SIAM J. Optim.
Vol. 28, No. 2, pp. 1049–1076

THE MULTILINEAR POLYTOPE FOR ACYCLIC HYPERGRAPHS*

ALBERTO DEL PIA† AND AIDA KHAJAVIRAD‡

Abstract. We consider the multilinear polytope defined as the convex hull of the set of binary points z satisfying a collection of equations of the form $z_e = \prod_{v \in e} z_v$, $e \in E$, where E denotes a family of subsets of $\{1, \ldots, n\}$ of cardinality at least two. Such sets are of fundamental importance in many types of mixed-integer nonlinear optimization problems, such as $0-1$ polynomial optimization. Utilizing an equivalent hypergraph representation, we study the facial structure of the multilinear polytope in conjunction with the acyclicity degree of the underlying hypergraph. We provide explicit characterizations of the multilinear polytopes corresponding to Berge-acylic and γ-acyclic hypergraphs. As the multilinear polytope for γ-acyclic hypergraphs may contain exponentially many facets in general, we present a strongly polynomial-time algorithm to solve the separation problem, implying polynomial solvability of the corresponding class of $0-1$ polynomial optimization problems. As an important byproduct, we present a new class of cutting planes for constructing tighter polyhedral relaxations of mixed-integer nonlinear optimization problems with multilinear subexpressions.

Key words. multilinear polytope, cutting planes, hypergraph acyclicity, separation algorithm, mixed-integer nonlinear optimization

AMS subject classifications. 90C10, 90C11, 90C26, 90C57

DOI. 10.1137/16M1095998

1. Introduction.

Consider a hypergraph $G = (V, E)$, where $V = V(G)$ is the set of nodes of G, and $E = E(G)$ is a set of subsets of V of cardinality at least two, called the edges of G. The *rank* r of G is defined as the maximum cardinality of an edge in E. With any hypergraph G, and cost vector $c \in \mathbb{R}^{V+E}$, we associate a $0-1$ multilinear optimization problem of the form

$$\text{(MO)} \qquad \begin{aligned} \max \quad & \sum_{v \in V} c_v z_v + \sum_{e \in E} c_e \prod_{v \in e} z_v \\ \text{s.t.} \quad & z_v \in \{0, 1\} \quad \forall v \in V. \end{aligned}$$

Without loss of generality we can assume that c_e is nonzero for every $e \in E$. We refer to the objective function of (MO) as a *multilinear function* and each product term $\prod_{v \in e} z_v$ as a *multilinear term*. Problem (MO) is a well-known $\mathcal{NP}$-hard optimization problem. Since $(z_v)^p = z_v$ for any $z_v \in \{0, 1\}$ and any positive integer p, problem (MO) is equivalent to unconstrained $0-1$ polynomial optimization. In particular, if $r = 2$, then we obtain the well-studied unconstrained $0-1$ quadratic optimization (QP) which is equivalent to the max-cut problem (see, e.g., [5, 25]). Moreover, it is simple to show that the maximum value of a multilinear function over a box is attained at a vertex of the box [28]. Clearly, multilinear functions are closed under scaling and shifting of variables. It then follows that (MO) is equivalent to

*Received by the editors September 27, 2016; accepted for publication (in revised form) December 4, 2017; published electronically April 10, 2018.
http://www.siam.org/journals/siopt/28-2/M109599.html
Funding: The research of the authors was supported in part by National Science Foundation award CMMI-1634768.
†Department of Industrial and Systems Engineering & Wisconsin Institute for Discovery, University of Wisconsin-Madison, Madison, WI 53706 (delpia@wisc.edu).
‡Department of Chemical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 (aida@cmu.edu).

maximizing a multilinear function over a box. The latter problem has been studied extensively by the global optimization community [1, 27, 31, 29, 24, 33, 23, 3].

It is common practice to linearize the objective function of problem (MO) by introducing a new variable for every multilinear term and obtain an equivalent optimization problem in a lifted space:

$$(\mathrm{MO}') \qquad \begin{aligned} \max \quad & \sum_{v \in V} c_v z_v + \sum_{e \in E} c_e z_e \\ \text{s.t.} \quad & z_e = \prod_{v \in e} z_v \quad \forall e \in E, \\ & z_v \in \{0,1\} \quad \forall v \in V. \end{aligned}$$

Subsequently, a convex relaxation of the feasible region of problem (MO′) is constructed and the resulting problem is solved to obtain an upper bound on the optimal value of problem (MO′). In this paper, we study the problem of constructing sharp polyhedral relaxations for the feasible region of problem (MO′). More precisely, we consider the *multilinear set* $\mathcal{S}_G$ defined as

$$(1) \qquad \mathcal{S}_G = \left\{ z \in \{0,1\}^{V+E} : z_e = \prod_{v \in e} z_v \ \ \forall e \in E \right\}.$$

Throughout the paper, we assume that each z_v, $v \in V$, appears in at least one multilinear term. In fact, if z_v does not appear in any multilinear term, then the set $\mathcal{S}_G$ can be written as the cartesian product $\mathcal{S}_{(V \setminus \{v\}, E)} \times \{0,1\}$. We refer to the convex hull of $\mathcal{S}_G$ as the *multilinear polytope* MP_G. Moreover, we refer to the rank r of the hypergraph G as the *degree* of the corresponding multilinear set $\mathcal{S}_G$. Building convex relaxations for multilinear sets has been a subject of extensive research by the mathematical programming community [1, 25, 13, 27, 31, 29, 24, 4, 23, 16, 15, 14, 7]. If all multilinear terms in $\mathcal{S}_G$ are bilinears, i.e., $r = 2$, the corresponding multilinear polytope coincides with the Boolean quadric polytope QP_G first defined by Padberg [25] in the context of unconstrained $0{-}1$ QPs. We should remark that for the Boolean quadric polytope, our hypergraph representation simplifies to the graph representation defined by Padberg [25]. In [16], we introduce the hypergraph representation framework for higher degree multilinear sets and study the facial structure of their convex hull. In particular, we develop the theory of various types of lifting operations, giving rise to many types of facet-defining inequalities in the space of the original variables. A great simplification in studying the facial structure of the multilinear polytope is possible when the corresponding multilinear set $\mathcal{S}_G$ is decomposable into simpler multilinear sets $\mathcal{S}_{G_j}$, $j \in J$; namely, the convex hull of $\mathcal{S}_G$ can be obtained by convexifying each $\mathcal{S}_{G_j}$, separately. In [15], we study the decomposability properties of multilinear sets.

1.1. Explicit characterization of MP_G and tractability of (MO). In this paper, we are interested in characterizing sufficient conditions under which the multilinear polytope admits a "desirable" explicit description. More precisely, for hypergraphs G with certain "degrees of acyclicity," we derive an explicit characterization of the polytope MP_G. In addition, we prove that for the same class of hypergraphs, this convex hull characterization enables us to solve problem (MO) in polynomial time.

In [25], Padberg derives a closed-form description of the Boolean quadric polytope QP_G, provided that the underlying graph G is acyclic or is series-parallel. Moreover, in those cases, given any objective function coefficient vector $c \in \mathbb{R}^{V+E}$, the corresponding unconstrained $0{-}1$ QP is polynomially solvable [5]. However, for higher degree multilinear optimization problems, similar tractability results are rather scarce. The explicit characterization of the multilinear polytope MP_G is available for the special case where $r = n$ and the edge set $E(G)$ contains all subsets of V of cardinality at least two (see, e.g., [32, 26]). In addition, in [14], the closed-from description of MP_G is given for a hypergraph G consisting of two edges that intersect in at least two nodes. Motivated by the existing results for the Boolean quadric polytope, in this paper we provide explicit characterizations of higher degree multilinear polytopes. Our new characterizations will be given in terms of easily verifiable assumptions on the structure of the corresponding hypergraph and serve as generalizations of those for unconstrained $0-1$ QPs.

We start by defining a well-known tractable relaxation of the multilinear polytope.

1.2. Standard linearization of multilinear sets. A valid polyhedral relaxation of the multilinear set $\mathcal{S}_G$ can be obtained by replacing each multilinear term $z_e = \prod_{v \in e} z_v$, by its convex hull over the unit hypercube:

$$\begin{aligned}\text{MP}_G^{\text{LP}} = \Big\{ z :\ & z_v \leq 1 \quad \forall v \in V,\\ & z_e \geq 0,\ z_e \geq \sum_{v \in e} z_v - |e| + 1 \quad \forall e \in E,\\ & z_e \leq z_v \quad \forall e \in E \quad \forall v \in e \Big\}.\end{aligned} \tag{2}$$

The above relaxation has been used extensively in the literature and is often referred to as the *standard linearization* of the multilinear set (see, e.g., [20, 13]). It is well-known that the Boolean quadric polytope QP_G coincides with its standard linearization QP_G^{LP} if and only if the graph G is acyclic [25]. To generalize this result to higher degree multilinear polytopes, it is natural to look into the notion acyclicity for hypergraphs. Interestingly, unlike graphs for which there is a single natural notion of acyclic graphs, there are several nonequivalent definitions of acyclicity for hypergraphs which collapse to graph acyclicity for the special case of ordinary graphs. In fact, the notion of graph acyclicity has been extended to several different degrees of acyclicity of hypergraphs [18]. Next, we briefly review the concept of cycles in hypergraphs, as it plays a crucial role in our subsequent developments.

1.3. Cycles in hypergraphs. Let $G = (V, E)$ be a hypergraph. The most restrictive type of acyclicity in hypergraphs is Berge-acyclicity. A hypergraph is *Berge-acyclic* when it contains no Berge-cycles, defined as follows (see [6, Chapter 5] for more details).

DEFINITION 1. *A* Berge-cycle *in G of length t for some $t \geq 2$ is a sequence $C = v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$ with the following properties:*
- *$v_1, v_2, \ldots, v_t$ are distinct nodes of G,*
- *$e_1, e_2, \ldots, e_t$ are distinct edges of G,*
- *$v_i, v_{i+1} \in e_i$ for $i = 1, \ldots, t-1$, and $v_t, v_1 \in e_t$.*

Note that Berge-cycles of length two are present only when two edges intersect in at least two nodes (see Figure 1(a)).

The next class of acyclic hypergraphs, in increasing order of generality, is the class of γ-acyclic hypergraphs. We first recall the notion of a γ-cycle (see, e.g., [17, 8] for more details).

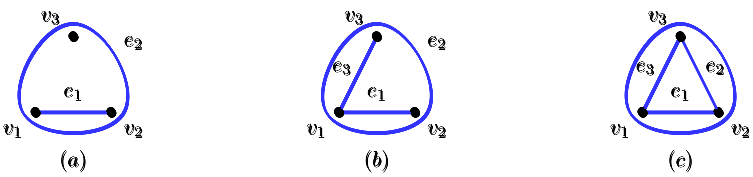

Fig. 1. *Examples of hypergraphs with different degrees of acyclicity: (a) γ-acyclic but not Berge-acyclic, since v_1, e_1, v_2, e_2, v_1 is a Berge-cycle; (b) β-acyclic but not γ-acyclic, since $v_1, e_1, v_2, e_2, v_3, e_3, v_1$ is a γ-cycle; (c) not β-acyclic since $v_1, e_1, v_2, e_2, v_3, e_3, v_1$ is a β-cycle.*

Fig. 2. *Examples of γ-acyclic hypergraphs containing Berge-cycles of length two and three.*

Definition 2. *A* γ-cycle *in G is a Berge-cycle $C = v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$ such that $t \geq 3$, and the node v_i belongs to e_{i-1}, e_i and no other e_j for all $i = 2, \ldots, t$.*

See Figure 1(b) for an example of a γ-cycle. A hypergraph is called *γ-acyclic* if it contains no γ-cycles. Clearly, a Berge-acyclic hypergraph is also γ-acyclic, while the converse is not true (see Figure 1(a)). In fact γ-acyclic hypergraphs are a significant generalization of Berge-acyclic hypergraphs and may contain Berge-cycles of arbitrary length, in general (see Figure 2). There exist several equivalent characterizations for γ-acyclic hypergraphs in the literature. We will be using the following characterization to prove our results in Section 4. First, we define a β-cycle [17].

Definition 3. *A β-cycle in G is a γ-cycle $C = v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$ such that the node v_1 belongs to e_1, e_t and no other e_j.*

See Figure 1(c) for an example of a β-cycle. A hypergraph is called *β-acyclic* if it does not contain any β-cycles. In the literature, β-acyclic hypergraphs have been also called *totally balanced hypergraphs* [22] and β-cycles have been also referred to as *special cycles* [2]. Using the notion of β-acyclicity, in [8] the author characterizes γ-acyclic hypergraphs as follows.

Proposition 4. *A hypergraph $G = (V, E)$ is γ-acyclic if and only if it satisfies the following properties:*

(i) *G is β-acyclic,*
(ii) *there do not exist distinct nodes v_1, v_2, v_3 such that $\{\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_2, v_3\}\} \subseteq \{e \cap \{v_1, v_2, v_3\} : e \in E\}$.*

Throughout this paper, given any cycle $C = v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$, we denote by $V(C) = \{v_1, \ldots, v_t\}$ the nodes of the cycle C, and by $E(C) = \{e_1, \ldots, e_t\}$ the edges of C.

1.4. Our contribution. In this paper, we present new explicit characterizations of multilinear polytopes corresponding to acyclic hypergraphs. As an important byproduct, we introduce a new class of cutting planes to construct tighter polyhedral relaxations of general multilinear sets. As we detail later, the separation problem for the proposed cutting planes can be solved efficiently for γ-acyclic hypergraphs, and for general hypergraphs with fixed rank.

Fig. 1. ***Examples of hypergraphs with different degrees of acyclicity: (a) γ-acyclic but not Berge-acyclic, since v_1,e_1,v_2,e_2,v_1 is a Berge-cycle; (b) β-acyclic but not γ-acyclic, since $v_1,e_1,v_2,e_2,v_3,e_3,v_1$ is a γ-cycle; (c) not β-acyclic, since $v_1,e_1,v_2,e_2,v_3,e_3,v_1$ is a β-cycle.***

Fig. 2. ***Examples of γ-acyclic hypergraphs containing Berge-cycles of length two and three.***

generalization of Berge-acyclic hypergraphs and may contain Berge-cycles of arbitrary length, in general (see Figure 2). There exist several equivalent characterizations for γ-acyclic hypergraphs. In the following we present an alternative characterization which will be used to prove our results in Section 4. First, we define a β-cycle [17].

Definition 3. ***A β-cycle in G is a γ-cycle $C = v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$ such that the node v_1 belongs to e_1, e_t and no other e_j.***

See Figure 1(c) for an example of a β-cycle. A hypergraph is called *β-acyclic* if it does not contain any β-cycles. In the literature, β-acyclic hypergraphs have been also called *totally balanced hypergraphs* [22] and β-cycles have been also referred to as *special cycles* [2]. Using the notion of β-acyclicity, in [8] the author characterizes γ-acyclic hypergraphs as follows:

Proposition 4. ***A hypergraph $G = (V, E)$ is γ-acyclic if and only if it satisfies the following properties:***

(i) G is β-acyclic,
(ii) there do not exist distinct nodes v_1, v_2, v_3 such that $\{\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_2, v_3\}\} \subseteq \{e \cap \{v_1, v_2, v_3\} : e \in E\}$.

Throughout this paper, given any cycle $C = v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$, we denote by $V(C) = \{v_1, \ldots, v_t\}$ the nodes of the cycle C, and by $E(C) = \{e_1, \ldots, e_t\}$ the edges of C.

1.4. Our contribution. In this paper, we present new explicit characterizations of Multilinear polytopes corresponding to acyclic hypergraphs. As an important byproduct, we introduce a new class of cutting planes to construct tighter polyhedral relaxations of general Multilinear sets. As we detail later, the separation problem for the proposed cutting planes can be solved efficiently for γ-acyclic hypergraphs, and for general hypergraphs with fixed rank.

