Binary polynomial optimization through a hypergraph theoretic lens

June 2025

Aida Khajavirad
Lehigh University

Joint work with: Alberto Del Pia (University of Wisconsin-Madison)

Problem definition

- Let $V = \{1, \dots, n\}$, let E be a set of subsets of cardinality at least two of V. Consider an unconstrained binary polynomial optimization problem:

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

- Problem BPO is NP-hard in general.

- Applications: Satisfiability problems (MAX-SAT), inference in higher-order graphical models, low-rank Boolean matrix/tensor factorization, ...

- Define $z_e := \prod_{v \in e} z_v$ for all $e \in E$:

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

Multilinear sets and polytope

- We define the multilinear set as:

$$\mathcal{S} = \Big\{ z \in \{0,1\}^{V \cup E} : z_e = \prod_{v \in e} z_v, \forall e \in E \Big\}.$$

- We define the multilinear polytope as the convex hull of the multilinear set:

$$\mathcal{P}_{\mathrm{m}} = \mathrm{conv}(\mathcal{S})$$

- If $|e| = 2$ for all $e \in E$, then the objective function of Problem (BPO) is quadratic and $\mathcal{P}_{\mathrm{m}}$ is the Boolean quadric polytope QP (Padberg, 89).

Hypergraph representation

- There is a bijection between multilinear sets and hypergraphs: one node v for each z_v, one edge e for each z_e, the edge e corresponding to z_e is $\{v : v \in e\}$

$$z_{12} = z_1 z_2$$
$$z_{24} = z_2 z_4$$
$$z_{123} = z_1 z_2 z_3$$
$$z_{134} = z_1 z_3 z_4$$

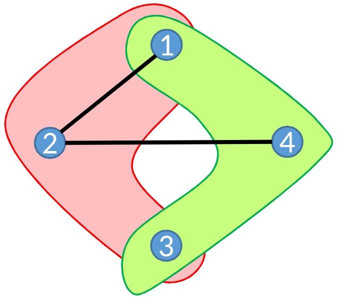

- $\mathcal{S} \to \mathcal{S}(G)$ and $\mathcal{P}_{\mathrm{m}} \to \mathcal{P}_{\mathrm{m}}(G)$.

- For quadratic sets, we obtain the graph representation of $\mathrm{QP}(G)$ (Padberg, 89).

$$z_{12} = z_1 z_2$$
$$z_{24} = z_2 z_4$$
$$z_{34} = z_3 z_4$$

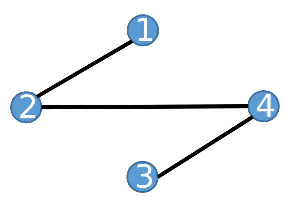

- The rank of G is the maximum cardinality of any edge in E.

The Boolean quadric polytope

- Replace each bilinear term $z_{ij} = z_i z_j$, by its convex hull over the unit hypercube and use $\bigcap_i \mathrm{conv}(\mathcal{S}_i) \supseteq \mathrm{conv}(\bigcap_i \mathcal{S}_i)$ to obtain the standard linearization, or the McCormick relaxation, of $\mathrm{QP}(G)$:

$$\mathrm{QP}^{\mathrm{LP}}(G) = \Big\{z : z_{ij} \geq 0,\ z_{ij} \geq z_i + z_j - 1,\ z_{ij} \leq z_i, z_{ij} \leq z_j,\ \forall (i,j) \in E\Big\}.$$

- $\mathrm{QP}(G) = \mathrm{QP}^{\mathrm{LP}}_G$ iff G is an acyclic graph (Padberg 89).

- Let $\mathrm{QP}^C(G)$ be polytope obtained by adding all odd cycle inequalities to $\mathrm{QP}^{\mathrm{LP}}(G)$; $\mathrm{QP}(G) = \mathrm{QP}^C(G)$ iff G is a series-parallel graph (Barahona 86, Padberg 89).

- Optimizing over $\mathrm{QP}^{\mathrm{LP}}(G)$ and $\mathrm{QP}^C(G)$ can be done in polynomial-time.

- Goal: obtaining similar results for higher degree multilinear sets in terms of easily verifiable conditions on the structure of underlying hypergraphs.

Cycles in hypergraphs

- Acyclic hypergraphs in increasing degree of generality:

$$\text{Berge} - \text{acyclic} \subset \gamma - \text{acyclic} \subset \beta - \text{acyclic} \subset \alpha - \text{acyclic}$$

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

- A hypergraph is Berge-acyclic iff it contains no Berge-cycles.

