

The pseudo-Boolean polytope and polynomial-size extended formulations for binary polynomial optimization

Alberto Del Pia ^{*} Aida Khajavirad [†]

July 1, 2024

Abstract

With the goal of obtaining strong relaxations for binary polynomial optimization problems, we introduce the pseudo-Boolean polytope defined as the set of binary points $z \in \{0, 1\}^{V \cup S}$ satisfying a collection of equalities of the form $z_s = \prod_{v \in s} \sigma_s(z_v)$, for all $s \in S$, where $\sigma_s(z_v) \in \{z_v, 1 - z_v\}$, and where S is a multiset of subsets of V . By representing the pseudo-Boolean polytope via a signed hypergraph, we obtain sufficient conditions under which this polytope has a polynomial-size extended formulation. Our new framework unifies and extends all prior results on the existence of polynomial-size extended formulations for the convex hull of the feasible region of binary polynomial optimization problems of degree at least three.

Key words: Binary polynomial optimization; Pseudo-Boolean optimization; Pseudo-Boolean polytope; Signed hypergraph; Polynomial-size extended formulation.

1 Introduction

We consider the problem of maximizing a multivariate polynomial function over the set of binary points, henceforth referred to as *binary polynomial optimization*. This problem class has numerous applications across science and engineering [8], and is NP-hard in general. Based on the encoding of the polynomial objective function, we obtain two popular optimization problems, which we refer to as “multilinear optimization” and “pseudo-Boolean optimization.” We first introduce these two optimization problems in Sections 1.1 and 1.2, and then compare and contrast them in Section 2. As binary quadratic optimization has been extensively studied in the literature (see for example [2, 6, 35]), in this paper, we focus on binary polynomial optimization problems of degree at least three.

1.1 Multilinear optimization

A *hypergraph* G is a pair (V, E) , where V is a finite set of nodes and E is a set of *edges*, which are subsets of V of cardinality at least two. Following the approach introduced in [16], with any hypergraph $G = (V, E)$, and cost vector $c \in \mathbb{R}^{V \cup E}$, we associate the following *multilinear optimization problem*:

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

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

[†]Department of Industrial and Systems Engineering, Lehigh University. E-mail: aida@lehigh.edu.

It can be checked that any binary polynomial optimization problem has a unique representation of the form Problem BPO_m . Next, we linearize the objective function of Problem BPO_m by introducing a new variable z_e for each product term $\prod_{v \in e} z_v$ to obtain an equivalent reformulation of this problem in a lifted space:

$$\begin{aligned} \max \quad & \sum_{v \in V} c_v z_v + \sum_{e \in E} c_e z_e \\ \text{s.t.} \quad & z_e = \prod_{v \in e} z_v \quad \forall e \in E \\ & z_v \in \{0, 1\} \quad \forall v \in V. \end{aligned} \tag{\ell \text{BPO}_m}$$

With the objective of understanding the facial structure of the convex hull of the feasible region of Problem ℓBPO_m , Del Pia and Khajavirad [16] introduced the *multilinear set*, defined as

$$\mathcal{S}_m(G) := \left\{ z \in \{0, 1\}^{V \cup E} : z_e = \prod_{v \in e} z_v, \forall e \in E \right\},$$

and its convex hull, which is called the *multilinear polytope* and is denoted by $\mathcal{P}_m(G)$.

In [17, 20], the authors show that the complexity of the facial structure of $\mathcal{P}_m(G)$ is closely related to the “acyclicity degree” of G . The most well-known types of acyclic hypergraphs, in increasing order of generality, are Berge-acyclic, γ -acyclic, β -acyclic, and α -acyclic hypergraphs [3, 9, 24, 25]. We next briefly review the existing results on the facial structure of the multilinear polytope of acyclic hypergraphs. Recall that the *rank* of a hypergraph G , denoted by r , is the maximum cardinality of any edge in E . In [10, 17], the authors prove that $\mathcal{P}_m(G)$ coincides with its standard linearization if and only if G is Berge-acyclic. This in turn implies that if G is Berge-acyclic, then $\mathcal{P}_m(G)$ is defined by $|V| + (r + 2)|E|$ inequalities in the original space. In [17], the authors prove that $\mathcal{P}_m(G)$ coincides with its flower relaxation if and only if G is γ -acyclic. This result implies that if G is γ -acyclic, then $\mathcal{P}_m(G)$ has a polynomial-size extended formulation with at most $|V| + 2|E|$ variables and at most $|V| + (r + 2)|E|$ inequalities. Subsequently, in [20], the authors present a polynomial-size extended formulation for the multilinear polytope of β -acyclic hypergraphs with at most $(r - 1)|V| + |E|$ variables and at most $(3r - 4)|V| + 4|E|$ inequalities.

On the other hand, in [15], the authors prove that Problem BPO_m is strongly NP-hard over α -acyclic hypergraphs. This result implies that, unless $\text{P} = \text{NP}$, one cannot construct, in polynomial time, a polynomial-size extended formulation for the multilinear polytope of α -acyclic hypergraphs. However, as we detail next, by making further assumptions on the rank of α -acyclic hypergraphs, one can construct a polynomial-size extended formulation for the multilinear polytope. In [5, 33, 38], the authors give extended formulations for the convex hull of the feasible set of (possibly constrained) multilinear optimization problems. The size of these extended formulations is parameterized in terms of the “tree-width” of their so-called intersection graphs. For the unconstrained case, as detailed in [19], their result can be equivalently stated as follows: If G is an α -acyclic hypergraph of rank r , with $r = O(\log \text{poly}(|V|, |E|))$, then $\mathcal{P}_m(G)$ has a polynomial-size extended formulation, where by $\text{poly}(|V|, |E|)$, we imply a polynomial function in $|V|, |E|$. Henceforth, for brevity, whenever for a hypergraph $G = (V, E)$ we have $r = O(\log \text{poly}(|V|, |E|))$, we say that G has *log-poly* rank.

For further results regarding polyhedral relaxations of multilinear sets of degree at least three, see [13, 14, 19, 21–23, 30, 32].

1.2 Pseudo-Boolean optimization

We define a *signed hypergraph* H as a pair (V, S) , where V is a finite set of nodes and S is a set of *signed edges*. A *signed edge* $s \in S$ is a pair (e, η_s) , where e is a subset of V of cardinality at least two,

and η_s is a map that assigns to each $v \in e$ a *sign* $\eta_s(v) \in \{-1, +1\}$. The *underlying edge* of a signed edge $s = (e, \eta_s)$ is e . Two signed edges $s = (e, \eta_s)$, $s' = (e', \eta_{s'}) \in S$ are said to be *parallel* if $e = e'$, and they are said to be *identical* if $e = e'$ and $\eta_s = \eta_{s'}$. Throughout this paper, we consider signed hypergraphs with no identical signed edges. However, our signed hypergraphs often contain parallel signed edges.

With any signed hypergraph $H = (V, S)$, and cost vector $c \in \mathbb{R}^{V \cup S}$, we associate the following *pseudo-Boolean optimization problem*:

$$\begin{aligned} \max \quad & \sum_{s \in S} c_s \prod_{v \in s} \sigma_s(z_v) \\ \text{s.t.} \quad & z \in \{0, 1\}^V, \end{aligned} \tag{BPO}_{\text{pB}}$$

where

$$\sigma_s(z_v) := \begin{cases} z_v & \text{if } \eta_s(v) = +1 \\ 1 - z_v & \text{if } \eta_s(v) = -1. \end{cases}$$

A variety of important applications such as maximum satisfiability problems [26], and inference in graphical models [29] can naturally be formulated as pseudo-Boolean optimization problems. Problem BPO_{pB} has been extensively studied in the literature; see [8] for a detailed survey of main results. Some of the main topics considered in these works are quadratic pseudo-Boolean optimization problems [6, 27], quadratization of general pseudo-Boolean optimization problems [7], and special problem classes such as super-modular pseudo-Boolean optimization problems [12].

As before, we linearize the objective function of Problem BPO_{pB} by introducing a new variable z_s , for each signed edge $s \in S$, to obtain an equivalent reformulation of Problem BPO_{pB} in a lifted space:

$$\begin{aligned} \max \quad & \sum_{s \in S} c_s z_s \\ \text{s.t.} \quad & z_s = \prod_{v \in s} \sigma_s(z_v) \quad \forall s \in S \\ & z \in \{0, 1\}^{V \cup S}. \end{aligned} \tag{\ell BPO}_{\text{pB}}$$

In this paper, we introduce the *pseudo-Boolean set* of the signed hypergraph $H = (V, S)$, as the feasible region of Problem $\ell \text{BPO}_{\text{pB}}$:

$$\mathcal{S}_{\text{pB}}(H) := \left\{ z \in \{0, 1\}^{V \cup S} : z_s = \prod_{v \in s} \sigma_s(z_v), \forall s \in S \right\},$$

and we refer to its convex hull as the *pseudo-Boolean polytope* and denote it by $\mathcal{P}_{\text{pB}}(H)$. With the objective of constructing strong linear programming (LP) relaxations for pseudo-Boolean optimization problems, in this paper, we study the facial structure of the pseudo-Boolean polytope.

It is important to note that, unlike the multilinear set/polytope, the pseudo-Boolean set/polytope may not be full-dimensional. For example, if $H = (V, S)$ contains three signed edges $s_1 = (e, \eta_{s_1})$, $s_2 = (e, \eta_{s_2})$, $s_3 = (e', \eta_{s_3})$, where $e' = e \setminus \{\bar{v}\}$ for some $\bar{v} \in V$, $\eta_{s_1}(v) = \eta_{s_2}(v) = \eta_{s_3}(v)$ for all $v \in e \setminus \{\bar{v}\}$, and $\eta_{s_1}(\bar{v}) = -\eta_{s_2}(\bar{v})$, then $\mathcal{S}_{\text{pB}}(H)$ and $\mathcal{P}_{\text{pB}}(H)$ are not full-dimensional as we have $z_{s_1} + z_{s_2} = z_{s_3}$. While the pseudo-Boolean polytope is significantly more complex than the multilinear polytope, as we demonstrate in this paper, by studying the facial structure of $\mathcal{P}_{\text{pB}}(H)$, one can obtain polynomial-size extended formulations for many instances for which such a formulation is not known for the corresponding multilinear polytope. In Section 2 we explore in detail the connections between the two polytopes.

1.3 Our contributions

In this paper, we demonstrate how the hypergraph framework pioneered in [16] for multilinear optimization is relevant and applicable to pseudo-Boolean optimization as well. Namely, using signed hypergraphs to represent pseudo-Boolean sets, we present sufficient conditions under which the pseudo-Boolean polytope admits a polynomial-size extended formulation. Our results unify and extend all prior results on polynomial-size representability of the multilinear polytope [5, 17, 19, 20, 33, 38].

We introduce a new technique, which we refer to as the “recursive inflate-and-decompose” framework to construct polynomial-size extended formulations for the pseudo-Boolean polytope. Our proposed framework relies on a recursive application of three key ingredients, each of which is of independent interest:

1. A sufficient condition for decomposability of pseudo-Boolean polytopes (see Theorem 1). This is the first result on decomposability of pseudo-Boolean polytopes and serves as a significant generalization of Theorem 4 in [20] (Section 3).
2. A polynomial-size extended formulation for the pseudo-Boolean polytope of pointed signed hypergraphs (see Theorem 2). The pseudo-Boolean polytope of pointed signed hypergraphs is the building block of our extended formulations, which appears as a result of applying our decomposition technique of Part 1 (Section 4).
3. An operation, which we refer to as “inflation of signed edges” (see Theorem 3) that we use to transform a large class of signed hypergraphs to those for which our results of Parts 1 and 2 are applicable (Section 5).

As we detailed in Section 1.1, at the time of this writing, the most general sufficient conditions under which one can obtain a polynomial-size extended formulation for the multilinear polytope $\mathcal{P}_m(G)$ are:

- (i) G is a β -acyclic hypergraph [20],
- (ii) G is an α -acyclic hypergraph with log-poly rank [33].

It is important to remark that neither of the above sufficient conditions implies the other one. Furthermore, the two results in [20] and [33] have been proven using entirely different techniques. In Section 6, we show that our recursive inflate-and-decompose framework implies as special cases both sufficient conditions (i) and (ii) above and extends to many more cases of interest. Below, we summarize these results.

Consider a signed hypergraph $H = (V, S)$. We define the *underlying hypergraph* of H as the hypergraph obtained from H by ignoring signs and dropping parallel edges. In Section 6.1 we prove that, if the underlying hypergraph of H is β -acyclic, then $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation (see Theorem 5). This is a significant generalization of case (i) above, as it only requires the β -acyclicity of the underlying hypergraph of H . In Section 6.2 we prove that if the underlying hypergraph of H is α -acyclic and has log-poly rank, then $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation (see Theorem 6). This result essentially coincides with case (ii) above. In Section 6.3, we introduce the notion of “gap” for hypergraphs, which roughly speaking, indicates if it is possible to inflate signed edges in an efficient manner. We then show that for signed hypergraphs, if certain gaps are not too large, by combining results of Theorems 3 and 5, one can obtain polynomial-size extended formulations for the pseudo-Boolean polytope (see Propositions 3 and 4). Finally, in Section 6.4 we outline some generalizations and directions of future research.

It is important to remark that the proofs of all our results regarding the existence of polynomial-size extended formulations for $\mathcal{P}_{\text{pB}}(H)$ are constructive and the proposed extended formulations can be constructed in polynomially many operations in $|V|, |S|$.

We would like to conclude this section by further emphasizing on the power of these extended formulations: Not only they serve as polynomial-size LP formulations for special classes of Problem BPO_m and Problem BPO_{PB} , they can also be used to construct strong LP *relaxations* for general mixed-integer nonlinear optimization problems whose factorable reformulations contain pseudo-Boolean sets (see for example [4, 21, 31, 34, 36, 37]).

1.4 Notations and preliminaries

In the following, we present all hypergraph terminology and notation that we will use throughout this paper.

Hypergraphs. Let $G = (V, E)$ be a hypergraph. We define the hypergraph obtained from G by *removing* a node $v \in V$ as the hypergraph $G - v$ with set of nodes $V \setminus \{v\}$ and set of edges $\{e - v : e \in E, |e - v| \geq 1\}$.

A node $v \in V$ is a β -leaf of G if the set of the edges of G containing v is totally ordered with respect to inclusion. A *sequence of β -leaves* of length t for some $1 \leq t \leq |V|$ of G is an ordering v_1, \dots, v_t of t distinct nodes of G , such that v_1 is a β -leaf of G , v_2 is a β -leaf of $G - v_1$, and so on, until v_t is a β -leaf of $G - v_1 - \dots - v_{t-1}$. The hypergraph G is said to be β -acyclic if it has a sequence of β -leaves of length $|V|$. An equivalent definition of β -acyclic hypergraphs can be obtained using the concept of β -cycle. A β -cycle of length q for some $q \geq 3$ in G is a sequence $v_1, e_1, v_2, e_2, \dots, v_q, e_q, v_1$ such that v_1, v_2, \dots, v_q are distinct nodes, e_1, e_2, \dots, e_q are distinct edges, and v_i belongs to e_{i-1}, e_i and no other edge among e_1, e_2, \dots, e_q , for all $i = 1, \dots, q$, where $e_0 := e_q$. A hypergraph is β -acyclic if and only if it does not contain any β -cycles [24].

A node $v \in V$ is an α -leaf of G if the set of edges of G containing v has a maximal element for inclusion. A *sequence of α -leaves* of length t for some $1 \leq t \leq |V|$ of G is an ordering v_1, \dots, v_t of t distinct nodes of G , such that v_1 is an α -leaf of G , v_2 is an α -leaf of $G - v_1$, and so on, until v_t is an α -leaf of $G - v_1 - \dots - v_{t-1}$. The hypergraph G is said to be α -acyclic if it has a sequence of α -leaves of length $|V|$. An equivalent definition of α -acyclic hypergraphs can be obtained using the concept of α -cycles introduced in [28].

We say that G is *connected* if for every $u, w \in V$, there exists a sequence $v_1, e_1, v_2, e_2, \dots, v_q, e_q, v_{q+1}$ such that $v_1 = u$, $v_{q+1} = w$, $v_2, v_3, \dots, v_q \in V$, $e_1, e_2, \dots, e_q \in E$, and e_i contains v_i, v_{i+1} , for all $i = 1, \dots, q$. The *connected components* of G are its maximal connected partial hypergraphs. A hypergraph $G' = (V', E')$ is a *partial hypergraph* of $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E$.

Signed hypergraphs. Let $H = (V, S)$ be a signed hypergraph. In the following, we define some useful operations on signed edges. Let $s = (e, \eta_s) \in S$. With a slight abuse of notation, we use set theoretical notation on s , with the understanding that s should be replaced with e . For example, we denote by $|s|$ the number $|e|$, we say that s is nonempty if e is nonempty, we write $v \in s$ meaning $v \in e$, and for $U \subseteq V$, we write $s \subseteq U$ (resp., $s \supseteq U$, $s = U$) meaning $e \subseteq U$ (resp., $e \supseteq U$, $e = U$). Similarly, we might write $s \subseteq s'$, for $s = (e, \eta_s), s' = (e', \eta_{s'}) \in S$, instead of $e \subseteq e'$. If $v \in s$, we denote by $s - v$ the signed edge $s' = (e', \eta_{s'})$, where $e' := e \setminus \{v\}$, and $\eta_{s'}$ is the restriction of η_s that assigns to each $v \in e'$ the sign $\eta_{s'}(v) = \eta_s(v)$. If $v \notin s$, we denote by $s + v^+$ the signed edge $s' = (e', \eta_{s'})$, where $e' := e \cup \{v\}$, and $\eta_{s'}$ is the extension of η_s that assigns to each $u \in e'$ the sign $\eta_{s'}(u) = \eta_s(u)$ and assigns to v the sign $\eta_{s'}(v) = +1$. Similarly, $s + v^-$ is defined as $s + v^+$ but with $\eta_{s'}(v) = -1$.

We define the signed hypergraph obtained from H by *removing* a node $v \in V$ as the signed hypergraph $H - v$ with set of nodes $V \setminus \{v\}$ and set of signed edges $\{s - v : s \in S, |s - v| \geq 1\}$.

1.5 Organization

The remainder of the paper is structured as follows. In Section 2 we detail on the connections between multilinear optimization and pseudo-Boolean optimization. In Section 3 we present a sufficient condition for decomposability of the pseudo-Boolean polytope. In Section 4 we obtain a polynomial-size extended formulation for the pseudo-Boolean polytope of pointed signed hypergraphs. In Section 5 we introduce the inflation operation. Combining the results of Sections 3-5, in Section 6, we introduce our new “recursive inflate-and-decompose” framework, which in turn enables us to present polynomial-size extended formulations for the pseudo-Boolean polytope of a large class of signed hypergraphs. Finally, Section 7 is the appendix and contains a technical proof omitted from Section 4.

2 Multilinear optimization versus pseudo-Boolean optimization

In this section we compare the two encodings of binary polynomial optimization introduced in Section 1; i.e., multilinear optimization defined by Problem BPO_m and pseudo-Boolean optimization defined by Problem BPO_{pB} . We also detail on the relationship between the multilinear set (resp. multilinear polytope) and the pseudo-Boolean set (resp. pseudo-Boolean polytope).