The remainder of this paper is organized as follows. In Section 2, we present a technical result regarding the decomposability of Multilinear sets, which will be used to prove our main results. Namely, we show that under certain assumptions the Multilinear set is decomposable into a collection of simpler Multilinear sets for which closed-form descriptions of convex hulls can be derived. In Section 3, we prove

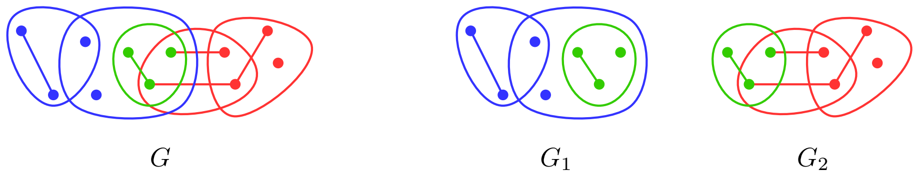

FIG. 3. *An example of a hypergraph G for which by Theorem 5 the Multilinear set $\mathcal{S}_G$ is decomposable into the sets $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$.*

that the standard linearization of a Multilinear set coincides with the Multilinear polytope if and only if the corresponding hypergraph is Berge-acyclic. Subsequently, in Section 4, we introduce a new class of valid inequalities for Multilinear sets, referred to as *flower inequalities*, and show that the polytope obtained by adding them to the standard linearization of $\mathcal{S}_G$ coincides with the Multilinear polytope MP_G if and only if G is γ-acyclic. Finally, as the number of facet-defining flower inequalities may grow exponentially in the rank of the hypergraph, in Section 5, we present a strongly polynomial-time separation algorithm for flower inequalities, when the underlying hypergraph is γ-acyclic. Polynomial-time algorithms for determining acyclicity degree of hypergraphs are available [18]. Together with our results in Sections 3–5, this implies that for Berge-acyclic and γ-acyclic hypergraphs, recognition and solution of problem (MO) can be done efficiently.

2. A sufficient condition for decomposability of Multilinear sets.

In this section, we derive a technical result on decomposability of Multilinear sets that will be used to obtain our convex hull characterizations in Sections 3 and 4. We refer the reader to [15] for an in-depth study of decomposability properties of Multilinear sets. However, the sufficient condition that we need in this section does not follow from our previous work in [15].

Let $G = (V, E)$ be a hypergraph. A hypergraph $G' = (V', E')$ is a *partial hypergraph* of G, if $V' \subseteq V$ and $E' \subseteq E$. Given a subset V' of V, the *section hypergraph of G induced by V'* is the partial hypergraph $G' = (V', E')$, where $E' = \{e \in E : e \subseteq V'\}$. Given hypergraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, we denote by $G_1 \cap G_2$ the hypergraph $(V_1 \cap V_2, E_1 \cap E_2)$, and we denote by $G_1 \cup G_2$, the hypergraph $(V_1 \cup V_2, E_1 \cup E_2)$.

Now consider the hypergraph G, and let G_1, G_2 be section hypergraphs of G such that $G_1 \cup G_2 = G$. We say that the set $\mathcal{S}_G$ is *decomposable into the sets $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$* if

$$\operatorname{conv} \mathcal{S}_G = \operatorname{conv} \bar{\mathcal{S}}_{G_1} \cap \operatorname{conv} \bar{\mathcal{S}}_{G_2},$$

where $\bar{\mathcal{S}}_{G_j}$, $j = 1, 2$ is the set of all points in the space of $\mathcal{S}_G$ whose projection in the space defined by G_j is $\mathcal{S}_{G_j}$. Next, in Theorem 5, we provide a sufficient condition for decomposability of Multilinear sets. Figure 3 illustrates a simple hypergraph G for which by Theorem 5 the set $\mathcal{S}_G$ is decomposable into $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$.

THEOREM 5. *Let G be a hypergraph, and let G_1, G_2 be section hypergraphs of G such that $G_1 \cup G_2 = G$. Denote by $\bar{p} := V(G_1) \cap V(G_2)$. Suppose that $\bar{p} \in V(G) \cup E(G)$, and that for every edge e of G containing nodes in $V(G_1) \setminus V(G_2)$ either $e \supset \bar{p}$, or $e \cap \bar{p} = \emptyset$. Then the set $\mathcal{S}_G$ is decomposable into $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$.*

Proof. Clearly the inclusion $\operatorname{conv} \mathcal{S}_G \subseteq \operatorname{conv} \bar{\mathcal{S}}_{G_1} \cap \operatorname{conv} \bar{\mathcal{S}}_{G_2}$ holds, since $\mathcal{S}_G \subseteq \bar{\mathcal{S}}_{G_1} \cap \bar{\mathcal{S}}_{G_2}$. Thus, it suffices to show the reverse inclusion. If either G_1 or G_2 coincides *and that for every edge e of G containing nodes in $V(G_1) \setminus V(G_2)$ either $e \supset \bar{p}$, or $e \cap \bar{p} = \emptyset$. Then the set $\mathcal{S}_G$ is decomposable into $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$.*

Proof. Clearly the inclusion $\text{conv}\mathcal{S}_G \subseteq \text{conv}\bar{\mathcal{S}}_{G_1} \cap \text{conv}\bar{\mathcal{S}}_{G_2}$ holds, since $\mathcal{S}_G \subseteq \bar{\mathcal{S}}_{G_1} \cap \bar{\mathcal{S}}_{G_2}$. Thus, it suffices to show the reverse inclusion. If either G_1 or G_2 coincides with G, then the statement is trivial. Henceforth, we assume that both G_1 and G_2 are different from G, implying $G \backslash G_1$ and $G \backslash G_2$ are nonempty. Let $\tilde{z} \in \text{conv}\bar{\mathcal{S}}_{G_1} \cap \text{conv}\bar{\mathcal{S}}_{G_2}$. We will show that $\tilde{z} \in \text{conv}\mathcal{S}_G$. Let $\bar{z}$ contain those components of $\tilde{z}$ corresponding to nodes and edges that are both in G_1 and in G_2. In particular, $\bar{z}_{\bar{p}}$ is a component of $\bar{z}$. Let z^1 be the vector containing the components of $\tilde{z}$ corresponding to nodes and edges in G_1 but not in G_2, and let z^2 be the vector containing the components of $\tilde{z}$ corresponding to nodes and edges in G_2 but not in G_1. Using these definitions, we can now write, up to reordering variables, $\tilde{z} = (z^1, \bar{z}, z^2)$.

By assumption, the vector $(z^1, \bar{z})$ is in $\text{conv}\mathcal{S}_{G_1}$. Thus, it can be written as a convex combination of points in $\mathcal{S}_{G_1}$; i.e., there exists $\mu \geq 0$ with $\sum_{(r,s) \in \mathcal{S}_{G_1}} \mu_{r,s} = 1$ such that

$$(3) \qquad \left(z^1, \bar{z}\right) = \sum_{(r,s) \in \mathcal{S}_{G_1}} \mu_{r,s}(r, s),$$

where the r vectors contain the components corresponding to nodes and edges in G_1 but not in G_2, and the s vectors contain the components corresponding to nodes and edges that are both in G_1 and in G_2.

Symmetrically, the vector $(\bar{z}, z^2)$ is in $\text{conv}\mathcal{S}_{G_2}$. Thus, it can be written as a convex combination of points in $\mathcal{S}_{G_2}$; i.e., there exists $\nu \geq 0$ with $\sum_{(s',t) \in \mathcal{S}_{G_2}} \nu_{s',t} = 1$ such that

$$(4) \qquad \left(\bar{z}, z^2\right) = \sum_{(s',t) \in \mathcal{S}_{G_2}} \nu_{s',t}(s', t),$$

where the s' vectors contain the components corresponding to nodes and edges that are both in G_1 and in G_2, and the t vectors contain the components corresponding to nodes and edges in G_2 but not in G_1.

By considering the component of (3) and of (4) corresponding to $\bar{p}$ we obtain

$$\bar{z}_{\bar{p}} = \sum_{(r,s) \in \mathcal{S}_{G_1} : s_{\bar{p}} = 1} \mu_{r,s} = \sum_{(s',t) \in \mathcal{S}_{G_2} : s'_{\bar{p}} = 1} \nu_{s',t},$$
$$1 - \bar{z}_{\bar{p}} = \sum_{(r,s) \in \mathcal{S}_{G_1} : s_{\bar{p}} = 0} \mu_{r,s} = \sum_{(s',t) \in \mathcal{S}_{G_2} : s'_{\bar{p}} = 0} \nu_{s',t}.$$

We claim that for every $(r, s) \in \mathcal{S}_{G_1}$ and every $(s', t) \in \mathcal{S}_{G_2}$ with $s_{\bar{p}} = s'_{\bar{p}}$ we have $(r, s', t) \in \mathcal{S}_G$. This is clearly true if $s_{\bar{p}} = s'_{\bar{p}} = 1$, as in this case all components of the two vectors s and s' are equal to one. Now, assume $s_{\bar{p}} = s'_{\bar{p}} = 0$. In this case we show that $(r, s') \in \mathcal{S}_{G_1}$, which implies $(r, s', t) \in \mathcal{S}_G$. Consider a component of r which corresponds to an edge, say, $\bar{e}$, of G_1 containing nodes in $V(G_2)$. By assumption $\bar{e} \supset \bar{p}$, and since $s_{\bar{p}} = 0$, if follows that $r_{\bar{e}} = 0$. Since $s'_{\bar{p}} = 0$ as well, we conclude that $(r, s') \in \mathcal{S}_{G_1}$.

Next, for every $(r, s) \in \mathcal{S}_{G_1}$ and $(s', t) \in \mathcal{S}_{G_2}$ with $s_{\bar{p}} = s'_{\bar{p}}$, we define

$$\tau_{r,s,s',t} := \begin{cases} \mu_{r,s} \cdot \nu_{s',t} / \bar{z}_{\bar{p}} & \text{if } s_{\bar{p}} = s'_{\bar{p}} = 1, \\ \mu_{r,s} \cdot \nu_{s',t} / (1 - \bar{z}_{\bar{p}}) & \text{if } s_{\bar{p}} = s'_{\bar{p}} = 0. \end{cases}$$

The multipliers $\tau_{r,s,s',t}$ are nonnegative and satisfy

$$
\begin{aligned}
&\sum_{(r,s)\in\mathcal{S}_{G_1},(s',t)\in\mathcal{S}_{G_2}:s_{\bar p}=s'_{\bar p}} \tau_{r,s,s',t} \\
&= \frac{\sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=1} \mu_{r,s} \sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=1} \nu_{s',t}}{\bar z_{\bar p}} + \frac{\sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=0} \mu_{r,s} \sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=0} \nu_{s',t}}{1-\bar z_{\bar p}} \\
&= \bar z_{\bar p} + (1-\bar z_{\bar p}) = 1.
\end{aligned}
$$

To prove that $\tilde z \in \mathrm{conv}\mathcal{S}_G$, it suffices to show that

$$
\sum_{(r,s)\in\mathcal{S}_{G_1},(s',t)\in\mathcal{S}_{G_2}:s_{\bar p}=s'_{\bar p}} \tau_{r,s,s',t}(r,s',t) = \left(z^1, \bar z, z^2\right). \tag{5}
$$

The restriction of (5) to the variables corresponding to nodes and edges in G_1 but not in G_2 can be shown as follows:

$$
\begin{aligned}
&\sum_{(r,s)\in\mathcal{S}_{G_1},(s',t)\in\mathcal{S}_{G_2}:s_{\bar p}=s'_{\bar p}} \tau_{r,s,s',t} r \\
&= \frac{\sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=1} \mu_{r,s} r \sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=1} \nu_{s',t}}{\bar z_{\bar p}} + \frac{\sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=0} \mu_{r,s} r \sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=0} \nu_{s',t}}{1-\bar z_{\bar p}} \\
&= \sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=1} \mu_{r,s} r + \sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=0} \mu_{r,s} r \\
&= \sum_{(r,s)\in\mathcal{S}_{G_1}} \mu_{r,s} r = z^1.
\end{aligned}
$$

The restriction of (5) to the remaining variables is shown below.

$$
\begin{aligned}
&\sum_{(r,s)\in\mathcal{S}_{G_1},(s',t)\in\mathcal{S}_{G_2}:s_{\bar p}=s'_{\bar p}} \tau_{r,s,s',t}(s',t) \\
&= \frac{\sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=1} \nu_{s',t}(s',t) \sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=1} \mu_{r,s}}{\bar z_{\bar p}} \\
&\quad + \frac{\sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=0} \nu_{s',t}(s',t) \sum_{(r,s)\in\mathcal{S}_{G_1}:s_{\bar p}=0} \mu_{r,s}}{1-\bar z_{\bar p}} \\
&= \sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=1} \nu_{s',t}(s',t) + \sum_{(s',t)\in\mathcal{S}_{G_2}:s'_{\bar p}=0} \nu_{s',t}(s',t) \\
&= \sum_{(s',t)\in\mathcal{S}_{G_2}} \nu_{s',t}(s',t) = \left(\bar z, z^2\right). \qquad \square
\end{aligned}
$$

3. The multilinear polytope for Berge-acyclic hypergraphs. In this section, we characterize multilinear sets for which the standard linearization defined by (2) is equivalent to the multilinear polytope. Namely, we show that $\mathrm{MP}^{\mathrm{LP}}_G = \mathrm{MP}_G$ if and only if the hypergraph G is Berge-acyclic. We start by establishing a property of MP_G and $\mathrm{MP}^{\mathrm{LP}}_G$ which enables us to identify conditions under which $\mathrm{MP}_G \subset \mathrm{MP}^{\mathrm{LP}}_G$ by examining the relative strength of such relaxations corresponding to hypergraphs with much simpler structures than G.

Given a hypergraph $G = (V, E)$ and $\bar{V} \subseteq V$, we define the *subhypergraph* of G induced by $\bar{V}$ as the hypergraph $G_{\bar{V}}$ with node set $\bar{V}$ and with edge set $\{e \cap \bar{V} : e \in E,\ |e \cap \bar{V}| \geq 2\}$. For every edge e of $G_{\bar{V}}$, there may exist several edges e' of G satisfying $e = e' \cap \bar{V}$; we denote by $e'(e)$ one such arbitrary edge of G. For ease of notation, we often identify an edge e of $G_{\bar{V}}$ with an edge $e'(e)$ of G. Denote by R a *relaxation* of the multilinear set; namely, R is a function that associates to each hypergraph G a set R_G containing all points in $\mathcal{S}_G$. Define

$$L_{\bar{V}} := \left\{ z \in \mathbb{R}^{V+E} : z_v = 1\ \forall v \in V \setminus \bar{V} \right\}. \tag{6}$$

Denote by $\mathrm{proj}_{G_{\bar{V}}}(R_G \cap L_{\bar{V}})$ the set obtained from $R_G \cap L_{\bar{V}}$ by projecting out all variables z_v for all $v \in V \setminus \bar{V}$, and z_f for all $f \in E \setminus \{e'(e) : e \in E(G_{\bar{V}})\}$. The following lemma establishes that for the multilinear polytope and the standard linearization the two sets $R_{G_{\bar{V}}}$ and $\mathrm{proj}_{G_{\bar{V}}}(R_G \cap L_{\bar{V}})$ are in fact identical.

LEMMA 6. *Let $G = (V, E)$ be a hypergraph and let $L_{\bar{V}}$ be a set defined by* (6) *for some $\bar{V} \subseteq V$. Then*
(i) $MP_{G_{\bar{V}}} = \mathrm{proj}_{G_{\bar{V}}}(MP_G \cap L_{\bar{V}})$,
(ii) $MP^{LP}_{G_{\bar{V}}} = \mathrm{proj}_{G_{\bar{V}}}(MP^{LP}_G \cap L_{\bar{V}})$.

