On decomposability of the multilinear polytope and its implications in mixed-integer nonlinear optimization

Alberto Del Pia * Aida Khajavirad †

May 17, 2018

This article is dedicated to the memory of Manfred Padberg whose work on the Boolean quadric polytope inspired us to start this line of research.

1 Introduction

Central to the efficiency of global optimization algorithms is their ability to construct sharp and cheaply computable convex relaxations. Factorable programming techniques are used widely in global optimization of mixed-integer nonlinear optimization problems (MINLPs) for bounding general nonconvex functions [9]. These techniques iteratively decompose a factorable function, through the introduction of variables and constraints for intermediate nonlinear expressions, until each intermediate expression can be convexified effectively.

Multilinear sets and polytopes. Factorable reformulations of many types of MINLPs, such as mixed-integer polynomial optimization problems, contain a collection of multilinear equations of the form $z_e = \prod_{v \in e} z_v$, $e \in E$, where E denotes a set of subsets of cardinality at least two of a ground set V. Let us define the set of points satisfying all multilinear equations present in a factorable reformulation of a MINLP as $\tilde{\mathcal{S}} = \{z : z_e = \prod_{v \in e} z_v \ \forall e \in E, \ z_v \in [0,1] \ \forall v \in V_1, \ z_v \in \{0,1\} \ \forall v \in V_2\}$, where V_1, V_2 forms a partition of V. It is well-known that the convex hull of $\tilde{\mathcal{S}}$ is a polytope and the projection of its vertices onto the space of the variables z_v, $v \in V$, is given by $\{0,1\}^V$. Hence, the facial structure of the convex hull of $\tilde{\mathcal{S}}$ can be equivalently studied by considering the following binary set:

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

In particular, this set represents the feasible region of a linearized unconstrained $0-1$ polynomial optimization problem. There is a one-to-one correspondence between sets of form (1) and hypergraphs $G = (V, E)$. Henceforth we refer to (1) as the *multilinear set* of the hypergraph G and denote it by $\mathcal{S}_G$, and refer to its convex hull as the *multilinear polytope* of G and denote it by MP_G. (See, e.g. [5])

If all multilinear equations defining $\mathcal{S}_G$ are bilinears, the multilinear polytope coincides with the Boolean quadric polytope defined by Padberg [10] in the context of $0-1$ quadratic optimization, in which case our hypergraph representation simplifies to the graph representation of Padberg. Indeed, a significant amount of research has been devoted to studying the facial structure of the Boolean quadric polytope and these theoretical developments have had a significant impact on the performance of branch-and-cut based algorithms for mixed-integer quadratic optimization problems. However, similar polyhedral studies for higher degree multilinear polytopes are quite scarce. Our ultimate goal is to bridge this gap by performing a systematic study of the facial structure of the multilinear polytope, and thus paving a way for devising novel optimization algorithms for nonconvex problems containing multilinear sub-expressions.

Decomposability. In this article, we provide an overview of some of our recent results [4, 6] on the facial structure of higher degree multilinear polytopes with a special focus on their "decomposability" properties. Namely, we demonstrate that for multilinear polytopes decomposability plays a key role from both theoretical and algorithmic viewpoints.

Let us start by introducing some hypergraph terminology that we need to formally define the notion of

*Department of Industrial and Systems Engineering & Wisconsin Institute for Discovery, University of Wisconsin-Madison. E-mail: delpia@wisc.edu.

†Department of Chemical Engineering, Carnegie Mellon University. E-mail: aida@cmu.edu.

decomposability for the multilinear polytope. Given a hypergraph $G = (V, E)$, and a subset V' of V, the *section hypergraph* of G induced by V' is the 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 a hypergraph G, and let G_j, $j \in J$, be distinct section hypergraphs of G such that $\cup_{j \in J} G_j = G$. Clearly, the system of all inequalities defining MP_{G_j} for all $j \in J$ provides a relaxation of MP_G as the convexification operation does not, in general, distribute over intersection. It is highly desirable to identify conditions under which these two sets coincide, as in such cases characterizing MP_G simplifies to characterizing each MP_{G_j} separately. More formally, we say that the polytope MP_G is *decomposable into polytopes MP_{G_j}, for $j \in J$*, if the following relation holds

$$\mathrm{MP}_G = \bigcap_{j \in J} \overline{\mathrm{MP}}_{G_j}, \tag{2}$$

where $\overline{\mathrm{MP}}_{G_j}$ is the set of all points in the space of MP_G whose projection in the space defined by G_j is MP_{G_j}.

