_D^\square(t)$ $(= T_m^\square(t))$. By applying Claim 2.6 to $\ddddot{B}_G^\square$, we obtain that

$$T_D^\square(t) = \begin{cases} \{t\} & \text{if } \varrho_D(t) = f(t), \\ T_D(t) - \{s : \varrho_D(s) = g(s)\} & \text{if } \varrho_D(t) > f(t), \end{cases} \tag{5.1}$$

implying that, in case $\varrho_D(t) > f(t)$, the set $T_D^\square(t)$ consists of those nodes s from which t is reachable and for which $\varrho_D(s) < g(s)$.

Formula (5.1) implies for distinct nodes s and t that the vector $m' := m + \chi_s - \chi_t$ belongs to $\ddddot{B}_G^\square$ precisely if there is an st-dipath (i.e. a dipath from s to t) for which $\varrho_D(s) < g(s)$ and $\varrho_D(t) > f(t)$. We call such a dipath P of D **reversible**. Note that the dipath P' of D' obtained by reorienting P is reversible in D'.

If P is a reversible st-dipath of D for which $\varrho_D(t) \geq \varrho_D(s) + 2$, then the orientation D' is decreasingly smaller than D. We call such a dipath **improving**. Therefore, reorienting an improving st-dipath corresponds to a 1-tightening step. Hence Theorem 2.1 implies the following extension of Theorem 4.5.

Theorem 5.1. *An (f, g)-bounded orientation D of G is dec-min if and only if there is no improving dipath, that is, a dipath from a node s to a node t for which $\varrho_D(t) \geq \varrho_D(s) + 2$, $\varrho_D(s) < g(s)$, and $\varrho_D(t) > f(t)$.* ∎

In Section 2.1 we have presented an algorithm that computes a dec-min element of an arbitrary M-convex set. By specializing it to $\ddddot{B}_G^\square$, we conclude that in order to construct a dec-min (f, g)-bounded orientation of G, one can start with an arbitrary (f, g)-bounded orientation, and then reorient (currently) improving dipaths one by one, as long as such a dipath exists. As we pointed out after

Theorem 4.5, after at most $|E|^2$ improving dipath reorientations, the algorithm terminates with a dec-min (f, g)-bounded orientation of G.

Canonical chain and essential value-sequence for (f, g)-bounded orientations   In Section 2.2, we indicated that Algorithm 2.3 can immediately be applied to compute the canonical chain, the canonical partition, and the essential value-sequence belonging to the intersection $\overset{....}{B}{}^{\square}$ of an arbitrary M-convex set $\overset{....}{B}$ with an integral box $T(f, g)$.

   This algorithm needs only the original subroutine to compute $T_m(u)$ since, by Claim 2.6, $T_m^{\square}(u)$ is easily computable from $T_m(u)$. As we indicated above, in the special case of orientations, the corresponding sets $T_D(t)$ and $T_D^{\square}(t)$ are immediately computable from D. Therefore this extended algorithm can be used in the special case when we are interested in dec-min (f, g)-bounded orientations of $G = (V, E)$. The algorithm starts with a dec-min (f, g)-bounded orientation D of G and outputs the canonical chain $C^{\square} = \{C_1^{\square}, \ldots, C_q^{\square}\}$, the canonical partition $\mathcal{P}^{\square} = \{S_1^{\square}, \ldots, S_q^{\square}\}$, and the essential value-sequence $\beta_1^{\square} > \cdots > \beta_q^{\square}$. In view of Theorem 2.4, we also define bounding functions f^* and g^* as

$$f^*(v) := \beta_i^{\square} - 1 \text{ if } v \in S_i \qquad (i = 1, \ldots, q),$$
$$g^*(v) := \beta_i^{\square} \qquad \text{if } v \in S_i \qquad (i = 1, \ldots, q).$$

We say that the small box

$$T^* := T(f^*, g^*) \tag{5.2}$$

belongs to $\overset{....}{B}{}_G^{\square}$. Clearly, $f \le f^*$ and $g^* \le g$, and hence $T(f^*, g^*) \subseteq T(f, g)$. In Section 5.2 below we assume that these data are available.

Remark 5.1. A special case of in-degree constrained orientations is when we have a prescribed subset T of V and a non-negative function $m_T : T \to \mathbf{Z}_+$ serving as an in-degree specification on T, and we are interested in orientations of G for which $\varrho(v) = m_T(v)$ holds for every $t \in T$. We call such an orientation T-**specified**. This notion will have applications in Section 5.4. ∎

5.2   Cheapest dec-min in-degree constrained orientations

We are given a cost-function c on the possible orientations of the edges of G and our goal is to find a cheapest dec-min (f, g)-bounded orientation of G. This will be done with the help of a purely graphical description of the set of all dec-min (f, g)-bounded orientations, which is given in Theorem 5.3.

   As a preparation, we derive the following claim as an immediate consequence of the structural result stated in Theorem 2.5. Let m be a dec-min element of an M-convex set $\overset{....}{B}$ on ground-set S. Suppose that $m' := m + \chi_s - \chi_t$ is in $\overset{....}{B}$ (that is, $s \in T_m(t)$). Since m is dec-min, $m(t) \le m(s) + 1$. If $m(t) = m(s) + 1$, then m' and m are value-equivalent and hence m' is also a dec-min element of $\overset{....}{B}$. We say that m' is obtained from m by an **elementary step**.

Claim 5.2. *Any dec-min element of $\overset{....}{B}$ can be obtained from a given dec-min element m by a sequence of at most $|S|$ elementary steps.*

Proof. By Theorem 2.5, the set of dec-min elements of $\overset{....}{B}$ is a matroidal M-convex set in the sense that it can be obtained from a matroid M^* by translating the incidence vectors of the bases of M^* by the same integral vector Δ^*. A simple property of matroids is that any basis can be obtained from a given basis through a sequence of at most $|S|$ bases such that each member of the series can be obtained from the preceding one by taking out one element and adding a new one. The corresponding change in the translated vector is exactly an elementary step. ∎

Theorem 5.3. *Let $G = (V, E)$ be an undirected graph admitting an (f, g)-bounded orientation. Let $\overset{\dots}{B}{}^{\square}_G$ denote the M-convex set consisting of the in-degree vectors of (f, g)-bounded orientations of G, and let T^* be the small box, belonging to $\overset{\dots}{B}{}^{\square}_G$, as defined in (5.2). There is a chain $\mathcal{Z}$ of subsets of V such that an (f, g)-bounded orientation D of G is a dec-min (f, g)-bounded orientation if and only if D is an orientation of G whose in-degree vector belongs to T^* and $\delta_D(Z) = 0$ holds for each $Z \in \mathcal{Z}$.*

Proof. Let D be a dec-min (f, g)-bounded orientation of G, and let m denote its in-degree vector. Consider the canonical chain $C^{\square} = \{C^{\square}_1, \dots, C^{\square}_q\}$, the canonical partition $\mathcal{P}^{\square} = \{S^{\square}_1, \dots, S^{\square}_q\}$, and the essential value-sequence $\beta^{\square}_1 > \cdots > \beta^{\square}_q$ belonging to $\overset{\dots}{B}{}^{\square}_G$.