Proof. (i) The set $\mathrm{MP}_G \cap L_{\bar{V}}$ is a face of MP_G; hence $\mathrm{MP}_G \cap L_{\bar{V}} = \mathrm{conv}(\mathcal{S}_G \cap L_{\bar{V}})$. Moreover, since the operations of taking the convex hull and taking the projection commute, we have $\mathrm{proj}_{G_{\bar{V}}}(\mathrm{conv}(\mathcal{S}_G \cap L_{\bar{V}})) = \mathrm{conv}(\mathrm{proj}_{G_{\bar{V}}}(\mathcal{S}_G \cap L_{\bar{V}}))$. Finally from the definition of the subhypergraph $G_{\bar{V}}$ it follows that $\mathrm{proj}_{G_{\bar{V}}}(\mathcal{S}_G \cap L_{\bar{V}}) = \mathcal{S}_{G_{\bar{V}}}$, which in turn implies that $\mathrm{proj}_{G_{\bar{V}}}(\mathrm{MP}_G \cap L_{\bar{V}}) = \mathrm{MP}_{G_{\bar{V}}}$.

(ii) From (2) and (6), it follows that

$$\begin{aligned} \mathrm{MP}^{\mathrm{LP}}_G \cap L_{\bar{V}} = \Big\{ z : {} & z_v \leq 1 \ \ \forall v \in \bar{V},\ z_v = 1,\ \forall v \in V \setminus \bar{V}, \\ & z_e \geq 0,\ z_e \geq \sum_{v \in e \cap \bar{V}} z_v - |e \cap \bar{V}| + 1 \ \ \forall e \in E, \\ & z_e \leq z_v \ \ \forall e \in E \ \ \forall v \in e \cap \bar{V},\ z_e \leq 1 \ \ \forall e \in E,\ \forall v \in e \setminus \bar{V} \Big\}. \end{aligned}$$

First notice that in the above system, each variable z_v, $v \in V \setminus \bar{V}$ only appears in the equality $z_v = 1$. Hence, projecting out these variables from $\mathrm{MP}^{\mathrm{LP}}_G \cap L_{\bar{V}}$ simply amounts to dropping the corresponding equalities from the above system. Now consider a variable z_f for some $f \in E \setminus \bar{E}$, where we define $\bar{E} := \{e'(e) : e \in E(G_{\bar{V}})\}$. The variable z_f appears in the following inequalities: $z_f \geq 0$, $z_f \geq \sum_{v \in f \cap \bar{V}} z_v - |f \cap \bar{V}| + 1$, $z_f \leq z_v$ for all $v \in f \cap \bar{V}$ and $z_f \leq 1$ for all $v \in f \setminus \bar{V}$. By projecting out z_f from these inequalities using Fourier–Motzkin elimination, we obtain a system of inequalities that are implied by the following system: $0 \leq z_v \leq 1$ for all $v \in f \cap \bar{V}$ and $z_v = 1$ for all $v \in f \setminus \bar{V}$. Consequently, by projecting out all variables z_v for $v \in V \setminus \bar{V}$, and

z_e for $e \in E \setminus \tilde{E}$, we obtain

$$\text{proj}_{\bar{V}}\left(\text{MP}_G^{\text{LP}} \cap L_{\bar{V}}\right) = \Big\{ z : z_v \leq 1 \quad \forall v \in \bar{V},$$
$$z_e \geq 0,\ z_e \geq \sum_{v \in e \cap \bar{V}} z_v - |e \cap \bar{V}| + 1 \quad \forall e \in \bar{E},$$
$$z_e \leq z_v \quad \forall e \in \bar{E} \quad \forall v \in e \cap \bar{V} \Big\}.$$

Finally, from the definitions of the standard linearization and the subhypergraph $G_{\bar{V}}$ it follows that $\text{proj}_{G_{\bar{V}}}(\text{MP}_G^{\text{LP}} \cap L_{\bar{V}}) = \text{MP}_{G_{\bar{V}}}^{\text{LP}}$. ∎

Now consider a hypergraph G for which we would like to show that $\text{MP}_G \subset \text{MP}_G^{\text{LP}}$. Since $\text{MP}_G \subseteq \text{MP}_G^{\text{LP}}$, it suffices to show that for some $\bar{V} \subseteq V(G)$ we have $\text{proj}_{G_{\bar{V}}}(\text{MP}_G \cap L_{\bar{V}}) \subset \text{proj}_{G_{\bar{V}}}(\text{MP}_G^{\text{LP}} \cap L_{\bar{V}})$. By Lemma 6, the latter inclusion can be established by showing that $\text{MP}_{G_{\bar{V}}} \subset \text{MP}_{G_{\bar{V}}}^{\text{LP}}$. Indeed, a careful selection of the subset $\bar{V}$ and employing the above technique is a key step in the proof of Theorem 7.

Let QP_G^{LP} denote the standard linearization of the Boolean quadric polytope QP_G. In [25], Padberg shows that $\text{QP}_G^{\text{LP}} = \text{QP}_G$ if and only if G is an acyclic graph. The following theorem generalizes the above result to higher degree multilinear sets using the notion of hypergraph acyclicity introduced in section 1. We remark that this result has been discovered independently in [9] using a different proof technique.

THEOREM 7. *$MP_G^{LP} = MP_G$ if and only if G is a Berge-acyclic hypergraph.*

Proof. "⇒" We first show that if the hypergraph G contains a Berge-cycle C of length two, then MP_G^{LP} does not coincide with MP_G. Let $E(C) = \{e_1, e_2\}$ with $|e_1 \cap e_2| \geq 2$. It then follows that the inequality

$$\sum_{v \in e_2 \setminus e_1} z_v + z_{e_1} - z_{e_2} \leq |e_2 \setminus e_1| \tag{7}$$

is valid for $\mathcal{S}_G$. To see this, observe that the value of $\sum_{v \in e_2 \setminus e_1} z_v + z_{e_1}$ does not exceed the right-hand side of inequality (7), unless $z_v = 1$ for all $v \in e_2 \setminus e_1$ and $z_{e_1} = 1$; however, this in turn implies that $z_{e_2} = 1$. Thus, inequality (7) is valid for $\mathcal{S}_G$. (See also [14], wherein the validity of inequalities (7) for $\mathcal{S}_G$ is established.) Now, consider the point $\tilde{z}$ defined as follows: $\tilde{z}_v = 1$ for all $v \in e_2 \setminus e_1$, $\tilde{z}_v = 1/2$ for all $v \in e_1$, $\tilde{z}_v = 0$ for the remaining nodes in G, $z_{e_1} = 1/2$, $z_{e_2} = 0$, $z_e = 1$ for all $e \subseteq e_2 \setminus e_1$, $z_e = 0$ for all $e \nsubseteq e_1 \cup e_2$, and $z_e = 1/2$ for all remaining edges in G. Clearly, this point does not satisfy inequality (7), as $|e_2 \setminus e_1| + 1/2 - 0 \nleq |e_2 \setminus e_1|$. However, it can be checked that $\tilde{z}$ belongs to MP_G^{LP}, provided that $|e_1 \cap e_2| \geq 2$. Hence, if the hypergraph G contains a Berge-cycle of length two, we have $\text{MP}_G \subset \text{MP}_G^{\text{LP}}$.

Now, consider a hypergraph G with $|e_1 \cap e_2| \leq 1$ for all $e_1, e_2 \in E(G)$; that is, G does not contain any Berge-cycle of length two. We show that if G contains a Berge-cycle of length greater than or equal to three, then $\text{MP}_G \subset \text{MP}_G^{\text{LP}}$. Denote by C a Berge-cycle of minimum length t, where $t \geq 3$. We claim that the subhypergraph $G_{V(C)}$ is a graph that consists of a chordless cycle of length t. To obtain a contradiction, suppose that $G_{V(C)}$ is not a chordless cycle. Since C is a Berge-cycle of minimum length, it follows that there exists an edge $\bar{e}$ in $E(G_{V(C)})$ containing at least three nodes in $V(C)$. Denote by $\tilde{e}$ an edge of G with $\bar{e} = \tilde{e} \cap V(C)$. Since by assumption $|e_i \cap e_j| \leq 1$ for all $e_i, e_j \in E(G)$, there exist no two nodes in $\bar{e}$ that are also present in another edge of G. Define $C = v_1, e_1, v_2, \ldots, v_t, e_t, v_1$. Without loss of generality, suppose that $v_1 \in \bar{e}$ and $v_2 \notin \bar{e}$. Let v_k be the next node of $V(C)$ after the first node v_1 that is present in $\bar{e}$. Clearly, $k < t$ since by assumption $\bar{e}$ contains at least three nodes of C. It then follows that the sequence $v_1, e_1, v_2, \ldots, e_{k-1}, v_k, \tilde{e}, v_1$ is a Berge-cycle of length k. However, this contradicts the assumption that C is Berge-cycle of minimum length. Therefore, the graph $G_{V(C)}$ consists of a chordless cycle. By Lemma 6, to prove $\mathrm{MP}_G \subset \mathrm{MP}^{\mathrm{LP}}_G$, it suffices to show that $\mathrm{MP}_{G_{V(C)}} \subset \mathrm{MP}^{\mathrm{LP}}_{G_{V(C)}}$. The polytope $\mathrm{MP}_{G_{V(C)}}$ is clearly integral. However, it is well-known that $\mathrm{MP}^{\mathrm{LP}}_{G_{V(C)}}$ is not integral, since the graph $G_{V(C)}$ consists of a chordless cycle [25]. Consequently, if the hypergraph G contains a Berge-cycle, we have $\mathrm{MP}_G \subset \mathrm{MP}^{\mathrm{LP}}_G$.

"$\Leftarrow$" Conversely, let G be a Berge-acyclic hypergraph. We show that $\mathrm{MP}^{\mathrm{LP}}_G = \mathrm{MP}_G$. The proof is by induction on the number of edges of G. In the base case G has only one edge and it is well-known that in this case $\mathrm{MP}^{\mathrm{LP}}_G = \mathrm{MP}_G$. To prove the inductive step, we assume that G has at least two edges. We first show that there exists at least one edge $\tilde{e}$ of G such that $\tilde{e} \cap (\cup_{e \in E(G) \setminus \tilde{e}} e) = \{\tilde{v}\}$ for some $\tilde{v} \in V(G)$. To obtain a contradiction, suppose that such an edge does not exist. By Berge-acyclicity, every two edges of G intersect in at most one node, as otherwise, they form a Berge-cycle of length two. It then follows that every edge of G intersects with at least two other edges in two distinct nodes. In particular, G has at least three edges. However, this implies that we can always find a Berge-cycle, which is in contradiction with the assumption that G is Berge-acyclic. Hence, G has an edge $\tilde{e}$ with $\tilde{e} \cap (\cup_{e \in E(G) \setminus \tilde{e}} e) = \{\tilde{v}\}$ for some $\tilde{v} \in V(G)$. We now define G_1 as the section hypergraph of G induced by $\tilde{e}$, and G_2 as the section hypergraph of G induced by $\cup_{e \in E(G) \setminus \tilde{e}} e$. Clearly, $G_1 \cup G_2 = G$ and $G_1 \cap G_2 = \{\tilde{v}\}$. Thus, by Theorem 5, the set $\mathcal{S}_G$ is decomposable into $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$. Both hypergraphs G_1 and G_2 have fewer edges than G and are Berge-acyclic since they are section hypergraphs of G. Therefore, by the induction hypothesis we have $\mathrm{MP}^{\mathrm{LP}}_{G_1} = \mathrm{MP}_{G_1}$ and $\mathrm{MP}^{\mathrm{LP}}_{G_2} = \mathrm{MP}_{G_2}$, implying $\mathrm{MP}^{\mathrm{LP}}_G = \mathrm{MP}_G$. ∎

Clearly, for a rank-r hypergraph $G = (V, E)$, the standard linearization $\mathrm{MP}^{\mathrm{LP}}_G$ has at most $|V| + (r + 2)|E|$ linear inequalities. Therefore, by Theorem 7, for a Berge-acyclic hypergraph G, problem (MO) can be solved via linear optimization in polynomial time, i.e., in a number of iterations bounded by a polynomial in $|V|$, $|E|$, and in the size of the vector c. (See [30] for more details.)

4. The multilinear polytope for γ-acyclic hypergraphs. As we detailed in section 1, Berge-acyclicity is the most restrictive type of hypergraph acyclicity. Indeed, by Theorem 7, the multilinear polytope for Berge-acyclic hypergraphs has a very simple structure; that is, $\mathrm{MP}_G = \mathrm{MP}^{\mathrm{LP}}_G$. In this section, we study the structure of the multilinear polytope for the next class of acyclic hypergraphs, in increasing order of generality, namely, the class of γ-acyclic hypergraphs. As we described in section 1, γ-acyclic hypergraphs represent a significant generalization of Berge-acyclic hypergraphs and may contain Berge-cycles of arbitrary lengths, in general.

We start by establishing a key connection between γ-acyclic and laminar hypergraphs. By building upon a result concerning balanced matrices and integral polyhedra, in section 4.1, we show that the multilinear polytope for laminar hypergraphs has a simple structure. Subsequently, in section 4.2, we introduce a generalization of the inequalities defined by (7), which we will refer to as flower inequalities. We introduce a new polyhedral relaxation of the multilinear set, obtained by addition of all flower inequalities to its standard linearization. Finally, using our decomposability results of section 2 together with our convex hull characterization for laminar hypergraphs, in section 4.3, we prove that this new relaxation coincides with the multilinear polytope if and only if the underlying hypergraph is γ-acyclic.

4.1. Laminar hypergraphs. Recall that a hypergraph G is *laminar* if for any two edges $e_1, e_2 \in E(G)$, one of the following is satisfied: (i) $e_1 \cap e_2 = \emptyset$, (ii) $e_1 \subset e_2$, (iii) $e_2 \subset e_1$. The following proposition establishes a key connection between laminar hypergraphs and γ-acyclic hypergraphs.

PROPOSITION 8. *Let $G = (V, E)$ be a γ-acyclic hypergraph, and let $e' \in E$. Then the subhypergraph $G_{e'}$ is laminar.*

Proof. Assume by contradiction that $G_{e'}$ is not laminar. Then there exist nodes $v_1, v_2, v_3 \in V(G_{e'})$ and edges $e_i, e_j \in E(G_{e'})$ such that $v_1, v_2 \in e_i$, $v_1, v_3 \in e_j$, $v_2 \notin e_j$, $v_3 \notin e_i$. Note that $e' \in E(G_{e'})$ contains all three nodes v_1, v_2, v_3. Let $\tilde{e}_i, \tilde{e}_j \in E$ such that $e_i = \tilde{e}_i \cap e', e_j = \tilde{e}_j \cap e'$. Then $\{\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_2, v_3\}\} = \{e \cap \{v_1, v_2, v_3\} : e \in \{\tilde{e}_i, \tilde{e}_j, e'\}\}$. As G is γ-acyclic, this contradicts property (ii) of Proposition 4. ∎

In particular, Proposition 8 implies that if a γ-acyclic hypergraph G has an edge that contains all nodes of G, then G is laminar. In our next result, we characterize the multilinear polytope for laminar hypergraphs. To do so, we make use of a fundamental result due to Conforti and Cornuéjols regarding the connection between integral polyhedra and balanced matrices. We recall that a $0, \pm 1$ matrix is *balanced* if, in every square submatrix with exactly two nonzero entries per row and per column, the sum of the entries is a multiple of 4.

THEOREM 9 (see [12, Theorem 6.13]). *Let A be a balanced $0, \pm 1$ matrix with rows a^i, $i \in S$, and let S_1, S_2, S_3 be a partition of S. For each a^i, let $n(a^i)$ denote the number of elements equal to -1. Then*

$$\begin{aligned} R(A) = \{x \in \mathbb{R}^n : a^i x &\geq 1 - n(a^i) \qquad &&\text{for } i \in S_1, \\ a^i x &= 1 - n(a^i) &&\text{for } i \in S_2, \\ a^i x &\leq 1 - n(a^i) &&\text{for } i \in S_3, \\ 0 \leq x &\leq 1\} \end{aligned}$$

is an integral polytope.