Organization. In Section 2 we provide a summary of our results in [4] regarding necessary and sufficient conditions for decomposability of multilinear polytopes based on the structure of their intersection hypergraphs. Subsequently, in Section 3 we present a polynomial-time algorithm to optimally decompose a multilinear polytope into a collection of nondecomposable multilinear polytopes. A detailed analysis of this algorithm can be found in [4]. In Section 4 we give a brief overview of our results in [6], wherein we study the complexity of the multilinear polytope in conjunction with the acyclicity degree of its hypergraph and show that for certain acyclic hypergraphs, the multilinear polytope is decomposable into a collection of simpler multilinear polytopes whose explicit description can be obtained directly.

2 Decomposability based on the intersection hypergraph

Suppose that G_1 and G_2 are section hypergraphs of G such that $G_1 \cup G_2 = G$. The following theorem provides a sufficient condition for decomposability of MP_G into MP_{G_1} and MP_{G_2}, based on the structure of the intersection hypergraph $G_1 \cap G_2$. In the following, we say that a hypergraph $\bar{G}$ is *complete* if all subsets of $V(\bar{G})$ of cardinality at least two are present in $E(\bar{G})$.

Theorem 1. *Let G be a hypergraph, and let G_1,G_2 be section hypergraphs of G such that $G_1 \cup G_2 = G$ and $G_1 \cap G_2$ is a complete hypergraph. Then the polytope MP_G is decomposable into MP_{G_1} and MP_{G_2}.*

Figure 1 illustrates some hypergraphs G for which MP_G is decomposable into MP_{G_1} and MP_{G_2}. To draw a hypergraph G, throughout this article, we represent the nodes in $V(G)$ by points, and the edges in $E(G)$ by closed curves enclosing the corresponding set of points.

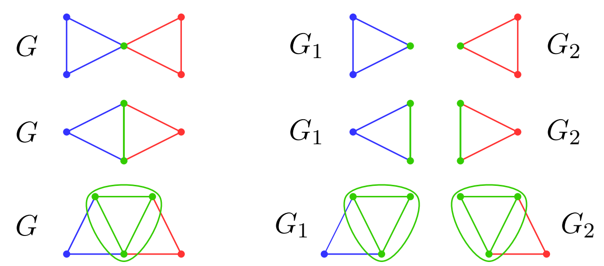

Figure 1

We now provide the proof sketch for Theorem 1. Given a vector z in the space defined by G, we denote by $z_\cap$ the vector that contains the components of z corresponding to nodes and edges that are in $G_1 \cap G_1$. Moreover, we denote by z_1 (resp. z_2) the vector that contains the components of z corresponding to nodes and edges that are in G_1 and not in G_2 (resp. in G_2 and not in G_1). The key step in proving Theorem 1 is to show that a vector $(\hat{z}_1, \hat{z}_\cap, \hat{z}_2)$ belongs to MP_G if $(\hat{z}_1, \hat{z}_\cap)$ can be written as a convex combination of vectors in $\mathcal{S}_{G_1}$ and $(\hat{z}_\cap, \hat{z}_2)$ can be written as a convex combination of vectors in $\mathcal{S}_{G_2}$. Clearly, given any two vectors $(z_1, z_\cap) \in \mathcal{S}_{G_1}$ and $(z'_\cap, z_2) \in \mathcal{S}_{G_2}$ with $z_\cap = z'_\cap$, we can combine them to obtain a vector $(z_1, z_\cap, z_2) \in \mathcal{S}_G$. Since by assumption the hypergraph $G_1 \cap G_2$ is complete, the polytope $\mathrm{MP}_{G_1 \cap G_2}$ is a simplex, implying that any vector $(\hat{z}_1, \hat{z}_\cap, \hat{z}_2)$ in MP_G can be written as a convex combination of the obtained vectors $(z_1, z_\cap, z_2)$ in $\mathcal{S}_G$. In partcIular, Theorem 1 unifies the existing decomposability results for the Boolean quadric polytope QP_G [10]:

Corollary 1. *Consider a graph $G = G_1 \cup G_2$, where G_1 and G_2 are induced subgraphs of G with*

$V(G_1) \cap V(G_2) = \{u\}$, *for some* $u \in V(G)$, *or* $V(G_1) \cap V(G_2) = \{u, v\}$, *for some* $\{u, v\} \in E(G)$. *Then* QP_G *is decomposable into* QP_{G_1} *and* QP_{G_2}.

The next theorem demonstrates the tightness of Theorem 1. We define the *rank* of a hypergraph G as the maximum cardinality of an edge in $E(G)$.

Theorem 2. *Let* $\bar{G}$ *be a rank-*r *hypergraph that is not complete. Then for any integer* $r' \geq \max\{r, 2\}$, *there exists a rank-*r' *hypergraph* $G = G_1 \cup G_2$, *where* G_1 *and* G_2 *are section hypergraphs of* G *with* $\bar{G} = G_1 \cap G_2$, *such that* MP_G *is not decomposable into* MP_{G_1} *and* MP_{G_2}.

Figure 2 illustrates some hypergraphs G for which MP_G is *not* decomposable into MP_{G_1} and MP_{G_2}.

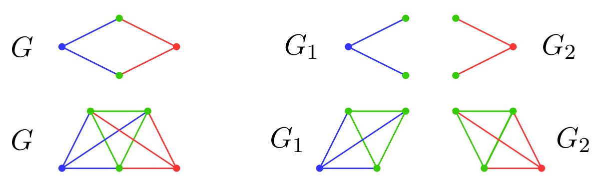

Figure 2

The proof of Theorem 2 is constructive. Namely, we search for a section hypergraph $\bar{H}$ of $\bar{G}$ with q nodes such that $E(\bar{H})$ consists all edges of cardinality between 2 and $q - 1$. Then we construct two hypergraphs H_1 and H_2 with $V(H_1) = V(\bar{H}) \cup \{u\}$, $E(H_1) = E(\bar{H}) \cup \{\{u, v\} : v \in V(\bar{H})\}$ and $V(H_2) = V(\bar{H}) \cup \{w\}$, $E(H_2) = E(\bar{H}) \cup \{\{w, v\} : v \in V(\bar{H})\}$. Subsequently, by letting $H = H_1 \cup H_2$, we provide a facet-defining inequality for MP_H with nonzero coefficients corresponding to some edges in $E(H_1) \setminus E(\bar{H})$ and in $E(H_2) \setminus E(\bar{H})$. This implies that MP_H is not decomposable into MP_{H_1} and MP_{H_2}. Next, we construct the hypergraph $G = H_1 \cup H_2 \cup \bar{G}$ and define G_1 and G_2 as the section hypergraphs of G induced by $V(H_1) \cup V(\bar{G})$ and $V(H_2) \cup V(\bar{G})$, respectively. We then show that since MP_H is not decomposable into MP_{H_1} and MP_{H_2}, the polytope MP_G is not decomposable into MP_{G_1} and MP_{G_2} either. It is simple to see that the rank of the hypergraph G constructed above is equal to $\max\{r, 2\}$. For any integer r' greater than $\max\{r, 2\}$, by adding a certain edge of cardinality r' to either G_1 or G_2, we can complete the proof.

In [10], Padberg poses a question regarding the decomposability of the Boolean quadric polytope when the intersection graph is a clique of cardinality greater than two. The proof of Theorem 2 implies that the answer to this question is negative for a clique with three or more nodes.

We conclude this section by remarking that in [4] we also present sufficient conditions for decomposability of multilinear polytopes with sparse intersection hypergraphs.