For $i \in \{1, \dots, q\}$, define
$$F_i := \{v : v \in S^{\square}_i, f(v) = \beta^{\square}_i\}.$$

Since $f(v) \leq m(v) \leq \beta^{\square}_i$ holds for every element v of $S^{\square}_i$, we obtain that $f(v) = m(v) = \beta^{\square}_i$ for $v \in F_i$. Note that F_i does not depend on D.

Claim 5.4. *For every $h = 1, \dots, i$, there is no dipath P from a node $s \in V - C^{\square}_h$ with $m(s) < g(s)$ to a node $t \in S^{\square}_h$ with $\beta^{\square}_h > f(t)$.*

Proof. Suppose indirectly that there is such a dipath P. If $m(t) = \beta^{\square}_h$, then P would be an improving dipath which is impossible since D is dec-min (f, g)-bounded. Therefore $m(t) = \beta^{\square}_h - 1$. But a property of the canonical partition is that there is an element t' of $S^{\square}_h - F_h$ for which $m(t') = \beta^{\square}_h$ and $t \in T^{\square}_D(t')$. This means that t' is reachable from t in D, and therefore there is a dipath from s to t' in D which is improving, a contradiction again. ∎

We are going to define a chain $\mathcal{Z}$ of subsets $Z_1 \supseteq Z_2 \supseteq \cdots \supseteq Z_q (= \emptyset)$ of V with the help of D, and will show that this chain actually does not depend on D. Let
$$Z_i := \{t : t \text{ is reachable in } D \text{ from a node } s \in V - C^{\square}_i \text{ with } \varrho_D(s) < g(s)\}. \tag{5.3}$$

Note that $Z_{i-1} \supseteq Z_i$ follows from the definition, where equality holds precisely if $\varrho_D(s) = g(s)$ for each $s \in S_i$.

Lemma 5.5. *Every dec-min (f, g)-bounded orientation defines the same family $\mathcal{Z}$.*

Proof. By Claim 5.2, it suffices to prove that a single elementary step does not change $\mathcal{Z}$. An elementary step in $\overset{\dots}{B}{}^{\square}_G$ corresponds to the reorientation of an st-dipath P in D where $s, t \in S^{\square}_h - F_h$, $m(t) = \beta^{\square}_h$ and $m(s) = \beta^{\square}_h - 1$ hold for some $h \in \{1, \dots, q\}$. We will show for $i \in \{1, \dots, q\}$ that the reorientation of P does not change Z_i.

If $h \leq i$, then Claim 5.4 implies that $Z_i \cap S^{\square}_h \subseteq F_h$. Since $\delta_D(Z_i) = 0$, the dipath P is disjoint from Z_i, implying that reorienting P does not affect Z_i.

Suppose now that $h \geq i + 1$. Since reorienting P results in a dec-min (f, g)-bounded orientation D', we get that $m(s) + 1 \leq g(s)$ and hence $s \in Z_i - C^{\square}_i$. Since $\delta_D(Z_i) = 0$, we obtain that $t \in Z_i - C^{\square}_i$. Since $\varrho_{D'}(t) = \varrho_D(t) - 1 < g(t)$ and the set of nodes reachable from s in D is equal to the set of nodes reachable from t in D', it follows that the reorientation of P does not change Z_i. ∎

Lemma 5.6. *The chain $\mathcal{Z}$ defined above meets the requirements of the theorem.*

Proof. Consider first an arbitrary dec-min (f, g)-bounded orientation D of G. Lemma 5.5 and the definition of Z_i in (5.3) imply that no arc of D can leave any $Z_i \in \mathcal{Z}$, that is, $\delta_D(Z_i) = 0$. Furthermore, by applying Theorem 2.4 to the present base-polyhedron $\overset{\dots}{B}{}^{\square}$, we obtain that the in-degree vector of D belongs to T^*.

Conversely, let D be an orientation of G for which $\delta_D(Z) = 0$ holds for every $Z \in \mathcal{Z}$ and the in-degree vector of D belongs to T^*, that is,
$$f^*(v) \leq \varrho_D(v) \leq g^*(v) \text{ for every } v \in V.$$

Then D is clearly (f, g)-bounded. The following claim makes complete the proof of the lemma.

22

Claim 5.7. *There is no improving dipath in D.*

Proof. Suppose, indirectly, that P is an improving st-dipath, that is, a dipath from s to t such that $\varrho_D(t) \geq \varrho_D(s) + 2$, $\varrho_D(t) > f(t)$, and $\varrho_D(s) < g(s)$. Suppose that t is in $S_i^\square$ for some $i \in \{1, \ldots, q\}$. If s is in $S_k^\square$ for some $k \in \{1, \ldots, q\}$, then

$$\beta_k^\square - 1 \leq \varrho_D(s) \leq \varrho_D(t) - 2 \leq \beta_i^\square - 2,$$

that is, $\beta_k^\square < \beta_i^\square$, and hence $k > i$, implying that s is in $V - C_i^\square$. This and $\varrho_D(s) < g(s)$ imply that s is in Z_i. Furthermore, $\beta_i^\square \geq \varrho_D(t) > f(t)$ implies that t is not in F_i, and since $S_i^\square \cap Z_i \subseteq F_i$, we obtain that t is not in Z_i. On the other hand, we must have $t \in Z_i$, since there is a dipath from $s \in V - C_i^\square$ to t and $\varrho_D(s) < g(s)$. This is a contradiction. ∎

By proving Claim 5.7, we have shown Lemma 5.6. Thus the proof of Theorem 5.3 is completed. ∎∎

Algorithm for computing a cheapest dec-min (f, g)-bounded orientation   First we compute a dec-min (f, g)-bounded orientation D of G with the help of the algorithm outlined in Section 5.1. Second, by applying the algorithm described in the same section, we compute the canonical chain and partition belonging to $\overset{....}{B_G^\square}$ along with the essential value-sequence. Once these data are available, the sets Z_i ($i = 1, \ldots, q$) defined in (5.3) are easily computable. Lemma 5.5 ensures that these sets Z_i do not depend on the starting dec-min (f, g)-bounded orientation D. Let E_0 be the union of the set of edges connecting some Z_i with $V - Z_i$, and define the orientation A_0 of E_0 by orienting each edge between Z_i and $V - Z_i$ toward Z_i.

Theorem 5.3 implies that, once E_0 and its orientation A_0 are available, the problem of computing a cheapest dec-min (f, g)-bounded orientation of G reduces to finding cheapest in-degree constrained (namely, (f^*, g^*)-bounded) orientation of a mixed graph. We indicated already in Section 4.4 that such a problem is easily solvable by the strongly polynomial min-cost flow algorithm of Ford and Fulkerson in a digraph with identically 1 capacities.

Remark 5.2. In Section 4.2 we have considered the capacitated dec-min orientation problem in the basic case where no in-degree constraints are imposed. With the technique presented there, we can cope with the capacitated, min-cost, in-degree constrained variants as well. Furthermore, the algorithms above can easily be extended, with a slight modification, to the case when one is interested in orientations of mixed graphs. ∎

5.3   Dec-min (f, g)-bounded orientations minimizing the in-degree of T

In this section, we discuss a rather specific orientation problem in full detail. The reason is that this framework is an indispensable tool for solving in Section 5.4 a common generalization of several previously investigated resource allocation problems. Our approach permits us to manage algorithmically even the minimum cost version of these problems.