Given a laminar hypergraph $G = (V, E)$, and an edge $e \in E$, we define $I(e) := \{p \in V \cup E : p \subset e,\ p \not\subset e' \text{ for } e' \in E,\ e' \subset e\}$. Given a $p \in V \cup E$ that is strictly contained in at least one edge of E, there exists a unique edge $\bar{e}$ of G that satisfies $p \in I(\bar{e})$. To obtain a contradiction, assume that there exist two distinct edges $\bar{e}_1, \bar{e}_2$ with $p \in I(\bar{e}_1) \cap I(\bar{e}_2)$. It then follows that $p \subset \bar{e}_1$ and $p \subset \bar{e}_2$. Since $p \in I(\bar{e}_1)$, we have $\bar{e}_2 \not\subset \bar{e}_1$. Symmetrically, since $p \in I(\bar{e}_2)$, we have $\bar{e}_1 \not\subset \bar{e}_2$. However, this contradicts the laminarity of G. The next theorem characterizes the multilinear polytope for laminar hypergraphs.

THEOREM 10. *Let $G = (V, E)$ be a laminar hypergraph. Then MP_G is described by the following system:*

$$z_v \leq 1 \qquad \forall v \in V, \tag{8}$$

$$-z_e \leq 0 \qquad \forall e \in E \text{ such that } e \not\subset f \text{ for } f \in E, \tag{9}$$

$$-z_p + z_e \leq 0 \qquad \forall e \in E,\ \forall p \in I(e), \tag{10}$$

$$\sum_{p \in I(e)} z_p - z_e \leq |I(e)| - 1 \qquad \forall e \in E. \tag{11}$$

Proof. Let Q be the polyhedron described by inequalities (8)–(11). In the following, we first show that the integer points in Q coincide with those of MP_G. To do so, it suffices to prove that $\text{MP}_G \subseteq Q \subseteq \text{MP}_G^{\text{LP}}$. Subsequently, we show that Q is an integral polytope, which together with the first claim implies $Q = \text{MP}_G$.

We start by showing that Q is a valid relaxation of $\mathcal{S}_G$, i.e., $\text{MP}_G \subseteq Q$. Clearly, inequalities (8) and (9) are present in the description of MP_G^{LP}. In addition, inequalities (10), if p is a node, and inequalities (11), if $I(e)$ only consists of nodes, are present in MP_G^{LP}. The validity of the remaining inequalities in (10) as well as inequalities (11) follows from the fact that for any $e \in E$, we have $z_e = 1$ if and only if $z_p = 1$ for all $p \in I(e)$. Hence, $\text{MP}_G \subseteq Q$.

We now show that $Q \subseteq \text{MP}_G^{\text{LP}}$. Let us consider the inequalities in the description of MP_G^{LP} given by (2). Inequalities $z_v \leq 1$ for every $v \in V$ are given by (8). Inequalities $z_e \geq 0$ for every e such that e is not contained in any other edge are given by (9). For every other edge e_0, let $e_1, \dots, e_t$ be a maximal sequence of edges such that $e_{i-1} \in I(e_i)$ for every $i = 1, \dots, t$. Then inequality $z_{e_0} \geq 0$ can be obtained by summing inequalities $z_{e_{i-1}} \geq z_{e_i}$ in (10) for every $i = 1, \dots, t$, and inequality $z_{e_t} \geq 0$ in (9). Inequalities $z_e \geq \sum_{v \in e} z_v - |e| + 1$ for every e such that e does not contain any other edge are given by (11). For every other edge e, inequality $z_e \geq \sum_{v \in e} z_v - |e| + 1$ can be obtained by summing inequalities $z_f \geq \sum_{p \in I(f)} z_p - |I(f)| + 1$ in (11) for every $f \subseteq e$. Inequalities $z_e \leq z_v$ for every edge $e \in E$ and node $v \in I(e)$ are given by (10). Now let e_0 be any edge and let v be a node not in $I(e_0)$. Let $e_1, \dots, e_t$ be a maximal sequence of edges such that $e_i \in I(e_{i-1})$ for every $i = 1, \dots, t$, and such that $v \in e_t$. Then inequality $z_{e_0} \leq z_v$ can be obtained by summing inequalities $z_{e_{i-1}} \leq z_{e_i}$ in (10) for every $i = 1, \dots, t$, and inequality $z_{e_t} \leq z_v$ in (10).

We now show that Q is an integral polytope. Clearly, inequalities (8)–(11) are of the form defined in the statement of Theorem 9. Thus by this theorem, it suffices to show that the constraint matrix of system (8)–(11) is balanced. In fact, by definition of a $0, \pm 1$ balanced matrix, we can equivalently show that the constraint matrix A corresponding to the system (10)–(11) is balanced as inequalities (8) and (9) introduce singleton rows in the constraint matrix. Assume by contradiction that there exists a square submatrix of A with exactly two nonzero entries per row and per column, such that the sum of the entries is congruous to 2 modulus 4. Let B be a square submatrix of this type with the minimum number of rows.

We show that no column of B corresponds to a node of G. By contradiction assume that a column of B corresponds to a node $\bar{v} \in V$. Let $\bar{e}$ be the unique edge of G that satisfies $\bar{v} \in I(\bar{e})$. Then $z_{\bar{v}}$ has a nonzero coefficient only in the following two inequalities from the system (10)–(11): $-z_{\bar{v}} + z_{\bar{e}} \leq 0$ defined by (10), and $\sum_{p \in I(\bar{e})} z_p - z_{\bar{e}} \leq |I(\bar{e})| - 1$ defined by (11). Since the column of B corresponding to $\bar{v}$ has two nonzero entries, these two inequalities must correspond to two rows of B. The first inequality has only one more nonzero coefficient, namely, the one corresponding to $\bar{e}$. Therefore, a column of B must correspond to $\bar{e}$. Now, let B' be obtained from B by removing the rows corresponding to the above two inequalities, and the columns corresponding to $\bar{v}$ and $\bar{e}$. The nonzero entries of B present in the removed rows and columns are a -1 and a $+1$ in the first inequality, and a $+1$ and a -1 in the second inequality, which implies that the sum of the entries of B' is congruous to 2 modulus 4. It follows that B' is a square submatrix of A with fewer rows than B, contradicting the minimality of B.

Since the sum of the entries of B is congruous to 2 modulus 4, there is at least one row of B with two entries of the same sign. This row then corresponds to an inequality in (11), say, the one corresponding to an edge $e_0 \in E$. Since no column of B corresponds to a node of G, the two entries of the same sign must correspond to two edges, say, e_1 and e_1' in $I(e_0)$. In particular, two columns of B correspond to e_1 and e_1'. Since each row contains only two nonzero entries, we also argue that no column of B corresponds to e_0.

We now show that there is at least one edge in $I(e_1)$, and that a column of B corresponds to it. As B has two nonzeros per column, there is another inequality among (10), (11) that corresponds to a row of B with a nonzero corresponding to e_1. If this inequality is in (10), then we claim that it is $-z_p + z_{e_1} \le 0$ for $p \in I(e_1)$. If not, since e_0 is the unique edge with $e_1 \in I(e_0)$, it must be $-z_{e_1} + z_{e_0} \le 0$. But then the corresponding row of B has only one nonzero entry since no column of B corresponds to e_0. If this inequality is in (11), since e_0 is the unique edge with $e_1 \in I(e_0)$, then it must be $\sum_{p \in I(e_1)} z_p - z_{e_1} \le |I(e_1)| - 1$. As no column of B corresponds to a node of G, in both cases we argue that a column of B must correspond to an edge, say, e_2, in $I(e_1)$.

Similarly, we show that there is at least one edge in $I(e_2)$, and that a column of B corresponds to it. There is another inequality among (10), (11) corresponding to a row of B with a nonzero corresponding to e_2. This inequality cannot have a nonzero coefficient corresponding to e_1, as otherwise we would obtain a column of B with three nonzero entries. Therefore, such inequality is either $-z_p + z_{e_2} \le 0$ in (10) for $p \in I(e_2)$, or $\sum_{p \in I(e_2)} z_p - z_{e_2} \le |I(e_2)| - 1$ in (11). In both cases we argue that a column of B must correspond to an edge, say, e_3, in $I(e_2)$.

By repeating the latter argument, we can show the existence of an edge $e_t \in I(e_{t-1})$ in G for any positive integer t, which contradicts the finiteness of G. ∎

Before proceeding further, we remark that our proof of Theorem 10 relies on the balancedness of the constraint matrix of the minimal system defining the polytope MP_G, which does not hold for γ-acyclic hypergraphs, in general. The following example demonstrates that if we relax the laminarity assumption of G, the constraint matrix of the minimal system defining MP_G is no longer balanced.

Example 1. Consider the γ-acyclic hypergraph G with $V(G) = \{v_1, v_2, v_3, v_4\}$ and $E(G) = \{e_{123}, e_{234}\}$, where edge e_I contains the nodes with indices in I. It can be checked that the following inequalities define facets of MP_G:

$$\begin{aligned} -z_{v_2} + z_{e_{123}} &\le 0, \\ -z_{v_3} + z_{e_{123}} &\le 0, \\ +z_{v_2} + z_{v_3} + z_{v_4} - z_{e_{234}} &\le 2. \end{aligned}$$

Let B be the constraint matrix of the above system. The square submatrix B' of B obtained by selecting columns corresponding to nodes v_2, v_3 and edge e_{123} has exactly two nonzero entries per row and per column, and the sum of the entries is congruous to 2 modulus 4. Therefore the constraint matrix of the minimal system defining MP_G is not balanced.

4.2. Flower inequalities. In what follows, we define the *support hypergraph* of a valid inequality $az \le \alpha$ for MP_G, as the hypergraph $G(a)$, where $V(G(a)) = \{v \in V : a_v \neq 0\} \cup (\cup_{e \in E: a_e \neq 0} e)$, and $E(G(a)) = \{e \in E : a_e \neq 0\}$. Let us revisit the valid inequalities for MP_G defined by (7). Clearly, the support hypergraph of these inequalities contains Berge-cycle of length two. In [14], the authors show that for a hypergraph G consisting of two edges intersecting in more than one node, the addition of these inequalities to the standard linearization yields the corresponding multilinear polytope. In this section, we present a significant generalization of this result.

Consider a hypergraph $G = (V, E)$. We say that two edges of G are *adjacent* if their intersection is not empty. Let e_0 be an edge of G and let e_k, $k \in K$, be the set of all edges adjacent to e_0 with $|e_0 \cap e_k| \geq 2$. Let T be a nonempty subset of K such that $e_i \cap e_j = \emptyset$ for all $i, j \in T$ with $i \neq j$. Then the *flower inequality* centered at e_0 with neighbors e_k, $k \in T$, is given by

$$\sum_{v \in e_0 \setminus \cup_{k \in T} e_k} z_v + \sum_{k \in T} z_{e_k} - z_{e_0} \leq |e_0 \setminus \cup_{k \in T} e_k| + |T| - 1. \tag{12}$$

It is simple to check that the support hypergraph of flower inequalities contains Berge-cycles of length two only. We first show that inequalities (12) are valid for MP_G. Clearly, for any given nonempty subset T of K, the left-hand side of these inequalities could exceed the right-hand side, only if $z_v = 1$ for all $v \in e_0 \setminus \cup_{k \in T} e_k$ and $z_{e_k} = 1$ for all $k \in T$. However, this in turn implies that $z_{e_0} = 1$. It then follows that inequalities (12) are valid for MP_G. We refer to the inequalities of the form (12) for all nonempty $T \subseteq K$ satisfying $e_i \cap e_j = \emptyset$ for all $i, j \in T$, as *the system of flower inequalities centered at* e_0. We define the *flower relaxation* MP^F_G as the polytope obtained by adding the system of flower inequalities centered at each edge of G to MP^{LP}_G.

Clearly, inequalities (8)–(11) in the statement of Theorem 10 are either flower inequalities or are present in (2): inequalities (8) and (9) are present in the description of MP^{LP}_G, and inequality (10) is present in the description of MP^{LP}_G for all $p \in V$, and is a flower inequality for all $p \in E$. Finally, inequality (11) corresponds to an inequality in MP^{LP}_G provided that $I(e)$ contains no edge of G; otherwise it is a flower inequality. Thus, we have the following result.

COROLLARY 11. *Let G be a laminar hypergraph. Then $MP_G = MP^F_G$.*

Consider a hypergraph G with $E(G) = \{e_k : k \in \{1, \ldots, K\}\}$ such that $e_1 \supset e_2 \supset \cdots \supset e_{K-1} \supset e_K$. Clearly, this hypergraph is laminar. The multilinear polytope for this special class of laminar hypergraphs is characterized in [19, 14].

4.3. γ-acyclic hypergraphs.

Our main result in this section states that MP^F_G coincides with MP_G if and only if the underlying hypergraph G is γ-acyclic. To this end, in the following two lemmata, we establish some basic properties of the polytope MP^F_G.

LEMMA 12. *Let $\tilde{G}$ be a partial hypergraph of the hypergraph G. Then all inequalities defining $MP^F_{\tilde{G}}$ are also present in the system defining MP^F_G.*

Proof. Clearly, the description of MP^{LP}_G contains all inequalities present in the description of $\text{MP}^{\text{LP}}_{\tilde{G}}$, since the latter is obtained by replacing each multilinear term $z_e = \prod_{v \in e} z_v$ by its convex hull over the unit hypercube for all $e \in E(\tilde{G})$ and we have $E(\tilde{G}) \subseteq E(G)$. In addition, from the definition of flower inequalities it follows that every flower inequality for $\mathcal{S}_{\tilde{G}}$ is also a flower inequality for $\mathcal{S}_G$, as again $E(\tilde{G}) \subseteq E(G)$. Consequently, all inequalities defining $\text{MP}^F_{\tilde{G}}$ are also present in the system defining MP^F_G. ∎

LEMMA 13. *Let $G = (V, E)$ be a hypergraph, let $\bar{V} \subseteq V$, and let $L_{\bar{V}}$ be a set defined by* (6). *Then*

$$MP^F_{G_{\bar{V}}} \subseteq \text{proj}_{G_{\bar{V}}} \left(MP^F_G \cap L_{\bar{V}} \right).$$

Proof. We prove the statement by showing that every nonredundant inequality in $\text{proj}_{G_{\bar{V}}}(\text{MP}^F_G \cap L_{\bar{V}})$ is also present in $\text{MP}^F_{G_{\bar{V}}}$. First, let us characterize the projection operation $\text{proj}_{G_{\bar{V}}}(\cdot)$ for the set $\text{MP}^F_G \cap L_{\bar{V}}$. As before we define $\bar{E} := \{e'(e) : e \in E(G_{\bar{V}})\}$. Consider an edge $e \in E \setminus \bar{E}$; based on the cardinality of $e \cap \bar{V}$, three cases arise:

(i) $e \cap \bar{V} = \emptyset$: since $z_v = 1$ for all $v \in e$, in this case the inequality $z_e \geq \sum_{v\in e} z_v - |e| + 1$ simplifies to $z_e \geq 1$, which together with $z_e \leq z_v$, $v \in e$ implies that $z_e = 1$.

(ii) $e \cap \bar{V} = \{\bar{v}\}$ for some $\bar{v} \in \bar{V}$: since $z_v = 1$ for all $v \in e \setminus \{\bar{v}\}$, in this case the inequality $z_e \geq \sum_{v\in e} z_v - |e| + 1$ simplifies to $z_e \geq z_{\bar{v}}$, which together with $z_e \leq z_v$, $v \in e$ implies that $z_e = z_{\bar{v}}$.