3 An optimal algorithm for decomposing the multilinear polytope

It is well-understood that branch-and-cut based MINLP solvers would highly benefit from our decomposition results as such techniques lead to significant reductions in CPU time during cut generation [1]. In this section, we present a simple and efficient algorithm for optimally decomposing a multilinear polytope into simpler and non-decomposable multilinear polytopes. Our proposed algorithm can be easily incorporated in MINLP solvers as a preprocessing step for cut generation. We start by presenting a sufficient condition for decomposability of MP_G into MP_{G_j}, for $j \in J$, which can be obtained by a recursive application of Theorem 1.

Theorem 3. *Let* G *be a hypergraph, and let* G_j, $j \in J$, *be section hypergraphs of* G *such that* $\cup_{j \in J} G_j = G$. *Suppose that for all* $j, j' \in J$ *with* $j \neq j'$, *the intersection* $G_j \cap G_{j'}$ *is the same complete hypergraph* $\bar{G}$. *Then* MP_G *is decomposable into* MP_{G_j}, *for* $j \in J$.

Now consider a hypergraph G and let $p \subset V(G)$. Denote by $\bar{G}$ the section hypergraph of G induced by p. We say that p *decomposes* G if

(a) the hypergraph $\bar{G}$ is complete,
(b) there exist at least two section hypergraphs G_j, $j \in J$, of G, with $V(G_j) \setminus V(G_{j'}) \neq \emptyset$ for all $j, j' \in J$ with $j \neq j'$, that together with $\bar{G}$ satisfy the hypothesis of Theorem 3.

If p does not decompose any G_j, $j \in J$, as defined in (b), then we refer to the family G_j, $j \in J$, as a p*-decomposition* of G. It can be shown that there exists a unique p-decomposition of G. The next result indicates that a p-decomposition test can be carried out efficiently.

Proposition 1. *Given a connected rank-*r *hypergraph* $G = (V, E)$ *and* $p \subset V$, *we can test if* p *decomposes* G, *and, if so, obtain the* p*-decomposition of* G *in* $O(r|E|)$ *time.*

Full Decompositions. In general, a multilinear polytope MP_G is decomposable into simpler polytopes via a series of p-decompositions of G until none of the newly generated multilinear polytopes are decomposable. In the following, whenever a polytope MP_G is decomposable into polytopes MP_{G_k}, $k \in K$, we refer to the family G_k, $k \in K$, as a *decomposition* of G. Given a hypergraph G, we define its *full-decomposition* as a decomposition of G given by a family G_k, $k \in K$, with the following properties:

(i) There exists no G_k, for some $k \in K$, and $p \subset V(G_k)$ such that p decomposes G_k.

(ii) No hypergraph G_s, for some $s \in K$, is a section hypergraph of another hypergraph G_t, for some $t \in K$ with $t \neq s$.

If G_s is a section hypergraph of G_t for some $s, t \in K$ with $s \neq t$, then MP_{G_s} corresponds to a face of MP_{G_t}. Thus, removing G_s from a decomposition of G amounts to removing redundant inequalities from the description of MP_G, which is computationally beneficial. It can be shown that the following algorithm gives a full-decomposition of G.

```
Gen_dec : General full-decomposition algorithm
  Input: A hypergraph G
  Output: A full-decomposition of G
  Initialize the family L = {G};
  while L does not satisfy property (i) of
   full-decomposition do
     select a hypergraph G̃ ∈ L and p ⊂ V(G̃);
     if p decomposes G̃ then
        let G_j, j ∈ J, be the p-decomposition of
          G̃;
        let J̃ be the subset of J such that each
          G_j, j ∈ J̃, is not a section hypergraph of
          any hypergraph in L different from G̃;
        in L, replace G̃ with G_j, j ∈ J̃;
  return L;
```

Decomposition orders. In `Gen_dec`, we have not specified which $\tilde{G} \in \mathcal{L}$ and $p \subset V(\tilde{G})$ to choose at each iteration. We refer to different choices of $\tilde{G}$ and p throughout the execution of `Gen_dec`, as *decomposition orders.* We denote a specific decomposition order by the sequence of choices that defines it, where each choice consists of a pair $(\tilde{G}, p)$, for some hypergraph $\tilde{G} \in \mathcal{L}$ and a set of nodes $p \subset V(\tilde{G})$ that is tested for p-decomposition of $\tilde{G}$. The next proposition indicates that a full-decomposition of G does not depend on the specific decomposition order used.