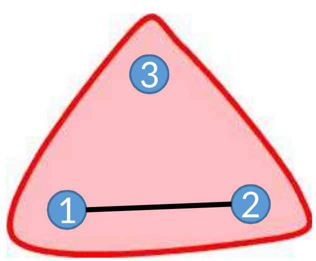

Berge-cycle:
$C = v_1, e_{12}, v_2, e_{123}, v_1$

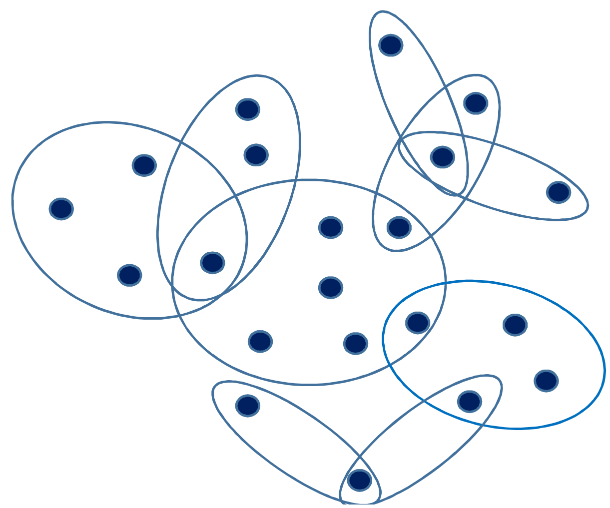

A Berge-acyclic Hypergraph

γ-acyclic hypergraphs

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

- A hypergraph is called γ-acyclic iff it contains no γ-cycles.

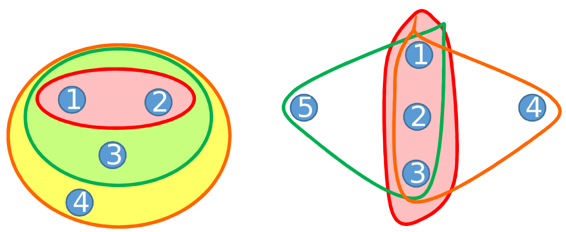

No γ-cycles; Berge-cycles of length two and three.

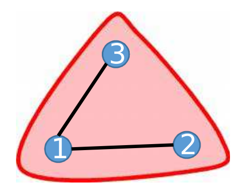

γ-cycle: $C = v_1, e_{12}, v_2, e_{123}, v_3, e_{13}, v_1$

A γ-acyclic hypergraph

β-acyclic hypergraphs

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

- A hypergraph is called β-acyclic iff it contains no β-cycles.

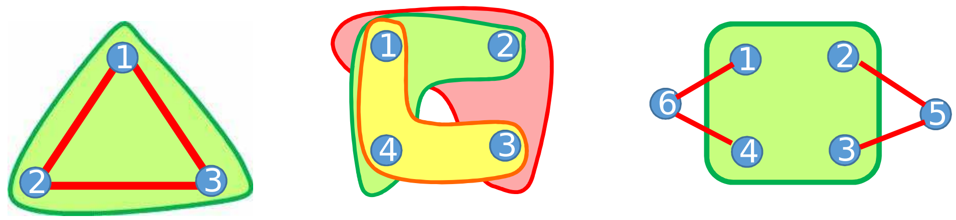

Examples of β-cycles

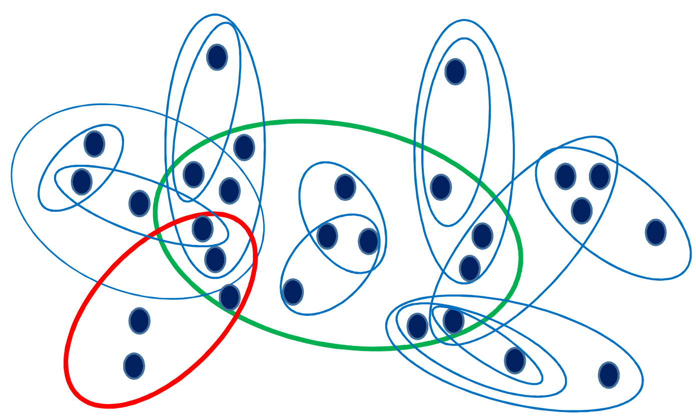

A β-acyclic hypergraph

α-acyclic hypergraphs

- Let $G = (V, E)$ be a hypergraph and let $C = e_1, e_2, \ldots, e_t, e_{t+1}$ with $e_{t+1} := e_1$ for some $t \geq 3$. Define $s_i := e_i \cap e_{i+1}$ for all $i \in [t]$. Then C is α-cycle of length t in G, if

$$(s_i \cup s_j \cup s_k) \setminus e \neq \emptyset \quad \forall 1 \leq i < j < k \leq \ell, \ \forall e \in E.$$

- A hypergraph is called α-acyclic iff it contains no α-cycles.