(iii) $|e\cap \bar{V}| \geq 2$: by definition of $\bar{E}$, for any $e \in E\setminus \bar{E}$ with $|e\cap \bar{V}| \geq 2$, there exists an edge $\bar{e} \in \bar{E}$ such that $\bar{e} \cap \bar{V} = e \cap \bar{V}$. Clearly, the following two flower inequalities are present in MP^F_G:

$$\sum_{v\in e\setminus \bar{e}} z_v + z_{\bar{e}} - z_e \leq |e \setminus \bar{e}|$$

and

$$\sum_{v\in \bar{e}\setminus e} z_v + z_e - z_{\bar{e}} \leq |\bar{e} \setminus e|.$$

Since by assumption $\bar{e}\cap\bar{V} = e\cap\bar{V}$, it follows that $z_v \in V \setminus \bar{V}$ for all $v \in (\bar{e}\setminus e)\cup(e\setminus \bar{e})$. Hence, substituting $z_v = 1$ for $v \in V \setminus \bar{V}$ in the above inequalities yields $z_e = z_{\bar{e}}$.

Hence, $\text{proj}_{G_{\bar{V}}}(\text{MP}^F_G\cap L_{\bar{V}})$ can be obtained from MP^F_G by substituting $z_v = 1$ for all $v \in V \setminus \bar{V}$, $z_e = 1$ for all $e \in E \setminus \bar{E}$ with $e \cap \bar{V} = \emptyset$, $z_e = z_{\bar{v}}$ for all $e \in E \setminus \bar{E}$ with $e \cap \bar{V} = \{\bar{v}\}$ for some $\bar{v} \in \bar{V}$, and $z_e = z_{\bar{e}}$ for all $e \in E \setminus \bar{E}$ with $e \cap \bar{V} = \bar{e} \cap \bar{V}$ for some $\bar{e} \in \bar{E}$, and dropping out all variables z_v, $v \in V \setminus \bar{V}$ and z_e for all $e \in E\setminus\bar{E}$ from the description of MP^F_G. It is then simple to verify that all nonredundant inequalities in $\text{proj}_{G_{\bar{V}}}(\text{MP}^F_G \cap L_{\bar{V}})$ corresponding to the standard linearization of $\mathcal{S}_G$ are present in $\text{MP}^{\text{LP}}_{G_{\bar{V}}}$ and hence are also present in $\text{MP}^F_{G_{\bar{V}}}$. Hence, it suffices to show that the same statement holds for the remaining inequalities in $\text{proj}_{G_{\bar{V}}}(\text{MP}^F_G\cap L_{\bar{V}})$, i.e., those corresponding to the flower inequalities in MP^F_G.

Consider a flower inequality for $\mathcal{S}_G$ centered at e_0 with neighbors e_k, $k \in T$, as defined by (12). Substituting $z_v = 1$ for all $v \in V \setminus \bar{V}$ in inequality (12), we obtain

$$\sum_{v\in(e_0\setminus\cup_{k\in T}e_k)\cap\bar{V}} z_v + \sum_{k\in T} z_{e_k} - z_{e_0} \leq |(e_0 \setminus \cup_{k\in T} e_k) \cap \bar{V}| + |T| - 1. \tag{13}$$

We would like to show that if inequality (13) is nonredundant for $\text{MP}^F_G\cap L_{\bar{V}}$, then the corresponding inequality in $\text{proj}_{G_{\bar{V}}}(\text{MP}^F_G \cap L_{\bar{V}})$ is also present in $\text{MP}^F_{G_{\bar{V}}}$. We claim that without loss of generality we can assume that $|e_0 \cap e_k \cap \bar{V}| \geq 2$ for all $k \in T$. To see this, suppose that $|e_0 \cap e_k \cap \bar{V}| \leq 1$ for all $k \in T'$, where T' is a nonempty subset of T. We now show that in this case inequality (13) is implied by a number of inequalities present in the description of $\text{MP}^F_G\cap L_{\bar{V}}$. Consider the following inequality:

$$\sum_{v\in e_0\cap(\cup_{k\in T'}e_k)} z_v + \sum_{v\in e_0\setminus\cup_{k\in T}e_k} z_v + \sum_{k\in T\setminus T'} z_{e_k} - z_{e_0} \leq |e_0\setminus\cup_{k\in T\setminus T'}e_k| + |T\setminus T'| - 1. \tag{14}$$

Clearly, the above inequality is a flower inequality for $\mathcal{S}_G$ centered at e_0 with neighbors

e_k, $k \in T \setminus T'$, if $T \setminus T'$ is nonempty and simplifies to $\sum_{v \in e_0} z_v - z_{e_0} \leq |e_0| - 1$, otherwise. Notice that the latter inequality is present in the description of $\mathrm{MP}^{\mathrm{LP}}_G$ and hence is present in MP^F_G. For each $k \in T'$, consider the inequality $z_{e_k} - z_{\bar{v}} \leq 0$ if $e_0 \cap e_k \cap \bar{V} = \{\bar{v}\}$, and $z_{e_k} \leq 1$ if $e_0 \cap e_k \cap \bar{V} = \emptyset$. Substituting $z_v = 1$ for all $v \in V \setminus \bar{V}$ in these inequalities together with inequality (14) and summing up the resulting inequalities, we obtain inequality (13). Henceforth, we can assume that in (13) we have $|e_0 \cap e_k \cap \bar{V}| \geq 2$ for all $k \in T$.

Denote by $\tilde{e}_0$ the edge of G in $\bar{E}$ with $\tilde{e}_0 \cap \bar{V} = e_0 \cap \bar{V}$. Note that if $e_0 \in \bar{E}$, then we have $\tilde{e}_0 = e_0$. Similarly, for each $k \in T$, denote by $\tilde{e}_k$ the edge of G in $\bar{E}$ with $\tilde{e}_k \cap \bar{V} = e_k \cap \bar{V}$. As we detailed before, by projecting out all variables z_v, $v \in V \setminus \bar{V}$, and z_e, $e \in E \setminus \bar{E}$, from inequality (13), we obtain

$$\sum_{v \in (\tilde{e}_0 \setminus \cup_{k \in T} \tilde{e}_k) \cap \bar{V}} z_v + \sum_{k \in T} z_{\tilde{e}_k} - z_{\tilde{e}_0} \leq |(\tilde{e}_0 \setminus \cup_{k \in T} \tilde{e}_k) \cap \bar{V}| + |T| - 1. \tag{15}$$

Now define $e'_0 = \tilde{e}_0 \cap \bar{V}$ and $e'_k = \tilde{e}_k \cap \bar{V}$ for all $k \in T$. Since $\tilde{e}_0 \in \bar{E}$ and $\tilde{e}_k \in \bar{E}$ for all $k \in T$, it follows that $e'_0 \in E(G_{\bar{V}})$ and $e'_k \in E(G_{\bar{V}})$ for all $k \in T$. Moreover, since $\tilde{e}_0 \cap \bar{V} = e_0 \cap \bar{V}$, $\tilde{e}_k \cap \bar{V} = e_k \cap \bar{V}$ for all $k \in T$, $|e_0 \cap e_k \cap \bar{V}| \geq 2$ for all $k \in T$, and $e_i \cap e_j = \emptyset$ for all $i, j \in T$, we have $|e'_0 \cap e'_k| \geq 2$ for all $k \in T$ and $e'_i \cap e'_j = \emptyset$ for all $i, j \in T$. Hence, the hypergraph $\tilde{G}$ with $E(\tilde{G}) = \cup_{k \in T} e'_k \cup e'_0$ is the support hypergraph of a flower inequality of the form (15). This implies that all nonredundant inequalities present in $\mathrm{proj}_{G_{\bar{V}}}(\mathrm{MP}^F_G \cap L_{\bar{V}})$ are also present in $\mathrm{MP}^F_{G_{\bar{V}}}$ and this completes the proof. ∎

Recall that in Theorem 7, in order to show that for certain hypergraphs G, we have $\mathrm{MP}_G \subset \mathrm{MP}^{\mathrm{LP}}_G$, by Lemma 6, we proved that $\mathrm{MP}_{G_{\bar{V}}} \subset \mathrm{MP}^{\mathrm{LP}}_{G_{\bar{V}}}$ for some $\bar{V} \subset V(G)$. The above lemma enables us to utilize a similar technique to prove that $\mathrm{MP}_G \subset \mathrm{MP}^F_G$ for some hypergraph G. That is, since by part (i) of Lemma 6, we have $\mathrm{MP}_{G_{\bar{V}}} = \mathrm{proj}_{\bar{V}}(\mathrm{MP}_G \cap L_{\bar{V}})$ and by Lemma 13, we have $\mathrm{MP}^F_{G_{\bar{V}}} \subseteq \mathrm{proj}_{\bar{V}}(\mathrm{MP}^F_G \cap L_{\bar{V}})$, it follows that if $\mathrm{MP}_{G_{\bar{V}}} \subset \mathrm{MP}^F_{G_{\bar{V}}}$ for some $\bar{V} \subset V$, then $\mathrm{proj}_{\bar{V}}(\mathrm{MP}_G \cap L_{\bar{V}}) \subset \mathrm{proj}_{\bar{V}}(\mathrm{MP}^F_G \cap L_{\bar{V}})$, which in turn implies that $\mathrm{MP}_G \subset \mathrm{MP}^F_G$. We are now ready to prove our main result.

THEOREM 14. *$MP_G = MP^F_G$ if and only if G is a γ-acyclic hypergraph.*

Proof. "⇒" First, we show that if G is not γ-acyclic, then we have $\mathrm{MP}_G \subset \mathrm{MP}^F_G$. To do so, we make use of Proposition 4. To obtain a contradiction, first suppose that G violates condition (ii) in Proposition 4. That is, suppose that there exist nodes $v_1, v_2, v_3 \in V(G)$ such that $\{\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_2, v_3\}\} \subseteq \{e \cap \{v_1, v_2, v_3\} : e \in E(G)\}$. Let $\bar{V} := \{v_1, v_2, v_3\}$, and denote by $e_{12} := \{v_1, v_2\}$, $e_{13} := \{v_1, v_3\}$, $e_{23} := \{v_2, v_3\}$, $e_{123} := \{v_1, v_2, v_3\}$. By part (i) of Lemma 6 and Lemma 13, to prove $\mathrm{MP}_G \subset \mathrm{MP}^F_G$, it suffices to show that $\mathrm{MP}_{G_{\bar{V}}} \subset \mathrm{MP}^F_{G_{\bar{V}}}$. Note that $E(G_{\bar{V}}) = \{e_{12}, e_{13}, e_{123}\}$ if $e_{23} \notin \{e \cap \{v_1, v_2, v_3\} : e \in E(G)\}$ and $E(G_{\bar{V}}) = \{e_{12}, e_{13}, e_{23}, e_{123}\}$, otherwise. It is simple to verify that the inequality $-z_{v_1} + z_{e_{12}} + z_{e_{13}} - z_{e_{123}} \leq 0$ defines a facet of $\mathrm{MP}_{G_{\bar{V}}}$. However, this inequality is not implied by the inequalities in $\mathrm{MP}^F_{G_{\bar{V}}}$, as its support hypergraph corresponds to a Berge-cycle of length three, while the support hypergraph of all inequalities in $\mathrm{MP}^F_{G_{\bar{V}}}$ correspond to a single edge, or a Berge-cycle of length two. Hence, if condition (ii) in Proposition 4 is violated, we have $\mathrm{MP}_G \subset \mathrm{MP}^F_G$.

Next, suppose that condition (ii) in Proposition 4 is satisfied but G contains at least one β-cycle. Denote by C a β-cycle of minimum length. We claim that the subhypergraph $G_{V(C)}$ is a graph that consists of a chordless cycle of length at least three. First note that by Definition 3, the set $\tilde{E} := \{e \cap V(C) : e \in E(C)\}$ is the edge set of a chordless cycle in $G_{V(C)}$. We would like to show that $E(G_{V(C)}) = \tilde{E}$. Observe that $\tilde{E} \subseteq E(G_{V(C)})$. To obtain a contradiction, assume that $\tilde{E} \subset E(G_{V(C)})$. Since by assumption C is a β-cycle of minimum length, it follows that $E(G_{V(C)})$ has an edge $\bar{e}$ with $|\bar{e}| \geq 3$. Denote by $\tilde{e}$ an edge of G with $\bar{e} = \tilde{e} \cap V(C)$. Two cases arise:

Case 1. If $\{v_1, v_2, v_3\} \subseteq \bar{e}$, where v_1, v_2, and v_3 are consecutive nodes in C, it follows that $\{\{v_1, v_2\}, \{v_1, v_3\}, \{v_1, v_2, v_3\}\} \subseteq \{e \cap \{v_1, v_2, v_3\} : e \in E(G)\}$, which contradicts the assumption that condition (ii) in Proposition 4 is satisfied.

Case 2. Suppose that $\bar{e}$ does not contain three consecutive nodes in C. Let the β-cycle C be given by the sequence $v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$. Suppose that $v_1 \in \bar{e}$ and $v_2 \notin \bar{e}$. Note that this assumption is without loss of generality, since $\bar{e}$ does not contain three consecutive nodes of C. Let v_k be the next node of $V(C)$ after the first node v_1 for which $v_k \in \bar{e}$. Clearly, $k \geq 3$ since by assumption $v_2 \notin \bar{e}$. In addition $k < t$, since by assumption $\bar{e}$ contains at least three nodes of C. Finally, by construction we have $\bar{e} \cap \{v_1, \ldots, v_k\} = \{v_1, v_k\}$. It then follows that the sequence $v_1, e_1, v_2, \ldots, e_{k-1}, v_k, \tilde{e}, v_1$ is a β-cycle of length k, where $k < t$. However, this contradicts the assumption that C is β-cycle of minimum length.

Hence, we conclude that $G_{V(C)}$ is a graph that consists of a chordless cycle. To show that $\mathrm{MP}_G \subset \mathrm{MP}^F_G$, by part (i) of Lemma 6 and Lemma 13, it is sufficient to prove that $\mathrm{MP}_{G_{V(C)}} \subset \mathrm{MP}^F_{G_{V(C)}}$. The latter inclusion is valid as the odd-cycle inequalities are facet-defining for $\mathrm{MP}_{G_{V(C)}}$ [25] and are clearly not present in $\mathrm{MP}^F_{G_{V(C)}}$. Consequently, if the hypergraph G contains a γ-cycle, we have $\mathrm{MP}_G \subset \mathrm{MP}^F_G$.

"$\Leftarrow$" Conversely, let G be a γ-acyclic hypergraph. We show that $\mathrm{MP}_G = \mathrm{MP}^F_G$. In the following, we say that an edge of a hypergraph G is *maximal* if it is not contained in any other edge of G. The proof is by induction on the number of maximal edges of G. First, consider the base case; that is, suppose that G has one maximal edge $e' = V(G)$. In this case, by Proposition 8, we conclude that G is a laminar hypergraph. Hence, by Corollary 11, we have $\mathrm{MP}_G = \mathrm{MP}^F_G$. We now proceed to the inductive step; namely, we assume that $\mathrm{MP}_G = \mathrm{MP}^F_G$, for any γ-acyclic hypergraph G with κ maximal edges. We would like to show that the same statement holds if G is a γ-acyclic hypergraph with $\kappa + 1$ maximal edges.