Proposition 2. *The full-decomposition of a hypergraph obtained by* `Gen_dec` *is independent of the decomposition order.*

Henceforth, we will speak of *the* full-decomposition of G. However, as we detail next, different decomposition orders result in different computational costs for `Gen_dec`. First, from the definition of `Gen_dec` it follows that the length of the decomposition order used is a reasonable measure for the overall cost of this algorithm and it can be shown that for a hypergraph G, every decomposition order contains at least $|V(G)| + |E(G)|$ pairs. Second, to ensure that property (ii) in the definition of the full-decomposition is satisfied, every time the p-decomposition of $\tilde{G}$ is generated, each new hypergraph G_j is compared with the existing ones and is added to $\mathcal{L}$ only if it is not a section hypergraph of another hypergraph in $\mathcal{L}$. Let us refer to the section hypergraphs not added to $\mathcal{L}$ as *redundant* hypergraphs. It can be shown that different decomposition orders in `Gen_dec` may result in distinct redundant hypergraphs. As the redundancy check is computationally expensive, it is beneficial to obtain a decomposition order that results in a minimum number of redundant hypergraphs.

The optimal decomposition algorithm. Next, we define a special sequence of choices $\bar{\mathcal{O}}$ in the execution of `Gen_dec` with highly desirable algorithmic properties. At a given iteration of `Gen_dec`, we say that $p \in V(\tilde{G}) \cup E(\tilde{G})$ is *tested* in $\tilde{G}$, if the pair $(\tilde{G}, p)$ has been already considered in an earlier iteration of `Gen_dec`. Moreover, we refer to the hypergraph $\tilde{G}$ in `Gen_dec` as the *parent* of each G_j. The *ancestors* of G_j are the parent of G_j, and the ancestors of the parent of G_j. At a given iteration, any hypergraph in the current family $\mathcal{L}$ can be chosen as $\tilde{G}$. Let the list $\{q_k, k \in K\}$ contain all nodes and edges of $\tilde{G}$ ordered by increasing cardinality. We define p to be the first element q_k in the above list that is *not* tested in $\tilde{G}$ or in any ancestor of $\tilde{G}$. The sequence $\bar{\mathcal{O}}$ ends when no such pair $(\tilde{G}, p)$ can be found.

Proposition 3. *The sequence $\bar{\mathcal{O}}$ is a decomposition order. Moreover, it creates no redundant hypergraphs. Consider a hypergraph G with n nodes and m edges. Let the decomposition order $\bar{\mathcal{O}}$ for G be given by $(G_1, p_1), (G_2, p_2), \ldots, (G_t, p_t)$. Then $t = n + m$.*

In [4], we present an optimal full-decomposition algorithm, referred to as `Opt_dec`, which is obtained by an efficient incorporation of the decomposition order $\bar{\mathcal{O}}$ in `Gen_dec`. We refer to this algorithm as optimal due to two reasons. First, `Opt_dec` applies the minimum number of p-decomposition tests needed to obtain the full-decomposition of any hypergraph. Second, no redundant hypergraph is generated in the course of `Opt_dec`, and hence the costly redundancy test (as described in `Gen_dec`) is not required. The following proposition gives the worst-case running time of `Opt_dec`.

Proposition 4. *Consider a connected rank-r hypergraph G with n nodes and m edges. Then, the running time of* `Opt_dec` *is $O(rm(n+m))$.*

In [4], we provide an example that demonstrates the significance of our optimal decomposition algorithm; namely we define a hypergraph G and a decomposition order $\tilde{O}$, such that when incorporated in `Gen_dec`, in comparison to $\bar{O}$, the decomposition order $\tilde{O}$ requires $n(m-1)/2$ additional p-decomposition tests to obtain a full-decomposition of G. In addition, a total number of $n(n-2)/4-1$ redundant hypergraphs are generated in the course of `Gen_dec`.