One may consider (f, g)-bounded orientations of G when the additional requirement is imposed that the in-degree of a specified subset T of nodes be as small as possible. We shall show that these orientations of G can be described as (f', g')-bounded orientations of a mixed graph arising from G by orienting the edges between a certain subset X_T of nodes and its complement $V - X_T$ toward $V - X_T$.

It is more comfortable, however, to show the analogous statement for a general M-convex set $\overset{....}{B}'(p) \subseteq \mathbf{Z}^V$ defined by a (fully) supermodular function p for which $\overset{....}{B}^\square := \overset{....}{B}'(p) \cap T(f, g)$ is non-empty. (Here, instead of the usual S, we use V to denote the ground-set of the general M-convex set. We are back at the special case of graph orientations when $p = i_G$.) We assume that each of p, f, and g is finite-valued.

23

Let $p^\square$ denote the unique (fully) supermodular function defining $B^\square$. This function can be expressed with the help of p, f, and g, as follows (see, for example, Theorem 14.3.9 in [16]):

$$p^\square(Y) = \max\{p(X) + \widetilde{f}(Y - X) - \widetilde{g}(X - Y) : X \subseteq V\} \qquad (Y \subseteq V). \qquad (5.4)$$

As $B^\square$ is defined by the supermodular function $p^\square$ (that is, $B^\square = B'(p^\square)$), we have

$$\min\{\widetilde{m}(T) : m \in \overset{....}{B^\square}\} = p^\square(T). \qquad (5.5)$$

This implies that the set of elements of $\overset{....}{B^\square}$ minimizing $\widetilde{m}(T)$ is itself an M-convex set. Namely, it is the set of integral elements of the base-polyhedron arising from $B^\square$ by taking its face defined by $\{m \in B^\square : \widetilde{m}(T) = p^\square(T)\}$. The next theorem shows how this M-convex set can be described in terms of f, g, and p, without referring to $p^\square$.

Theorem 5.8. *There is a box $T(f', g') \subseteq T(f, g)$ and a subset $X_T \subseteq V$ such that an element $m \in \overset{....}{B^\square}$ minimizes $\widetilde{m}(T)$ if and only if $\widetilde{m}(X_T) = p(X_T)$ and $m \in \overset{....}{B} \cap T(f', g')$.*

Proof. Let X_T be a set maximizing the right-hand side of (5.4) for the definition of $p^\square(Y)$.

Claim 5.9. *An element $m \in \overset{....}{B^\square}$ is a minimizer of the left-hand side of (5.5) if and only if the following three optimality criteria hold:*

$$\widetilde{m}(X_T) = p(X_T),$$
$$v \in T - X_T \quad \text{implies} \quad m(v) = f(v),$$
$$v \in X_T - T \quad \text{implies} \quad m(v) = g(v).$$

Proof. For any $m \in \overset{....}{B^\square}$ and $X \subseteq V$, we have $\widetilde{m}(T) = \widetilde{m}(X) + \widetilde{m}(T - X) - \widetilde{m}(X - T) \geq p(X) + \widetilde{f}(T - X) - \widetilde{g}(X - T)$. Here we have equality if and only if $\widetilde{m}(X) = p(X)$, $\widetilde{m}(T - X) = \widetilde{f}(T - X)$, and $\widetilde{m}(X - T) = \widetilde{g}(X - T)$, implying the claim. ∎

Define f' and g' as follows:

$$f'(v) := \begin{cases} g(v) & \text{if } v \in X_T - T, \\ f(v) & \text{if } v \in V - (X_T - T), \end{cases} \qquad (5.6)$$

$$g'(v) := \begin{cases} f(v) & \text{if } v \in T - X_T, \\ g(v) & \text{if } v \in V - (T - X_T). \end{cases} \qquad (5.7)$$

The claim implies that $T(f', g')$ and X_T meet the requirement of the theorem. ∎∎

As the set of elements of $\overset{....}{B^\square}$ minimizing $\widetilde{m}(T)$ is itself an M-convex set, all the algorithms developed earlier can be applied once we are able to compute set X_T occurring in Theorem 5.8. (By definitions (5.6) and (5.7), X_T immediately determines f' and g').

The following straightforward algorithm computes an element $m \in \overset{....}{B^\square}$ minimizing the left-hand side of (5.5) and a subset X_T maximizing the right-hand side of (5.4). Start with an arbitrary element $m \in \overset{....}{B^\square}$. By an improving step we mean the change of m to $m' := m + \chi_s - \chi_t$ for some elements $s \in V - T, t \in T$ for which $m(s) < g(s)$, $m(t) > f(t)$, and $s \in T_m(t)$, where $T_m(t)$ is the smallest m-tight set (with respect to p) containing t. Clearly, $m' \in \overset{....}{B^\square}$, and $\widetilde{m}'(T) = \widetilde{m}(T) - 1$. The algorithm applies improving steps as long as possible. When no more improving step exists, the set $X_T := \cup(T_m(t) : t \in T, m(t) > f(t))$ meets the three optimality criteria. The algorithm is polynomial if $|p(X)|$ is bounded by a polynomial of $|V|$.

By applying Theorem 5.8 to the special case of $p = i_G$, we obtain the following.

Corollary 5.10. *Let $G = (V, E)$ be a graph admitting an (f, g)-bounded orientation. There is a box $T(f', g') \subseteq T(f, g)$ and a subset $X_T \subseteq V$ such that an (f, g)-bounded orientation of G minimizes the in-degree of T if and only if D is an (f', g')-bounded orientation for which $\varrho_D(X_T) = 0$.* ∎

In this case, the algorithm above to compute X_T starts with an (f, g)-bounded orientation D of G, whose in-degree vector is denoted by m. As long as there is an st-dipath P with $s \in V - T, t \in T, m(s) < g(s)$, and $m(t) > f(t)$, reorient P. When no such a dipath exists anymore, the set X_T of nodes from which a node $t \in T$ with $m(t) > f(t)$ is reachable in D, along with the bounding functions f' and g' defined in (5.6) and in (5.7), meet the requirement in the corollary.

Minimum cost version It follows that, in order to compute a minimum cost dec-min (f, g)-bounded orientation for which the in-degree of T is minimum, we can apply the algorithm described in Section 5.2 for the mixed graph obtained from G by orienting each edge between X_T and $V - X_T$ toward $V - X_T$.

Remark 5.3. Instead of a single subset T of V, we may consider a chain $\mathcal{T}$ of subsets $T_1 \subset T_2 \subset \cdots \subset T_h$ of V. Then $\mathcal{T}$ defines a face $B_{\text{face}}^\square$ of the base-polyhedron $B^\square$. Namely, an element m of $B^\square$ belongs to $B_{\text{face}}^\square$ precisely if $\widetilde{m}(T_i) = p^\square(T_i)$ for each $i \in \{1, \ldots, h\}$. This implies that the integral elements of $B_{\text{face}}^\square$ simultaneously minimize $\widetilde{m}(T_i)$ for each $i \in \{1, \ldots, h\}$ (over the elements of $\overset{....}{B^\square}$). Therefore, we can consider (f, g)-bounded orientations of G with the additional requirement that each of the in-degrees of $T_1, T_2, \ldots, T_h$ is (simultaneously) minimum. Corollary 5.10 can be extended to this case, implying that we have an algorithm to compute a minimum cost dec-min (f, g)-bounded orientation of G that simultaneously minimizes the in-degree of each member of the chain $\{T_1, T_2, \ldots, T_h\}$. ∎

5.4 Application in resource allocation: semi-matchings

For a general M-convex set $\overset{....}{B}$, it is known (see Theorem 2.2 in Section 2.1) that an element m of $\overset{....}{B}$ is dec-min if and only if m is a minimizer of the square-sum of the components over $\overset{....}{B}$. Therefore the corresponding equivalences hold in the special case of in-degree constrained (in particular, T-specified) orientations of undirected graphs.