Lifting and decomposition. Consider a maximal edge e' of G, and define E' to be the set of edges contained in e', and $\bar{V} := e' \cap (\cup_{e \in E \setminus E'} e)$. Clearly, $E \setminus E' \neq \emptyset$, as by assumption G contains at least two maximal edges. We say that e' is a *leaf* of G if $\bar{V} \subset \tilde{e}$ for some, $\tilde{e} \in E \setminus E'$. We claim that G contains a leaf e'. To obtain a contradiction, suppose that G does not contain any leaves. It then follows that for every maximal edge e', and every maximal edge e'' adjacent to e', there exists another maximal edge adjacent to e', say, e''', such that neither of the two sets $e' \cap e''$ and $e' \cap e'''$ is a subset of another. From Proposition 8, it follows that the sets $e' \cap e''$ and $e' \cap e'''$ are disjoint. We now show that G contains a β-cycle, which violates property (i) of Proposition 4. Let e_1 denote a maximal edge of G. Denote by e_2 a maximal edge of G adjacent to e_1 and let e_3 denote a maximal edge of G adjacent to e_2 such that $e_2 \cap e_3$ is disjoint from $e_1 \cap e_2$. Recursively, let e_i be a maximal edge of G adjacent to e_{i-1} such that $e_{i-1} \cap e_i$ is disjoint from $e_{i-2} \cap e_{i-1}$. Eventually, there exists an index i such that e_i intersects some e_j for $j \leq i - 1$. Let t be the first such index, and let $s \leq t - 2$ be the largest index such that e_s intersects e_t. Now let v_s be a node in

$e_s \cap e_t$, and, for every $i = s+1, \ldots, t$, let v_i be a node in $e_{i-1} \cap e_i$. Then the sequence $v_s, e_s, v_{s+1}, e_{s+1}, \ldots, v_t, e_t, v_s$ is a β-cycle of length $t - s + 1 \geq 3$.

Now, let e' be a leaf of G and, as before, let $\bar{V} := e' \cap (\cup_{e \in E \setminus E'} e)$. We define G^+ as the hypergraph obtained by adding the edge $\bar{V}$ to G if $\bar{V} \notin V \cup E$, and $G^+ := G$, otherwise. Subsequently, we define G_1 as the section hypergraph of G^+ induced by e', and G_2 as the section hypergraph of G^+ induced by $\cup_{e \in E \setminus E(G_1)} e$. Clearly, both G_1 and G_2 are different from G^+. In addition, we have $G_1 \cup G_2 = G^+$ and $G_1 \cap G_2 = \bar{V}$. By Proposition 8, the subhypergraph $G_{e'}$ of G is laminar. Moreover, the hypergraph G_1 is a partial hypergraph of $G_{e'}$, and thus G_1 is laminar as well. As G_1 contains the edge $\bar{V}$, this implies that every edge e'' of G_1 containing nodes in $V(G_1) \setminus V(G_2)$ satisfies either $e'' \supset \bar{V}$ or $e'' \cap \bar{V} = \emptyset$. Thus all assumptions of Theorem 5 are satisfied and the set $\mathcal{S}_{G^+}$ is decomposable into $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$.

Since G_1 is laminar, by Corollary 11 we have $\mathrm{MP}_{G_1} = \mathrm{MP}^F_{G_1}$. Now, consider the hypergraph G_2. This hypergraph has κ maximal edges which are the κ maximal edges of G that are different from e'. In addition, the hypergraph G_2 is γ-acyclic. To see this, suppose that G_2 contains a γ-cycle C. Then $\bar{V}$ must be an edge of G_2 and $E(C)$ must contain the edge $\bar{V}$, as otherwise C is a γ-cycle of G as well. Since $e' \cap V(G_2) = \bar{V}$, it follows that by replacing $\bar{V}$ with e' in C, we obtain a γ-cycle of G, which is in contradiction with the assumption that G is γ-acyclic. Therefore, by the induction hypothesis we have $\mathrm{MP}_{G_2} = \mathrm{MP}^F_{G_2}$, which together with $\mathrm{MP}_{G_1} = \mathrm{MP}^F_{G_1}$ and the decomposability of $\mathcal{S}_{G^+}$ into $\mathcal{S}_{G_1}$ and $\mathcal{S}_{G_2}$ implies $\mathrm{MP}_{G^+} = \mathrm{MP}^F_{G^+}$.

If $G = G^+$, that is, if $\bar{V} \in V(G) \cup E(G)$, we obtain $\mathrm{MP}_G = \mathrm{MP}^F_G$ and this completes the proof. Henceforth, we assume that $\bar{V} \notin V(G) \cup E(G)$. To obtain MP_G, it suffices to project out the auxiliary variable $z_{\bar{V}}$ from the facet-description of MP_{G^+}. In the following, we perform this projection using Fourier–Motzkin elimination.

Projection. First consider an inequality in the description $\mathrm{MP}^F_{G^+}$ that does not contain $z_{\bar{V}}$. Clearly, the support hypergraph of such an inequality is a partial hypergraph of G. Thus, by Lemma 12, this inequality is also present in the description MP^F_G. Thus to complete the proof, we need to show that by projecting out $z_{\bar{V}}$ from the remaining inequalities of MP_{G^+}, we obtain valid inequalities for MP^F_G.

First, consider MP_{G_1}; denote by $\bar{e}$ the edge of G_1 containing $\bar{V}$ such that there exists no other edge $e \in E(G_1)$ with $e \supset \bar{V}$ and $e \subset \bar{e}$. Note that the edge $\bar{e}$ is well-defined by the laminarity of G_1. Then, by Theorem 10, the auxiliary variable $z_{\bar{V}}$ appears in the following inequalities, which we will refer to as system (I) in the rest of the proof:

$$-z_p + z_{\bar{V}} \leq 0 \qquad \forall p \in I(\bar{V}), \tag{16}$$

$$-z_{\bar{V}} + z_{\bar{e}} \leq 0, \tag{17}$$

$$\sum_{p \in I(\bar{V})} z_p - z_{\bar{V}} \leq |I(\bar{V})| - 1, \tag{18}$$

$$\sum_{p \in I(\bar{e})} z_p - z_{\bar{e}} \leq |I(\bar{e})| - 1. \tag{19}$$

Note that by definition of $\bar{e}$ we have $\bar{V} \in I(\bar{e})$.

Now consider the polytope $\mathrm{MP}_{G_2} = \mathrm{MP}^F_{G_2}$. Let $\bar{E}$ contain all edges of G_2 that are adjacent to $\bar{V}$ and let $\tilde{\mathcal{E}}$ be the set containing all subsets $\tilde{E}$ of $\bar{E}$ with $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in \tilde{E}$. Observe that $\tilde{\mathcal{E}}$ contains the empty set. For each $\hat{e} \in \bar{E}$, let $\mathcal{U}_{\hat{e}}$ be the set containing all subsets of adjacent edges to $\hat{e}$ denoted by $U_{\hat{e}}$ such that $\bar{V} \in U_{\hat{e}}$

and $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in U_{\hat{e}}$. Then, the inequalities in the description of $\mathrm{MP}^F_{G_2}$ containing the auxiliary variable $z_{\bar{V}}$ are the following:

$$-z_p + z_{\bar{V}} \leq 0 \qquad \forall p \in I(\bar{V}), \tag{20}$$

$$\sum_{v \in \bar{V} \setminus \cup_{e \in \tilde{E}} e} z_v + \sum_{e \in \tilde{E}} z_e - z_{\bar{V}} \leq |\bar{V} \setminus \cup_{e \in \tilde{E}} e| + |\tilde{E}| - 1 \qquad \forall \tilde{E} \in \tilde{\mathcal{E}}, \tag{21}$$

$$\sum_{v \in \hat{e} \setminus \cup_{e \in U_{\hat{e}}} e} z_v + \sum_{e \in U_{\hat{e}}} z_e - z_{\hat{e}} \leq |\hat{e} \setminus \cup_{e \in U_{\hat{e}}} e| + |U_{\hat{e}}| - 1 \qquad \forall \hat{e} \in \bar{E}, \forall U_{\hat{e}} \in \mathcal{U}_{\hat{e}}. \tag{22}$$

We should remark that inequalities (21) are flower inequalities provided that $\tilde{E} \neq \emptyset$ and amount to the inequality $\sum_{v \in \bar{V}} z_v - z_{\bar{V}} \leq |\bar{V}| - 1$ present in the standard linearization of $\mathcal{S}_{G_2}$, otherwise. In the remainder of the proof, we will refer to the inequalities (20)–(22) as system (II).

Now consider the system of linear inequalities (I)–(II). We eliminate $z_{\bar{V}}$ from this system using Fourier–Motzkin elimination. First consider the case where we select two inequalities from system (I). Denote by G'_1 the hypergraph obtained by removing the edge $\bar{V}$ from G_1. It then follows that the inequality $az \leq \alpha$ obtained as a result of such projection is valid for the multilinear polytope $\mathrm{MP}_{G'_1}$. Since G'_1 is a laminar hypergraph, by Corollary 11, we have $\mathrm{MP}_{G'_1} = \mathrm{MP}^F_{G'_1}$. Finally, since G'_1 is a partial hypergraph of G, by Lemma 12, $az \leq \alpha$ is a valid inequality for MP^F_G. Similarly, we can show that by projecting out $z_{\bar{V}}$ from two inequalities present in system (II), we obtain an inequality that is valid for MP^F_G . This is due to the fact that the hypergraph G'_2 obtained by removing $\bar{V}$ from G_2 is a γ-acyclic hypergraph with κ maximal edges for which by the induction hypothesis we have $\mathrm{MP}_{G'_2} = \mathrm{MP}^F_{G'_2}$. Note that G'_2 is γ-acyclic as it is a partial hypergraph of the γ-acyclic hypergraph G_2. Hence, it suffices to show that the remaining inequalities obtained by projecting out $z_{\bar{V}}$ are valid for MP^F_G as well. Therefore, it suffices to examine inequalities obtained by projecting out $z_{\bar{V}}$ starting from two inequalities one of which is only present in system (I) while the other one is only present in system (II).

We start by selecting one inequality in (16) from system (I). Clearly, this inequality is identical to inequality (20) present in system (II). Hence, by the above discussion, we do not need to consider inequalities (16). Next, consider inequality (17) from system (I). Since the coefficient of $z_{\bar{V}}$ in (17) is negative, it suffices to consider inequalities (20) and (22) from system (II). In addition, we do not need to consider (20) since it is already present system (I). By summing inequalities (17) and (22), we obtain

$$\sum_{v \in \hat{e} \setminus \cup_{e \in U_{\hat{e}}} e} z_v + \sum_{e \in U_{\hat{e}} \setminus \bar{V}} z_e + z_{\bar{e}} - z_{\hat{e}} \leq |\hat{e} \setminus \cup_{e \in U_{\hat{e}}} e| + |U_{\hat{e}}| - 1 \;\; \forall \hat{e} \in \bar{E},\ \forall U_{\hat{e}} \in \mathcal{U}_{\hat{e}}.$$

To show that the above system represents a system of flower inequalities for MP_G, it suffices to show that the set $(U_{\hat{e}} \setminus \bar{V}) \cup \bar{e}$ satisfies two properties: (i) all edges in $(U_{\hat{e}} \setminus \bar{V}) \cup \bar{e}$ are adjacent to $\hat{e}$ and (ii) $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in (U_{\hat{e}} \setminus \bar{V}) \cup \bar{e}$. By construction, all edges in $U_{\hat{e}}$ are adjacent to $\hat{e}$, and $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in U_{\hat{e}}$. It in addition, we have $\hat{e} \cap \bar{V} = \hat{e} \cap \bar{e}$ for all $\hat{e} \in \bar{E}$. It then follows that for each $\hat{e} \in \bar{E}$ the above system is contained in the system of flower inequalities for MP_G, centered at $\hat{e}$.

Next, we select inequalities (18) from system (I). Define a partition of $I(\bar{V}) = I_v(\bar{V}) \cup I_e(\bar{V})$, where $I_v(\cdot)$ and $I_e(\cdot)$ contain the nodes and edges of $I(\cdot)$, respectively. It then follows that $I_e(\bar{V}) \in \tilde{\mathcal{E}}$, as the section hypergraph of G induced by $\bar{V}$ is laminar. Consequently, inequalities (18) are implied by inequalities (21), which in turn imply that we do not need to consider these inequalities and proceed with inequalities (19) from system (I). Since the coefficient of $z_{\bar{V}}$ in (19) is positive, it suffices to consider inequalities (21) from system (II). By summing inequalities (19) and (21), we get

$$\text{(23)} \quad \sum_{p \in I(\bar{e}) \setminus \bar{V}} z_p + \sum_{v \in \bar{V} \setminus \cup_{e \in \tilde{E}} e} z_v + \sum_{e \in \tilde{E}} z_e - z_{\bar{e}} \leq |\bar{V} \setminus \cup_{e \in \tilde{E}}| + |\tilde{E}| + |I(\bar{e})| - 2 \quad \forall \tilde{E} \in \tilde{\mathcal{E}}.$$

As before, define a partition of $I(\bar{e}) = I_v(\bar{e}) \cup I_e(\bar{e})$. Consider the set of edges defined as $\tilde{E}' = \tilde{E} \cup (I_e(\bar{e}) \setminus \bar{V})$. Clearly, all edges in $\tilde{E}'$ are adjacent to $\bar{e}$ as $\tilde{E}$ represents a set of edges adjacent to $\bar{V}$ and by definition all edges in $I_e(\bar{e})$ are contained in $\bar{e}$. Also, we have $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in \tilde{E}'$ since (i) G_1 is a laminar hypergraph, which implies $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in I(\bar{e})$, and in particular $e_i \cap \bar{V} = \emptyset$ for all $e_i \in I(\bar{e}) \setminus \bar{V}$, (ii) by definition $e_i \cap e_j = \emptyset$ for all $e_i, e_j \in \tilde{E}$, and (iii) by definition $e_i \cap \bar{e} \subseteq \bar{V}$ for all $e_i \in \tilde{E}$. It is simple to check that $\bar{e} \setminus \cup_{e \in \tilde{E}'} e = (\bar{V} \setminus \cup_{e \in \tilde{E}} e) \cup I_v(\bar{e})$. Moreover, we have $|\tilde{E}'| = |\tilde{E}| + |I_e(\bar{e})| - 1$. Define $\tilde{\mathcal{E}}' = \{\tilde{E} \cup (I_e(\bar{e}) \setminus \bar{V}) : \tilde{E} \in \tilde{\mathcal{E}}\}$. Hence, inequality (23) can be equivalently written as

$$\sum_{v \in \bar{e} \setminus \cup_{e \in \tilde{E}'} e} z_v + \sum_{e \in \tilde{E}'} z_e - z_{\bar{e}} \leq |\bar{e} \setminus \cup_{e \in \tilde{E}'} e| + |\tilde{E}'| - 1 \quad \forall \tilde{E}' \in \tilde{\mathcal{E}}'.$$

Now it is simple to verify that for each $\tilde{E}' \in \tilde{\mathcal{E}}'$, the above inequality is a flower inequality for MP_G centered at $\bar{e}$ with the neighbors $e \in \tilde{E}'$. Hence, we have shown that all inequalities obtained by projecting out $z_{\bar{V}}$ from the facet description of MP_{G^+} are implied MP_G^F. It then follows that $\mathrm{MP}_G = \mathrm{MP}_G^F$ and this completes the proof. □

5. Separation of flower inequalities. The following example demonstrates that, even for γ-acyclic hypergraphs, the number of facets of MP_G^F may not be bounded by a polynomial in $|V(G)|, |E(G)|$.