4 The multilinear polytope of acyclic hypergraphs

In this section, we demonstrate the key role of decomposition in obtaining explicit descriptions for the multilinear polytope of certain acyclic hypergraphs. Moreover, these convex hull characterizations enable us to optimize a linear function over MP_G in polynomial time. We start by providing a sufficient condition for decomposability of multilinear polytopes that will be used for the subsequent developments.

Theorem 4. *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 MP_G is decomposable into MP_{G_1} and MP_{G_2}.*

Figure 3 illustrates a hypergraph G for which by Theorem 4 the polytope MP_G is decomposable into MP_{G_1} and MP_{G_2}.

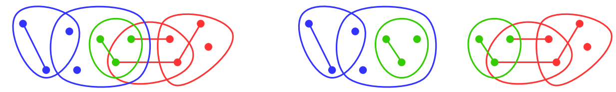
G G_1 G_2

Figure 3

As in Theorem 1, to prove Theorem 4, we need to show that a vector $(\hat{z}_1, \hat{z}_\cap, \hat{z}_2)$ belongs to MP_G if $(\hat{z}_1, \hat{z}_\cap)$ can be written as a convex combination of vectors in $\mathcal{S}_{G_1}$ and $(\hat{z}_\cap, \hat{z}_2)$ can be written as a convex combination of vectors in $\mathcal{S}_{G_2}$. Moreover, as before, it is sufficient to consider vectors in $\mathcal{S}_G$ obtained by combining one vector $(z_1, z_\cap)$ in $\mathcal{S}_{G_1}$ with one vector $(z'_\cap, z_2)$ in $\mathcal{S}_{G_2}$. However, since in Theorem 4 the intersection hypergraph is not complete, it is no longer sufficient to only combine vectors with $z_\cap = z'_\cap$. In this case, we need to consider all vectors $(z_1, z'_\cap, z_2)$ obtained by combining a vector $(z_1, z_\cap) \in \mathcal{S}_{G_1}$ and a vector $(z'_\cap, z_2) \in \mathcal{S}_{G_2}$ with $z_{\bar{e}} = z'_{\bar{e}}$. The presence of $(z_1, z'_\cap, z_2)$ in $\mathcal{S}_G$ follows from the assumption that every edge that is only in G_1 either contains $\bar{e}$ or is disjoint from it. Moreover, the existence of the edge $\bar{e}$ implies that we can write the vector $(\hat{z}_1, \hat{z}_\cap, \hat{z}_2)$ as a convex combination of the obtained vectors $(z_1, z'_\cap, z_2)$ in $\mathcal{S}_G$.

Acyclic hypergraphs. Padberg [10] shows that for an acyclic graph, the Boolean quadric polytope admits a simple and compact description. This result can be obtained by showing that the Boolean quadric polytope of an acyclic graph is decomposable into a collection of Boolean quadric polytopes whose graphs consists of a single edge. To obtain similar characterizations for higher degree multilinear polytopes, it is then natural to look into the notion of acyclicity for hypergraphs. Interestingly, unlike graphs for which there is a single natural notion of acyclicity, for hypergraphs several different degrees of acyclicity have been defined [8]. In the following, we present two types of hypergraph acyclicity which will be used for the subsequent developments.

The most restrictive class of acyclic hypergraphs is the class of Berge-acyclic hypergraphs. A *Berge-cycle* in G is a sequence $v_1, e_1, v_2, e_2, \ldots, v_t, e_t, v_1$ with $t \geq 2$, such that (i) $v_1, v_2, \ldots, v_t$ are distinct nodes of G, (ii) $e_1, e_2, \ldots, e_t$ are distinct edges of G, (iii) $v_i, v_{i+1} \in e_i$ for $i = 1, \ldots, t-1$, and $v_t, v_1 \in e_t$. A hypergraph is *Berge-acyclic* if it contains no Berge-cycles. The next class of acyclic hypergraphs in increasing order of generality, is the class of γ-acyclic hypergraphs. A γ-*cycle* in G is a Berge-cycle such that $t \geq 3$, and for each $i \in \{2, \ldots, t\}$, the node v_i belongs to e_{i-1},

e_i and no other e_j. A hypergraph is γ-*acyclic* if it contains no γ-cycles.