As an application of this equivalence, we show first how a result of Harvey et al. [31] concerning a resource allocation problem follows immediately. They introduced the notion of a semi-matching of a simple bipartite graph $G = (S, T; E)$ as a subset F of edges for which $d_F(t) = 1$ holds for every node $t \in T$, and solved the problem of finding a semi-matching F for which $\sum[d_F(s)(d_F(s) + 1) : s \in S]$ is minimum. (Here $d_F(v)$ denotes the number of edges in F for which v is an end-node.) The problem was motivated by practical applications in the area of resource allocation in computer science. Note that

$$\sum[d_F(s)(d_F(s) + 1) : s \in S] = \sum[d_F(s)^2 : s \in S] + \sum[d_F(s) : s \in S]$$
$$= \sum[d_F(s)^2 : s \in S] + |F| = \sum[d_F(s)^2 : s \in S] + |T|,$$

and therefore the problem of Harvey et al. is equivalent to finding a semi-matching F of G that minimizes the square-sum of degrees in S.

By orienting each edge in F toward S and each edge in $E - F$ toward T, a semi-matching can be identified with the set of arcs directed toward S in an orientation of $G = (S, T; E)$ in which the out-degree of every node $t \in T$ is 1 (that is, $\varrho(t) = d_G(t) - 1$), and $d_F(s) = \varrho(s)$ for each $s \in S$. Since $\varrho(t)$ for $t \in T$ is the same in these orientations, it follows that the total sum of $\varrho(v)^2$ over $S \cup T$ is minimized precisely if $\sum[\varrho(s)^2 : s \in S] = \sum[d_F(s)^2 : s \in S]$ is minimized. Therefore the semi-matching problem of Harvey et al. is nothing but a special dec-min T-specified orientation problem. Note that not only semi-matching problems can be managed with graph orientations, but conversely, an orientation of a graph $G = (V, E)$ can also be interpreted as a semi-matching of the bipartite graph obtained from

G by subdividing each edge by a new node (where subdividing an edge $e = uv$ formally means that we replace e by a path $(uz_e, z_e v)$ of length two with a new node z_e). This implies, for example, that the algorithm of Harvey et al. to compute a semi-matching minimizing $\sum[d_F(v)^2 : v \in S]$ is able to compute an orientation of a graph G for which $\sum[\varrho(v)^2 : v \in S]$ is minimum. Furthermore, an orientation of a hypergraph means that we assign an element of each hyper-edge Z to Z as its head. In this sense, semi-matchings of bipartite graphs and orientations of hypergraphs are exactly the same. Several graph orientation results have been extended to hypergraph orientation, for an overview, see, e.g. [16].

Bokal et al. [6] extended the results to subgraphs of G meeting a more general degree-specification on T when, rather than the identically 1 function, one imposes an arbitrary degree-specification m_T on T satisfying $0 \leq m_T(t) \leq d_G(t)$ $(t \in T)$. The same orientation approach applies in this more general setting. We may call a subset F of edges an m_T-**semi-matching** if $d_F(t) = m_T(t)$ for each $t \in T$. The extended resource allocation problem is to find an m_T-semi-matching F that minimizes $\sum[d_F(s)^2 : s \in S]$. This is equivalent to finding a T-specified orientation of G for which the square-sum of the in-degrees is minimum and the in-degree specification in $t \in T$ is $m'_T(t) := d_G(t) - m_T(t)$. Therefore this extended resource allocation problem is equivalent to finding a dec-min T-specified orientation of G.

The same orientation approach, when applied to in-degree constrained orientations, allows us to extend the m_T-semi-matching problem when we have upper and lower bounds imposed on the nodes in S. This may be a natural requirement in practical applications where the elements of S correspond to available resources (e.g. computers), the elements of T correspond to users, and we are interested in a fair (= dec-min = square-sum minimizer) distribution (=m_T-semi-matchings) of the resources when the load (or burden) of each resource is requested to meet a specified upper and/or lower bound. Note that in the resource allocation framework, the degree $d_F(s)$ of node $s \in S$ may be interpreted as the burden of s, and hence a difference-sum minimizer semi-matching minimizes the total sum of burden-differences.

Katrenič and Semanišin [40] investigated the problem of finding a dec-min "maximum (f, g)-semi-matching" problem where there is a lower-bound function f_T on T and an upper bound function g_S on S (in the present notation) and one is interested in maximum cardinality subgraphs of G meeting these bounds. They describe an algorithm to compute a dec-min subgraph of this type. With the help of the orientation model discussed in Section 5.3 (where, besides the in-degree bounds on the nodes, the in-degree of a specified subset T was requested to be minimum), we have a strongly polynomial algorithm for an extension of the model of [40] when there may be upper and lower bounds on both S and T. Actually, even the minimum cost version of this problem was solved in Section 5.3.

In another variation, we also have degree bounds (f_S, g_S) on S and (f_T, g_T) on T, but we impose an arbitrary positive integer γ for the cardinality of F. We consider degree-constrained subgraphs $(S, T; F)$ of G for which $|F| = \gamma$, and want to find such a subgraph for which $\sum[d_F(s)^2 : s \in S]$ is minimum. (Notice the asymmetric role of S and T.) This is equivalent to finding an in-degree constrained orientation D of G for which $\varrho_D(S) = \gamma$ and $\sum[\varrho_D(s)^2 : s \in S]$ is minimum. Here the corresponding in-degree bound (f, g) on S is the given (f_S, g_S) while (f, g) on T is defined for $t \in T$ by

$$f(t) := d_G(t) - g_T(t) \quad \text{and} \quad g(t) := d_G(t) - f_T(t).$$

Let B denote the base-polyhedron spanned by the in-degree vectors of the degree-constrained orientations of G. Then the restriction of B to S is a g-polymatroid Q. By intersecting Q with the hyperplane $\{x : \widetilde{x}(S) = \gamma\}$, we obtain an integral base-polyhedron B_S in $\mathbf{R}^S$, and then the elements of $\ddot{B}_S$ are exactly the in-degree vectors of the requested orientations restricted to S. That is, the elements of $\ddot{B}_S$ are the restriction of the degree-vectors of the requested subgraphs of G to S. Since B_S is a base-polyhedron, a dec-min element of $\ddot{B}_S$ will be a solution to our minimum degree-square sum problem.