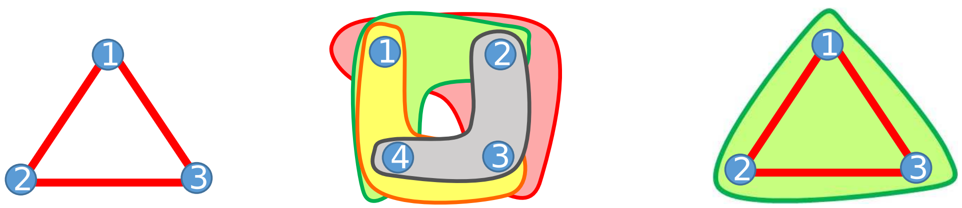

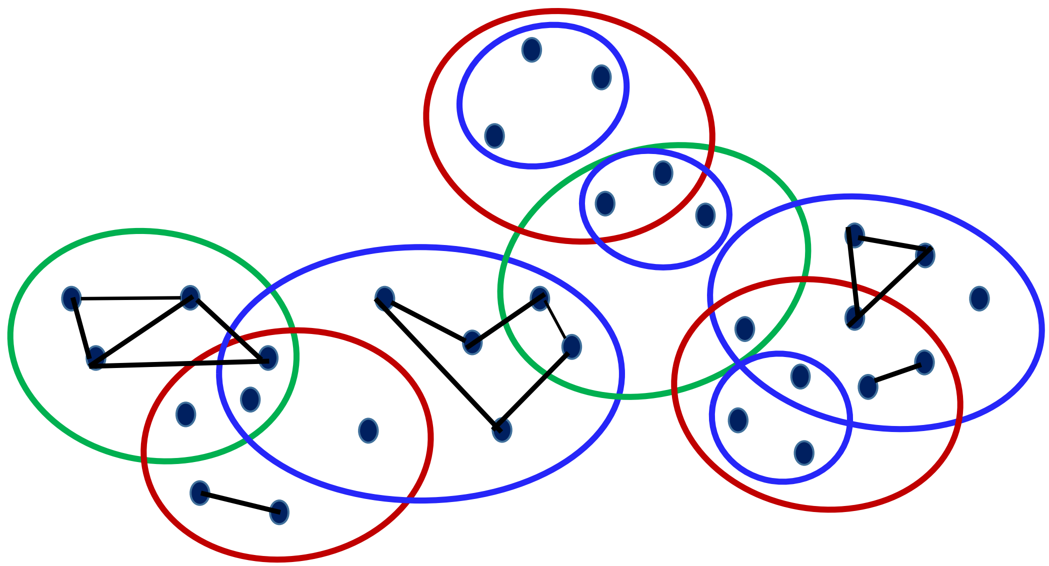

Extended formulations

- When does the multilinear polytope of acyclic hypergraphs admit a polynomial-size extended formulation which can be constructed in polynomial-time?

Decomposability of multilinear sets

- Consider hypergraphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ such that $V_1 \cap V_2 \neq \emptyset$. Denote by $G_1 \cup G_2$, the hypergraph $(V_1 \cup V_2, E_1 \cup E_2)$, and by $G_1 \cap G_2$, the hypergraph $(V_1 \cap V_2, E_1 \cap E_2)$.

- Let $G := G_1 \cup G_2$. Then $\mathcal{P}_{\mathrm{m}}(G)$ is decomposable into $\mathcal{P}_{\mathrm{m}}(G_1)$ and $\mathcal{P}_{\mathrm{m}}(G_2)$, if

$$\mathcal{P}_{\mathrm{m}}(G) = \mathcal{P}_{\mathrm{m}}(G_1) \cap \mathcal{P}_{\mathrm{m}}(G_2),$$

  i.e., the system comprised a description of $\mathcal{P}_{\mathrm{m}}(G_1)$ and a description of $\mathcal{P}_{\mathrm{m}}(G_2)$, is a description of $\mathcal{P}_{\mathrm{m}}(G)$.