Acyclicity and decomposability. The link between hypergraph acyclicity and decomposability is given by the concept of leaf of a hypergraph. Consider a hypergraph $G = (V, E)$. An edge of G is *maximal* if it is not contained in any other edge of G. We say that an edge e' is a *leaf* of G if it is a maximal edge and $e' \cap (\cup_{e \in E \setminus E'} e) \subset \tilde{e}$ for some $\tilde{e} \in E \setminus E'$, where E' is the set of edges contained in e'. It can be shown that every γ-acyclic hypergraph contains a leaf. The existence of a leaf, together with the special structure of Berge-acyclic and γ-acyclic hypergraphs enables us to employ our decomposition results and derive an explicit description of MP_G by induction on the number of maximal edges of G.

Berge-acyclic hypergraphs. The *standard linearization* $\mathrm{MP}_G^{\mathrm{LP}}$ is a widely-used relaxation of $\mathcal{S}_G$ and is obtained by replacing each multilinear equation $z_e = \prod_{v \in e} z_v$ by its convex hull over the unit hypercube (see, e.g., [3]):

$$\begin{aligned}
& z_v \leq 1 && \forall v \in V, \\
& z_e \geq 0 && \\
& z_e \geq \textstyle\sum_{v \in e} z_v - |e| + 1 && \forall e \in E, \\
& z_e \leq z_v && \forall e \in E, \ \forall v \in e.
\end{aligned}$$

We now show that for a Berge-acyclic hypergraph, we have $\mathrm{MP}_G = \mathrm{MP}_G^{\mathrm{LP}}$.

Any two edges of a Berge-acyclic hypergraph intersect in at most one node. It then follows that the hypergraph considered in the base case of the induction consists of a single edge. Hence, the corresponding multilinear polytope coincides with the standard linearization.

In the inductive step, we construct MP_G in two steps:

1. Decompose the polytope MP_G into MP_{G_1} and MP_{G_2}, where G_1 is the section hypergraph of G induced by the leaf e' and G_2 is the section hypergraph of G induced by $\cup_{e \in E \setminus E'} e$.
2. Obtain MP_G by juxtaposing the description of MP_{G_1} and of MP_{G_2} given by the induction hypothesis.

For a Berge-acyclic hypergraph, the intersection of the leaf e' with the hypergraph G_2, i.e., the set $\bar{p}$ defined in Theorem 4, consists of at most one node. Hence, all assumptions of Theorem 4 are trivially satisfied and we can utilize this result to perform the decomposition described in Step 1. Hence, if G is a Berge-acyclic hypergraph, we have $\mathrm{MP}_G = \mathrm{MP}_G^{\mathrm{LP}}$. In fact, we have proved that the converse holds as well. More precisely, we have shown the following:

Theorem 5. *$MP_G = MP_G^{LP}$ if and only if G is a Berge-acyclic hypergraph.*

It follows directly from Theorem 5 that for a Berge-acyclic hypergraph G, we can optimize a linear function over MP_G via linear optimization in polynomial time.

In [10], Padberg shows that the standard linearization coincides with the Boolean quadric polytope if and only if G is an acyclic graph. Therefore Theorem 5 generalizes Padberg's result to higher degree multilinear polytopes.

γ-acyclic hypergraphs. To characterize the multilinear polytope of γ-acyclic hypergraphs, we introduce a class of valid inequalities for MP_G which we will refer to as flower inequalities. Let e_0 be an edge of G and let e_k, $k \in K$, be a collection of edges such that $|e_0 \cap e_k| \geq 2$ for every $k \in K$, and $e_i \cap e_j = \emptyset$ for all $i, j \in K$ with $i \neq j$. Then a *flower inequality* for MP_G is given by:

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

We define the *flower relaxation* MP_G^F as the relaxation of the multilinear set obtained by adding all flower inequalities to its standard linearization $\mathrm{MP}_G^{\mathrm{LP}}$. We now show that for a γ-acyclic hypergraph, we have $\mathrm{MP}_G = \mathrm{MP}_G^F$.