We briefly indicate that a capacitated version of the semi-matching problem can also be formu-

lated as a dec-min in-degree constrained and capacitated orientation problem (cf., Section 4.2 and Remark 5.2). Let $G = (S, T; E)$ be again a bipartite graph, γ a positive integer, and f_V and g_V integer-valued bounding functions on $V := S \cup T$ for which $f_V \leq g_V$. In addition, an integer-valued capacity function g_E is also given on the edge-set E, and we are interested in finding a non-negative integral vector $z : E \rightarrow \mathbf{Z}_+$ for which $\widetilde{z}(E) = \gamma$, $z \leq g_E$ and $f_V(v) \leq d_z(v) \leq g_V(v)$ for every $v \in V$. (Here $d_z(v) := \sum[z(uv) : uv \in E]$.) We call such a vector **feasible**. The problem is to find a feasible vector z whose degree vector restricted to S (that is, the vector $(d_z(s) : s \in S)$ is decreasingly minimal.

By replacing each edge e with $g_E(e)$ parallel edges, it follows from the uncapacitated case above that the vectors $\{(d_z(s) : s \in S) : z$ is a feasible integral vector$\}$ form an M-convex set. In this case, however, the basic algorithm is not necessarily polynomial since the values of g_E may be large. Therefore we need the general strongly polynomial algorithm described in Section 2.4. In this case the general Subroutine (2.4) can be realized via max-flow min-cut computations.

Minimum edge-cost dec-min semi-matchings　Harada et al. [30] developed an algorithm to solve the minimum edge-cost version of the original semi-matching problem of Harvey et al. [31]. As the dec-min in-degree bounded orientation problem covers all the extensions of semi-matching problems mentioned above, the minimum edge-cost version of these extensions can also be solved with the strongly polynomial algorithms developed in Section 5.2 for minimum cost dec-min in-degree bounded orientations.

We close this section with some historical remarks. The problem of Harvey et al. is closely related to earlier investigations in the context of minimizing a separable convex function over (integral elements of) a base-polyhedron. For example, Federgruen and Groenevelt [12] provided a polynomial time algorithm in 1986. Hochbaum and Hong [33] in 1995 developed a strongly polynomial algorithm; their proof, however, included a technical gap, which was fixed by Moriguchi, Shioura, and Tsuchimura [45] in 2011. For an early book on resource allocation, see the one by Ibaraki and Katoh [35] while three more recent surveys are due to Katoh and Ibaraki [38] from 1998, to Hochbaum [32] from 2007, and to Katoh, Shioura, and Ibaraki [39] from 2013. Algorithmic aspects of minimum degree square-sum problems for general graphs were discussed by Apollonio and Sebő [1].

6　Orientations of graphs with edge-connectivity requirements

In this section, we investigate various edge-connectivity requirements for the orientations of G. The main motivation behind these investigations is a conjecture of Borradaile et al. [7] on decreasingly minimal strongly connected orientations. Our goal is to prove their conjecture in a more general form.

6.1　Strongly connected orientations

Suppose that G is 2-edge-connected, implying that it has a strong orientation by a theorem of Robbins [51]. We are interested in dec-min strong orientations, meaning that the in-degree vector is decreasingly minimal over the strong orientations of G. This problem of Borradaile et al. was motivated by a practical application concerning optimal interval routing schemes.

Analogously to Theorem 4.5, they described a natural way to improve a strong orientation D to another one whose in-degree vector is decreasingly smaller. Suppose that there are two nodes s and t for which $\varrho(t) \geq \varrho(s) + 2$ and there are two edge-disjoint dipaths from s to t in D. Then reorienting an arbitrary st-dipath of D results in another strongly connected orientation of D which is clearly decreasingly smaller than D.

Borradaile et al. [7] conjectured the truth of the converse (and this conjecture was the starting point of our investigations). The next theorem states that the conjecture is true.

Theorem 6.1. *A strongly connected orientation D of $G = (V, E)$ is decreasingly minimal if and only if there are no two arc-disjoint st-dipaths in D for nodes s and t with $\varrho(t) \geq \varrho(s) + 2$.*

Proof. Suppose first that there are nodes s and t with $\varrho(t) \geq \varrho(s) + 2$ such that there are two arc-disjoint st-dipaths of D. Let P be any st-dipath in D and let D' denote the digraph arising from D by reorienting P. Then D' is strongly connected, since if it had a node-set Z ($\emptyset \subset Z \subset V$) with no entering arcs, then Z must be a $t\bar{s}$-set and P enters Z exactly once. But then $0 = \varrho_{D'}(Z) = \varrho_D(Z) - 1 \geq 2 - 1 = 1$, a contradiction. Therefore D' is indeed strongly connected and its in-degree vector is decreasingly smaller than that of D.

To see the non-trivial part, define a set-function p_1 as follows:

$$p_1(X) := \begin{cases} 0 & \text{if } X = \emptyset, \\ |E| & \text{if } X = V, \\ i_G(X) + 1 & \text{if } \emptyset \subset X \subset V. \end{cases}$$

Then p_1 is crossing supermodular and hence $B_1 := B'(p_1)$ is a base-polyhedron.

Claim 6.2. *An integral vector m is the in-degree vector of a strong orientation of G if and only if m is in $\dddot{B}_1$.*

Proof. If m is the in-degree vector of a strong orientation of G, then $\widetilde{m}(V) = |E| = p_1(V)$, $\widetilde{m}(\emptyset) = 0 = p_1(\emptyset)$, and

$$\widetilde{m}(Z) = \sum[\varrho(v) : v \in Z] = \varrho(Z) + i_G(Z) \geq 1 + i_G(Z) = p_1(Z)$$

for $\emptyset \subset Z \subset V$, that is, $m \in \dddot{B}_1$.

Conversely, let $m \in \dddot{B}_1$. Then $m \in B_G$ and hence by Claim 4.2, G has an orientation D with in-degree vector m. We claim that D is strongly connected. Indeed,

$$\varrho(Z) = \sum[\varrho(v) : v \in Z] - i_G(Z) = \widetilde{m}(Z) - i_G(Z) \geq p_1(Z) - i_G(Z) = 1$$

whenever $\emptyset \subset Z \subset V$. ■

Claim 6.3. *Let D be a strong orientation of G with in-degree vector m. Let t and s be nodes of G. The vector $m' := m + \chi_s - \chi_t$ is in B_1 if and only if D admits two arc-disjoint dipaths from s to t.*

Proof. $m' \in B_1$ holds precisely if there is no $t\bar{s}$-set X which is m-tight with respect to p_1, that is, $\widetilde{m}(X) = i_G(X) + 1$. Since $\varrho(Y) + i_G(Y) = \sum[\varrho(v) : v \in Y] = \widetilde{m}(Y)$ holds for any set $Y \subset V$, the tightness of X (that is, $\widetilde{m}(X) = i_G(X) + 1$) is equivalent to requiring that $\varrho(X) = 1$. Therefore $m' \in B_1$ if and only if $\varrho(Y) > 1$ holds for every $t\bar{s}$-set Y, which is, by Menger's theorem, equivalent to the existence of two arc-disjoint st-dipaths of D. ■

By Theorem 2.1, m is a dec-min element of $\dddot{B}_1$ if and only if there is no 1-tightening step for m. By Claim 6.3 this is just equivalent to the condition in the theorem that there are no two arc-disjoint st-dipaths in D for nodes s and t for which $\varrho(t) \geq \varrho(s) + 2$. ■■

We remark that a recent work of Zhou and Hou [56] describes a proof of the conjecture of Borradaile et al. that uses purely graph-theoretical concepts and does not rely on any knowledge of base-polyhedra or supermodular functions. An immediate consequence of Claim 6.2 and Theorem 2.1 is the following.