Example 2. Consider the γ-acyclic hypergraph G with $E(G) = \{e_0, e_1, \ldots, e_m\}$, such that $e_j \cap e_{j'} = \emptyset$ for all $j, j' \in J = \{1, \ldots, m\}$, $|e_0 \cap e_j| \geq 2$ and, $e_j \setminus e_0 \neq \emptyset$ for all $j \in J$. In this example, the number of flower inequalities present in MP_G^F grows exponentially with the number of edges of G; to see this, note that we can write $2^m - 1$ flower inequalities centered at e_0, while there exists exactly one flower inequality centered at each e_j, $j \in J$. Hence, the total number of flower inequalities in MP_G^F is $2^m + m - 1$. We show that for this example, all flower inequalities centered at e_0 are facet-defining for MP_G, implying that this polytope has exponentially many facets. By (12), any flower inequality centered at e_0 can be written as

$$\text{(24)} \quad \sum_{v \in e_0 \setminus \cup_{j \in T} e_j} z_v + \sum_{j \in T} z_{e_j} - z_{e_0} \leq |e_0 \setminus \cup_{j \in T} e_j| + |T| - 1,$$

where T denotes a nonempty subset of J. We start by characterizing the sets of points in $\mathcal{S}_G$ that satisfy the above inequality tightly. Subsequently, we show that any nontrivial valid inequality $az \leq \alpha$ for $\mathcal{S}_G$ that is satisfied tightly at all such points coincides with (24) up to a positive scaling. Since MP_G is full dimensional [16], this in turn implies that inequality (24) is facet-defining for MP_G. It is simple to verify that inequality (24) is satisfied tightly by the following sets of points in $\mathcal{S}_G$:

(i) any point $z \in \mathcal{S}_G$ with $z_{e_0} = 1$ and $z_{e_j} = 1$ for all $j \in T$;
(ii) any point $z \in \mathcal{S}_G$ with $z_v = 1$ for all $v \in e_0 \setminus \{v'\}$ and $z_{v'} = 0$, where $v' \in e_0 \setminus \cup_{j \in T} e_j$ and $z_{e_j} = 1$ for all $j \in T$;
(iii) any point $z \in \mathcal{S}_G$ with $z_v = 1$ for all $v \in (\cup_{j \in J} e_j \cup e_0) \setminus e_{j''}$ for some $j'' \in T$ and $z_v = 0$ for all $v \in V'' \subseteq e_0 \cap e_{j''}$ with $V'' \neq \emptyset$.

Consider the case where $J \setminus T \neq \emptyset$ and construct a tight point of type (ii) defined above with $v' \in e_0 \cap e_{j'}$ for some $j' \in J \setminus T$ and $z_v = 0$ for all $v \in (\cup_{j \in J \setminus T} e_j) \setminus e_0$. Substituting this point in $az \leq \alpha$ gives

$$\sum_{v \in e_0 \setminus (\cup_{j \in T} e_j \cup \{v'\})} a_v z_v + \sum_{j \in T} a_{e_j} z_{e_j} = \alpha. \tag{25}$$

Now consider another tight point of type (ii) obtained by letting $z_{\tilde{v}} = 1$ for some $\tilde{v} \in e_{j'} \setminus e_0$ in the tight point defined above. Note that if $J \setminus T \neq \emptyset$, then a node of the form $\tilde{v}$ always exists since by assumption $e_j \setminus e_0 \neq \emptyset$ for all $j \in J$. Substituting this point in $az \leq \alpha$ yields

$$\sum_{v \in e_0 \setminus (\cup_{j \in T} e_j \cup \{v'\})} a_v z_v + a_{\tilde{v}} z_{\tilde{v}} + \sum_{j \in T} a_{e_j} z_{e_j} = \alpha. \tag{26}$$

From (25) and (26) it follows that

$$a_v = 0 \quad \forall v \in e_j \setminus e_0,\ \forall j \in J \setminus T. \tag{27}$$

Construct a tight point of type (i) with $z_v = 0$ for all $v \in e_j \setminus e_0$, for all $j \in J \setminus T$. Subsequently, construct a new tight point of type (i) by letting $z_{e'_j} = 1$ for some $j' \in J \setminus T$ in the previous point. Substituting these points in $az = \alpha$ and using (27), we obtain

$$a_{e_j} = 0 \quad \forall j \in J \setminus T. \tag{28}$$

Next consider a point in $\mathcal{S}_G$ of type (iii) defined above with $z_v = 0$ for all $v \in e_{j''}$, where $j'' \in T$. Clearly in this case we have $V'' = e_0 \cap e_{j''}$. Subsequently, construct a second tight point by letting $z_{\bar{v}} = 1$ for some $\bar{v} \in e_{j''}$ in the previous tight point. Note that the second point is also a tight point of type (iii) for any $\bar{v} \in e_{j''}$, since by assumption $|e_0 \cap e_{j''}| \geq 2$, which in turn implies $V'' \setminus \{\bar{v}\} \neq \emptyset$ for all $\bar{v} \in e_{j''}$. Substituting these two points in $az = \alpha$ and subtracting the resulting relations, we obtain

$$a_v = 0 \quad \forall v \in e_j,\ \forall j \in T. \tag{29}$$

Now, consider a tight point of type (i) with $z_v = 0$ for all $v \in e_j \setminus e_0$, $j \in J \setminus T$ and construct a tight point of type (iii) by letting $z_v = 0$ for all $v \in V'' \subseteq e_0 \cap e_{j''}$ for some $j'' \in T$ in the first point. Substituting these points in $az = \alpha$ and using (29), we obtain

$$a_{e_j} + a_{e_0} = 0 \quad \forall j \in T. \tag{30}$$

Consider the case $e_0 \setminus \cup_{j \in T} e_j \neq \emptyset$. Construct a tight point of type (i) defined above with $z_v = 0$ for all $v \in e_j \setminus e_0$, $j \in J \setminus T$. Now construct a new point by letting

$z_{v'} = 0$ for some $v' \in e_0 \setminus \cup_{j\in T} e_j$ in the first point. Clearly, the second point is a tight point of type (ii). Substituting the two points in $az = \alpha$ and subtracting the resulting equalities, we obtain

$$a_v + a_{e_0} = 0 \quad \forall v \in e_0 \setminus \cup_{j\in T} e_j. \tag{31}$$

From (27), (28), (29), (30), and (31), it follows that the inequality $az \leq \alpha$, up to a positive scaling, can be equivalently written as

$$\sum_{v\in e_0 \setminus \cup_{j\in T} e_j} z_v + \sum_{j\in T} z_{e_j} - z_{e_0} \leq \alpha.$$

Moreover by substituting a tight point of type (i) in this inequality we obtain $\alpha = |e_0 \setminus \cup_{j\in T} e_j| + |T| - 1$. Hence, $az \leq \alpha$ coincides with inequality (24) up to a positive scaling, implying that (24) defines a facet of MP_G for any nonempty $T \subseteq J$. We have shown that all flower inequalities centered at e_0 are facet-defining for MP_G. Since there are a total number of $2^m - 1$ such inequalities present in MP^F_G, we conclude that for a γ-acyclic hypergraph G, the polytope MP^F_G may have exponentially many facets.

5.1. Separation problem. We start by defining the separation problem for flower inequalities as follows. (See [30] more details.)

DEFINITION 15. *Given a hypergraph G and a vector $\bar{z} \in \mathbb{R}^{V+E}$, decide whether $\bar{z}$ satisfies all flower inequalities or not, and in the latter case, find a flower inequality that is violated by $\bar{z}$.*

Given a γ-acyclic hypergraph G, we are interested in solving the separation problem over all flower inequalities in strongly polynomial time, i.e., in a number of iterations bounded by a polynomial in $|V|$ and $|E|$. This in turn implies that the optimization problem (MO) is polynomially solvable over γ-acyclic hypergraphs.

We show that the separation problem over all flower inequalities centered at a given edge of a γ-acyclic hypergraph can be equivalently stated as a minimum-weight perfect matching problem over a related laminar hypergraph. Subsequently, we present a strongly polynomial-time combinatorial algorithm to solve this matching problem. Recall that a *matching* in a hypergraph is a set of edges M with the property that $e \cap f = \emptyset$ for all $e, f \in M$ with $e \neq f$. A matching is called *perfect* if each node is contained in exactly one edge of the matching. Finding a minimum-weight perfect matching in a general hypergraph is $\mathcal{NP}$-hard [21]. However, for balanced hypergraphs, this problem can be solved in polynomial time by solving a linear optimization problem [10]. A hypergraph is said to be *balanced* if every Berge-cycle of odd length has an edge containing three vertices of the cycle; that is, a hypergraph is balanced if and only if it does not contain any β-cycle of odd length. As laminar hypergraphs are balanced, this result in particular implies that finding a minimum-weight perfect matching in a laminar hypergraph can be done in polynomial time. Consequently, the separation problem over flower inequalities for γ-acyclic hypergraphs can be done in polynomial time. In order to attain a strongly polynomial-time separation algorithm, in the following, we present a strongly polynomial-time combinatorial algorithm to solve the matching subproblems.

THEOREM 16. *Given a γ-acyclic hypergraph $G = (V, E)$ and a vector $\bar{z} \in \mathbb{R}^{V+E}$, there exists a strongly polynomial-time algorithm that solves the separation problem over all flower inequalities.*

Proof. We show how to solve the separation problem over the flower inequalities centered at an edge e_0 of G. By applying the algorithm $|E|$ times, we can then solve the separation problem over all the flower inequalities.

Let e_k, $k \in K$, be the set of all edges adjacent to e_0 with $|e_0 \cap e_k| \geq 2$ for all $k \in K$. There exists a flower inequality violated by the vector $\bar{z}$ if and only if there exists a nonempty subset T of K with $e_i \cap e_j = \emptyset$ for all $i, j \in T$ with $i \neq j$, such that

$$\sum_{v \in e_0 \setminus \cup_{k \in T} e_k} \bar{z}_v + \sum_{k \in T} \bar{z}_{e_k} - \bar{z}_{e_0} > |e_0 \setminus \cup_{k \in T} e_k| + |T| - 1,$$

or equivalently

$$\sum_{v \in e_0 \setminus \cup_{k \in T} e_k} (1 - \bar{z}_v) + \sum_{k \in T} (1 - \bar{z}_{e_k}) < 1 - \bar{z}_{e_0}. \tag{32}$$

Since the right-hand side of inequality (32) does not depend on T, it suffices to show how to minimize its left-hand side over all possible sets T. More precisely, if the minimum of the left-hand side of inequality (32) is greater than or equal to $1 - \bar{z}_{e_0}$, then the vector $\bar{z}$ satisfies all flower inequalities centered at e_0. Otherwise, any subset T realizing the minimum value yields a flower inequality violated by $\bar{z}$.

Let $\bar{V} := e_0$, $\bar{L} := \{\{v\} : v \in \bar{V}\}$, $\bar{E} := \bar{L} \cup \{e' \cap e_0 : e' \in E \setminus \{e_0\},\ |e' \cap e_0| \geq 2\}$, and define the hypergraph $\bar{G} := (\bar{V}, \bar{E})$. By Proposition 8, the hypergraph $\bar{G}$ is laminar. Note that unlike G, the hypergraph $\bar{G}$ has *loops*, i.e., edges containing only one node. We associate a *weight* to each loop $\{v\} \in \bar{L}$, defined as $w_{\{v\}} := 1 - \bar{z}_v$. For every edge $e \in \bar{E} \setminus \bar{L}$, there may exist several edges $e' \in E$ satisfying $e = e' \cap e_0$. We denote by $e'(e)$ an edge that maximizes $\bar{z}_{e'(e)}$. We associate a weight to each edge $e \in \bar{E} \setminus \bar{L}$ defined as $w_e := 1 - \bar{z}_{e'(e)}$.

We now show that the problem of minimizing the left-hand side of (32) over all possible sets T can be solved by finding a perfect matching M of $\bar{G}$ of minimum weight. Indeed, given a perfect matching M of $\bar{G}$, the set $T := \{e'(e) : e \in M \setminus \bar{L}\}$ yields a left-hand side of (32) whose value equals the weight of the matching. Conversely, given a subset T, the set $M := \{e' \cap e_0 : e' \in T\} \cup \{\{v\} : v \in \bar{V} \setminus (\cup_{e' \in T} e')\}$ is a perfect matching of $\bar{G}$ whose weight is no greater than the value of the left-hand side of (32).

Next, we present a strongly polynomial-time combinatorial algorithm that finds a minimum weight perfect matching of the laminar hypergraph $\bar{G} = (\bar{V}, \bar{E})$. At iteration t of this algorithm, we start with a laminar hypergraph $\bar{G}^t = (\bar{V}, \bar{E}^t)$, which is a partial hypergraph of $\bar{G}$, and with a perfect matching M^t of $\bar{G}^t$ with the additional property that for every edge $e \in M^t$, no other edge $e' \in \bar{E}^t$ is contained in e. If $M^t = \bar{E}^t$, then M^t is a minimum weight perfect matching of $\bar{G}^t$, as $\bar{G}^t$ has no other perfect matching and the algorithm terminates. Otherwise, we construct a laminar partial hypergraph of $\bar{G}^t$ denoted by $\bar{G}^{t+1} = (\bar{V}, \bar{E}^{t+1})$ and a perfect matching M^{t+1} of $\bar{G}^{t+1}$ with the same property with respect to $\bar{G}^{t+1}$; i.e., for every edge $e \in M^{t+1}$, no other edge $e' \in \bar{E}^{t+1}$ is contained in e.

We initialize the algorithm by setting $\bar{G}^0 = \bar{G}$ and by setting M^0 to be the trivial perfect matching of $\bar{G}^0$ that consists of all the loops of $\bar{G}$. By construction, all hypergraphs $\bar{G}^t$ are partial hypergraphs of $\bar{G}$ with the same node set $\bar{V}$. As a result all intermediate perfect matchings M^t of $\bar{G}^t$ correspond to perfect matchings of $\bar{G}$ as well. In addition, as we detail in the following, the proposed algorithm is a greedy algorithm in the sense that the weight of these perfect matchings decreases at every iteration until a minimum weight perfect matching of $\bar{G}$ is found, that is $M^s = \bar{E}^s$ for some $s \geq 0$.

We now describe the tth iteration of the proposed algorithm. We start by selecting a *minimal* edge f of $\bar{G}^t$ that is not in M^t; that is, we select an edge f in $\bar{E}^t \setminus M^t$ that does not contain any other edge in $\bar{E}^t \setminus M^t$. Note that the special property of M^t implies that f contains edges in M^t. Moreover, laminarity of $\bar{G}^t$ implies that the edges $e \in M^t$ with $e \subset f$ partition the nodes in f. We construct the hypergraph $\bar{G}^{t+1}$ and its perfect matching M^{t+1} as follows:

Case A. If $w_f \geq \sum_{e \in M^t : e \subset f} w_e$, we define $\bar{E}^{t+1} := \bar{E}^t \setminus \{f\}$ and $M^{t+1} := M^t$.

Case B. Otherwise, if $w_f < \sum_{e \in M^t : e \subset f} w_e$, we define $\bar{E}^{t+1} := \bar{E}^t \setminus \{e \in M^t : e \subset f\}$ and $M^{t+1} := M^t \setminus \{e \in M^t : e \subset f\} \cup f$.

We now show that M^{t+1} is a perfect matching of $\bar{G}^{t+1}$ with the property that for every edge $e \in M^{t+1}$, no other edge $e' \in \bar{E}^{t+1}$ is contained in e. In Case A, this follows from the fact that $\bar{G}^{t+1}$ is obtained from $\bar{G}$ by removing an edge that is not present in M^t. In Case B, M^{t+1} is obtained from M^t by adding the new edge f, and by removing all edges $e \in M^t$ with $e \subset f$. Since the edges $e \in M^t$ with $e \subset f$ partition the nodes in f, the set M^{t+1} is a perfect matching of $\bar{G}^{t+1}$. Moreover, since f does not contain any other edge in $\bar{E}^{t+1}$, the matching M^{t+1} satisfies the aforementioned property.