First, it is simple to see that Problem BPO_m is a special case of Problem BPO_{pB} obtained by letting $\eta_s(v) = +1$ for every $s \in S$ and $v \in s$. Similarly, the multilinear set $\mathcal{S}_m(G)$ and the multilinear polytope $\mathcal{P}_m(G)$ are special cases of the pseudo-Boolean set $\mathcal{S}_{pB}(H)$ and of the pseudo-Boolean polytope $\mathcal{P}_{pB}(H)$, respectively. Vice versa, it is also possible to go from the pseudo-Boolean setting to the multilinear setting. In the remainder of the section we discuss this direction, which requires more work and may result in an exponential blowup of the size of the objects considered.

Reformulating Problem BPO_{pB} as Problem BPO_m . Each product term in the objective function of Problem BPO_{pB} can be expanded and written as:

$$\prod_{v \in s} \sigma_s(z_v) = \sum_{e \in E} d_e \prod_{v \in e} z_v + d_0, \quad (1)$$

where the coefficients d_0 and d_e , for $e \in E$, are $0, \pm 1$ and can be easily computed by expanding the product on the left-hand side of (1). To rewrite (1) with explicit coefficients, let $p(s) := \{v \in s : \eta_s(v) = +1\}$ and, for $i = 0, 1, 2, \dots, |s|$, let $N_i(s) := \{t \subseteq s : \eta_s(v) = -1 \ \forall v \in t, |t| = i\}$. We then obtain

$$\prod_{v \in s} \sigma_s(z_v) = \sum_{i=0}^{|s|} \sum_{t \in N_i(s)} (-1)^{|t|} \prod_{v \in p(s) \cup t} z_v. \quad (2)$$

It then follows that Problem BPO_{pB} over a signed hypergraph $H = (V, S)$ can be reformulated (up to a constant in the objective) as Problem BPO_m over a hypergraph, which we call the “multilinear hypergraph” of H and that we define next.

Given a signed hypergraph $H = (V, S)$, the *multilinear hypergraph* of H , denoted by $\text{mh}(H)$, is the hypergraph (V, E) , where E is constructed as follows: for each $s \in S$, and every $t \subseteq s$ with $\eta_s(v) = -1$ for all $v \in t$, the set E contains $\{v \in s : \eta_s(v) = +1\} \cup t$ if it has cardinality at least two. It then follows from the definition of $\text{mh}(H)$ that all sets $p(s) \cup t$ that appear in (2) are edges of $\text{mh}(H)$. We summarize the above discussion in the next observation:

Observation 1. *Problem BPO_{pB} over a signed hypergraph $H = (V, S)$ can be reformulated (up to a constant in the objective) as Problem BPO_m over $\text{mh}(H)$.*

It is important to note that the number of edges in the definition of $\text{mh}(H)$ that arise from a single signed edge $s \in S$ can be *exponential* in the number of nodes contained in s . In fact, the reformulation of Problem BPO_{pB} as Problem BPO_{m} described above may lead to an exponential increase in the size of the problem. One easy example of this is given by the following special case of Problem BPO_{pB} :

$$\begin{aligned} \max \quad & \prod_{v \in V} (1 - z_v) \\ \text{s.t.} \quad & z \in \{0, 1\}^V. \end{aligned}$$

In this problem, S contains only one signed edge $s = (V, \eta_s)$ with $\eta_s(v) = -1$ for all $v \in V$. Once reformulated as Problem BPO_{m} , the resulting objective function contains all possible $2^{|V|} - 1$ monomials. However, one can observe that the change of variable $z'_v := 1 - z_v$, for all $v \in V$, directly transforms this special case of Problem BPO_{pB} to Problem BPO_{m} with only one monomial in the objective function, namely $\prod_{v \in V} z'_v$.

Next, we provide a more involved example that is insensitive to changes of variables of the form $z'_v := 1 - z_v$, for some $v \in V$. Consider a signed hypergraph $H = (V, S)$, where S contains the two signed edges $s = (V, \eta_s)$ and $t = (V, \eta_t)$ such that $\eta_s(v) = -\eta_t(v)$ for every $v \in V$. We then consider the special case of Problem BPO_{pB} given by

$$\begin{aligned} \max \quad & 2 \prod_{v \in s} \sigma_s(z_v) + \prod_{v \in t} \sigma_t(z_v) \\ \text{s.t.} \quad & z \in \{0, 1\}^V. \end{aligned}$$

It is then simple to check that, once reformulated as Problem BPO_{m} , the number of monomials in the resulting objective function is at least $\sum_{i=1}^m \binom{m}{i} = 2^m - 1$, where $m := \lceil |V|/2 \rceil$. This bound holds even if we apply any change of variables of the form $z'_v := 1 - z_v$, for $v \in U \subseteq V$.

The pseudo-Boolean set as a projection of the multilinear set. The pseudo-Boolean set $\mathcal{S}_{\text{pB}}(H)$ of a signed hypergraph $H = (V, S)$ can be obtained as a projection of the multilinear set of $\text{mh}(H)$. Let $\text{mh}(H) = (V, E)$. In the remainder of the section, we denote by x_v , for $v \in V$, the variables corresponding to nodes of H , which coincide with nodes of $\text{mh}(H)$. These variables are used both for the pseudo-Boolean set and the multilinear set. We denote the variables of the multilinear set corresponding to edges in E by y_e , for $e \in E$. The variables of the pseudo-Boolean set corresponding to signed edges in S are instead denoted by z_s , for $s \in S$.

To obtain the pseudo-Boolean set as a projection of the multilinear set, we need to construct the z variables from the x, y variables. This can be done by linearizing equality (2), for every $s \in S$. Specifically, in (2), the left-hand side is replaced with z_s , and each product $\prod_{v \in p(s) \cup t} z_v$ on the right-hand side is replaced with $y_{p(s) \cup t}$ if $|p(s) \cup t| \geq 2$, and with $x_{p(s) \cup t}$ if $|p(s) \cup t| = 1$. We denote by \mathcal{I} this system of linear equalities obtained by linearizing equality (2), for every $s \in S$. The following projection result is then easy to check:

Observation 2. *The pseudo-Boolean set $\mathcal{S}_{\text{pB}}(H)$ is the projection onto the space of the x, z variables of the sets of points $(x, y, z) \in \mathbb{R}^{V \cup E \cup S}$ satisfying the system of linear equalities \mathcal{I} and such that (x, y) lies in the multilinear set of $\text{mh}(H)$.*

The pseudo-Boolean polytope as a projection of the multilinear polytope. The pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ of a signed hypergraph $H = (V, S)$ can be obtained as a projection of the multilinear polytope of $\text{mh}(H)$.

Observation 3. *The pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ is the projection onto the space of the x, z variables of the set of points $(x, y, z) \in \mathbb{R}^{V \cup E \cup S}$ satisfying the system of linear equalities \mathcal{I} and such that (x, y) lies in the multilinear polytope of $\text{mh}(H)$.*

Proof. Denote by Q the set of points $(x, y, z) \in \mathbb{R}^{V \cup E \cup S}$ satisfying the system of linear equalities \mathcal{I} and such that (x, y) lies in the multilinear polytope of $\text{mh}(H)$. Denote by Q' the projection of Q onto the space of variables x, z . The inequalities defining the multilinear polytope of $\text{mh}(H)$ only involve variables x, y , while the linear equalities in \mathcal{I} involve variables x, y, z . In particular, we obtain that Q is a polytope, and therefore Q' is a polytope too. Let (\bar{x}, \bar{z}) be a vertex of Q' . It suffices to show that (\bar{x}, \bar{z}) is in the pseudo-Boolean set $\mathcal{S}_{\text{pB}}(H)$. By definition of Q' , there exists \bar{y} such that $(\bar{x}, \bar{y}, \bar{z})$ is a vertex of Q . Therefore, there exist $|V| + |E| + |S|$ linearly independent constraints among the linear inequalities defining Q such that $(\bar{x}, \bar{y}, \bar{z})$ is the unique vector that satisfies these constraints at equality. Note that the constraints in \mathcal{I} are the only ones containing variables z , each such constraint contains exactly one z variable with nonzero coefficient, and there is precisely one constraint for each z variable. Thus, $|S|$ of the linearly independent constraints defining $(\bar{x}, \bar{y}, \bar{z})$ must be the $|S|$ constraints in \mathcal{I} . The remaining $|V| + |E|$ linearly independent constraints defining $(\bar{x}, \bar{y}, \bar{z})$ must be inequalities defining the multilinear polytope of $\text{mh}(H)$, and these constraints only involve x, y variables. Hence, (\bar{x}, \bar{y}) is a vertex of the multilinear polytope of $\text{mh}(H)$, and hence it is in the multilinear set of $\text{mh}(H)$. It then follows from our previous discussion that (\bar{x}, \bar{z}) is in the pseudo-Boolean set $\mathcal{S}_{\text{pB}}(H)$. \square

3 Decomposability of pseudo-Boolean polytopes

In this section, we present a sufficient condition for decomposability of pseudo-Boolean polytopes that we will use to obtain our extended formulations. Our decomposition result is the first known sufficient condition for decomposability of the pseudo-Boolean polytope. Existing decomposability results for the multilinear polytope are Theorem 1 in [18], Theorem 5 in [17], Theorem 1 in [19], Theorem 4 in [14], and Theorem 4 in [20]. Our decomposition result serves as a significant generalization of Theorem 4 in [20].

Consider a signed hypergraph $H = (V, S)$, let $V_1, V_2 \subseteq V$ such that $V = V_1 \cup V_2$, let $S_1 \subseteq \{s \in S : s \subseteq V_1\}$, $S_2 \subseteq \{s \in S : s \subseteq V_2\}$ such that $S = S_1 \cup S_2$. Let $H_1 := (V_1, S_1)$ and $H_2 := (V_2, S_2)$. In order to define the concept of decomposability we now explain how we can write a vector $z \in \mathbb{R}^{V \cup S}$ by partitioning its components into three sub-vectors. The vector z_\cap contains the components of z corresponding to nodes and signed edges that are both in H_1 and in H_2 . The vector z_1 contains the components of z corresponding to nodes and signed edges in H_1 but not in H_2 . Finally, the vector z_2 contains the components of z corresponding to nodes and signed edges in H_2 but not in H_1 . Using these definitions, we can now write a vector z in the space defined by H as $z = (z_1, z_\cap, z_2)$, up to reordering variables. Similarly, we can write a vector z in the space defined by H_1 as (z_1, z_\cap) , and a vector z in the space defined by H_2 as (z_\cap, z_2) . We now define the concept of decomposability.

Definition 1 (Decomposability). *We say that the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ is decomposable into the pseudo-Boolean polytopes $\mathcal{P}_{\text{pB}}(H_1)$ and $\mathcal{P}_{\text{pB}}(H_2)$ if the following relation holds*

$$(z^1, z^\cap) \in \mathcal{P}_{\text{pB}}(H_1), (z^\cap, z^2) \in \mathcal{P}_{\text{pB}}(H_2) \quad \Rightarrow \quad (z^1, z^\cap, z^2) \in \mathcal{P}_{\text{pB}}(H). \quad (3)$$

We observe that the opposite implication in (3), from right to left, always holds. Furthermore, if $\mathcal{P}_{\text{pB}}(H)$ is decomposable into $\mathcal{P}_{\text{pB}}(H_1)$ and $\mathcal{P}_{\text{pB}}(H_2)$, then the system comprised of a description of $\mathcal{P}_{\text{pB}}(H_1)$ and a description of $\mathcal{P}_{\text{pB}}(H_2)$, is a description of $\mathcal{P}_{\text{pB}}(H)$. We are now ready to present our decomposability result, in Theorem 1 below. An example of signed hypergraphs H, H_1, H_2 that follow the construction of Theorem 1 is given in Figure 1. Throughout this paper, for an integer k , we define $[k] := \{1, \dots, k\}$.




Figure 1: Example of signed hypergraphs H, H_1, H_2 that follow the construction of Theorem 1. The figure depicts only the underlying hypergraphs of H, H_1, H_2 , denoted by G, G_1, G_2 , respectively. The decomposition result holds for any possible signing of the edges of H . Note that if in the edge of cardinality three containing v , all nodes have sign -1 , then v is not a nest point in the multilinear hypergraph of H , and Theorem 4 in [20] cannot be applied.

Theorem 1. *Let $H = (V, S)$ be a signed hypergraph, and assume that the underlying hypergraph of H has a β -leaf v . Let $s_1 \subseteq s_2 \subseteq \dots \subseteq s_k$ be the signed edges of H containing v , and assume that S contains the signed edges $s_i - v$ such that $|s_i - v| \geq 2$, for every $i \in [k]$. Then $\mathcal{P}_{\text{pB}}(H)$ is decomposable into $\mathcal{P}_{\text{pB}}(H_1)$ and $\mathcal{P}_{\text{pB}}(H_2)$, where H_1 and H_2 are defined as follows. $H_1 := (V_1, S_v \cup P_v)$, where V_1 is the underlying edge of s_k , $S_v := \{s_1, \dots, s_k\}$, $P_v := \{s_i - v : |s_i - v| \geq 2, i \in [k]\}$, and $H_2 := H - v$.*

Proof. We assume $k \geq 1$, as otherwise the result is obvious. Let z_1, z_\cap , and z_2 be defined as above, so that we can write a vector z in the space defined by H as $z = (z_1, z_\cap, z_2)$, a vector z in the space defined by H_1 as (z_1, z_\cap) , and a vector z in the space defined by H_2 as $z = (z_\cap, z_2)$. In particular, the vector z_\cap contains the components of z corresponding to nodes in $V_1 \setminus \{v\}$ and signed edges in P_v , while the vector z_1 contains the components of z corresponding to node v and signed edges in S_v .

Let \tilde{z} be such that $(\tilde{z}_1, \tilde{z}_\cap) \in \mathcal{P}_{\text{pB}}(H_1)$ and $(\tilde{z}_\cap, \tilde{z}_2) \in \mathcal{P}_{\text{pB}}(H_2)$. We show $\tilde{z} \in \mathcal{P}_{\text{pB}}(H)$. To do so, we will write \tilde{z} explicitly as a convex combinations of vectors in $\mathcal{S}_{\text{pB}}(H)$. In the next claim, we show how a vector in $\mathcal{S}_{\text{pB}}(H_1)$ and a vector in $\mathcal{S}_{\text{pB}}(H_2)$ can be combined to obtain a vector in $\mathcal{S}_{\text{pB}}(H)$.

Define $p_i := s_i - v$ for every $i \in [k]$. Note that, for some $i = 1, \dots, k$, we might have $|p_i| = 1$, and so $p_i \notin P_v$, meaning that p_i is not a signed edge of H . In these cases, we have that p_i contains a single node, say u , and we call p_i a *signed loop*. Consistently with our notation for signed edges, we will write z_{p_i} and $\sigma_{p_i}(z_u)$ to denote z_u if $\eta_{p_i}(u) = +1$, or to denote $1 - z_u$ if $\eta_{p_i}(u) = -1$.

Claim 1. *Let $(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)$ and $(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2)$ be such that $z_{p_i} = z'_{p_i}$ for every $i \in [k]$. Then, $(z_1, z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H)$.*

Proof of claim. We have $(z_1, z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H)$ if and only if $(z_1, z'_\cap) \in \mathcal{S}_{\text{pB}}(H_1)$ and $(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2)$, thus it suffices to show $(z_1, z'_\cap) \in \mathcal{S}_{\text{pB}}(H_1)$. The signed edges of H_1 are $S_v \cup P_v$. Consider first a signed edge $p_i \in P_v$. In (z_1, z'_\cap) , the component corresponding to the signed edge p_i is in z'_\cap , and the components corresponding to the nodes in p_i are in z'_\cap , thus we need to show the equality $z'_{p_i} = \prod_{u \in p_i} \sigma_{p_i}(z'_u)$, which follows directly from $(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2)$. Consider now a signed edge $s_i \in S_v$. In (z_1, z'_\cap) , the component corresponding to the signed edge s_i is in z_1 , and the components corresponding to the nodes in s_i are in z'_\cap , except for the component corresponding to node v that is in z_1 . Thus, we need to show the equality

$$z_{s_i} = \sigma_{s_i}(z_v) \prod_{u \in s_i \setminus \{v\}} \sigma_{s_i}(z'_u).$$

Since s_i is a signed edge of H_1 and $(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)$, we know

$$z_{s_i} = \sigma_{s_i}(z_v) \prod_{u \in s_i \setminus \{v\}} \sigma_{s_i}(z_u).$$

We then obtain

$$\prod_{u \in s_i \setminus \{v\}} \sigma_{s_i}(z_u) = \prod_{u \in p_i} \sigma_{p_i}(z_u) = z_{p_i} = z'_{p_i} = \prod_{u \in p_i} \sigma_{p_i}(z'_u) = \prod_{u \in s_i \setminus \{v\}} \sigma_{s_i}(z'_u).$$

Here, the first equality holds by definition of p_i for $i \in [k]$; the second equality holds since p_i is a signed edge of H_1 or a signed loop containing a node of H_1 ; in the third equality we use the assumption $z_{p_i} = z'_{p_i}$; the fourth equality holds because p_i is a signed edge of H_2 or a signed loop containing a node of H_2 , and $(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2)$; the last equality holds by definition of p_i for $i \in [k]$. \diamond

In the remainder of the proof, we show how to write explicitly \tilde{z} as a convex combination of the vectors in $\mathcal{S}_{\text{pB}}(H)$ obtained in Claim 1. By assumption, the vector $(\tilde{z}_1, \tilde{z}_\cap)$ is in $\text{conv } \mathcal{S}_{\text{pB}}(H_1)$. Thus, it can be written as a convex combination of points in $\mathcal{S}_{\text{pB}}(H_1)$; i.e., there exists $\mu \geq 0$ such that

$$\begin{aligned} \sum_{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)} \mu_{(z_1, z_\cap)} &= 1 \\ \sum_{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)} \mu_{(z_1, z_\cap)} (z_1, z_\cap) &= (\tilde{z}_1, \tilde{z}_\cap). \end{aligned} \quad (4)$$

Similarly, the vector $(\tilde{z}_\cap, \tilde{z}_2)$ is in $\text{conv } \mathcal{S}_{\text{pB}}(H_2)$ and it can be written as a convex combination of points in $\mathcal{S}_{\text{pB}}(H_2)$; i.e., there exists $\nu \geq 0$ such that

$$\begin{aligned} \sum_{(z_\cap, z_2) \in \mathcal{S}_{\text{pB}}(H_2)} \nu_{(z_\cap, z_2)} &= 1 \\ \sum_{(z_\cap, z_2) \in \mathcal{S}_{\text{pB}}(H_2)} \nu_{(z_\cap, z_2)} (z_\cap, z_2) &= (\tilde{z}_\cap, \tilde{z}_2). \end{aligned} \quad (5)$$

In the remainder of the proof, given $z_{p_1}, \dots, z_{p_k} \in \{0, 1\}$, we will consider the number $\max\{j \in [k] : z_{p_j} = 1\} \in \{0, \dots, k\}$, with the understanding that this number equals 0 when $z_{p_1} = \dots = z_{p_k} = 0$. In the next technical claim, we study the sums of the multipliers μ and ν corresponding to binary vectors z with a fixed $\max\{j \in [k] : z_{p_j} = 1\}$. To do so, we need to introduce the rather technical notation $\text{ext}(i)$, for $i = 0, \dots, k$. For $j, t \in \{1, \dots, k\}$ with $j < t$, we say that p_t is an *extension* of p_j if the common nodes of p_j, p_t have the same sign in p_j and in p_t ; formally, for every $v \in p_j$ we have $\eta_{p_j}(v) = \eta_{p_t}(v)$. We can now define $\text{ext}(0)$ as the set of indices $t \in \{1, \dots, k\}$, such that there is no $j \in \{1, \dots, t-1\}$ such that p_t is an extension of p_j . For $i = 1, \dots, k$, we define $\text{ext}(i)$ as the set of indices $t \in \{i+1, \dots, k\}$, such that p_t is a *minimal extension* of p_i , that is, p_t is an extension of p_i , and there is no $j \in \{i+1, \dots, t-1\}$ such that p_t is an extension of p_j . Note that, if instead there exists $j \in \{i+1, \dots, t-1\}$ such that p_t is an extension of p_j , then p_j is also an extension of p_i . Observe that $\text{ext}(k) = \emptyset$. To familiarize with the definition, note that in the special case where all signed edges s_1, \dots, s_k have all nodes with a +1 sign, i.e., $\eta_{s_i}(v) = +1$ for every $v \in s_i$ and $i = 1, \dots, k$, we have $\text{ext}(0) = \{1\}$ and $\text{ext}(i) = \{i+1\}$, for $i = 1, \dots, k$. We are now ready to present our technical claim.