Corollary 6.4. *A strong orientation of G is dec-min if and only if it is inc-max.* ■

We indicated in Section 5.1 how in-degree constrained dec-min orientations can be managed due to the fact that the intersection of an integral base-polyhedron B with an integral box T is an integral base-polyhedron. The same approach works for degree-constrained strong orientations. For example, in this case dec-min and inc-max again coincide and one can formulate the in-degree constrained version of Theorem 6.1. In Section 6.3, we overview more general cases.

Given mixed graph   Dec-min in-degree $(3, 1, 3, 3)$ Inc-max in-degree $(2, 2, 2, 4)$

Figure 1: Dec-min and inc-max strong orientations of a mixed graph may differ

Figure 2: A mixed graph for the second variant

6.2 Counterexample for mixed graphs

Although Robbins' theorem on strong orientability of undirected graphs easily extends to mixed graphs, as was pointed out by Boesch and Tindell [5], it is not true anymore that a decreasingly minimal strong orientation of a mixed graph is always increasingly maximal. Actually, one may consider two natural variants.

In the first one, decreasing minimality and increasing maximality are defined for the total in-degree vector of the directed graph obtained from the initial mixed graph orienting its undirected edges. Consider the mixed graph $M = (V, A + E)$ in the left of Figure 1. Here $V = \{a, b, c, d\}$ while the set E of undirected edges (to be oriented) has just two elements: ab and dc. There are two strong orientations of M. In the first one (the middle digraph in Figure 1), these are the arcs ba and dc, in which case the total in-degree vector is $(3, 1, 3, 3)$. In the second one (the digraph in the right of Figure 1), the orientations of the elements of E are ab and cd, in which case the total in-degree vector is $(2, 2, 2, 4)$. Now $(3, 1, 3, 3)$ is dec-min while $(2, 2, 2, 4)$ is inc-max.

In the second variant, we are interested in the in-degree vector of the digraph obtained by orienting the originally undirected part E. For this version, the counterexample mixed graph is demonstrated in Figure 2. In any strong orientations of M, parallel edges in E must be oriented oppositely. Therefore the contribution of the actual orientations of these four pairs of parallel undirected edges (in the order of a, b, c, d, u, v, x, y) is $(2, 1, 0, 1, 1, 1, 1, 1)$.

Therefore there are essentially two distinct strong orientations of M. In the first one, the undirected edges ab, cd are oriented as ba, dc, while in the second one the undirected edges ab, cd are oriented as ab, cd. Hence the in-degree vector of the first strong orientation corresponding to the orientation of G (in the order of a, b, c, d, u, v, x, y) is $(3, 1, 1, 1, 1, 1, 1, 1)$. The in-degree vector of second strong orientation corresponding to the orientation of G is $(2, 2, 0, 2, 1, 1, 1, 1)$. The first vector is inc-max while the second vector is dec-min.

These examples give rise to the question: what is behind the phenomenon that while dec-min and inc-max coincide for strong orientations of undirected graphs, they differ for strong orientations of mixed graph? The explanation is, as we pointed out earlier, that for an M-convex set the two notions coincide and the set of in-degree vectors of strong orientations of an undirected graph is an M-convex set, while the corresponding set for a mixed graph is, in general, not an M-convex set. It is actually the

intersection of two M-convex sets. An algorithm for computing a dec-min element of the intersection of two M-convex sets is described in [23].

6.3 Higher edge-connectivity

An analogous approach works in a much more general setting. We say that a digraph **covers** a set-function h if $\varrho(X) \geq h(X)$ holds for every set $X \subseteq V$. The following result was proved in [14].

Theorem 6.5 ([14]). *Let h be a finite-valued, non-negative crossing supermodular function with $h(\emptyset) = h(V) = 0$. A graph $G = (V, E)$ has an orientation covering h if and only if*

$$e_{\mathcal{P}} \geq \sum_{i=1}^{q} h(V_i) \quad \text{and} \quad e_{\mathcal{P}} \geq \sum_{i=1}^{q} h(V - V_i)$$

hold for every partition $\mathcal{P} = \{V_1, \ldots, V_q\}$ of V, where $e_{\mathcal{P}}$ denotes the number of edges connecting distinct parts of $\mathcal{P}$. ∎

This theorem easily implies the classic orientation result of Nash-Williams [48] stating that a graph G has a k-edge-connected orientation precisely if G is $2k$-edge-connected. Namely, by defining h to be identically equal to k on non-empty proper subsets of V while $h(\emptyset) := h(V) := 0$, and observing that $e_{\mathcal{P}} = \sum[d_G(X)/2 : X \in \mathcal{P}]$, we can see that both inequalities in the theorem are equivalent to $\sum[d_G(X)/2 : X \in \mathcal{P}] \geq kq$. But this latter inequality follows from the $(2k)$-edge-connectivity of G since $\sum[d_G(X)/2 : X \in \mathcal{P}] \geq \sum[k : X \in \mathcal{P}] = kq$.

Even more, call a digraph (k, ℓ)-edge-connected ($\ell \leq k$) (with respect to a root-node r_0) if $\varrho(X) \geq k$ whenever $\emptyset \subset X \subseteq V - r_0$ and $\varrho(X) \geq \ell$ whenever $r_0 \in X \subset V$. (By Menger's theorem, (k, ℓ)-edge-connectedness is equivalent to requiring that there are k arc-disjoint dipaths from r_0 to every node and there are ℓ arc-disjoint dipaths from every node to r_0.) Then Theorem 6.5 implies:

Theorem 6.6. *A graph $G = (V, E)$ has a (k, ℓ)-edge-connected orientation if and only if*

$$e_{\mathcal{P}} \geq k(q - 1) + \ell$$

holds for every q-partite partition $\mathcal{P}$ of V. ∎

Note that an even more general special case of Theorem 6.5 can be formulated to characterize graphs admitting in-degree constrained and (k, ℓ)-edge-connected orientations.

It is important to emphasize that however general Theorem 6.5 is, it does not say anything about strong orientations of mixed graphs. In particular, it does not imply the pretty but easily provable theorem of Boesch and Tindell [5]. The problem of finding decreasingly minimal in-degree constrained k-edge-connected orientation of mixed graphs can be solved as a special case of decreasing minimization over the intersection of two M-convex sets.

The next lemma shows why the set of in-degree vectors of orientations of G covering the set-function h appearing in Theorem 6.5 is an M-convex set, ensuring in this way the possibility of applying the results on decreasing minimization over M-convex sets to general graph orientation problems.

Lemma 6.7. *An orientation D of G covers h if and only if its in-degree vector m is in the base-polyhedron $B = B'(p)$, where $p := h + i_G$ is a crossing supermodular function.*

Proof. Suppose first that m is the in-degree vector of a digraph covering h. Then $h(X) \leq \varrho(X) = \widetilde{m}(X) - i_G(X)$ for $X \subset V$ and $h(V) = 0 = \varrho(V) = \widetilde{m}(V) - i_G(V)$, that is, m is indeed in B.