- A hypergraph $G = (V, E)$ is complete if all subsets of V of cardinality at least two are in E.

- Given $V' \subset V$, the section hypergraph of G induced by V' is $G' = (V', E')$, where $E' = \{e \in E : e \subseteq V'\}$.

- Theorem: Let G_1, G_2 be section hypergraphs of G such that $G_1 \cup G_2 = G$ and $G_1 \cap G_2$ is a complete hypergraph. Then the set $\mathcal{P}_{\mathrm{m}}(G)$ is decomposable into $\mathcal{P}_{\mathrm{m}}(G_1)$ and $\mathcal{P}_{\mathrm{m}}(G_2)$.

Acyclic graphs

- Theorem (Padberg 89): The Boolean quadric polytope QP(G) of an acyclic graph $G = (V, E)$ has a formulation with $4|E|$ inequalities.

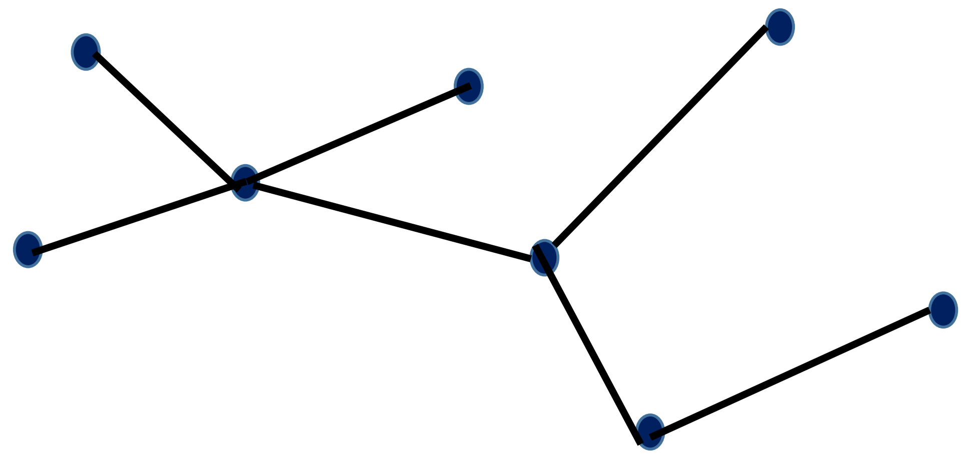

Acyclic graphs

- Theorem (Padberg 89): The Boolean quadric polytope $\mathrm{QP}(G)$ of an acyclic graph $G = (V, E)$ has a formulation with $4|E|$ inequalities.

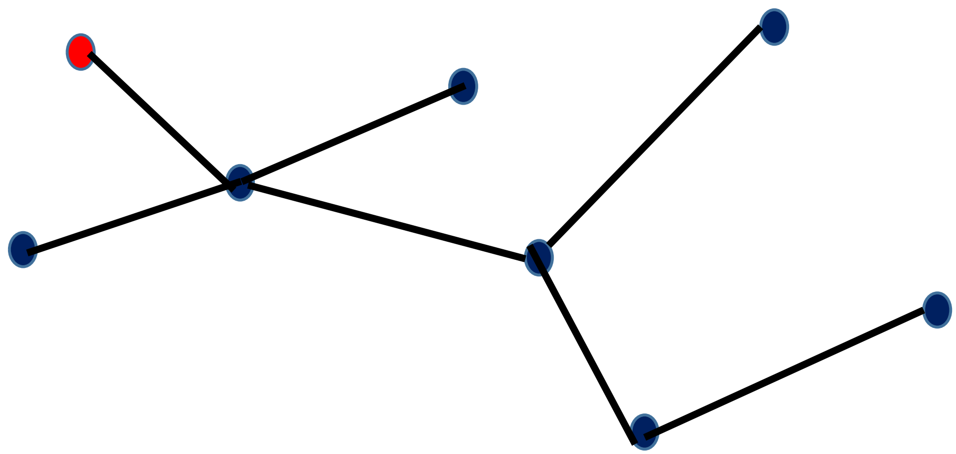

Acyclic graphs

- Theorem (Padberg 89): The Boolean quadric polytope $\mathrm{QP}(G)$ of an acyclic graph $G = (V, E)$ has a formulation with $4|E|$ inequalities.