Claim 2. For $i \in \{0, \dots, k\}$, we have

$$\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \mu_{(z_1, z_\cap)} = \sum_{\substack{(z_\cap, z_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \nu_{(z_\cap, z_2)} = \begin{cases} 1 - \sum_{t \in \text{ext}(0)} \tilde{z}_{p_t} & \text{if } i = 0 \\ \tilde{z}_{p_i} - \sum_{t \in \text{ext}(i)} \tilde{z}_{p_t} & \text{if } i \in \{1, \dots, k\}. \end{cases}$$

Proof of claim. By considering the component of (4) corresponding to p_i , for $i \in [k]$, we obtain

$$\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_i} = 1}} \mu_{(z_1, z_\cap)} = \tilde{z}_{p_i}.$$

We first consider the case $i = k$. We have

$$\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = k}} \mu_{(z_1, z_\cap)} = \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_k} = 1}} \mu_{(z_1, z_\cap)} = \tilde{z}_{p_k} = \tilde{z}_{p_k} - \sum_{t \in \text{ext}(k)} \tilde{z}_{p_t}.$$

Next, we consider the case $i \in \{1, \dots, k-1\}$. We have

$$\begin{aligned} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \mu_{(z_1, z_\cap)} &= \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_i} = 1 \\ z_{p_t} = 0 \ \forall t \in \{i+1, \dots, k\}}} \mu_{(z_1, z_\cap)} \\ &= \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_i} = 1 \\ z_{p_t} = 0 \ \forall t \in \text{ext}(i)}} \mu_{(z_1, z_\cap)} \\ &= \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_i} = 1}} \mu_{(z_1, z_\cap)} - \sum_{t \in \text{ext}(i)} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_t} = 1}} \mu_{(z_1, z_\cap)} \\ &= \tilde{z}_{p_i} - \sum_{t \in \text{ext}(i)} \tilde{z}_{p_t}. \end{aligned}$$

To check the second equality, we assume $(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1)$ with $z_{p_i} = 1$ and $z_{p_t} = 0$ for all $t \in \text{ext}(i)$ and we show that $z_{p_t} = 0$ for all $t \in \{i+1, \dots, k\}$. For a contradiction, assume that there exists $t \in \{i+1, \dots, k\}$ with $z_{p_t} = 1$. Assume that t is the smallest index of this type. Then $t \notin \text{ext}(i)$. Now p_t must be an extension of p_i , as otherwise $z_{p_i} = 1$ implies $z_{p_t} = 0$, a contradiction. Hence, p_t is not a minimal extension of p_i , thus there exists $j \in \{i+1, \dots, t-1\}$ such that p_t is an extension of p_j . By the choice of t we have $z_{p_j} = 0$, which implies $z_{p_t} = 0$, since p_t is an extension of p_j . We have obtained a contradiction.

To check the third equality, we assume $(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1)$ with $z_{p_i} = 1$ and $z_{p_t} = 1$ for some $t \in \text{ext}(i)$ and we show that $z_{p_t} = 1$ for precisely one index $t \in \text{ext}(i)$. For a contradiction, assume there exist two distinct $t, j \in \text{ext}(i)$ with $z_{p_t} = z_{p_j} = 1$. Without loss of generality, assume $j < t$. By definition of $\text{ext}(i)$, both p_j and p_t are extensions of p_i , and p_t is not an extension of p_j . But then $z_{p_j} = 1$ implies $z_{p_t} = 0$, a contradiction.

Lastly, we consider the case $i = 0$. We have

$$\begin{aligned} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = 0}} \mu_{(z_1, z_\cap)} &= \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_t} = 0 \ \forall t \in \{1, \dots, k\}}} \mu_{(z_1, z_\cap)} \\ &= \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_t} = 0 \ \forall t \in \text{ext}(0)}} \mu_{(z_1, z_\cap)} \\ &= \sum_{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1)} \mu_{(z_1, z_\cap)} - \sum_{t \in \text{ext}(0)} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ z_{p_t} = 1}} \mu_{(z_1, z_\cap)} \end{aligned}$$

$$= 1 - \sum_{t \in \text{ext}(0)} \tilde{z}_{pt}.$$

The equalities above follow with arguments similar to those in the previous case.

The statement for ν follows symmetrically, starting with (5) rather than (4). \diamond

For ease of notation, we define, for $i \in \{0, \dots, k\}$, the quantity $m(i)$ to be the sum obtained in Claim 2, i.e.,

$$m(i) := \begin{cases} 1 - \sum_{t \in \text{ext}(0)} \tilde{z}_{pt} & \text{if } i = 0 \\ \tilde{z}_{p_i} - \sum_{t \in \text{ext}(i)} \tilde{z}_{pt} & \text{if } i \in \{1, \dots, k\}. \end{cases}$$

We are now ready to define the multipliers λ that will allow us to write \tilde{z} as a convex combination of vectors in $\mathcal{S}_{\text{pB}}(H)$ obtained in Claim 1. For every $(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)$ and $(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2)$ such that $z_{p_j} = z'_{p_j}$ for every $j \in [k]$, we define

$$\lambda_{(z_1, z'_\cap, z'_2)} := \begin{cases} \frac{\mu_{(z_1, z_\cap)} \nu_{(z'_\cap, z'_2)}}{m(i)} & \text{if } m(i) \neq 0 \\ 0 & \text{if } m(i) = 0, \end{cases}$$

where $i := \max\{j \in [k] : z_{p_j} = 1\} = \max\{j \in [k] : z'_{p_j} = 1\}$. In the next claim, we show that the multipliers λ are nonnegative and sum to one.

Claim 3. *We have $\lambda \geq 0$ and*

$$\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ (z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ z_{p_j} = z'_{p_j} \quad \forall j \in [k]}} \lambda_{(z_1, z'_\cap, z'_2)} = 1.$$

Proof of claim. It follows from Claim 2 that $m(i) \geq 0$ for all $i \in \{0, \dots, k\}$. Thus, using the definition of λ , we obtain $\lambda \geq 0$. Using Claim 2, we obtain

$$\begin{aligned} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ (z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ z_{p_j} = z'_{p_j} \quad \forall j \in [k]}} \lambda_{(z_1, z'_\cap, z'_2)} &= \sum_{i=0, \dots, k} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \lambda_{(z_1, z'_\cap, z'_2)} \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \lambda_{(z_1, z'_\cap, z'_2)} \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \frac{\mu_{(z_1, z_\cap)} \nu_{(z'_\cap, z'_2)}}{m(i)} \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \frac{1}{m(i)} \left(\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \mu_{(z_1, z_\cap)} \right) \left(\sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \nu_{(z'_\cap, z'_2)} \right) \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \frac{(m(i))^2}{m(i)} = \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} m(i) = \sum_{i=0, \dots, k} m(i) = 1. \end{aligned}$$

The last equality $\sum_{i=0, \dots, k} m(i) = 1$ can be seen using the definition of $m(i)$, because each index in $\{1, \dots, k\}$ is exactly in one set among $\text{ext}(0), \dots, \text{ext}(k)$. \diamond

Our last claim, which concludes the proof of the theorem, shows that the multipliers λ yield \tilde{z} as a convex combination of the vectors in $\mathcal{S}_{\text{pB}}(H)$ obtained in Claim 1.

Claim 4. *We have*

$$(\tilde{z}_1, \tilde{z}_\cap, \tilde{z}_2) = \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ (z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ z_{p_j} = z'_{p_j} \quad \forall j \in [k]}} \lambda_{(z_1, z'_\cap, z'_2)}(z_1, z'_\cap, z'_2), \quad (6)$$

Proof of claim. Using the definition of λ , we rewrite (6) in the form

$$\begin{aligned} (\tilde{z}_1, \tilde{z}_\cap, \tilde{z}_2) &= \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ (z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ z_{p_j} = z'_{p_j} \quad \forall j \in [k]}} \lambda_{(z_1, z'_\cap, z'_2)}(z_1, z'_\cap, z'_2) \\ &= \sum_{i=0, \dots, k} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \lambda_{(z_1, z'_\cap, z'_2)}(z_1, z'_\cap, z'_2) \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \lambda_{(z_1, z'_\cap, z'_2)}(z_1, z'_\cap, z'_2) \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \frac{\mu_{(z_1, z_\cap)} \nu_{(z'_\cap, z'_2)}}{m(i)}(z_1, z'_\cap, z'_2). \end{aligned}$$

We now verify the obtained equality, first for components \tilde{z}_1 , and then for components $\tilde{z}_\cap, \tilde{z}_2$. We start with components \tilde{z}_1 . We obtain

$$\begin{aligned} \tilde{z}_1 &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \frac{\mu_{(z_1, z_\cap)} \nu_{(z'_\cap, z'_2)}}{m(i)} z_1 \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \frac{1}{m(i)} \left(\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \mu_{(z_1, z_\cap)} z_1 \right) \left(\sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{pB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \nu_{(z'_\cap, z'_2)} \right) \\ &= \sum_{\substack{i=0, \dots, k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \mu_{(z_1, z_\cap)} z_1 \\ &= \sum_{\substack{i=0, \dots, k \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)} \mu_{(z_1, z_\cap)} z_1 \\ &= \sum_{(z_1, z_\cap) \in \mathcal{S}_{\text{pB}}(H_1)} \mu_{(z_1, z_\cap)} z_1. \end{aligned}$$

Here, the third equality holds due to Claim 2. The fourth equality holds because, from Claim 2, if $m(i) = 0$, then all $\mu_{(z_1, z_\cap)}$ in the second sum equal zero. The obtained equality is implied by (4).

Next, we consider components $\tilde{z}_\cap, \tilde{z}_2$. We obtain

$$\begin{aligned}
(\tilde{z}_\cap, \tilde{z}_2) &= \sum_{\substack{i=0,\dots,k \\ m(i) \neq 0}} \sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{PB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \frac{\mu_{(z_1, z_\cap)} \nu_{(z'_\cap, z'_2)}(z'_\cap, z'_2)}{m(i)} \\
&= \sum_{\substack{i=0,\dots,k \\ m(i) \neq 0}} \frac{1}{m(i)} \left(\sum_{\substack{(z_1, z_\cap) \in \mathcal{S}_{\text{PB}}(H_1) \\ \max\{j \in [k] : z_{p_j} = 1\} = i}} \mu_{(z_1, z_\cap)} \right) \left(\sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{PB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \nu_{(z'_\cap, z'_2)}(z'_\cap, z'_2) \right) \\
&= \sum_{\substack{i=0,\dots,k \\ m(i) \neq 0}} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{PB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \nu_{(z'_\cap, z'_2)}(z'_\cap, z'_2) \\
&= \sum_{i=0,\dots,k} \sum_{\substack{(z'_\cap, z'_2) \in \mathcal{S}_{\text{PB}}(H_2) \\ \max\{j \in [k] : z'_{p_j} = 1\} = i}} \nu_{(z'_\cap, z'_2)}(z'_\cap, z'_2) \\
&= \sum_{(z'_\cap, z'_2) \in \mathcal{S}_{\text{PB}}(H_2)} \nu_{(z'_\cap, z'_2)}(z'_\cap, z'_2).
\end{aligned}$$

Here, the third equality holds due to Claim 2. The fourth equality holds because, from Claim 2, if $m(i) = 0$, then all $\nu_{(z'_\cap, z'_2)}$ in the second sum equal zero. The obtained equality is implied by (5). \diamond

We should remark that the overall structure of the proof of Theorem 1 is similar to that of Theorem 4 in [20]. The difference lies in the key construction of the proof, that is, the multipliers introduced to construct the convex combination. Due to the presence of the signed edges, these new multipliers are significantly more involved than the ones presented in the proof of Theorem 4 in [20]. In the special case where all signs are positive, the new multipliers simplify to the ones in the proof of Theorem 4 in [20].

4 Pointed signed hypergraphs

The signed hypergraph H_1 defined in the statement of Theorem 1 plays a key role in our convex hull characterizations. In this section, we prove that the pseudo-Boolean polytope $\mathcal{P}_{\text{PB}}(H_1)$ has a polynomial-size extended formulation. Together with the decomposition technique of Theorem 1, this result enables us to obtain polynomial-size extended formulations for the pseudo-Boolean polytope of a large family of signed hypergraphs.

Definition 2 (Pointed signed hypergraph). *Consider a signed hypergraph $H = (V, S)$ and let $v \in V$ be a β -leaf of the underlying hypergraph of H . Denote by S_v the set of all signed edges in S containing v . Define $P_v := \{s - v : s \in S_v, |s| \geq 3\}$. We say that H is pointed at v (or is a pointed signed hypergraph) if V coincides with the underlying edge of the signed edge of maximum cardinality in S_v and $S = S_v \cup P_v$.*

From this definition it follows that the signed hypergraph H_1 in the statement of Theorem 1 is pointed at v (the hypergraph G_1 in Figure 1 is an example of the underlying hypergraph of a pointed signed hypergraph). In order to characterize the pseudo-Boolean polytope of a signed hypergraph H pointed at v , we first characterize the pseudo-Boolean polytope of a simpler type of signed hypergraphs corresponding to faces of $\mathcal{P}_{\text{PB}}(H)$ defined by $z_v = 0$ or $z_v = 1$. An extended formulation for $\mathcal{P}_{\text{PB}}(H)$ can then be obtained using disjunctive programming.

4.1 The pseudo-Boolean polytope of nested signed hypergraphs

In this section, we characterize the pseudo-Boolean polytope of nested signed hypergraphs in the original space. This result will then enable us to obtain a polynomial-size extended formulation for the pseudo-Boolean polytope of pointed signed hypergraphs.

Definition 3 (Nested signed hypergraph). *Let $H = (V, S)$ be a signed hypergraph with $V = \{v_1, \dots, v_n\}$. Denote by E the set of underlying edges of S (note that several signed edges may collapse to the same underlying edge). Define $e_k := \{v_1, \dots, v_{k+1}\}$ for all $k \in [n-1]$ and $\bar{E} := \{e_1, \dots, e_{n-1}\}$. We say that H is a nested signed hypergraph if it satisfies the following conditions:*

(N1) $E = \bar{E}$,

(N2) for any signed edge $s \in S$ with underlying edge e_k for some $k \in [n-1]$, the following two signed edges are also present in S : $\ell(s)$, obtained from s by flipping the sign of v_{k+1} , and $p(s) := s - v_{k+1}$.

Remark 1. *Nested signed hypergraphs have two important properties that we will use to obtain our extended formulations:*

1. *As we will prove in Proposition 1, the pseudo-Boolean polytope of a nested signed hypergraph $H = (V, S)$ in the original space is defined by $\frac{|S|}{2} + 1$ equalities, and $|S| + 2|V| - 4$ inequalities; i.e., $2(|S| + |V| - 1)$ inequalities.*
2. *Let $H = (V, S)$ be a signed hypergraph such that for any $s, s' \in S$ we either have $s \subseteq s'$ or $s' \subseteq s$. It then follows that by adding at most $2|S|(|V| - 2) + |S|$ signed edges to H we obtain a nested signed hypergraph. To see this, note that to satisfy property (N2) of nested signed hypergraphs, for each edge $s \in S$ we need to add at most $2(|s| - 2) + 1$ signed edges to H . Together with case 1 above, this in turn implies that $\mathcal{P}_{\text{PB}}(H)$ has a polynomial-size extended formulation with at most $2|S|(|V| - 1) + |V|$ variables and at most $4|S|(|V| - 1) + 2|V|$ inequalities.*

We now proceed with characterizing the pseudo-Boolean polytope of nested signed hypergraphs. To this end, we first introduce some notation. We denote by \mathcal{E}_k , $k \in [n-1]$, the set consisting of signed edges in S whose underlying edges are e_k . For each $k \in [n-1] \setminus \{1\}$, we define

$$\mathcal{E}_k^+ := \left\{ s \in \mathcal{E}_k : \eta_s(v_{k+1}) = +1 \right\}, \quad \text{and} \quad \mathcal{E}_k^- := \left\{ s \in \mathcal{E}_k : \eta_s(v_{k+1}) = -1 \right\}.$$

For any $s \in S$, we define

$$N(s) := \{s' \in S : s = p(s')\},$$

where as before for any $s' \in S$ with underlying edge e_k , we define $p(s') = s' - v_{k+1}$. Notice that by property (N2) of nested signed hypergraphs, $|N(s)|$ equals zero or two.

The following proposition characterizes the pseudo-Boolean polytope of nested signed hypergraphs in the original space. Due to its length and technical nature, we have deferred the proof to the appendix; i.e., Section 7. While our proof relies on a standard disjunctive programming technique [1] followed by a projection step using Fourier-Motzkin elimination, the novelty of the proof lies in the manner Fourier-Motzkin elimination is implemented. Namely, it is well-understood that a generic application of Fourier-Motzkin elimination leads to a rapid increase in the number of inequalities defining the polyhedron. In our proof, the auxiliary variables are projected out in a specific order so that the projection does not contain redundant inequalities.