Conversely, suppose that $m \in B$. Since h is finite-valued and non-negative, we have $\widetilde{m}(X) \geq p(X) \geq i_G(X)$ for $X \subset V$ and $\widetilde{m}(V) = i_G(V)$ and hence, by the Orientation lemma, there is an orientation D of G with in-degree vector m. Moreover, this digraph D covers h since $\varrho_D(X) = \widetilde{m}(X) - i_G(X) \geq p(X) - i_G(X) = h(X)$ holds for $X \subset V$. ∎

By Lemma 6.7, Theorem 2.1 can be applied again to the general orientation problem covering a non-negative and crossing supermodular set-function h in the same way as it was applied in the special case of strong orientation above, but we formulate the result only for the special case of in-degree constrained and k-edge-connected orientations.

Theorem 6.8. *Let $G = (V, E)$ be an undirected graph endowed with a lower bound function $f : V \to \mathbf{Z} \cup \{-\infty\}$ and an upper bound function $g : V \to \mathbf{Z} \cup \{+\infty\}$ with $f \le g$. A k-edge-connected and in-degree constrained orientation D of G is decreasingly minimal if and only if there are no two nodes s and t for which $\varrho(t) \ge \varrho(s) + 2$, $\varrho(t) > f(t)$, $\varrho(s) < g(s)$, and there are $k + 1$ arc-disjoint st-dipaths.* ∎

The theorem can be extended even further to in-degree constrained and (k, ℓ)-edge-connected orientations ($\ell \le k$).

Locally k-edge-connected orientations For a subset S of nodes of a digraph $D = (V, A)$, we say that D is **locally k-edge-connected** in S if there are k arc-disjoint dipaths in D from any node s of S to any other node t of S. This is equivalent to requiring (by Menger's theorem) that $\varrho_D(X) \ge k$ holds for each subset $X \subset V$ for which neither $X \cap S$ nor $S - X$ is empty. Clearly, when $S = V$, we are back at ordinary k-edge-connectivity. Király and Lau [41] considered a closely related concept when the requirement for D is that there are k edge-disjoint dipaths from a specified root-node r of S to every other node of S. Such a digraph may be called locally rooted k-edge-connected in S.

By relying on a base-polyhedral approach, Király and Lau [41] proved (Theorem 1.3) that there exists a polynomial algorithm for finding an orientation of G which is locally rooted k-edge-connected in S *and* the orientation meets an in-degree specification $m_{\overline{S}}$ on the nodes in $V - S$. By using an analogous approach, we can solve the orientation problem in which, beside the in-degree specification on $V - S$, the requirement is local k-edge-connectivity in S. To this end, define a set-function h on S as follows:

$$h(Z) := \begin{cases} 0 & \text{if } Z = \emptyset, \\ |E| - \widetilde{m_{\overline{S}}}(V - S) & \text{if } Z = S, \\ \max\{k + i_G(Z \cup X) - \widetilde{m_{\overline{S}}}(X) : X \subseteq V - S\} & \text{if } \emptyset \subset Z \subset S. \end{cases}$$

It is not difficult to see that h is a non-negative, crossing supermodular function on S. Let $B_S := B'(h)$ denote the integral base-polyhedron defined by h. By relying on Lemma 6.7, one can derive the following.

Theorem 6.9. *Let $G = (V, E)$ be an undirected graph with a specified subset S of V. Let $m_{\overline{S}}$ be an in-degree specification on $V - S$. The set of in-degree vectors of those orientations of G which are locally k-edge-connected in S and in-degree specified in $V - S$ is an M-convex set B, namely, $B = \{(m_S, m_{\overline{S}}) : m_S \in B_S\}$.* ∎

By this theorem, we can determine a decreasingly minimal orientation among those which are k-edge-connected in S and in-degree specified in $V - S$. Even additional in-degree constraints can be imposed on the elements of S.

Remark 6.1. One may consider the orientation problem $(*)$ where, beside local k-edge-connectivity in S, we impose upper and lower bounds on the in-degrees of nodes outside S (instead of exact values). Bernáth et al. [4] proved that the in-degree constrained version of Nash-Willams' well-balanced orientation problem [48] is NP-complete. By using a similar method, Bernáth [3] recently proved that already the problem $(*)$ is NP-complete, even in the special case with $k = 1$. ∎

Hypergraph orientation   Let $H = (V, \mathcal{E})$ be a hypergraph for which we assume that each hyperedge has at least 2 nodes. Orienting a hyperedge Z means that we designate an element z of Z as its head-node. A hyperedge Z with a designated head-node $z \in Z$ is a directed hyperedge denoted by (Z, z). Orienting a hypergraph means the operation of orienting each of its hyperedges. We say that a directed hyperedge (Z, z) enters a subset X of nodes if $z \in X$ and $Z - X \neq \emptyset$. A directed hypergraph is called k-edge-connected if the in-degree of every non-empty proper subset of nodes is at least k.

The following result was proved in [19] (see, also Theorem 2.22 in the survey paper [18]).

Theorem 6.10. *The set of in-degree vectors of k-edge-connected and in-degree constrained orientations of a hypergraph forms an M-convex set.* ∎

Therefore we can apply the general results obtained for decreasing minimization over M-convex sets.

References

[1] Apollonio, N., Sebő, A.: Minconvex factors of prescribed size in graphs. SIAM Journal on Discrete Mathematics **23**, 1297–1310 (2009)

[2] Bérczi, K., Frank, A.: Supermodularity in unweighted graph optimization, Part I: Branchings and matchings. Mathematics of Operations Research **43**, 726–753 (2018)

[3] Bernáth, A.: Personal communication (2021)

[4] Bernáth, A., Iwata, S., Király, T., Király, Z, Szigeti, Z.: Recent results on well-balanced orientations. Discrete Optimization **5**, 663–676 (2008)

[5] Boesch, F., Tindell, R.: Robbins's theorem for mixed multigraphs. American Mathematical Monthly **87**, 716–719 (1980)

[6] Bokal, D., Brešar, B., Jerebic, J.: A generalization of Hungarian method and Hall's theorem with applications in wireless sensor networks. Discrete Applied Mathematics **160**, 460–470 (2012)

[7] Borradaile, G., Iglesias, J., Migler, T., Ochoa, A., Wilfong, G., Zhang, L.: Egalitarian graph orientations. Journal of Graph Algorithms and Applications **21**, 687–708 (2017)

[8] Borradaile, G., Migler, T., Wilfong, G.: Density decompositions of networks. Journal of Graph Algorithms and Applications **23**, 625–651 (2019)

[9] Edmonds, J.: Submodular functions, matroids and certain polyhedra. In: Guy, R., Hanani, H., Sauer, N., Schönheim, J. (eds.) Combinatorial Structures and Their Applications, pp. 69–87. Gordon and Breach, New York (1970)

[10] Edmonds, J.: Edge-disjoint branchings. In: Rustin, B. (ed.) Combinatorial Algorithms, pp. 91–96. Academic Press, New York (1973)

[11] Edmonds, J.: Matroid intersection. Annals of Discrete Mathematics **14**, 39–49 (1979)

[12] Federgruen, A., Groenevelt, H.: The greedy procedure for resource allocation problems: necessary and sufficient conditions for optimality. Operations Research **34**, 909–918 (1986)

[13] Ford, L. R., Jr., Fulkerson, D.R.: Flows in Networks. Princeton University Press, Princeton NJ (1962)

[14] Frank, A.: On the orientation of graphs. Journal of Combinatorial Theory, Ser. B **28**, 251–261 (1980)