To establish the base case of the induction, we make use of the fact that a γ-acyclic hypergraph with one maximal edge is a laminar hypergraph. The multilinear polytope of a laminar hypergraph can be characterized using a fundamental result due to Conforti and Cornuéjols regarding the connection between integral polyhedra and balanced matrices [2]. This characterization in turn implies that for a laminar hypergraph the multilinear polytope coincides with its flower relaxation.

In the inductive step, Theorem 5 cannot be directly applied to MP_G as was the case for Berge-acyclic hypergraphs. However, we can utilize this result after the addition of one extra edge to G. In more detail, we construct MP_G in four steps:

1. Define the hypergraph G^+ obtained from G by adding the edge $\bar{p} := e' \cap (\cup_{e \in E \setminus E'} e)$.

2. Decompose the polytope MP_{G^+} into MP_{G_1} and MP_{G_2}, where G_1 is the section hypergraph of G^+ induced by e' and G_2 is the section hypergraph of G^+ induced by $\cup_{e \in E \setminus E'} e$.
3. Obtain MP_{G^+} by juxtaposing the description of MP_{G_1} given by the base case, and of MP_{G_2} given by the induction hypothesis.
4. Obtain MP_G by projecting out the variable $\bar{p}$ from the description of MP_{G^+}.

The section hypergraph induced by an edge of a γ-acyclic hypergraph is laminar. This in particular implies that for every edge e of G containing nodes in $V(G_1) \setminus V(G_2)$ either $e \supset \bar{p}$, or $e \cap \bar{p} = \emptyset$. Hence, we can employ Theorem 4 to perform the decomposition described in Step 2. Finally, by projecting out the variable $z_{\bar{p}}$ from the description of MP_{G^+} using Fourier-Motzkin elimination, we conclude that $\text{MP}_G = \text{MP}_G^F$.

In fact, we have shown that the converse holds as well. More precisely, we have shown the following:

Theorem 6. *$MP_G = MP_G^F$ if and only if G is a γ-acyclic hypergraph.*

For γ-acyclic hypergraphs, the number of facets of MP_G^F may not be bounded by a polynomial in $|V(G)|, |E(G)|$. However, flower inequalities can be separated in strongly polynomial time, and this allows us to optimize a linear function over MP_G in polynomial time.

We conclude this article by remarking that in [7] we extend the above decomposition based technique to characterize the multilinear polytope for a more general class of acyclic hypergraphs.

Acknowledgments

We would like to thank the Young Researcher Prize Committee consisting of Michael Ferris (chair), Jonathan Eckstein, Simge Küçükyavuz, and Suvrajeet Sen; it is a great honor to have been selected as a recipient of this prize. This research was supported in part by National Science Foundation award CMMI-1634768.

References

[1] X. Bao, A. Khajavirad, N.V. Sahinidis, and M. Tawarmalani. Global optimization of nonconvex problems with multilinear intermediates. *Mathematical Programming Computation*, pages 1–37, 2014.

[2] G. Cornuéjols. *Combinatorial Optimization: Packing and Covering*, volume 74 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 2001.

[3] Y. Crama. Concave extensions for non-linear 0−1 maximization problems. *Mathematical Programming*, 61:53–60, 1993.

[4] A. Del Pia and A. Khajavirad. On decomposability of multilinear sets. *Mathematical Programming*, 2017.

[5] A. Del Pia and A. Khajavirad. A polyhedral study of binary polynomial programs. *Mathematics of Operations Research*, 42(2):389–410, 2017.

[6] A. Del Pia and A. Khajavirad. The multilinear polytope for acyclic hypergraphs. *SIAM Journal on Optimization*, 28:1049–1076, 2018.

[7] A. Del Pia and A. Khajavirad. The running intersection relaxation of the multilinear polytope. *Optimization Online manuscript 2018/05/6618*, 2018.

[8] Ronald Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *Journal of the ACM (JACM)*, 30(3):514–550, 1983.

[9] G.P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I—Convex underestimating problems. *Mathematical Programming*, 10:147–175, 1976.

[10] M. Padberg. The boolean quadric polytope: Some characteristics, facets and relatives. *Mathematical Programming*, 45:139–172, 1989.