Proposition 1. *Let $H = (V, S)$ with $n := |V|$ be a nested signed hypergraph. Suppose that $\mathcal{E}_1 = \{q_1, q_2, q_3, q_4\}$ with $\eta_{q_1}(v_1) = \eta_{q_1}(v_2) = \eta_{q_2}(v_1) = \eta_{q_3}(v_2) = +1$ and $\eta_{q_2}(v_2) = \eta_{q_3}(v_1) = \eta_{q_4}(v_1) = \eta_{q_4}(v_2) = -1$. Then the pseudo-Boolean polytope $\mathcal{P}_{\text{PB}}(H)$ is given by:*

$$z_s + z_{\ell(s)} = z_{p(s)}, \quad \forall s \in \mathcal{E}_k^+, k \in [n-1] \setminus \{1\} \quad (7)$$

$$z_s \geq 0, \quad \forall s \in S \quad (8)$$

$$\sum_{s \in \mathcal{E}_k^+} z_s \leq z_{v_{k+1}}, \quad \forall k \in [n-1] \setminus \{1\} \quad (9)$$

$$\sum_{s \in \mathcal{E}_k^-} z_s \leq 1 - z_{v_{k+1}}, \quad \forall k \in [n-1] \setminus \{1\} \quad (10)$$

$$z_{q_1} + z_{q_2} = z_{v_1}, \quad z_{q_1} + z_{q_3} = z_{v_2}, \quad z_{q_3} + z_{q_4} = 1 - z_{v_1}. \quad (11)$$

Remark 2. In Proposition 1, the assumption on the structure of \mathcal{E}_1 is not restrictive. Recall that by property (N2) of a nested signed hypergraph we must either have $|\mathcal{E}_1| = 2$ or $|\mathcal{E}_1| = 4$. Moreover, if $|\mathcal{E}_1| = 2$, then we must either have $\mathcal{E}_1 = \{q_1, q_3\}$ or $\mathcal{E}_1 = \{q_2, q_4\}$. Without loss of generality, suppose that $\mathcal{E}_1 = \{q_1, q_3\}$. Then to obtain a description of $\mathcal{P}_{\text{pB}}(H)$ in the original space, it suffices to project out z_{q_2}, z_{q_4} from system (7)–(11). To do so, notice that z_{q_2}, z_{q_4} appear only in the following constraints of system (7)–(11):

$$\begin{aligned} z_{q_2} &\geq 0, \quad z_{q_4} \geq 0 \\ z_{q_1} + z_{q_2} &= z_{v_1}, \quad z_{q_3} + z_{q_4} = 1 - z_{v_1}. \end{aligned}$$

Projecting out z_{q_2}, z_{q_4} , we obtain:

$$z_{q_1} \leq z_{v_1}, \quad z_{q_3} \leq 1 - z_{v_1},$$

which together with the remaining equalities and inequalities of system (7)–(11) gives a description of $\mathcal{P}_{\text{pB}}(H)$ in the original space.

It is important to note that, in spite of its simple structure, the constraint matrix of the pseudo-Boolean polytope of a nested signed hypergraph is *not* totally unimodular. The following example demonstrates this fact. In this example, for ease of notation, we write a signed edge in a compact form by listing its nodes with their signs as superscripts. For example, the signed edge $s = (e, \eta_s)$, where $e = \{v_1, v_2, v_3\}$ and $\eta_s(v_1) = +1$, $\eta_s(v_2) = -1$, $\eta_s(v_3) = +1$, will be written compactly as $s = \{v_1^+, v_2^-, v_3^+\}$.

Example 1. Consider the nested signed hypergraph $H = (V, S)$, with $V = \{v_1, v_2, v_3, v_4\}$, $\mathcal{E}_1 = \{s_1, s_2, s_3, s_4\}$, $s_1 = \{v_1^+, v_2^+\}$, $s_2 = \{v_1^-, v_2^+\}$, $s_3 = \{v_1^+, v_2^-\}$, $s_4 = \{v_1^-, v_2^-\}$, $\mathcal{E}_2 = \{s_5, s_6, s_7, s_8\}$, $s_5 = \{v_1^+, v_2^+, v_3^+\}$, $s_6 = \{v_1^+, v_2^+, v_3^-\}$, $s_7 = \{v_1^-, v_2^+, v_3^-\}$, $s_8 = \{v_1^-, v_2^+, v_3^+\}$, $\mathcal{E}_3 = \{s_9, s_{10}, s_{11}, s_{12}\}$, $s_9 = \{v_1^+, v_2^+, v_3^-, v_4^+\}$, $s_{10} = \{v_1^+, v_2^+, v_3^-, v_4^-\}$, $s_{11} = \{v_1^-, v_2^+, v_3^+, v_4^+\}$, $s_{12} = \{v_1^-, v_2^+, v_3^+, v_4^-\}$. Then by Proposition 1, the following constraints are present in the description of $\mathcal{P}_{\text{pB}}(H)$:

$$\begin{aligned} z_{s_5} + z_{s_6} &= z_{s_1} \\ z_{s_9} + z_{s_{10}} &= z_{s_6} \\ z_{s_{11}} + z_{s_{12}} &= z_{s_8} \\ z_{s_5} + z_{s_8} &\leq z_{v_3} \\ z_{s_9} + z_{s_{11}} &\leq z_{v_4} \end{aligned}$$

Consider the submatrix of the above equalities and inequalities corresponding to $z_{s_5}, z_{s_6}, z_{s_8}, z_{s_9}, z_{s_{11}}$. It can be checked that the determinant of this submatrix equals 2, implying the constraint matrix of $\mathcal{P}_{\text{pB}}(H)$ is not totally unimodular. In addition, each row and each column of this submatrix has two non-zero entries, implying it is not a balanced matrix either [11].

4.2 The pseudo-Boolean polytope of pointed signed hypergraphs

Thanks to the compact description for the pseudo-Boolean polytope of nested signed hypergraphs given by Proposition 1, we now provide a polynomial-size extended formulation for the pseudo-Boolean polytope of pointed signed hypergraphs. Recall that in a signed hypergraph $H = (V, S)$ pointed at \bar{v} for some $\bar{v} \in V$, we have $S = S_{\bar{v}} \cup P_{\bar{v}}$; the set $S_{\bar{v}}$ consists of all signed edges containing \bar{v} , and the set $P_{\bar{v}}$ has the form $P_{\bar{v}} = \{s - \bar{v} : s \in S_{\bar{v}}, |s| \geq 3\}$.

Remark 3. Consider a pointed signed hypergraph $H = (V, S)$. There are two special cases for which a polynomial-size description of $\mathcal{P}_{\text{pB}}(H)$ follows from previously known results:

- (i) Suppose that for each $v \in V$ we have $\eta_s(v) = \eta_{s'}(v)$ for all signed edges s, s' in H containing v . Then (possibly after a one-to-one linear transformation) the explicit description of $\mathcal{P}_{\text{pB}}(H)$ in the original space consisting of at most $5|V| + 2$ inequalities is given in Theorem 5 of [20].
- (ii) Suppose that all signed edges in $S_{\bar{v}}$ are parallel. It then follows that for any $s, s' \in S$ we either have $s \subseteq s'$ or $s' \subseteq s$. Then by Part 2 of Remark 1 and definition of $P_{\bar{v}}, S_{\bar{v}}$, the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation with at most $|S|(|V| - 1) + |V|$ variables and at most $2(|S|(|V| - 1) + |V|)$ inequalities.

The following theorem gives a polynomial-size extended formulation for the pseudo-Boolean polytope of general pointed signed hypergraphs.

Theorem 2. Let $H = (V, S)$ be a pointed signed hypergraph. Then the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation with at most $2|V|(|S| + 1)$ variables and at most $4(|S|(|V| - 2) + |V|)$ inequalities. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(H)$ are $0, \pm 1$.

Proof. Suppose that $H = (V, S)$ is pointed at \bar{v} for some $\bar{v} \in V$, implying that \bar{v} is a β -leaf of the underlying hypergraph of H and $S = S_{\bar{v}} \cup P_{\bar{v}}$. Define the signed hypergraph $L = (V \setminus \{\bar{v}\}, P_{\bar{v}})$. By definition of $P_{\bar{v}}$, for any two edges $s, s' \in P_{\bar{v}}$ we either have $s \subseteq s'$ or $s' \subseteq s$. Hence after the addition of at most $2|P_{\bar{v}}|(|V| - 3) + |P_{\bar{v}}|$ edges to L we can construct a nested signed hypergraph denoted by L' . Denote by $P'_{\bar{v}}$ the set of signed edges of L' . Note that $|P'_{\bar{v}}| = 2|P_{\bar{v}}|(|V| - 2) \leq |S|(|V| - 2)$. Define $H' := (V, S_{\bar{v}} \cup P'_{\bar{v}})$. Since a description of $\mathcal{P}_{\text{pB}}(H')$ serves as an extended formulation for $\mathcal{P}_{\text{pB}}(H)$, to complete the proof, it suffices to show that $\mathcal{P}_{\text{pB}}(H')$ has a polynomial-size extended formulation. Denote by H^0 (resp. H^1) the signed hypergraph corresponding to the face of $\mathcal{P}_{\text{pB}}(H')$ with $z_{\bar{v}} = 0$ (resp. $z_{\bar{v}} = 1$). We then have:

$$\mathcal{P}_{\text{pB}}(H') = \text{conv} \left(\mathcal{P}_{\text{pB}}(H^0) \cup \mathcal{P}_{\text{pB}}(H^1) \right).$$

Denote by \bar{z} the vector consisting of $z_v, v \in V \setminus \{\bar{v}\}$, and $z_s, s \in P'_{\bar{v}}$. It then follows that:

$$\begin{aligned} \mathcal{P}_{\text{pB}}(H^0) &= \left\{ z \in \mathbb{R}^{V \cup S_{\bar{v}} \cup P'_{\bar{v}}} : z_{\bar{v}} = 0, z_s = 0, \forall s \in S_{\bar{v}} \text{ with } \eta_s(\bar{v}) = +1, z_s = z_{p(s)}, \forall s \in S_{\bar{v}} \text{ with} \right. \\ &\quad \left. \eta_s(\bar{v}) = -1, \bar{z} \in \mathcal{P}_{\text{pB}}(L'_{\bar{v}}) \right\} \\ \mathcal{P}_{\text{pB}}(H^1) &= \left\{ z \in \mathbb{R}^{V \cup S_{\bar{v}} \cup P'_{\bar{v}}} : z_{\bar{v}} = 1, z_s = z_{p(s)}, \forall s \in S_{\bar{v}} \text{ with } \eta_s(\bar{v}) = +1, z_s = 0, \forall s \in S_{\bar{v}} \text{ with} \right. \\ &\quad \left. \eta_s(\bar{v}) = -1, \bar{z} \in \mathcal{P}_{\text{pB}}(L'_{\bar{v}}) \right\}. \end{aligned}$$

By Proposition 1, the polytope $\mathcal{P}_{\text{pB}}(L'_{\bar{v}})$ is given by system (7)-(11) and this description has at most $2(|S|(|V| - 2) + |V|)$ inequalities. For notational simplicity, let us write $\mathcal{P}_{\text{pB}}(L'_{\bar{v}})$ compactly as

$$\mathcal{P}_{\text{pB}}(L'_{\bar{v}}) = \{\bar{z} : A\bar{z} \leq b, C\bar{z} = d\}.$$

Then, using Balas' formulation for the union of polytopes [1], we obtain a polynomial-size extended formulation for $\mathcal{P}_{\text{pB}}(H')$:

$$\mathcal{P}_{\text{pB}}(H') = \left\{ z \in \mathbb{R}^{V \cup S_{\bar{v}} \cup P'_{\bar{v}}} : \exists (z^0, z^1, z, \lambda) \text{ s.t. } z = z^0 + z^1, \lambda_0 + \lambda_1 = 1, z_{\bar{v}}^0 = 0, \right. \\ z_s^0 = 0, \forall s \in S_{\bar{v}} \text{ with } \eta_s(\bar{v}) = +1, z_s^0 = z_{p(s)}^0, \forall s \in S_{\bar{v}} \text{ with } \eta_s(\bar{v}) = -1, \\ z_{\bar{v}}^1 = \lambda_1, z_s^1 = z_{p(s)}^1, \forall s \in S_{\bar{v}} \text{ with } \eta_s(\bar{v}) = +1, z_s^1 = 0, \forall s \in S_{\bar{v}} \text{ with } \eta_s(\bar{v}) = -1, \\ \left. Az^0 \leq b\lambda_0, Cz^0 = d\lambda_0, Az^1 \leq b\lambda_1, Cz^1 = d\lambda_1, \lambda_0, \lambda_1 \geq 0 \right\}.$$

The size of the above extended formulation can be further reduced by projecting out variables $\lambda_0, \lambda_1, z^0$ using the equalities $z_{\bar{v}} = \lambda_1, \lambda_0 + \lambda_1 = 1$, and $z = z^0 + z^1$. Hence, we obtain an extended formulation for $\mathcal{P}_{\text{pB}}(H')$ with at most

$$2(|V| + |S_{\bar{v}}| + |P'_{\bar{v}}|) = 2(|V| + |S_{\bar{v}}| + 2|P_{\bar{v}}|(|V| - 2)) = 2(|V| + |S| + |P_{\bar{v}}|(2|V| - 5)) \\ \leq 2(|V| + |S|(|V| - \frac{3}{2})) \leq 2|V|(|S| + 1), \quad (12)$$

variables and at most

$$2(2(|S|(|V| - 2) + |V|)) \leq 4(|S|(|V| - 2) + |V|), \quad (13)$$

inequalities. The second equality in (12) follows from $|S| = |S_{\bar{v}}| + |P_{\bar{v}}|$ and the inequality in (12) follows from $|P_{\bar{v}}| \leq \frac{|S|}{2}$.

By Proposition 1, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(I'_{\bar{v}})$ are $0, \pm 1$. This together with the fact that $\lambda_0 = 1 - z_{\bar{v}}$ and $\lambda_1 = z_{\bar{v}}$ implies that the same statement holds for the extended formulation of $\mathcal{P}_{\text{pB}}(H')$ and as a result for the extended formulation of $\mathcal{P}_{\text{pB}}(H)$ as well. \square

5 Inflation of signed edges

Our decomposition result stated in Theorem 1 relies on the key assumption that the underlying hypergraph of H has at least one β -leaf. We argue that in many cases of interest, this restrictive assumption can be removed. To this end, in the following, we introduce an operation on signed hypergraphs, which we refer to as the *inflation of signed edges*. Starting from a signed hypergraph H whose underlying hypergraph does not contain any β -leaves, by inflating certain edges, we obtain a new signed hypergraph H' whose underlying hypergraph has a sequence of β -leaves. By relating the extended formulations of $\mathcal{P}_{\text{pB}}(H)$ and $\mathcal{P}_{\text{pB}}(H')$, we are able to obtain polynomial-size extended formulations for the pseudo-Boolean polytope of various classes of signed hypergraphs whose underlying hypergraphs *contain* β -cycles.

Definition 4 (Inflation operation). *Let $H = (V, S)$ be a signed hypergraph, let $s \in S$, and let $e \subseteq V$ such that $s \subset e$. Denote by $I(s, e)$ the set of all possible signed edges s' parallel to e such that $\eta_s(v) = \eta_{s'}(v)$ for every $v \in s$. We say that $H' = (V, S')$ is obtained from H by inflating s to e if $S' = S \cup I(s, e) \setminus \{s\}$. We also say that H' is obtained from H via an inflation operation.*

We refer the reader to Figure 2 for an example of the inflation operation. The next theorem indicates that if an extended formulation for $\mathcal{P}_{\text{pB}}(H')$ is available, one can obtain an extended formulation for $\mathcal{P}_{\text{pB}}(H)$ as well.




Figure 2: The underlying hypergraph of a signed hypergraph $H = (V, S)$ is shown on the left. The underlying hypergraph of the signed hypergraph H' obtained from H by inflating all signed edges in S to V is shown on the right.

Theorem 3. *Let $H = (V, S)$ be a signed hypergraph, let $s \in S$, and let $e \subseteq V$ such that $s \subseteq e$. Let $H' = (V, S')$ be obtained from H by inflating s to e . Then an extended formulation of $\mathcal{P}_{\text{pB}}(H)$ can be obtained by juxtaposing an extended formulation of $\mathcal{P}_{\text{pB}}(H')$ and the equality constraint*

$$z_s = \sum_{s' \in I(s,e)} z_{s'}. \quad (14)$$

Moreover, if $\mathcal{P}_{\text{pB}}(H')$ has a polynomial-size extended formulation and $|e| - |s| = O(\log \text{poly}(|V|, |S|))$, then $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation as well.

Proof. Given a signed hypergraph H and a set C , we denote by $\text{proj}_H(C)$ the orthogonal projection of the set C onto the subspace of variables corresponding to nodes and signed edges of H . Let $A(z, y) \leq b$ be an extended formulation of $\mathcal{P}_{\text{pB}}(H')$. Let $H'' = (V, S'')$, where $S'' := S \cup I(s, e) = S' \cup \{s\}$. It suffices to prove that $A(z, y) \leq b$ together with (14) is an extended formulation of $\mathcal{P}_{\text{pB}}(H'')$. That is, we need to show that

$$Q := \text{proj}_{H''}\{(z, y) : A(z, y) \leq b, (14)\} = \mathcal{P}_{\text{pB}}(H'').$$

Since $S \subseteq S''$, it follows that an extended formulation of $\mathcal{P}_{\text{pB}}(H'')$ is also an extended formulation of $\mathcal{P}_{\text{pB}}(H)$. For every binary z , we have $z \in \mathcal{P}_{\text{pB}}(H'')$ if and only if

$$z_s = \prod_{v \in s} \eta_s(z_v) = \prod_{v \in s} \eta_s(z_v) \prod_{v \in e \setminus s} (z_v + (1 - z_v)) = \sum_{s' \in I(s,e)} z_{s'}. \quad (15)$$

Thus, it suffices to show that the vertices of Q are binary. We have

$$\begin{aligned} Q &= \text{proj}_{H''}\{(z, y) : A(z, y) \leq b\} \cap \{z : (14)\} \\ &= \text{proj}_{H'}\{(z, y) : A(z, y) \leq b\} \cap \{z : (14)\} \\ &= \mathcal{P}_{\text{pB}}(H') \cap \{z : (14)\}, \end{aligned} \quad (16)$$

where the first equality holds since all variables that appear with nonzero coefficients in (14) correspond to signed edges in H'' , the second equality holds because $\{s\} = S'' \setminus S'$ and so the variable z_s does not appear in the system $A(z, y) \leq b$, and the third equality holds since $A(z, y) \leq b$ is an the extended formulation of $\mathcal{P}_{\text{pB}}(H')$.

Let \bar{z} be a vertex of Q . The variable \bar{z}_s is only present in (14), thus we have $\bar{z}_s = \sum_{s' \in I(s,e)} \bar{z}_{s'}$. The other constraints of the system defining Q that are active at \bar{z} are all from the projection of the system $A(z, y) \leq b$, thus they determine a vertex of $\mathcal{P}_{\text{pB}}(H')$, which is binary. This implies that all components of \bar{z} , except for \bar{z}_s , are binary. It then follows from (15) that \bar{z}_s is binary as well.