[15] Frank, A.: Augmenting graphs to meet edge-connectivity requirements. SIAM Journal on Discrete Mathematics **5**, 22–53 (1992)

[16] Frank, A.: Connections in Combinatorial Optimization. Oxford University Press, Oxford (2011)

[17] Frank, A., Jordán, T.: Minimal edge-coverings of pairs of sets. Journal of Combinatorial Theory, Ser. B **65**, 73–110 (1995)

[18] Frank, A., Király, T.: A survey on covering supermodular functions. In: Cook, W., Lovász, L., Vygen, J. (eds.) Research Trends in Combinatorial Optimization, pp. 87–126. Springer, Berlin (2009)

[19] Frank, A., Király, T., Király, Z.: On the orientation of graphs and hypergraphs. Discrete Applied Mathematics **131**, 385–400 (2003)

[20] Frank, A., Murota, K.: Discrete decreasing minimization, Part II: Views from discrete convex analysis. arXiv: 1808.08477 (August 2018)

[21] Frank, A., Murota, K.: Decreasing minimization on M-convex sets: Background and structures. Submitted for publication. arXiv: 2007.09616 (July 2020)

[22] Frank, A., Murota, K.: Fair integral flows. Submitted for publication. arXiv: 1907.02673 (September 2020)

[23] Frank, A., Murota, K.: Fair integral submodular flows. arXiv: 2012.07325 (December 2020)

[24] Frank, A., Tardos, É.: Generalized polymatroids and submodular flows. Mathematical Programming **42**, 489–563 (1988)

[25] Fujishige, S.: Lexicographically optimal base of a polymatroid with respect to a weight vector. Mathematics of Operations Research **5**, 186–196 (1980)

[26] Fujishige, S.: Structures of polyhedra determined by submodular functions on crossing families. Mathematical Programming **29**, 125–141 (1984)

[27] Fujishige, S.: Submodular Functions and Optimization, 2nd edn. Annals of Discrete Mathematics **58**, Elsevier, Amsterdam (2005)

[28] Goemans, M.X., Gupta, S., Jaillet, P.: Discrete Newton's algorithm for parametric submodular function minimization. In: Eisenbrand, F., Koenemann, J. (eds.) Integer Programming and Combinatorial Optimization. Lecture Notes in Computer Science, vol. 10328, pp. 212–227 (2017)

[29] Hakimi, S.L.: On the degrees of the vertices of a directed graph. Journal of The Franklin Institute **279**, 290–308 (1965)

[30] Harada, Y., Ono, H., Sadakane, K., Yamashita, M.: Optimal balanced semi-matchings for weighted bipartite graphs. IPSJ Digital Courier **3**, 693–702 (2007)

[31] Harvey, N.J.A., Ladner, R.E., Lovász, L., Tamir, T.: Semi-matchings for bipartite graphs and load balancing. Journal of Algorithms **59**, 53–78 (2006)

[32] Hochbaum, D.S.: Complexity and algorithms for nonlinear optimization problems. Annals of Operations Research **153**, 257–296 (2007)

[33] Hochbaum, D.S., Hong, S.-P.: About strongly polynomial time algorithms for quadratic optimization over submodular constraints. Mathematical Programming **69**, 269–309 (1995)

[34] Hoffman, A.J.: Some recent applications of the theory of linear inequalities to extremal combinatorial analysis. In: Bellman, R., Hall, M., Jr. (eds.) Combinatorial Analysis (Proceedings of the Symposia of Applied Mathematics **10**) pp. 113–127. American Mathematical Society, Providence, Rhode Island (1960)

[35] Ibaraki, T., Katoh, N.: Resource Allocation Problems: Algorithmic Approaches. MIT Press, Boston (1988)

[36] Iwata, S., Fleischer, L., Fujishige, S.: A combinatorial, strongly polynomial-time algorithm for minimizing submodular functions. Journal of the ACM **48**, 761–777 (2001)

[37] Jiang, H.: Minimizing convex functions with integral minimizers. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, pp. 976–985 (2021)

[38] Katoh, N., Ibaraki, T.: Resource allocation problems. In: Du, D.-Z., Pardalos, P.M. (eds.) Handbook of Combinatorial Optimization, Vol.2, pp. 159–260. Kluwer Academic Publishers, Boston (1998)

[39] Katoh, N., Shioura, A., Ibaraki, T.: Resource allocation problems. In: Pardalos, P.M., Du, D.-Z., Graham, R.L. (eds.) Handbook of Combinatorial Optimization, 2nd ed., Vol. 5, pp. 2897-2988, Springer, Berlin (2013)

[40] Katrenič, J., Semanišin, G.: Maximum semi-matching problem in bipartite graphs. Discussiones Mathematicae, Graph Theory **33**, 559–569 (2013)

[41] Király, T., Lau, L.-C.: Approximate min-max theorems for Steiner rooted-orientations of graphs and hypergraphs. Journal of Combinatorial Theory, Ser. B **98**, 1233–1252 (2008)

[42] Levin, A., Onn, S.: Shifted matroid optimization. Operations Research Letters **44**, 535–539 (2016)

[43] Megiddo, N.: Optimal flows in networks with multiple sources and sinks. Mathematical Programming **7**, 97–107 (1974)

[44] Megiddo, N.: A good algorithm for lexicographically optimal flows in multi-terminal networks. Bulletin of the American Mathematical Society **83**, 407–409 (1977)

[45] Moriguchi, S., Shioura, A., Tsuchimura, N.: M-convex function minimization by continuous relaxation approach—Proximity theorem and algorithm. SIAM Journal on Optimization **21**, 633–668 (2011)

[46] Murota, K.: Discrete convex analysis. Mathematical Programming **83**, 313–371 (1998)

[47] Murota, K.: Discrete Convex Analysis. Society for Industrial and Applied Mathematics, Philadelphia (2003)

[48] Nash-Williams, C.St.J.A.: On orientations, connectivity and odd vertex pairings in finite graphs. Canadian Journal of Mathematics **12**, 555–567 (1960)

[49] Orlin, J.B.: A faster strongly polynomial time algorithm for submodular function minimization. Mathematical Programming, Series A **118**, 237–251 (2009)

[50] Radzik, T.: Fractional combinatorial optimization. In: Pardalos, P.M., Du, D.-Z., Graham, R.L. (eds.) Handbook of Combinatorial Optimization, 2nd edn., pp. 1311–1355. Springer Science+Business Media, New York (2013)

[51] Robbins, H.E.: A theorem on graphs with an application to a problem of traffic control. American Mathematical Monthly **46**, 281–283 (1939)

[52] Schrijver, A.: A combinatorial algorithm minimizing submodular functions in strongly polynomial time. Journal of Combinatorial Theory, Series B **80**, 346–355 (2000)

[53] Schrijver, A.: Combinatorial Optimization: Polyhedra and Efficiency. Springer, Heidelberg (2003)

[54] Tardos, É.: A strongly polynomial minimum cost circulation algorithm. Combinatorica **5**, 247–255 (1985)

[55] Topkis, D.M.: Minimizing a submodular function on a lattice. Operations Research **26**, 305–321 (1978)

[56] Zhou, H., Hou, X.: Strongly connected orientation with minimum lexicographic order of indegrees. arXiv:2103.01017 (March 2021)