In the following, we show that there exists a minimum weight perfect matching of $\bar{G}^t$ that is also a perfect matching of $\bar{G}^{t+1}$. Since every perfect matching of $\bar{G}^{t+1}$ is also a perfect matching of $\bar{G}^t$ with the same weight, the above claim implies that any minimum weight perfect matching of $\bar{G}^t$ for all $t \geq 0$ is also a minimum weight perfect matching of $\bar{G}$. This in turn completes the proof of the correctness of the proposed algorithm as upon termination, this algorithm returns a minimum weight perfect matching of $\bar{G}^s$ for some $s \geq 0$. In Case A defined above, this amounts to showing that $\bar{G}^t$ contains a minimum weight perfect matching that does not include f. Let $\tilde{M}$ be a minimum weight perfect matching of $\bar{G}^t$. If $\tilde{M}$ does not contain f, we are done. Thus, assume that $\tilde{M}$ contains f. Let $\tilde{M}'$ be obtained from $\tilde{M}$ by replacing f with the edges $e \in M^t$ with $e \subset f$. The set $\tilde{M}'$ is a perfect matching of $\bar{G}^t$ that does not contain f, and it is of minimum weight because $w_f \geq \sum_{e \in M^t : e \subset f} w_e$. Therefore, the hypergraph $\bar{G}^t$ has a minimum weight perfect matching that does not contain f. In Case B, we can show that a stronger property is satisfied; that is, each minimum weight perfect matching of $\bar{G}^t$ does not contain any of the edges $e \in M^t$ with $e \subset f$. To obtain a contradiction, assume that $\tilde{M}$ is a minimum weight perfect matching of $\bar{G}^t$ that contains at least one of these edges. This in turn implies that $\tilde{M}$ does not contain f. Since f does not contain any edge in $\bar{E}^t \setminus M^t$, and $\tilde{M}$ is a perfect matching of the laminar hypergraph $\bar{G}^t$, it must contain all the edges $e \in M^t$ with $e \subset f$. Now let $\tilde{M}'$ be obtained from $\tilde{M}$ by replacing all the edges $e \in M^t$ with $e \subset f$ with the edge f. The set $\tilde{M}'$ is a perfect matching of $\bar{G}^t$, and its weight is strictly smaller than that of $\tilde{M}$ because $w_f < \sum_{e \in M^t : e \subset f} w_e$. This contradicts the assumption that $\tilde{M}$ is a minimum weight perfect matching. Hence, no minimum weight perfect matching of G^t contains an edge $e \in M^t$ with $e \subset f$.

Hence, the separation problem over all flower inequalities consists of solving $|E|$ minimum-weight perfect matching problems for laminar hypergraphs. Since at iteration t of the proposed matching algorithm (as described by Case A and Case B), at least one edge is removed from $\bar{E}^t$, we conclude that the algorithm terminates after at most $|V| + |E|$ iterations. It then follows that the separation problem over all flower inequalities can be solved in strongly polynomial time. ∎

We now analyze the computational complexity of the separation algorithm described in the proof of Theorem 16. For brevity, we make use of the notation introduced in this proof without redefining it. In the following, we assume that a hypergraph is represented by an incidence-list in which edges are stored as objects, and every edge stores its incident vertices. In order to use efficient searching algorithms, we assume that the vertex list for each edge is sorted. Otherwise, such a sorted data structure for a rank-r hypergraph can be obtained in $O(r|E|)$ time by using some integer sorting algorithm such as counting sort [11]. In addition, we assume that the edges of E are sorted in increasing cardinality, and edges of the same cardinality are sorted lexicographically. For a rank-r hypergraph, such a sorting order can be obtained using the least significant digit radix sort in $O(r|E|)$ operations (see, e.g., [11]).

PROPOSITION 17. *Given a rank-r γ-acyclic hypergraph $G = (V, E)$, the separation problem over all flower inequalities can be solved in $O(r|E|^2(|V| + |E|))$ operations.*

Proof. Let us first consider the separation problem over all flower inequalities centered at $e_0 \in E$. As we described in the proof of Theorem 16, this problem can be equivalently solved by finding a minimum-weight perfect matching of a laminar hypergraph $\bar{G} = (\bar{V}, \bar{E})$ defined before. We argue that the matching algorithm proposed in the proof of Theorem 16 terminates after at most $|\bar{E} \setminus \bar{L}|$ iterations. To see this, consider an iteration of this algorithm in which we select a minimal edge f of $\bar{G}^t$ that is not in M^t. If the condition in Case A is satisfied, then $\bar{G}^{t+1}$ is obtained by removing f from $\bar{G}^t$. Since all subsequent hypergraphs $\bar{G}^s$, $s > t + 1$, are partial hypergraphs of $\bar{G}^{t+1}$, the edge f will never be selected again. Now, consider Case B; in this case, the edge f is added to M^t and it will not be reselected unless it is removed from M^s for some $s > t$. However, if f is removed from M^s, then it is also removed from $\bar{E}^s$ and hence by the above argument it will not be selected in the subsequent iterations of the proposed algorithm. Recall that all loops $e \in \bar{L}$ are initially present in M^0 and hence by the above argument will not be selected in the following iterations. Again, by the above argument, each edge $f \in \bar{E} \setminus \bar{L}$ is selected at most once throughout the matching algorithm. Since $|\bar{E} \setminus \bar{L}| \leq |E|$ for all $e_0 \in E$, we conclude that each minimum weight perfect matching problem is solved in at most $|E|$ iterations.

Next, we analyze the cost of each iteration in the matching algorithm. The first step is to construct the laminar hypergraph $\bar{G} = (\bar{V}, \bar{E})$. As we detailed before, we represent the hypergraph G by an incidence-list in which edges are stored as objects, every edge stores its incident vertices, and these vertices are sorted. In addition, edges are sorted in increased cardinality and edges of the same cardinality are sorted lexicographically. Thus, to obtain $\bar{G}$ from G, it suffices to construct the set $\{e' \cap e_0 : e' \in E \setminus \{e_0\},\ |e_0 \cap e'| \geq 2\}$. Since the set of vertices contained in each edge $e \in E$ are sorted, for each $e' \in E \setminus \{e_0\}$, we can obtain $e' \cap e_0$ in $O(\max(|e_0|, |e'|))$ time. It then follows that $\bar{E} \setminus \bar{L}$ can be obtained in $O(r|E|)$ time. Subsequently, we sort the edges $\bar{E} \setminus \bar{L}$ similar to those of G in $O(r|E|)$ time. Finally, we append $\bar{L}$ to the adjacency-list in a sorted order, which can be done in $O(|V|)$ operations. For each $e \in \bar{E}$, we compute and store the weight w_e, as defined before, which can be done in $O(|V| + |E|)$ operations. Finally, we remove parallel edges based on the values of these weights in $O(|E|)$ operations. In subsequent iterations of the algorithm, all hypergraphs $\bar{G}^t$ are obtained from $\bar{G}$ by removing certain edges from this data structure. In addition, with each edge $e \in \bar{E}^t$, we associate a label $m(e)$ defined as follows: if $e \in M^t$, then we let $m(e) = 1$; otherwise, we set $m(e) = 0$. We initialize M^0 by letting $m(e) = 1$ for all $e \in \bar{L}$ and $m(e) = 0$ for all $e \in \bar{E} \setminus \bar{L}$, which takes $O(|V| + |E|)$ operations. As described in the proof of Theorem 16, at iteration t of this algorithm, we select a minimal edge $f \in \bar{E}^t \setminus M^t$. Using the aforementioned data structure for $\bar{G}^t$, this can be done in $O(1)$ time by selecting the first edge $f \in \bar{E}^t$ with $m(f) = 0$, in the order the edges are sorted in $\bar{E}^t$. This is due to the fact that at iteration t, edges are sorted in increased cardinality in $\bar{E}^t$, and each time we select a new edge f from $\bar{E}^t$ all edges $e' \subset f$ have been already considered in a previous iteration; that is, all such edges are either added to M^s or are removed from E^s for some $s < t$ and as we described before, none of such edges will be present in $\bar{E}^t \setminus M^t$. Now suppose that we select a minimal edge $f \in \bar{E}^t \setminus M^t$. We need to identify all edges $e' \in \bar{E}^t$ with $e' \subset f$ and $m(e') = 1$. This can be done by scanning all edges $e \in \bar{E}^t$ that are listed before f and for each of them test whether $e \subset f$; the latter can be solved in $O(r)$ operations as the vertices corresponding to each edge are sorted. As a result, we can identify all edges with $e \subset f$ and $e \in M^t$ in $O(r(|V|+|E|))$ operations. Subsequently, we compute $\tilde{w} = \sum_{e \subset f,\, e \in M^t} w_e$ in $O(r)$ time and compare it against w_f. Two cases arise:

Case A. If $w_f \geq \tilde{w}$, then we remove the edge f from $\bar{E}^t$, which can be done in constant time using a proper data structure.

*Cas*e B. Otherwise, if $w_f < \tilde{w}$, we set $m(f) = 1$ and we remove the edges $e \subset f$ and $e \in M^t$ from $\bar{E}^t$, which for a rank-r hypergraph can be done in $O(r)$ operations.

It then follows that the cost of separation problem over all flower inequalities centered at e_0 is $O(r|E|(|V| + |E|))$, which in turn implies the overall cost of solving the separation problem over all flower inequalities for a rank-r γ-acyclic hypergraph G is $O(r|E|^2(|V| + |E|))$. ∎

As we detailed before, the polytope MP_G^{LP} consists of at most $|V| + (r + 2)|E|$ inequalities. By polynomial equivalence of separation and optimization (see, e.g., [30]) and Theorem 14, the following holds.

Corollary 18. *Problem* (MO) *is polynomially solvable over γ-acyclic hypergraphs.*

As we mentioned before, Conforti, Cornuéjols, and Vušković [10] proved that a minimum-weight perfect matching in balanced hypergraphs can be obtained in polynomial time via solving a linear optimization problem. It is simple to show that if G is a balanced hypergraph, then the hypergraph $\bar{G}$ defined in the proof of Theorem 16 is balanced as well. Our proposed separation algorithm over all flower inequalities consists of solving $|E(G)|$ minimum-weight perfect matching problems for hypergraphs of the form $\bar{G}$. Consequently, we have the following result.

Theorem 19. *Given a balanced hypergraph $G = (V, E)$ and a vector $\bar{z} \in \mathbb{R}^{V+E}$, the separation problem over all flower inequalities can be solved in polynomial time, i.e., in a number of iterations bounded by a polynomial in $|V|$, $|E|$, and in the size of the vector $\bar{z}$.*

It can be shown that a naive implementation of the separation problem over flower inequalities for general hypergraphs has a time complexity of $O(r^3|E|^{\lfloor r/2 \rfloor + 1})$. For the multilinear sets that appear in MINLPs, we often have $r \ll \min\{|V|, |E|\}$. In fact, for all practical purposes we can assume that $r < 10$ and therefore it is reasonable to assume that r is a fixed parameter, in which case we conclude that the separation of flower inequalities over general multilinear sets can be done efficiently. In fact, in Example 2, in which the polytope MP_G^F has exponentially many facets, we have

$r \geq 2|E|$. Hence, the proposed flower inequalities can be effectively incorporated in a branch-and-cut framework to construct tighter polyhedral relaxations of general MINLPs containing a collection of multilinear subexpressions.

REFERENCES

[1] F. Al-Khayyal and J. Falk, *Jointly constrained biconvex programming*, Math. Oper. Res., 8 (1983), pp. 273–286.

[2] R. Anstee, *Hypergraphs with no special cycles*, Combinatorica, 3 (1983), pp. 141–146.

[3] X. Bao, A. Khajavirad, N. Sahinidis, and M. Tawarmalani, *Global optimization of nonconvex problems with multilinear intermediates*, Math. Program. Comput., 7 (2015), pp. 1–37.

[4] X. Bao, N. Sahinidis, and M. Tawarmalani, *Multiterm polyhedral relaxations for nonconvex, quadratically-constrained quadratic programs*, Optim. Methods Softw., 24 (2009), pp. 485–504.

[5] F. Barahona and A. Mahjoub, *On the cut polytope*, Math. Program., 36 (1986), pp. 157–173.

[6] C. Berge, *Hypergraphs: Combinatorics of Finite Sets*, North-Holland Math. Libr., Elsevier, New York, 1984.

[7] E. Boros and P. Hammer, *Pseudo-boolean optimization*, Discrete Appl. Math., 123 (2002), pp. 155–225.

[8] J. Brault-Baron, *Hypergraph acyclicity revisited*, ACM Comput. Surv., 49 (2016), pp. 54:1–54:26.

[9] C. Buchheim, Y. Crama, and E. Rodríguez-Heck, *Berge-Cyclic Multilinear* $0-1$ *Optimization Problems*, manuscript, 2016.

[10] M. Conforti, G. Cornuéjols, and K. Vušković, *Balanced matrices*, Discrete Math., 306 (2006), pp. 2411–2437.

[11] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms*, Vol. 6, MIT Press, Cambridge, 2001.

[12] G. Cornuéjols, *Combinatorial Optimization: Packing and Covering*, CBMS-NSF Regional Conf. Ser. in Appl. Math. 74, SIAM, Philadelphia, 2001.

[13] Y. Crama, *Concave extensions for non-linear* $0-1$ *maximization problems*, Math. Program., 61 (1993), pp. 53–60.

[14] Y. Crama and E. Rodríguez-Heck, *A class of valid inequalities for multilinear* $0-1$ *optimization problems*, Discrete Optim., 25 (2017), pp. 28–47.

[15] A. Del Pia and A. Khajavirad, *On decomposability of multilinear sets*, Math. Program., 2017.

[16] A. Del Pia and A. Khajavirad, *A polyhedral study of binary polynomial programs*, Math. Oper. Res., 42 (2017), pp. 389–410.

[17] D. Duris, *Some characterizations of* γ *and* β*-acyclicity of hypergraphs*, Inform. Process. Lett., 112 (2012), pp. 617–620.

[18] R. Fagin, *Degrees of acyclicity for hypergraphs and relational database schemes*, J. ACM, 30 (1983), pp. 514–550.

[19] A. Fischer, F. Fischer, and S. McCormick, *Matroid optimisation problems with nested non-linear monomials in the objective function*, Math. Program., 2017.

[20] R. Fortet, *Applications de l'algèbre de boole en recherche opérationelle*, Revue Française d'Automatique, Informatique et Recherche Opérationnelle, 4 (1960), pp. 17–26.

[21] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, 1979.

[22] L. Lovász, *Combinatorial Problems and Exercises*, AMS, Providence, RI, 1993.

[23] J. Luedtke, M. Namazifar, and J. Linderoth, *Some results on the strength of relaxations of multilinear functions*, Math. Program., 136 (2012), pp. 325–351.

[24] C. Meyer and C. Floudas, *Convex envelopes for edge-concave functions*, Math. Program., 103 (2005), pp. 207–224.

[25] M. Padberg, *The boolean quadric polytope: Some characteristics, facets and relatives*, Math. Program., 45 (1989), pp. 139–172.

[26] M. Padberg and M. Wilczak, *Boolean polynomials and set functions*, Math. Comput. Model., 17 (1993), pp. 3–6.

[27] A. Rikun, *A convex envelope formula for multilinear functions*, J. Global Optim., 10 (1997), pp. 425–437.

[28] I. Rosenberg, $0-1$ *optimization and non-linear programming*, Revue Française d'Automatique, Informatique et Recherche Opérationnelle, 6 (1972), pp. 95–97.

[29] H. Ryoo and N. Sahinidis, *Analysis of bounds for multilinear functions*, J. Global Optim., 19 (2001), pp. 403–424.

[30] A. Schrijver, *Theory of Linear and Integer Programming*, Wiley, Chichester, 1986.
[31] H. Sherali, *Convex envelopes of multilinear functions over a unit hypercube and over special discrete sets*, Acta Math. Vietnam., 22 (1997), pp. 245–270.
[32] H. Sherali and W. Adams, *A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems*, SIAM J. Discrete Math., 3 (1990), pp. 411–430.
[33] M. Tawarmalani, J.-P. Richard, and C. Xiong, *Explicit convex and concave envelopes through polyhedral subdivisions*, Math. Program., 138 (2013), pp. 531–577.