Finally, since $|I(s, e)| \leq 2^{|e| - |s|}$, we conclude that, $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation if $\mathcal{P}_{\text{pB}}(H')$ has a polynomial-size extended formulation and $|e| - |s| = O(\log \text{poly}(|V|, |S|))$. \square

The inflation operation is of independent interest as it enables us to obtain polynomial-size extended formulations for the pseudo-Boolean polytope of certain signed hypergraphs. The following proposition is an illustration of this fact:

Proposition 2. *Consider a signed hypergraph $H = (V, S)$. Suppose that each $s \in S$ contains at least $|V| - k$ nodes. Then $\mathcal{P}_{\text{pB}}(H)$ has an extended formulation with $O(2^k |V| |S|)$ variables and inequalities. In particular, if $k = O(\log \text{poly}(|V|, |S|))$, then $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(H)$ are $0, \pm 1$.*

Proof. Consider the signed hypergraph $H' = (V, S')$ obtained from H by inflating every $s \in S$ with $|s| < |V|$ to V . We then have $|S'| \leq 2^k |S|$. By Part 2 of Remark 1, the polytope $\mathcal{P}_{\text{pB}}(H')$ has a polynomial-size extended formulation with at most $2^{k+1} |S| (|V| - 1) + |V|$ variables and at most $2^{k+2} |S| (|V| - 1) + 2|V|$ inequalities. Therefore, from Lemma 3 we deduce that an extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is obtained by adding at most $|S|$ equality constraints (containing $|S|$ additional variables) to the extended formulation of $\mathcal{P}_{\text{pB}}(H')$. Since all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(H')$ as well as in equalities (14) are $0, \pm 1$, the statement holds for the system defining $\mathcal{P}_{\text{pB}}(H)$. \square

Figure 2 illustrates the proof idea in Proposition 2. That is, starting from a signed hypergraph H whose underlying edges have large cardinalities, one can inflate all edges to their union, and subsequently use Part (ii) of Remark 3 to obtain a polynomial-size extended formulation for the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$.

6 The recursive inflate-and-decompose framework

With the objective of constructing extended formulations for the pseudo-Boolean polytope, by combining our decomposition scheme of Theorem 1, our convex hull characterization of Theorem 2, and the inflation operation of Theorem 3, we introduce a new framework, which we refer to as the *recursive inflate-and-decompose* framework. We show that this framework enables us to obtain polynomial-size extended formulations for the pseudo-Boolean polytope of large families of signed hypergraphs.

As we detail in this section, our new framework unifies and extends all existing results on polynomial-size extended formulations for the convex hull of the feasible region of Problem (ℓ BPO_m) [5, 17, 19, 20, 33, 38]. In the following, we present our recursive inflate-and-decompose framework.

The Recursive inflate-and-decompose (RID) framework

Input. A signed hypergraph $H = (V, S)$.

Output. An extended formulation for $\mathcal{P}_{\text{pB}}(H)$.

Step 0. Set $H^{(0)} := H$, $i := 0$.

Step 1. If we can obtain $\bar{H}^{(i)}$ from $H^{(i)}$ via a number of inflation operations, such that a suitable extended formulation for $\mathcal{P}_{\text{pB}}(\bar{H}^{(i)})$ is available, then by Theorem 3 we are done. Otherwise, proceed to Step 2.

Step 2. Choose a node \bar{v} of $H^{(i)}$. If \bar{v} is a β -leaf of the underlying hypergraph of $H^{(i)}$, then set $\bar{H}^{(i)} := H^{(i)}$ and proceed to Step 3. Otherwise, construct $\bar{H}^{(i)}$ from $H^{(i)}$ via inflation operations, such that \bar{v} is a β -leaf of the underlying hypergraph of $\bar{H}^{(i)}$. By Theorem 3, it suffices to find an extended formulation for $\mathcal{P}_{\text{pB}}(\bar{H}^{(i)})$.

Step 3. Using Theorem 1, decompose $\mathcal{P}_{\text{pB}}(\bar{H}^{(i)})$ into $\mathcal{P}_{\text{pB}}(\bar{H}_1^{(i)})$ and $\mathcal{P}_{\text{pB}}(\bar{H}_2^{(i)})$, where $\bar{H}_1^{(i)}$ denotes the signed hypergraph containing node \bar{v} . Since an extended formulation for $\mathcal{P}_{\text{pB}}(\bar{H}_1^{(i)})$ is given by Theorem 2, it suffices to find an extended formulation for $\mathcal{P}_{\text{pB}}(\bar{H}_2^{(i)})$. Set $H^{(i+1)} := \bar{H}_2^{(i)}$, increment i by one, and go to Step 1.

Clearly, Steps 1 and 2 of the RID framework can be performed in many different ways. That is, we have not specified how node \bar{v} in Step 2 should be selected or how the inflation operations of Step 1 and Step 2 should be performed. A simple way to obtain a β -leaf in Step 2 is to inflate each signed edge containing \bar{v} to the union of all signed edges containing \bar{v} . While the RID framework is fairly general, we are naturally interested in specifying conditions under which this framework results in a polynomial-size extended formulation for the pseudo-Boolean polytope. To this end, first note that at each iteration, one node of the signed hypergraph $H^{(i)}$ is removed. Hence, the number of iterations of the RID framework is upper bounded by the number of nodes of H . It then follows that RID provides a polynomial-size extended formulation for $\mathcal{P}_{\text{pB}}(H)$ if the following conditions are satisfied:

- (I) In Step 1, the algorithm should terminate, only if a polynomial-size extended formulation for $\mathcal{P}_{\text{pB}}(\bar{H}^{(i)})$ is available.
- (II) The total number of new edges introduced as a result of inflation operations in Steps 1 and 2 is upper bounded by a polynomial in $|V|, |S|$.

In the remainder of this section, we consider various types of signed hypergraphs for which one can customize the RID framework so that conditions (I)–(II) above are satisfied and hence polynomial-size extended formulations for the corresponding pseudo-Boolean polytope can be constructed.

Our results are stated in terms of easily verifiable conditions on the structure of the underlying hypergraphs, namely, β -acyclic hypergraphs, α -acyclic hypergraphs with log-poly ranks, and certain classes of hypergraphs with log-poly “gaps”. In each case, we first provide a self contained proof without any reference to the RID framework. Subsequently, we explain how each proof can be obtained via a specific implementation of RID.

6.1 β -acyclic hypergraphs

Consider a signed hypergraph $H = (V, S)$. In this section, we show that if the underlying hypergraph of H has a sequence of β -leaves, then $\mathcal{P}_{\text{pB}}(H)$ can be described in terms of pseudo-Boolean polytopes of simpler signed hypergraphs. In particular, if the underlying hypergraph of H is β -acyclic, then $\mathcal{P}_{\text{pB}}(H)$ has an extended formulation of size polynomial in $|V|, |S|$. These theorems serve as significant generalizations of the main results in [20].

Theorem 4. *Let $H = (V, S)$ be a signed hypergraph of rank r , and let v_1, \dots, v_t for some $t \geq 1$ be a sequence of β -leaves of the underlying hypergraph of H . Then an extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is given by a description of $\mathcal{P}_{\text{pB}}(H - v_1 - \dots - v_t)$ together with a system of at most $4rt(2|S| + 1)$ inequalities consisting of at most $2rt(2|S| + 1)$ variables. Moreover, all coefficients and right-hand side constants in the latter are $0, \pm 1$.*

Proof. Denote by S_{v_1} the set of all signed edges of H containing v_1 . Since v_1 is a β -leaf of the underlying hypergraph of H , the set S_{v_1} is totally ordered. Define the signed hypergraph $H'_1 := (V, S'_1)$ with $S'_1 = S \cup P_{v_1}$, where $P_{v_1} := \{s - v_1 : s \in S_{v_1}, |s| \geq 3\}$. Clearly, an extended formulation for $\mathcal{P}_{\text{pB}}(H'_1)$ serves as an extended formulation for $\mathcal{P}_{\text{pB}}(H)$ as well. Now define the pointed signed hypergraph $H_{v_1} := (V_1, S_{v_1} \cup P_{v_1})$, where V_1 denotes the underlying edge of a signed edge of maximum cardinality in S_{v_1} . We then have $H'_1 = H_{v_1} \cup (H - v_1)$, where we used the fact that $H'_1 - v_1 = H - v_1$. Hence by Theorem 1, the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H'_1)$ is decomposable into pseudo-Boolean polytopes $\mathcal{P}_{\text{pB}}(H_{v_1})$ and $\mathcal{P}_{\text{pB}}(H - v_1)$.

Next consider the signed hypergraph $H - v_1$. By definition v_2 is a β -leaf of the underlying hypergraph of $H - v_1$. Denote by S_{v_2} the set of signed edges $H - v_1$ containing v_2 . Define $S'_2 := (S \setminus S_{v_1}) \cup P_{v_2}$, where $P_{v_2} := \{s - v_2 : s \in S_{v_2}, |s| \geq 3\}$, and define the signed hypergraph $H'_2 := (V \setminus \{v_1\}, S'_2)$. Again, an extended formulation for $\mathcal{P}_{\text{pB}}(H'_2)$ serves as an extended formulation for $\mathcal{P}_{\text{pB}}(H - v_1)$ as well.

Define the pointed signed hypergraph $H_{v_2} := (V_2, S_{v_2} \cup P_{v_2})$, where V_2 denotes the underlying edge of a signed edge of maximum cardinality in S_{v_2} . Then by Theorem 1, $\mathcal{P}_{\text{pB}}(H_2)$ is decomposable into $\mathcal{P}_{\text{pB}}(H_{v_2})$ and $\mathcal{P}_{\text{pB}}(H - v_1 - v_2)$.

By a recursive application of the above argument after t times, we conclude that an extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is given by a description of $\mathcal{P}_{\text{pB}}(H - v_1 - \dots - v_t)$ together with a system of inequalities defining $\mathcal{P}_{\text{pB}}(H_{v_i})$ for all $i \in \{1, \dots, t\}$, where H_{v_i} is a signed hypergraph pointed at v_i defined as $H_{v_i} := (V_i, S_{v_i} \cup P_{v_i})$, where S_{v_i} denotes the set of signed edge of $H - v_1 - \dots - v_{i-1}$ (we define $H - v_1 - v_0 = H$) containing v_i , $P_{v_i} := \{s - v_i : s \in S_{v_i}, |s| \geq 3\}$, and V_i denotes the underlying edge of a signed edge of maximal cardinality in S_{v_i} .

By Theorem 2 the polytopes $\mathcal{P}_{\text{pB}}(H_{v_i})$, $i \in \{1, \dots, t\}$ have polynomial-size extended formulations each of which consists of at most $2|V_i|(|S_{v_i}| + |P_{v_i}| + 1) \leq 2r(2|S| + 1)$ variables and at most $4(|S_{v_i}| + |P_{v_i}|)(|V_i| - 2) + 4|V_i| \leq 8(r - 2)|S| + 4r \leq 4r(2|S| + 1)$ inequalities, where the inequalities follow from $|V_i| \leq r$ and $|P_{v_i}| \leq |S_{v_i}| \leq |S|$. \square

Remark 4. *The proof of Theorem 4 follows from the RID framework. To see this, first note that the node \bar{v} chosen in Step 2 of the i th iteration of the RID framework is the β -leaf v_i for all $i \in [t]$. Moreover, at iteration $t + 1$, the algorithm terminates at Step 1 with a description of $\mathcal{P}_{\text{pB}}(H - v_1 - \dots - v_t)$. In this case, no inflation operation at any step of the RID is needed.*

Theorem 5. *Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is β -acyclic. Then the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation with at most $2r(2|S| + 1)|V|$ variables, and at most $4r(2|S| + 1)|V|$ inequalities. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(H)$ are $0, \pm 1$.*

Proof. Since the underlying hypergraph of H is β -acyclic, it has a sequence of β -leaves of length $|V|$. The proof then follows directly from Theorem 4. \square

Remark 5. *Let $G = (V, E)$ be a β -acyclic hypergraph of rank r . Theorem 1 in [20] gives an extended formulation for the multilinear polytope of G with at most $(r - 1)|V|$ variables and at most $(3r - 4)|V| + 4|E|$ inequalities. Theorem 5 is a significant generalization of this result as it only requires the β -acyclicity of the underlying hypergraph of H . Indeed the multilinear hypergraph $\text{mh}(H)$ may contain many β -cycles. Nonetheless, this generalization has a cost: while the size of the extended formulation for $\mathcal{P}_{\text{m}}(G)$ is quadratic in $|V|, |E|$, the size of the extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is cubic in $|V|, |S|$.*

The next example illustrates the power of Theorem 5.

Example 2. *Consider the signed hypergraph $H = (V, S)$ with $V = \{v_1, \dots, v_n\}$ for some $n \geq 4$, and $S = \{s_1, s_2, s_3, s_4\}$, where $s_i = (e_i, \eta_{s_i})$ for $i \in \{1, \dots, 4\}$, $e_1 = e_2 = \{v_1, \dots, v_{n-1}\}$, $e_3 = e_4 = \{v_2, \dots, v_n\}$, $\eta_{s_1}(v) = 1$ for $v \in \{v_1, \dots, v_{n-2}\}$, $\eta_{s_1}(v_{n-1}) = -1$, $\eta_{s_2}(v) = 1$ for $v \in \{v_1, v_3, \dots, v_{n-1}\}$, $\eta_{s_2}(v_2) = -1$, $\eta_{s_3}(v) = 1$ for $v \in \{v_2, \dots, v_{n-2}, v_n\}$, $\eta_{s_3}(v_{n-1}) = -1$, $\eta_{s_4}(v) = 1$ for $v \in \{v_3, \dots, v_n\}$, and $\eta_{s_1}(v_2) = -1$. It then follows that the underlying hypergraph of H is β -acyclic and hence a polynomial-size extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is given by Theorem 5. It can be checked that the multilinear hypergraph of H is given by $\text{mh}(H) = (V, E)$ with $E = \{f_1, \dots, f_6\}$, where $f_1 = f_2 = V$, $f_3 = \{v_1, \dots, v_{n-2}\}$, $f_4 = \{v_1, v_3, \dots, v_{n-1}\}$, $f_5 = \{v_2, \dots, v_{n-2}, v_n\}$, and $f_6 = \{v_3, \dots, v_n\}$. It then follows that $\text{mh}(H)$ contains a β -cycle of length four given by $C = v_1 f_3 v_2 f_5 v_n f_6 v_{n-1} f_4 v_1$. Therefore, Theorem 1 in [20] cannot be used to obtain an extended formulation for the multilinear polytope of $\text{mh}(H)$. See Figure 3 for an illustration.*




Figure 3: The hypergraphs of Example 2 with $n = 4$. The underlying hypergraph of the signed hypergraph H , which is β -acyclic is shown on the left. The multilinear hypergraph of H that contains a β -cycle of length four is shown on the right.

6.2 α -acyclic hypergraphs with log-poly ranks

As we mentioned before, α -acyclic hypergraphs are the most general type of acyclic hypergraphs. In [15], the authors prove that unless $P = NP$, one cannot construct, in polynomial time, a polynomial-size extended formulation for the multilinear polytope of α -acyclic hypergraphs. In [5, 33, 38], the authors give extended formulations for the convex hull of the feasible set of (possibly constrained) multilinear optimization problems. For the unconstrained case, their result can be equivalently stated as follows (see Theorem 5 in [19]): If $G = (V, E)$ is an α -acyclic hypergraph of rank r such that $r = O(\log \text{poly}(|V|, |E|))$, then $\mathcal{P}_m(G)$ has a polynomial-size extended formulation. In the following we show that this result follows from Theorems 1–3; that is, it is a special case of the RID framework.

Theorem 6. *Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is α -acyclic. Then $\mathcal{P}_{\text{pB}}(H)$ has an extended formulation with at most $(\frac{2}{3}3^r + 2(r-1)(2^r + 1))|V|$ variables and inequalities. In particular, if $k = O(\log \text{poly}(|V|, |S|))$, then $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(H)$ are $0, \pm 1$.*

Proof. By definition of α -acyclic hypergraphs, the underlying hypergraph of H has a sequence of α -leaves of length $n := |V|$. Let us denote this sequence by v_1, v_2, \dots, v_n . Since v_1 is an α -leaf of the underlying hypergraph of H , there exists a maximal (for inclusion) signed edge \bar{s} of H containing v_1 . Let H' be obtained from H by first inflating each signed edge s' containing v_1 such that $s' \subset \bar{s}$ to \bar{s} , and then by adding all signed edges $s - v_1$ for all signed edges containing v_1 . Denote by N_{v_1} the number of signed edges $s \in S \setminus \{\bar{s}\}$ containing v_1 such that $s \subset \bar{s}$. By Theorem 3, it suffices to find an extended formulation of $\mathcal{P}_{\text{pB}}(H')$. An extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is then obtained by juxtaposing the extended formulation for $\mathcal{P}_{\text{pB}}(H')$ and N_{v_1} equalities containing at most N_{v_1} additional variables. Since H is a rank- r signed hypergraph we have:

$$N_{v_1} \leq N_{\max} = 2 \left(\sum_{k=1}^{r-2} \binom{r-1}{k} 2^{r-1-k} \right) = 2(3^{r-1} - 2^{r-1} - 1).$$

Moreover, by construction, v_1 is a β -leaf of the underlying hypergraph of H' , implying we can apply the decomposition result of Theorem 1. Namely, let V_1 be the underlying edge of \bar{s} , let S_{v_1} be the set of signed edges of H' parallel to \bar{s} , and let $P_{v_1} := \{s - v_1 : |s - v_1| \geq 2\}$. Define the signed hypergraphs $H_{v_1} := (V_1, S_{v_1} \cup P_{v_1})$ and $H'_1 = H' - v_1$. Then, by Theorem 1, $\mathcal{P}_{\text{pB}}(H')$ is decomposable into $\mathcal{P}_{\text{pB}}(H_{v_1})$ and $\mathcal{P}_{\text{pB}}(H'_1)$. An extended formulation for the pseudo-Boolean polytope of H_{v_1} can then be obtained from Part 2 of Remark 1.

Next, we show that we can apply the above construction recursively to $H' - v_1$ with sequence of α -leaves v_2, \dots, v_n . Note that the underlying hypergraph of $H' - v_1$ is obtained from the underlying

hypergraph of $H - v_1$ by possibly removing some edges contained in $\bar{s} - v_1$. First, note that the rank of $H' - v_1$ is at most r , as removing nodes and edges from a hypergraph cannot increase its rank. Next, we show that v_2, \dots, v_n is a sequence of α -leaves of the underlying hypergraph of $H' - v_1$. For $k \in \{1, \dots, n\}$, we show that v_k is an α -leaf of the underlying hypergraph of $H'_k := H' - v_1 - v_2 - \dots - v_{k-1}$. Since v_k is an α -leaf of the underlying hypergraph of $H_k := H - v_1 - v_2 - \dots - v_{k-1}$, let \hat{s} be a signed edge of H_k containing v_k that is maximal for inclusion. Let \tilde{s} be a signed edge of H such that $\hat{s} = \tilde{s} - v_1 - v_2 - \dots - v_{k-1}$. If \tilde{s} is also a signed edge of H' we are done. Otherwise, by definition of H' , we have $\tilde{s} \subseteq \bar{s}$. But then $\bar{s}' := \bar{s} - v_1 - v_2 - \dots - v_{k-1}$ is a signed edge of H'_k and we have $\hat{s} \subseteq \bar{s}'$. Therefore, \bar{s}' is a signed edge of H'_k containing v_k that is maximal for inclusion. Hence, we can apply the above construction recursively. The proof follows by induction on n .

Define the signed hypergraph $H_{v_k} = (V_k, S_{v_k} \cup P_{v_k})$, $k \in \{1, \dots, n\}$, where S_{v_k} is a set of parallel signed edges containing v_k , and $P_{v_k} := \{s - v_k : |s - v_k| \geq 2\}$. Moreover, V_k is the underlying edge of a signed edge in S_{v_k} . By the above argument an extended formulation for $\mathcal{P}_{\text{pB}}(H)$ is obtained by juxtaposing extended formulations of $\mathcal{P}_{\text{pB}}(H_{v_k})$ for all $k \in \{1, \dots, n\}$ and at most $N_{\max}|V|$ equalities containing at most $N_{\max}|V|$ additional variables.

Finally, consider the pointed signed hypergraph H_{v_k} for some $k \in \{1, \dots, n\}$. Since the rank of H is r , we have $|S_{v_k}| \leq 2^r$. Moreover, since by construction all signed edges in S_{v_k} are parallel, by Part (ii) of Remark 3 we conclude that $\mathcal{P}_{\text{pB}}(H_{v_k})$ has an extended formulation with at most $(r-1)2^r + r$ variables and at most $2((r-1)2^r + r)$ inequalities. Therefore, the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ has an extended formulation with at most

$$((r-1)2^r + r)|V| + 2(3^{r-1} - 2^{r-1} - 1)|V| = \left(\frac{2}{3}3^r + (r-2)(2^r + 1)\right)|V|,$$

variables, and at most

$$(2(r-1)2^r + 2r)|V| + 2(3^{r-1} - 2^{r-1} - 1)|V| \leq \left(\frac{2}{3}3^r + 2(r-1)(2^r + 1)\right)|V|$$

inequalities. □

Remark 6. *The proof of Theorem 6 follows from the RID framework. To see this, first note that the node \bar{v} chosen in Step 2 of the i th iteration of the RID framework is the α -leaf v_i for all $i \in [n]$. Since by definition of α -leaves each v_i is contained in a maximal edge \bar{s} , the inflation operation of Step 2 is defined as follows: inflate all edges containing v_i to \bar{s} . Since H is a rank- r hypergraph, from the proof of Theorem 6 it follows that the total number of edges added due to inflation operations is upper bounded by $2^r|V|$, which is a polynomial in $|V|, |S|$ if we have $r = O(\log \text{poly}(|V|, |S|))$.*

6.3 Hypergraphs with log-poly gaps

Consider a hypergraph $G = (V, E)$, and let $E' \subseteq E$; we define the *gap* of E' as

$$\text{gap}(E') := \max \left\{ \left| \bigcup_{f \in E'} f \right| - |e| : e \in E' \right\}.$$

Moreover, for a signed hypergraph $H = (V, S)$, and $S' \subseteq S$, the *gap* of S' , denoted by $\text{gap}(S')$ is defined as the gap of the set of the underlying edges of $s \in S'$.

Consider the signed hypergraph $H = (V, S)$ defined in the statement of Proposition 2. It can be checked that the gap of S is upper bounded by k . Then, from the proof of Proposition 2 it follows that, if $\text{gap}(S) = O(\log \text{poly}(|V|, |S|))$, the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation. In this section, we strengthen this result by providing polynomial-size extended

formulations for the pseudo-Boolean polytope of signed hypergraphs $H = (V, S)$ with $\text{gap}(S') = O(\log \text{poly}(|V|, |S|))$ for only certain (small) subsets $S' \subset S$.

In the following, given a signed hypergraph $H = (V, S)$, we say that $s \in S$ is a *maximal signed edge* if there is no $s' \in S$ with $s' \supset s$.

Proposition 3. *Consider a signed hypergraph $H = (V, S)$ of rank r . For each $s \in S$, among all maximal signed edges of H containing s , denote by f_s one with minimum cardinality. Let k be such that*

$$\text{gap}(\{s, f_s\}) \leq k, \quad \forall s \in S. \quad (17)$$

Denote by \bar{S} the set of maximal signed edges of H . If the underlying hypergraph of (V, \bar{S}) is β -acyclic, then $\mathcal{P}_{\text{PB}}(H)$ has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities. In particular, if $k = O(\log \text{poly}(|V|, |S|))$, then $\mathcal{P}_{\text{PB}}(H)$ has a polynomial-size extended formulation. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{PB}}(H)$ are $0, \pm 1$.

Proof. Denote by $H' = (V, S')$ the signed hypergraph obtained from H by inflating every non-maximal signed edge $s \in S$ to f_s , where f_s is defined in the statement. From (17) it follows that $|S'| \leq 2^k(|S| - |\bar{S}|) + |\bar{S}| \leq 2^k|S|$. Notice that the underlying hypergraph of H' coincides with the underlying hypergraph of (V, \bar{S}) , which by assumption is β -acyclic. Hence by Theorem 5, $\mathcal{P}_{\text{PB}}(H')$ has an extended formulation with at most $2r|V|(2^{k+1}|S| + 1)$ variables, and at most $4r|V|(2^{k+1}|S| + 1)$ inequalities. Therefore, from Lemma 3 we deduce that an extended formulation for $\mathcal{P}_{\text{PB}}(H)$ is obtained by adding at most $|S| - |\bar{S}|$ equalities consisting of at most $|S| - |\bar{S}|$ additional variables to $\mathcal{P}_{\text{PB}}(H')$, and this completes the proof. \square

The next example illustrates the power of Proposition 3.

Example 3. *Consider the hypergraph $G = (V, E)$ with $V = \{v_1, \dots, v_{2n}\}$ for some $n \geq 3$ and $E = \{e_1, \dots, e_n\} \cup \{f_1, \dots, f_n\} \cup \{g_1, \dots, g_n\}$, where $e_1 = \{v_1, v_2, v_3\}$, $e_i = \{v_{2i-2}, v_{2i-1}, v_{2i}, v_{2i+1}\}$ for all $i \in \{2, \dots, n-1\}$, $e_n = \{v_{2n-2}, v_{2n-1}, v_{2n}\}$, $f_1 = \{v_1, v_2\}$, $f_i = \{v_{2i-2}, v_{2i}\}$ for all $i \in \{2, \dots, n\}$, $g_i = \{v_{2i-1}, v_{2i+1}\}$ for all $i \in \{1, \dots, n-1\}$ and $g_n = \{v_{2n-1}, v_{2n}\}$. See Figure 4 for an illustration. It can be checked that G has a β -cycle of length $2n$ given by*

$$C = v_1, f_1, v_2, f_2, v_4, \dots, v_{2n-2}, f_n, v_{2n}, g_n, v_{2n-1}, \dots, v_5, g_2, v_3, g_1, v_1.$$

We now use Proposition 3 to obtain a polynomial-size extended formulation for $\mathcal{P}_m(G)$. To this end, notice that $\text{gap}(\{f_1, e_1\}) = \text{gap}(\{g_1, e_1\}) = \text{gap}(\{f_n, e_n\}) = \text{gap}(\{g_n, e_n\}) = 1$ and $\text{gap}(\{f_i, e_i\}) = \text{gap}(\{g_i, e_i\}) = 2$ for all $i \in \{2, \dots, n-1\}$. Moreover hypergraph (V, \bar{E}) , where $\bar{E} = \{e_1, \dots, e_n\}$ is β -acyclic. Therefore, by proof of Proposition 3, by inflating each f_i, g_i to e_i for all $i \in \{1, \dots, n\}$, one obtains a signed hypergraph $H = (V, S)$ whose underlying hypergraph is β -acyclic and $|S| \leq 4|E|$. Therefore, by Theorem 5, $\mathcal{P}_{\text{PB}}(H)$ and as a result $\mathcal{P}_m(G)$ have polynomial-size extended formulations.

We should remark that while G is not β -acyclic, it is an α -acyclic hypergraph with rank $r = 4$. Hence, Theorem 6 also provides (a different) polynomial-size extended formulation for $\mathcal{P}_m(G)$. However, it is simple to modify this example so that G will no longer have a log-poly rank and hence Theorem 6 will no longer be applicable. Denote by U_i a set of n nodes such that $U_i \cap V = \emptyset$ for all $i \in [n-1]$ and $U_i \cap U_j = \emptyset$ for all $i \neq j \in [n-1]$. Define the hypergraph $G' = (V', E')$, with $V' = \bigcup_{i \in [n-1]} U_i \cup V$ and $E' = \{e'_1, \dots, e'_n\} \cup \{f'_1, \dots, f'_n\} \cup \{g'_1, \dots, g'_n\}$, where $e'_1 = e_1 \cup U_1$, $f'_1 = f_1 \cup U_1$, $g'_1 = g_1 \cup U_1$, $e'_i = e_i \cup U_{i-1} \cup U_i$, $f'_i = f_i \cup U_{i-1} \cup U_i$, $g'_i = g_i \cup U_{i-1} \cup U_i$ for $i \in \{2, \dots, n-1\}$ and $e'_n = e_n \cup U_{n-1}$, $f'_n = f_n \cup U_{n-1}$, $g'_n = g_n \cup U_{n-1}$. Then the rank of G' is larger than $2n$ and hence Theorem 6 is no longer applicable. However, the gaps remain unchanged; i.e., $\text{gap}(\{f_i, e_i\}) = \text{gap}(\{f'_i, e'_i\})$ and $\text{gap}(\{g_i, e_i\}) = \text{gap}(\{g'_i, e'_i\})$ for all $i \in [n]$ and the hypergraph (V', \bar{E}) is β -acyclic, implying that by Proposition 3 one can obtain a polynomial-size extended formulation for $\mathcal{P}_m(G')$.




Figure 4: The hypergraph $G = (V, E)$ of Example 3; G contains a long β -cycle of length $|V|$.

Let $G = (V, E)$ be a hypergraph and let $C = v_1, e_1, v_2, e_2, \dots, v_q, e_q, v_1$ for some $q \geq 3$ be a β -cycle of G . Then the *support hypergraph* of C is the hypergraph $G[C] = (V[C], E[C])$, where $V[C] = \bigcup_{i=1}^q e_i$ and $E[C] = \{e_1, \dots, e_q\}$. The next proposition essentially indicates that if edge sets of the support hypergraphs of β -cycles of the underlying hypergraph of H have log-poly gaps, then one can *remove* these β -cycles by inflating the edges of the cycle, and obtain a polynomial-size extended formulation for $\mathcal{P}_{\text{PB}}(H)$.

Proposition 4. *Consider a signed hypergraph $H = (V, S)$ of rank r , and denote by $G = (V, E)$ its underlying hypergraph. Let \mathcal{C} denote the set of all β -cycles in G , and for each $C \in \mathcal{C}$, let $G[C] = (V[C], E[C])$ be the support hypergraph of C . Let \tilde{G} be the hypergraph $(\bigcup_{C \in \mathcal{C}} V[C], \bigcup_{C \in \mathcal{C}} E[C])$, let $(V_1, E_1), (V_2, E_2), \dots, (V_\omega, E_\omega)$ for some $\omega \geq 1$ be the connected components of \tilde{G} , and let k be such that*

$$\text{gap}(E_i) \leq k, \quad \forall i \in [\omega]. \quad (18)$$

Then, $\mathcal{P}_{\text{PB}}(H)$ has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities. In particular, if $k = O(\log \text{poly}(|V|, |S|))$, then $\mathcal{P}_{\text{PB}}(H)$ has a polynomial-size extended formulation. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{PB}}(H)$ are $0, \pm 1$.

Proof. Each β -cycle in G is contained in precisely one of the connected components of \tilde{G} . Let H' be obtained from H by inflating, for each $j \in \{1, 2, \dots, \omega\}$, every signed edge s contained in V_j , to V_j . Let G' be the underlying hypergraph of H' . We show that G' contains no β -cycle. Assume for a contradiction that G' contains a β -cycle C' . Consider first the case where C' does not contain any set V_j as an edge. Then C' is a β -cycle in G not contained in any connected component of \tilde{G} , which gives us a contradiction. Next, consider the case where C' contains at least one set V_j as an edge. We now show how we can modify C' to obtain a β -cycle C in G , not contained in any connected component of \tilde{G} , which gives us again a contradiction. Assume that C' contains the edge V_j , and let u, v be the nodes in C' before and after V_j . We then define C by replacing, in C' , the edge V_j with a minimal sequence $f_1, w_2, \dots, w_t, f_t$, where f_j are edges in E_j and w_j are nodes in V_j , such that $u \in f_1$ and $v \in f_t$. Note that this sequence exists because (V_j, E_j) is connected. Clearly, nodes w_2, \dots, w_t and edges f_1, \dots, f_t did not belong to C' . It is simple to check that applying recursively the above construction to all sets among V_1, \dots, V_ω contained in C' yields a β -cycle in C not contained in any connected component of \tilde{G} , which gives us a contradiction. We have therefore shown that G' is β -acyclic.

By Lemma 3 and assumption (18), an extended formulation for $\mathcal{P}_{\text{PB}}(H)$ can be obtained by juxtaposing an extended formulation for $\mathcal{P}_{\text{PB}}(H')$ together with at most $\bigcup_{C \in \mathcal{C}} |E[C]| \leq |S|$ equalities, containing at most $\bigcup_{C \in \mathcal{C}} |E[C]| \leq |S|$ additional variables. Now consider the signed hypergraph $H' = (V, S')$. We then have $|S'| \leq |S| - \sum_{i \in [\omega]} |E_i| + \sum_{i \in [\omega]} 2^k |E_i| \leq 2^k |S|$. Since the underlying hypergraph of H' is β -acyclic, by Corollary 5, $\mathcal{P}_{\text{PB}}(H')$ has an extended formulation with at most $2r|V|(2^{k+1}|S| + 1)$ variables, and at most $4r|V|(2^{k+1}|S| + 1)$ inequalities. \square

Remark 7. Under more restrictive assumptions, the removal of β -cycles can be performed in a more efficient manner than the technique in the proof of Proposition 4. Two β -cycles are equivalent if one can be obtained from the other via a circular permutation and/or reversing the order of the edges. Now consider the case where all β -cycles in each (V_i, E_i) , $i \in [\omega]$, are equivalent. Also notice that each β -cycle is contained in precisely one of the connected components (V_i, E_i) . For each $i \in [\omega]$, denote by f_i an edge of maximum cardinality in E_i . By the definition of a β -cycle, each node of the cycle should be present in exactly two edges of the cycle. Therefore, by inflating f_i to V_i , we can eliminate all β -cycles in this connected component. Let k be such that

$$\text{gap}(\{f_i, V_i\}) \leq k, \quad \forall C \in \mathcal{C}. \quad (19)$$

Then one can replace assumption (18) by the weaker assumption (19)

Let $C = v_1, e_1, v_2, e_2, \dots, v_q, e_q, v_1$ be a β -cycle of length q for some $q \geq 3$ in G . Define $\delta = \max_{i \in \{1, \dots, q\}} |e_i|$. We then have $\text{gap}(E[C]) \geq q - \delta$. Therefore, Proposition 4 (or its enhancement described in Remark 7) is not applicable for hypergraphs containing long β -cycles consisting of edges of small cardinality. For example, the hypergraph G of Example 3 does not satisfy the assumptions of Proposition 4, since the gap of the edge set of its β -cycle equals $2n - 2$, which in turn equals $|V| - 2$. The following example illustrates the usefulness of Proposition 4.

Example 4. Consider a hypergraph $G = (V, E)$ with $V = \{v_1, \dots, v_n\}$ for some $n \geq 12$ and

$$E = \left\{ V, \{v_1, v_2\}, \{v_2, v_3\}, \{v_4, v_5\}, \{v_5, v_6\}, \{v_4, v_7\}, \{v_6, v_7\}, \{v_4, v_5, v_6\}, \{v_8, v_9\}, \{v_9, v_{10}\}, \{v_8, v_{10}\}, \right. \\ \left. \{v_9, v_{11}, v_{12}\}, \{v_{10}, v_{11}, v_{12}\} \right\}.$$

See Figure 5 for an illustration. It then follows that the hypergraph \tilde{G} defined in Proposition 4 has two connected components: (V_1, E_1) where $V_1 = \{v_4, v_5, v_6, v_7\}$, $E_1 = \{e \in E : e \subseteq V_1\}$, and (V_2, E_2) where $V_2 = \{v_8, v_9, v_{10}, v_{11}, v_{12}\}$, $E_2 = \{e \in E : e \subseteq V_2\}$. We then have $\text{gap}(E_1) = 2$ and $\text{gap}(E_2) = 3$. Hence the assumptions of Proposition 4 are satisfied and by inflating every edge in E_1 to V_1 and by inflating every edge in E_2 to V_2 , we obtain a signed hypergraph H whose underlying hypergraph is β -acyclic. Therefore, by Theorem 5, $\mathcal{P}_{\text{pB}}(H)$ and as a result $\mathcal{P}_{\text{m}}(G)$ have polynomial-size extended formulations.

Notice that the hypergraph of Example 4 does not satisfy the assumptions of Proposition 3 as we have $\text{gap}(\{s, f_s\}) = |V| - 2$ for $s = \{v_1, v_2\}$. That is, Propositions 3 and 4 are rather complementary and an effective combination of the two can be used to obtain polynomial-size extended formulations for many signed hypergraphs whose underlying hypergraphs contain β -cycles.




Figure 5: The hypergraph G of Example 4 with $n = 12$ is shown on the left. The β -acyclic hypergraph obtained from G after inflation operations is shown on the right.

Remark 8. *Proofs of Propositions 3 and 4 follow from Step 1 of the RID framework: given a signed hypergraph $H = (V, S)$, via a number of inflation operations, we obtain the hypergraph $H' = (V, S')$ such that:*

- (i) *the underlying hypergraph of H' is β -acyclic, hence, by Theorem 5, the polytope $\mathcal{P}_{\text{pB}}(H')$ has a polynomial-size extended formulation in $|V|, |S'|$,*
- (ii) *$|S'| \leq |S| + \text{poly}(|V|, |S|)$, where this inequality is satisfied because of the log-poly gaps.*

6.4 A more general framework

So far, all inflation operations we employed to obtain polynomial-size extended formulations (i.e., Theorem 6, and Propositions 2–4) have a simple form: we inflated a number of signed edges to the same set. This in turn enabled us to use Theorems 1 and 2 to obtain our extended formulations. As we detail next, the inflation of signed edges can be useful in a more general setting. Consider a signed hypergraph $H = (V, S)$. Assume that S is not totally ordered. Let \mathcal{E} denote a set of subsets of V that is totally ordered. Moreover suppose that for each $s \in S$, there exists some $e \in \mathcal{E}$ such that $e \supseteq s$. Now inflate each $s \in S$ to some element in \mathcal{E} , which we denote by $e(s)$. Define

$$I(S, \mathcal{E}) := \bigcup_{s \in S} I(s, e(s)).$$

Define the signed hypergraph $H' = (V, I(S, \mathcal{E}))$. By Part 2 of Remark 1, the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(H')$ has a polynomial-size extended formulation in $|V|, |I(S, \mathcal{E})|$. Therefore, from Lemma 3 we deduce that if $|I(S, \mathcal{E})|$ is a polynomial in $|V|, |S|$, we can obtain a polynomial-size extended formulation for $\mathcal{P}_{\text{pB}}(H)$ as well. The following proposition is an illustration of how this general setting can be used to obtain a polynomial-size extended formulation for the pseudo-Boolean polytope.

Proposition 5. *Consider a signed hypergraph $H = (V, S)$. Let $S = S_1 \cup S_2$, such that each $s \in S_1$ contains at least $|V| - k_1$ nodes, whereas each $s \in S_2$ is contained in $U \subset V$, where $|U| = k_2$. Then $\mathcal{P}_{\text{pB}}(H)$ has an extended formulation with $O(2^{\max\{k_1, k_2\}} |S| |V|)$ variables and inequalities. In particular, if $\max\{k_1, k_2\} = O(\log \text{poly}(|V|, |S|))$, then $\mathcal{P}_{\text{pB}}(H)$ has a polynomial-size extended formulation. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\text{pB}}(H)$ are $0, \pm 1$.*

Proof. Consider the signed hypergraph $H' = (V, S')$ obtained from H by inflating every $s \in S_1$ with $|s| < |V|$ to V , and by inflating every $s \in S_2$ with $|s| < |U|$ to U . We then have $|S'| \leq 2^{k_1} |S_1| + 2^{k_2} |S_2| \leq 2^{\max\{k_1, k_2\}} |S|$. The remainder of the proof is identical to the proof of Proposition 2. \square

We conclude the paper by remarking that while we presented several classes of signed hypergraphs for which our proposed RID framework provides polynomial-size extended formulations for the corresponding pseudo-Boolean polytopes, a complete characterization of such signed hypergraphs remains an open question and is a subject of future research.

7 Appendix

Proof of Proposition 1. The proof is by induction on the number of nodes n of H . The base case is $n = 2$; in this case we have $S = \mathcal{E}_1 = \{q_1, q_2, q_3, q_4\}$. It is simple to check that $\mathcal{P}_{\text{pB}}(H)$ is given by:

$$\begin{aligned} z_{q_1} + z_{q_2} &= z_{v_1}, \quad z_{q_1} + z_{q_3} = z_{v_2}, \quad z_{q_3} + z_{q_4} = 1 - z_{v_1} \\ z_{q_1}, z_{q_2}, z_{q_3}, z_{q_4} &\geq 0. \end{aligned}$$

Henceforth, let $n \geq 3$. Denote by $\mathcal{P}_{\text{pB}}^0(H)$ (resp. $\mathcal{P}_{\text{pB}}^1(H)$) the face of $\mathcal{P}_{\text{pB}}(H)$ defined by $z_{v_n} = 0$ (resp. $z_{v_n} = 1$). We then have:

$$\mathcal{P}_{\text{pB}}(H) = \text{conv} \left(\mathcal{P}_{\text{pB}}^0(H) \cup \mathcal{P}_{\text{pB}}^1(H) \right). \quad (20)$$

Define the signed hypergraph $\bar{H} := (\bar{V}, \bar{S})$ with $\bar{V} := V \setminus \{v_n\}$ and $\bar{S} := S \setminus \mathcal{E}_{n-1}$. Since \bar{H} has one fewer node than H and is a nested signed hypergraph, by the induction hypothesis, the pseudo-Boolean polytope $\mathcal{P}_{\text{pB}}(\bar{H})$ is given by:

$$\begin{aligned} z_s + z_{\ell(s)} &= z_{p(s)}, \quad \forall s \in \mathcal{E}_k^+, k \in [n-2] \setminus \{1\} \\ z_s &\geq 0, \quad \forall s \in S \setminus \mathcal{E}_{n-1} \\ \sum_{s \in \mathcal{E}_k^+} z_s &\leq z_{v_{k+1}}, \quad \forall k \in [n-2] \setminus \{1\} \\ \sum_{s \in \mathcal{E}_k^-} z_s &\leq 1 - z_{v_{k+1}}, \quad \forall k \in [n-2] \setminus \{1\} \\ z_{q_1} + z_{q_2} &= z_{v_1}, \quad z_{q_1} + z_{q_3} = z_{v_2}, \quad z_{q_3} + z_{q_4} = 1 - z_{v_1}. \end{aligned}$$

Denote by \bar{z} the vector consisting of z_v for all $v \in \bar{V}$, and z_s for all $s \in \bar{S}$. It then follows that

$$\begin{aligned} \mathcal{P}_{\text{pB}}^0(H) &= \left\{ z \in \mathbb{R}^{V \cup S} : z_{v_n} = 0, z_s = 0, \forall s \in \mathcal{E}_{n-1}^+, z_s = z_{p(s)}, \forall e \in \mathcal{E}_{n-1}^-, \bar{z} \in \mathcal{P}_{\text{pB}}(\bar{H}) \right\} \\ \mathcal{P}_{\text{pB}}^1(H) &= \left\{ z \in \mathbb{R}^{V \cup S} : z_{v_n} = 1, z_s = z_{p(s)}, \forall e \in \mathcal{E}_{n-1}^+, z_s = 0, \forall s \in \mathcal{E}_{n-1}^-, \bar{z} \in \mathcal{P}_{\text{pB}}(\bar{H}) \right\}. \end{aligned}$$

By (20) and using Balas' formulation for the union of polytopes [1], it follows that $\mathcal{P}_{\text{pB}}(H)$ is the projection onto the space of the $z_v, v \in V$, and $z_s, s \in S$, variables of the polyhedron defined by the following system (21)–(23):

$$\begin{aligned} \lambda_0 + \lambda_1 &= 1, \quad \lambda_0 \geq 0, \quad \lambda_1 \geq 0 \\ z_v &= z_v^0 + z_v^1, \quad \forall v \in V \\ z_s &= z_s^0 + z_s^1, \quad \forall s \in S \end{aligned} \quad (21)$$

$$\begin{aligned} z_{v_n}^0 &= 0 \\ z_s^0 &= 0, \quad \forall s \in \mathcal{E}_{n-1}^+ \\ z_s^0 &= z_{p(s)}^0, \quad \forall s \in \mathcal{E}_{n-1}^- \\ z_s^0 + z_{\ell(s)}^0 &= z_{p(s)}^0, \quad \forall s \in \mathcal{E}_k^+, k \in [n-2] \setminus \{1\} \\ z_s^0 &\geq 0, \quad \forall s \in S \setminus \mathcal{E}_{n-1} \\ \sum_{s \in \mathcal{E}_k^+} z_s^0 &\leq z_{v_{k+1}}^0, \quad \forall k \in [n-2] \setminus \{1\} \\ \sum_{s \in \mathcal{E}_k^-} z_s^0 &\leq \lambda_0 - z_{v_{k+1}}^0, \quad \forall k \in [n-2] \setminus \{1\} \\ z_{q_1}^0 + z_{q_2}^0 &= z_{v_1}^0, \quad z_{q_1}^0 + z_{q_3}^0 = z_{v_2}^0, \quad z_{q_3}^0 + z_{q_4}^0 = \lambda_0 - z_{v_1}^0 \end{aligned} \quad (22)$$

$$\begin{aligned}
z_{v_n}^1 &= \lambda_1 \\
z_s^1 &= z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-1}^+ \\
z_s^1 &= 0, \quad \forall s \in \mathcal{E}_{n-1}^- \\
z_s^1 + z_{\ell(s)}^1 &= z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+, k \in [n-2] \setminus \{1\} \\
z_s^1 &\geq 0, \quad \forall s \in S \setminus \mathcal{E}_{n-1} \\
\sum_{s \in \mathcal{E}_k^+} z_s^1 &\leq z_{v_{k+1}}^1, \quad \forall k \in [n-2] \setminus \{1\} \\
\sum_{s \in \mathcal{E}_k^-} z_s^1 &\leq \lambda_1 - z_{v_{k+1}}^1, \quad \forall k \in [n-2] \setminus \{1\} \\
z_{q_1}^1 + z_{q_2}^1 &= z_{v_1}^1, \quad z_{q_1}^1 + z_{q_3}^1 = z_{v_2}^1, \quad z_{q_3}^1 + z_{q_4}^1 = \lambda_1 - z_{v_1}^1.
\end{aligned} \tag{23}$$

In the remainder of this proof, we project out $z^0, z^1, \lambda_0, \lambda_1$ from system (21)–(23) and obtain a description of $\mathcal{P}_{\text{PB}}(H)$ in the original space. From $z_{v_n} = z_{v_n}^0 + z_{v_n}^1$, $z_{v_n}^0 = 0$, $z_{v_n}^1 = \lambda_1$, and $\lambda_0 + \lambda_1 = 1$, it follows that

$$\lambda_0 = 1 - z_{v_n}, \quad \lambda_1 = z_{v_n}. \tag{24}$$

By $z_s = z_s^0 + z_s^1$ for all $s \in S$, $z_s^0 = 0$, and $z_s^1 = z_{p(s)}^1$ for all $s \in \mathcal{E}_{n-1}^+$, we get

$$z_s^1 = z_{p(s)}^1 = z_s, \quad z_{p(s)}^0 = z_{p(s)} - z_s, \quad \forall s \in \mathcal{E}_{n-1}^+. \tag{25}$$

Similarly, by $z_{\ell(s)}^1 = 0$ and $z_{\ell(s)}^0 = z_{p(s)}^0$ for all $s \in \mathcal{E}_{n-1}^+$, we get

$$z_{\ell(s)}^0 = z_{p(s)}^0 = z_{\ell(s)}, \quad z_{p(s)}^1 = z_{p(s)} - z_{\ell(s)}, \quad \forall s \in \mathcal{E}_{n-1}^+. \tag{26}$$

Using (21) to project out z_v^0 , $v \in V$, z_s^0 , $s \in S$, using (24) to project out λ_0, λ_1 , and using (25) and (26) to project out z_s^0, z_s^1 , $s \in \mathcal{E}_{n-1}$ with $N(s) \neq \emptyset$, we deduce that system (21)–(23) simplifies to:

$$z_{s+v_n^-} + z_{\ell(s)+v_n^-} = z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2}^+ : s + v_n^-, \ell(s) + v_n^- \in \mathcal{E}_{n-1}^- \tag{27}$$

$$z_{s+v_n^-} + z_{\ell(s)} - z_{\ell(s)}^1 = z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2} : s + v_n^- \in \mathcal{E}_{n-1}^-, N(\ell(s)) = \emptyset \tag{28}$$

$$z_s - z_s^1 + z_{\ell(s)} - z_{\ell(s)}^1 = z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+ : k \in [n-3] \setminus \{1\}, \text{ or} \\ k = n-2, N(s) = N(\ell(s)) = \emptyset \tag{29}$$

$$z_s - z_s^1 \geq 0, \quad \forall s \in \mathcal{E}_k : k \in [n-3] \text{ or } k = n-2, N(s) = \emptyset \tag{30}$$

$$\sum_{\substack{s \in \mathcal{E}_{n-2}^+ \\ s+v_n^- \in \mathcal{E}_{n-1}^-}} z_{s+v_n^-} + \sum_{\substack{s \in \mathcal{E}_{n-2}^+ \\ N(s)=\emptyset}} (z_s - z_s^1) \leq z_{v_{n-1}} - z_{v_{n-1}}^1, \tag{31}$$

$$\sum_{\substack{s \in \mathcal{E}_{n-2}^- \\ s+v_n^- \in \mathcal{E}_{n-1}^-}} z_{s+v_n^-} + \sum_{\substack{s \in \mathcal{E}_{n-2}^- \\ N(s)=\emptyset}} (z_s - z_s^1) \leq 1 - z_{v_n} - z_{v_{n-1}} + z_{v_{n-1}}^1, \tag{32}$$

$$\sum_{s \in \mathcal{E}_k^+} (z_s - z_s^1) \leq z_{v_{k+1}} - z_{v_{k+1}}^1, \quad \forall k \in [n-3] \setminus \{1\} \tag{33}$$

$$\sum_{s \in \mathcal{E}_k^-} (z_s - z_s^1) \leq 1 - z_{v_n} - z_{v_{k+1}} + z_{v_{k+1}}^1, \quad \forall k \in [n-3] \setminus \{1\} \tag{34}$$

$$\begin{aligned}
z_{q_1} - z_{q_1}^1 + z_{q_2} - z_{q_2}^1 &= z_{v_1} - z_{v_1}^1, \quad z_{q_1} - z_{q_1}^1 + z_{q_3} - z_{q_3}^1 = z_{v_2} - z_{v_2}^1, \\
z_{q_3} - z_{q_3}^1 + z_{q_4} - z_{q_4}^1 &= 1 - z_{v_n} - z_{v_1} + z_{v_1}^1,
\end{aligned} \tag{35}$$

and

$$z_{s+v_n^+} + z_{\ell(s)+v_n^+} = z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2}^+ : s + v_n^+, \ell(s) + v_n^+ \in \mathcal{E}_{n-1}^+ \tag{36}$$

$$z_{s+v_n^+} + z_{\ell(s)}^1 = z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2} : s + v_n^+ \in \mathcal{E}_{n-1}^+, N(\ell(s)) = \emptyset \tag{37}$$

$$z_s^1 + z_{\ell(s)}^1 = z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+ : k \in [n-3] \setminus \{1\}, \text{ or } k = 2, N(s) = N(\ell(s)) = \emptyset \tag{38}$$

$$z_s^1 \geq 0, \quad \forall s \in \mathcal{E}_k : k \in [n-3], \text{ or } k = n-2, N(s) = \emptyset \tag{39}$$

$$\sum_{\substack{s \in \mathcal{E}_{n-2}^+ \\ s+v_n^+ \in \mathcal{E}_{n-1}^+}} z_{s+v_n^+} + \sum_{\substack{s \in \mathcal{E}_{n-2}^+ \\ N(s)=\emptyset}} z_s^1 \leq z_{v_{n-1}}^1 \tag{40}$$

$$\sum_{\substack{s \in \mathcal{E}_{n-2}^- \\ s+v_n^+ \in \mathcal{E}_{n-1}^+}} z_{s+v_n^+} + \sum_{\substack{s \in \mathcal{E}_{n-2}^- \\ N(s)=\emptyset}} z_s^1 \leq z_{v_n} - z_{v_{n-1}}^1 \tag{41}$$

$$\sum_{s \in \mathcal{E}_k^+} z_s^1 \leq z_{v_{k+1}}^1, \quad \forall k \in [n-3] \setminus \{1\} \tag{42}$$

$$\sum_{s \in \mathcal{E}_k^-} z_s^1 \leq z_{v_n} - z_{v_{k+1}}^1, \quad \forall k \in [n-3] \setminus \{1\} \tag{43}$$

$$z_{q_1}^1 + z_{q_2}^1 = z_{v_1}^1, \quad z_{q_1}^1 + z_{q_3}^1 = z_{v_2}^1, \quad z_{q_3}^1 + z_{q_4}^1 = z_{v_n} - z_{v_1}^1, \tag{44}$$

together with

$$z_s + z_{\ell(s)} = z_{p(s)}, \quad \forall s \in \mathcal{E}_{n-1}^+ \tag{45}$$

$$z_s \geq 0, \quad \forall s \in \mathcal{E}_{n-1}. \tag{46}$$

Note that equalities (45) are present among equalities (7) and inequalities (46) are present among inequalities (8).

In the remainder of the proof, we project out variables $z_v^1, v \in V \setminus \{v_n\}, z_s^1, s \in S \setminus \{s \in \mathcal{E}_{n-1} : N(s) \neq \emptyset\}$ in a specific order:

Projecting out $z_{v_k}^1, k \in \{1, \dots, n-1\}$:

- The variable $z_{v_1}^1$ appears only in equalities (35) and (44). Hence projecting out $z_{v_1}^1$ we obtain:

$$z_{q_1} + z_{q_2} = z_{v_1} \tag{47}$$

$$z_{q_3} + z_{q_4} = 1 - z_{v_1} \tag{48}$$

$$z_{q_1}^1 + z_{q_2}^1 + z_{q_3}^1 + z_{q_4}^1 = z_{v_n}. \tag{49}$$

Equalities (47) and (48) are present among equalities (11). Moreover, the above equalities imply:

$$z_{q_1} - z_{q_1}^1 + z_{q_2} - z_{q_2}^1 + z_{q_3} - z_{q_3}^1 + z_{q_4} - z_{q_4}^1 = 1 - z_{v_n} \tag{50}$$

We will use equality (50) to simplify our derivations.

- The variable $z_{v_2}^1$ appears only in equalities (35) and (44). Hence projecting out this variable we obtain:

$$z_{q_1} + z_{q_3} = z_{v_2}, \quad (51)$$

which is present among equalities (11).

- The variable $z_{v_{n-1}}^1$ appears only in inequalities (31), (32), (40) and (41). Projecting out $z_{v_{n-1}}^1$ from (31) and (32) gives:

$$\sum_{s \in \mathcal{E}_{n-1}^-} z_s + \sum_{s \in \mathcal{E}_{n-2}: N(s)=\emptyset} (z_s - z_s^1) \leq 1 - z_{v_n}, \quad (52)$$

while projecting out $z_{v_{n-1}}^1$ from (40) and (41) gives:

$$\sum_{s \in \mathcal{E}_{n-1}^+} z_s + \sum_{s \in \mathcal{E}_{n-2}: N(s)=\emptyset} z_s^1 \leq z_{v_n}. \quad (53)$$

As we argue shortly, inequality (52) (resp. (53)) is implied by equality (50) (resp. (49)) and inequality (30) (resp. (39)). Projecting out $z_{v_{n-1}}^1$ from (31) and (40) gives:

$$\sum_{s \in \mathcal{E}_{n-2}^+} z_s \leq z_{v_{n-1}}, \quad (54)$$

which coincides with inequalities (9) for $k = n - 2$. Lastly, projecting out $z_{v_{n-1}}^1$ from (32) and (41) we obtain:

$$\sum_{s \in \mathcal{E}_{n-2}^-} z_s \leq 1 - z_{v_{n-1}}, \quad (55)$$

which coincides with inequalities (10) for $k = n - 2$.

- Variables $z_{v_{k+1}}^1$, $k \in \{2, \dots, n - 3\}$, are only present in inequalities (33), (34), (42), and (43). Projecting out $z_{v_{k+1}}^1$, $k \in \{2, \dots, n - 3\}$, from (33) and (34) gives:

$$\sum_{s \in \mathcal{E}_k} (z_s - z_s^1) \leq 1 - z_{v_n}, \quad \forall k \in [n - 3] \setminus \{1\},$$

which by equalities (29) can be equivalently written as:

$$\sum_{s \in \mathcal{E}_k^+} (z_{p(s)} - z_{p(s)}^1) \leq 1 - z_{v_n}, \quad \forall k \in [n - 3] \setminus \{1\}. \quad (56)$$

By (27)–(29), for each $k \in [n - 1] \setminus \{1\}$, we have $\sum_{s \in \mathcal{E}_{k+1}} (z_s - z_s^1) \leq \sum_{s \in \mathcal{E}_k} (z_s - z_s^1)$, implying that inequalities (52) and (56) are implied by the following inequality:

$$\sum_{s \in \mathcal{E}_2^+} (z_{p(s)} - z_{p(s)}^1) \leq 1 - z_{v_n},$$

which in turn is implied by inequalities (30) and equality (50). Similarly, projecting out $z_{v_{k+1}}^1$, $k \in \{2, \dots, n - 3\}$, from (42) and (43) gives:

$$\sum_{s \in \mathcal{E}_k} z_s^1 \leq z_{v_n}, \quad \forall k \in [n - 3] \setminus \{1\},$$

which by equalities (38) can be equivalently written as:

$$\sum_{s \in \mathcal{E}_k^+} z_{p(s)}^1 \leq z_{v_n}, \quad \forall k \in [n-3] \setminus \{1\}. \quad (57)$$

Since for each $k \in [n-1] \setminus \{1\}$, by (36)–(38), we have $\sum_{s \in \mathcal{E}_{k+1}} z_s^1 \leq \sum_{s \in \mathcal{E}_k} z_s^1$, inequalities (53) and (57) are implied by:

$$\sum_{s \in \mathcal{E}_2^+} z_{p(s)}^1 \leq z_{v_n},$$

which in turn is implied by inequalities (39) and equality (49). Projecting out $z_{v_{k+1}}^1$, $k \in \{2, \dots, n-3\}$, from (33) and (42) we obtain:

$$\sum_{s \in \mathcal{E}_k^+} z_s \leq z_{v_{k+1}}, \quad \forall k \in [n-3] \setminus \{1\}, \quad (58)$$

which are present among inequalities (9). Similarly, projecting out $z_{v_{k+1}}^1$, $k \in \{2, \dots, n-3\}$, from (34) and (43), we obtain

$$\sum_{s \in \mathcal{E}_k^-} z_s \leq 1 - z_{v_{k+1}}, \quad \forall k \in [n-3] \setminus \{1\}, \quad (59)$$

which are present among inequalities (10).

Hence, we have shown that projecting out variables $z_{v_k}^1$, $k \in \{1, \dots, n-1\}$, from (31)–(35) and (40)–(44), we obtain inequalities (9)–(10) when $k \neq n-1$, equalities (11), and equalities (49)–(50).

Projecting out z_s^1 , $s \in \mathcal{E}_{n-2}$ with $N(s) = \emptyset$: Consider $z_{\bar{s}}^1$ for some $\bar{s} \in \mathcal{E}_{n-2}^+$ with $N(\bar{s}) = \emptyset$. This variable is only present in (28)–(30) and (37)–(39). Two cases arise:

- If $N(\ell(\bar{s})) \neq \emptyset$, then $z_{\bar{s}}^1$ is only present in (28), (30), (37), (39); i.e., the following system:

$$\begin{aligned} z_{\bar{s}} - z_{\bar{s}}^1 + z_{\ell(\bar{s})+v_n^-} &= z_{p(\bar{s})} - z_{p(\bar{s})}^1 \\ z_{\bar{s}} - z_{\bar{s}}^1 &\geq 0 \\ z_{\bar{s}}^1 + z_{\ell(\bar{s})+v_n^+} &= z_{p(\bar{s})}^1 \\ z_{\bar{s}}^1 &\geq 0. \end{aligned}$$

Projecting out $z_{\bar{s}}^1$ from the above system yields:

$$\begin{aligned} z_{\bar{s}} + z_{\ell(\bar{s})} &= z_{p(\bar{s})} \\ z_{\bar{s}} &\geq 0 \\ z_{\ell(\bar{s})+v_n^-} &\leq z_{p(\bar{s})} - z_{p(\bar{s})}^1 \\ z_{\ell(\bar{s})+v_n^+} &\leq z_{p(\bar{s})}^1 \end{aligned}$$

where to obtain the first equality we made use of equalities (45).

- If $N(\ell(\bar{s})) = \emptyset$, then $z_{\bar{s}}^1$ and $z_{\ell(\bar{s})}^1$ are present only in (29)–(30), and (38)–(39); i.e., the following system:

$$z_{\bar{s}} - z_{\bar{s}}^1 + z_{\ell(\bar{s})} - z_{\ell(\bar{s})}^1 = z_{p(\bar{s})} - z_{p(\bar{s})}^1,$$

$$\begin{aligned}
z_{\bar{s}} - z_{\bar{s}}^1 &\geq 0 \\
z_{\ell(\bar{s})} - z_{\ell(\bar{s})}^1 &\geq 0 \\
z_{\bar{s}}^1 + z_{\ell(\bar{s})}^1 &= z_{p(\bar{s})}^1, \\
z_{\bar{s}}^1 &\geq 0 \\
z_{\ell(\bar{s})}^1 &\geq 0.
\end{aligned}$$

Projecting out $z_{\bar{s}}^1$ and $z_{\ell(\bar{s})}^1$ from the above system yields:

$$\begin{aligned}
z_{\bar{s}} + z_{\ell(\bar{s})} &= z_{p(\bar{s})} \\
z_{\bar{s}} &\geq 0 \\
z_{\ell(\bar{s})} &\geq 0 \\
z_{p(\bar{s})}^1 &\geq 0 \\
z_{p(\bar{s})} - z_{p(\bar{s})}^1 &\geq 0.
\end{aligned}$$

By a recursive application of the two steps detailed above to project out all z_s^1 , $s \in \mathcal{E}_{n-2}$ with $N(s) = \emptyset$, from (28)–(30) and (37)–(39), we obtain:

$$\begin{aligned}
z_{s+v_n^+} &\leq z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2} : s + v_n^+ \in \mathcal{E}_{n-1}^+, N(\ell(s)) = \emptyset, \\
z_{s+v_n^-} &\leq z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2} : s + v_n^- \in \mathcal{E}_{n-1}^-, N(\ell(s)) = \emptyset, \\
z_{p(s)}^1 &\geq 0, \quad z_{p(s)} - z_{p(s)}^1 \geq 0, \quad \forall s \in \mathcal{E}_{n-2} : N(s) = N(\ell(s)) = \emptyset
\end{aligned}$$

together with

$$\begin{aligned}
z_s + z_{\ell(s)} &= z_{p(s)}, \quad \forall s \in \mathcal{E}_{n-2}^+ \\
z_s &\geq 0, \quad \forall s \in \mathcal{E}_{n-2}.
\end{aligned} \tag{60}$$

Hence, to complete the projection, it suffices to project out variables z_s^1 , $s \in \mathcal{E}_k$, $k \in [n-3]$, from the following system:

$$\begin{aligned}
z_{s+v_n^-} + z_{\ell(s)+v_n^-} &= z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2}^+ : s + v_n^-, \ell(s) + v_n^- \in \mathcal{E}_{n-1}^- \\
z_{s+v_n^+} + z_{\ell(s)+v_n^+} &= z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2}^+ : s + v_n^+, \ell(s) + v_n^+ \in \mathcal{E}_{n-1}^+ \\
z_{s+v_n^-} &\leq z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2} : s + v_n^- \in \mathcal{E}_{n-1}^-, N(\ell(s)) = \emptyset \\
z_{s+v_n^+} &\leq z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-2} : s + v_n^+ \in \mathcal{E}_{n-1}^+, N(\ell(s)) = \emptyset \\
z_s - z_s^1 + z_{\ell(s)} - z_{\ell(s)}^1 &= z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+, k \in [n-3] \setminus \{1\} \\
z_s^1 + z_{\ell(s)}^1 &= z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+, k \in [n-3] \setminus \{1\} \\
z_s - z_s^1 &\geq 0, \quad \forall s \in \mathcal{E}_k : k \in [n-3] \\
z_s^1 &\geq 0, \quad \forall s \in \mathcal{E}_k : k \in [n-3] \\
z_{q_1} - z_{q_1}^1 + z_{q_2} - z_{q_2}^1 + z_{q_3} - z_{q_3}^1 + z_{q_4} - z_{q_4}^1 &= 1 - z_{v_n} \\
z_{q_1}^1 + z_{q_2}^1 + z_{q_3}^1 + z_{q_4}^1 &= z_{v_n}.
\end{aligned} \tag{61}$$

For each $s \in \mathcal{E}_{n-3}^+$ define

$$N_2^+(p(s)) := \{s' \in \mathcal{E}_{n-1}^+ : p(p(s')) = s\}$$

and

$$N_2^-(p(s)) := \{s' \in \mathcal{E}_{n-1}^- : p(p(s')) = s\}.$$

Note that $0 \leq |N_2^+(p(s))| \leq 4$ and $0 \leq |N_2^-(p(s))| \leq 4$ for all $s \in \mathcal{E}_{n-3}^+$. Projecting out z_s^1 , $s \in \mathcal{E}_{n-3}$, from system (61) yields:

$$\begin{aligned} z_s &\geq 0, \quad \forall s \in \mathcal{E}_{n-3} \\ z_s + z_{\ell(s)} &= z_{p(s)}, \quad \forall s \in \mathcal{E}_{n-3}^+, \end{aligned}$$

together with

$$\begin{aligned} \sum_{s' \in N_2^-(p(s))} z_{s'} &= z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-3}^+ : |N_2^-(p(s))| = 4 \\ \sum_{s' \in N_2^+(p(s))} z_{s'} &= z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-3}^+ : |N_2^+(p(s))| = 4 \\ \sum_{s' \in N_2^-(p(s))} z_{s'} &\leq z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-3}^+ : |N_2^-(p(s))| < 4 \\ \sum_{s' \in N_2^+(p(s))} z_{s'} &\leq z_{p(s)}^1, \quad \forall s \in \mathcal{E}_{n-3}^+ : |N_2^+(p(s))| < 4 \\ z_s - z_s^1 + z_{\ell(s)} - z_{\ell(s)}^1 &= z_{p(s)} - z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+, k \in [n-4] \setminus \{1\} \\ z_s^1 + z_{\ell(s)}^1 &= z_{p(s)}^1, \quad \forall s \in \mathcal{E}_k^+, k \in [n-4] \setminus \{1\} \\ z_s - z_s^1 &\geq 0, \quad \forall k \in [n-4] \\ z_s^1 &\geq 0, \quad \forall k \in [n-4] \\ z_{q_1}^1 + z_{q_2}^1 + z_{q_3}^1 + z_{q_4}^1 &= z_{v_n} \\ z_{q_1} - z_{q_1}^1 + z_{q_2} - z_{q_2}^1 + z_{q_3} - z_{q_3}^1 + z_{q_4} - z_{q_4}^1 &= 1 - z_{v_n} \end{aligned}$$

By a recursive application of the above argument $n-5$ times to project out variables z_s^1 , $s \in \mathcal{E}_k$, $k \in [n-4] \setminus \{1\}$, from the above system we obtain:

$$\begin{aligned} \sum_{s \in \mathcal{E}_{n-1}^- : s \supset q_i} z_s &\leq z_{q_i} - z_{q_i}^1, \quad \forall i \in \{1, \dots, 4\} \\ \sum_{s \in \mathcal{E}_{n-1}^+ : s \supset q_i} z_s &\leq z_{q_i}^1, \quad \forall i \in \{1, \dots, 4\} \\ z_{q_i} - z_{q_i}^1 &\geq 0, \quad \forall i \in \{1, \dots, 4\} \\ z_{q_i}^1 &\geq 0, \quad i \in \{1, \dots, 4\} \\ z_{q_1}^1 + z_{q_2}^1 + z_{q_3}^1 + z_{q_4}^1 &= z_{v_n} \\ z_{q_1} - z_{q_1}^1 + z_{q_2} - z_{q_2}^1 + z_{q_3} - z_{q_3}^1 + z_{q_4} - z_{q_4}^1 &= 1 - z_{v_n}, \end{aligned} \tag{62}$$

together with

$$\begin{aligned} z_s &\geq 0, \quad \forall s \in \mathcal{E}_k, k \in [n-3] \setminus \{1\} \\ z_s + z_{\ell(s)} &= z_{p(s)}, \quad \forall e \in \mathcal{E}_k^+, k \in [n-3] \setminus \{1\}. \end{aligned} \tag{63}$$

Hence it remains to project out $z_{q_i}^1$, $i \in \{1, \dots, 4\}$, from system (62). Since $\bigcup_{i \in \{1, \dots, 4\}} \{s \in \mathcal{E}_{n-1}^- : s \supset q_i\} = \mathcal{E}_{n-1}^-$ and $\bigcup_{i \in \{1, \dots, 4\}} \{s \in \mathcal{E}_{n-1}^+ : s \supset q_i\} = \mathcal{E}_{n-1}^+$, it follows that the projection of system (62) onto the

space of z is given by:

$$\begin{aligned}
\sum_{s \in \mathcal{E}_{n-1}^-} z_s &\leq 1 - z_{v_n} \\
\sum_{s \in \mathcal{E}_{n-1}^+} z_s &\leq z_{v_n} \\
z_{q_i} &\geq 0, \quad \forall i \in \{1, \dots, 4\} \\
z_{q_1} + z_{q_2} + z_{q_3} + z_{q_4} &= 1,
\end{aligned} \tag{64}$$

where the last equality is implied by equalities (11). From (45)–(48), (51), (54), (55), (58)–(60), (63), and (64), it follows that the pseudo-Boolean polytope $\mathcal{P}_{\text{PB}}(H)$ is defined by system (7)–(11) and this completes the proof. \square

Statements and Declarations

Competing interests: The authors have no financial or non-financial conflicts of interests.

Availability of data and materials: The manuscript has not used any data.

Funding: A. Del Pia is partially funded by AFOSR grant FA9550-23-1-0433. A. Khajavirad is in part supported by AFOSR grant FA9550-23-1-0123. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Air Force Office of Scientific Research.

Acknowledgments: The authors would like to thank two anonymous referees for comments and suggestions that improved the quality of this manuscript.

References

- [1] E. Balas. Disjunctive programming: properties of the convex hull of feasible points. *Discrete Applied Mathematics*, 89(1–3):3–44, 1998.
- [2] F. Barahona and A. R. Mahjoub. On the cut polytope. *Mathematical programming*, 36:157–173, 1986.
- [3] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30:479–513, 1983.
- [4] P. Belotti, J. Lee, L. Liberti, F. Margot, and A. Wächter. Branching and bounds tightening techniques for non-convex MINLP. *Optimization Methods & Software*, 24(4-5):597–634, 2009.
- [5] D. Bienstock and G. Munoz. LP formulations for polynomial optimization problems. *SIAM Journal on Optimization*, 28(2):1121–1150, 2018.
- [6] E. Boros, Y. Crama, and P. L. Hammer. Upper-bounds for quadratic 0–1 maximization. *Operations Research Letters*, 9(2):73–79, 1990.
- [7] E. Boros, Y. Crama, and E. Rodríguez-Heck. Compact quadratizations for pseudo-boolean functions. *Journal of combinatorial optimization*, 39:687–707, 2020.
- [8] E. Boros and P. L. Hammer. Pseudo-Boolean optimization. *Discrete applied mathematics*, 123(1-3):155–225, 2002.

- [9] J. Brault-Baron. Hypergraph acyclicity revisited. *ACM Computing Surveys*, 49(3):54:1–54:26, 2016.
- [10] C. Buchheim, Y. Crama, and E. Rodríguez-Heck. Berge-acyclic multilinear 0–1 optimization problems. *European Journal of Operational Research*, 273(1):102–107, 2019.
- [11] G. Cornuéjols. *Combinatorial Optimization: Packing and Covering*, volume 74 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM, 2001.
- [12] Y. Crama. Recognition problems for special classes of polynomials in 0–1 variables. *Mathematical programming*, 44:139–155, 1989.
- [13] Y. Crama and E. Rodríguez-Heck. A class of valid inequalities for multilinear 0 – 1 optimization problems. *Discrete Optimization*, 25:28–47, 2017.
- [14] A. Del Pia and S. Di Gregorio. Chvátal rank in binary polynomial optimization. *INFORMS Journal on Optimization*, 3(4):315–349, 2021.
- [15] A. Del Pia and S. Di Gregorio. On the complexity of binary polynomial optimization over acyclic hypergraphs. *Algorithmica*, 85:2189–2213, 2023.
- [16] A. Del Pia and A. Khajavirad. A polyhedral study of binary polynomial programs. *Mathematics of Operations Research*, 42(2):389–410, 2017.
- [17] A. Del Pia and A. Khajavirad. The multilinear polytope for acyclic hypergraphs. *SIAM Journal on Optimization*, 28(2):1049–1076, 2018.
- [18] A. Del Pia and A. Khajavirad. On decomposability of multilinear sets. *Mathematical Programming, Series A*, 170(2):387–415, 2018.
- [19] A. Del Pia and A. Khajavirad. The running intersection relaxation of the multilinear polytope. *Mathematics of Operations Research*, 46(3):1008–1037, 2021.
- [20] A. Del Pia and A. Khajavirad. A polynomial-size extended formulation for the multilinear polytope of beta-acyclic hypergraphs. *Mathematical Programming, Series A*, pages 1–33, 2023.
- [21] A. Del Pia, A. Khajavirad, and N. Sahinidis. On the impact of running-intersection inequalities for globally solving polynomial optimization problems. *Mathematical Programming Computation*, 12:165–191, 2020.
- [22] A. Del Pia and M. Walter. Simple odd β -cycle inequalities for binary polynomial optimization. In *Proceedings of IPCO 2022*, volume 13265 of *Lecture Notes in Computer Science*, pages 181–194. Springer, 2022.
- [23] A. Del Pia and M. Walter. Simple odd β -cycle inequalities for binary polynomial optimization. *Mathematical Programming, Series B*, 2023.
- [24] D. Duris. Some characterizations of γ and β -acyclicity of hypergraphs. *Information Processing Letters*, 112:617–620, 2012.
- [25] R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *Journal of the ACM (JACM)*, 30(3):514–550, 1983.
- [26] M. Goemans and D. P. Williamson. New 3/4-approximation algorithms for the maximum satisfiability problem. *SIAM Journal on Discrete Mathematics*, 7(4):656–666, 1994.

- [27] P. L. Hammer, P. Hansen, and B. Simeone. Roof duality, complementation and persistency in quadratic 0–1 optimization. *Mathematical programming*, 28:121–155, 1984.
- [28] P. Jégou and S. N. Ndiaye. On the notion of cycles in hypergraphs. *Discrete Mathematics*, 309:6535–6543, 2009.
- [29] M. I. Jordan. Graphical Models. *Statistical Science*, 19(1):140 – 155, 2004.
- [30] A. Khajavirad. On the strength of recursive McCormick relaxations for binary polynomial optimization. *Operations Research Letters*, 51(2):146–152, 2023.
- [31] A. Khajavirad and N. V. Sahinidis. A hybrid LP/NLP paradigm for global optimization relaxations. *Mathematical Programming Computation*, 10(3):383–421, May 2018.
- [32] J. Kim, J. P. Richard, and M. Tawarmalani. A reciprocity between tree ensemble optimization and multilinear optimization. *Optimization Online*, <https://optimization-online.org/2022/03/8828/>, 2022.
- [33] M. Laurent. Sums of squares, moment matrices and optimization over polynomials. In *Emerging Applications of Algebraic Geometry*, volume 149 of *The IMA Volumes in Mathematics and its Applications*, pages 157–270. Springer, 2009.
- [34] R. Misener and C. A. Floudas. ANTIGONE: algorithms for continuous/integer global optimization of nonlinear equations. *Journal of Global Optimization*, 59(2):503–526, 2014.
- [35] M. Padberg. The Boolean quadric polytope: Some characteristics, facets and relatives. *Mathematical Programming*, 45(1–3):139–172, 1989.
- [36] N. Sahinidis and M. Tawarmalani. Convexification and global optimization in continuous and mixed–integer nonlinear programming: theory, algorithms, software, and applications, 2003.
- [37] S. Vigerske and A. Gleixner. SCIP: Global optimization of mixed-integer nonlinear programs in a branch-and-cut framework. *Optimization Methods and Software*, 33(3):563–593, 2018.
- [38] M.J. Wainwright and M.I. Jordan. Treewidth-based conditions for exactness of the Sherali-Adams and Lasserre relaxations. Technical Report 671, University of California, 2004.