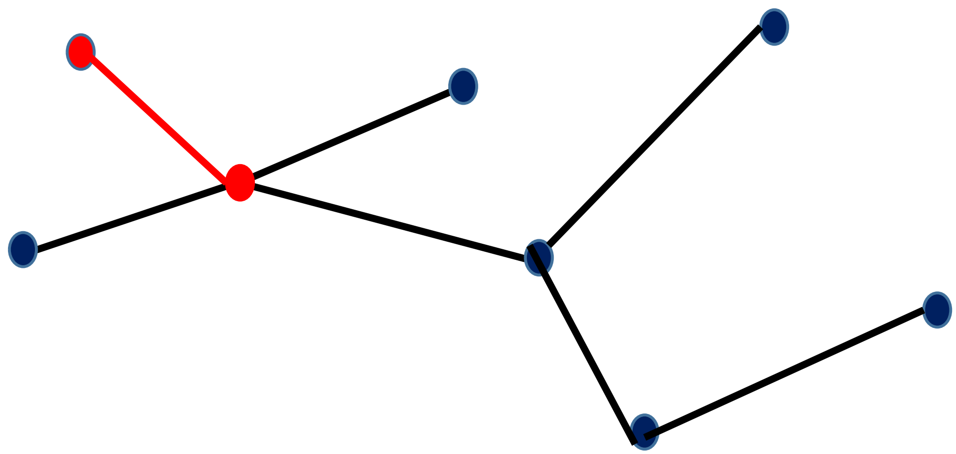

Acyclic graphs

- Theorem (Padberg 89): The Boolean quadric polytope $\mathrm{QP}(G)$ of an acyclic graph $G = (V, E)$ has a formulation with $4|E|$ inequalities.

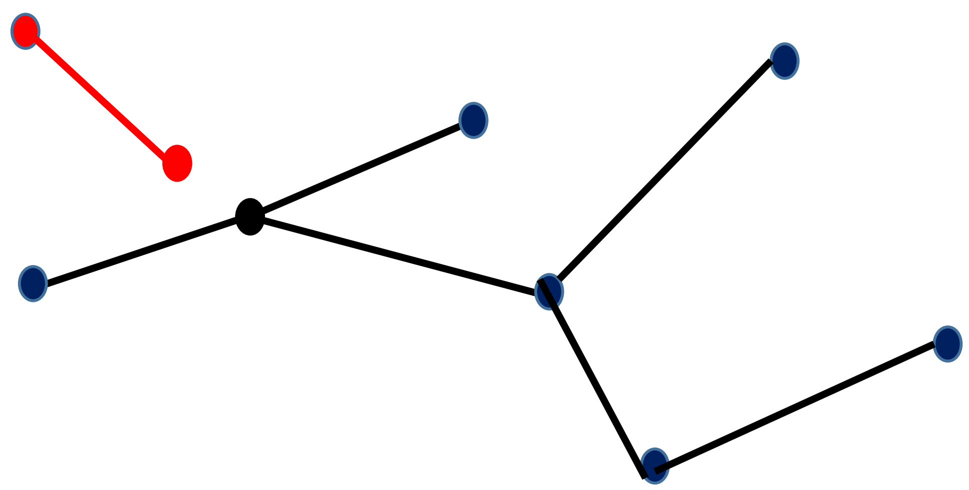

The multilinear polytope of acyclic hypergraphs

- Acyclic hypergraphs in increasing degree of generality:

$$\text{Berge} - \text{acyclic} \;\subset\; \gamma - \text{acyclic} \;\subset\; \beta - \text{acyclic} \;\subset\; \alpha - \text{acyclic}$$

- Optimizing over the multilinear polytope of α-acyclic hypergraphs is NP-hard in general.

- Theorem: The multilinear polytope of an α-acyclic hypergraph of rank r has an extended formulation with at most $O(2^r|V|)$ variables and inequalities.

- If r is upper bounded by the log of a polynomial in the size of the hypergraph, then the multilinear polytope has a polynomial-size extended formulation.

- Equivalent to assuming bounded treewidth for the intersection graph: Wainwright-Jordan 2004, Laurent 2009, Bienstock-Munoz 2018.

α-acyclic hypergraphs

- Theorem: The multilinear polytope of an α-acyclic hypergraph of rank r has an extended formulation with at most $O(2^r|V|)$ variables and inequalities.

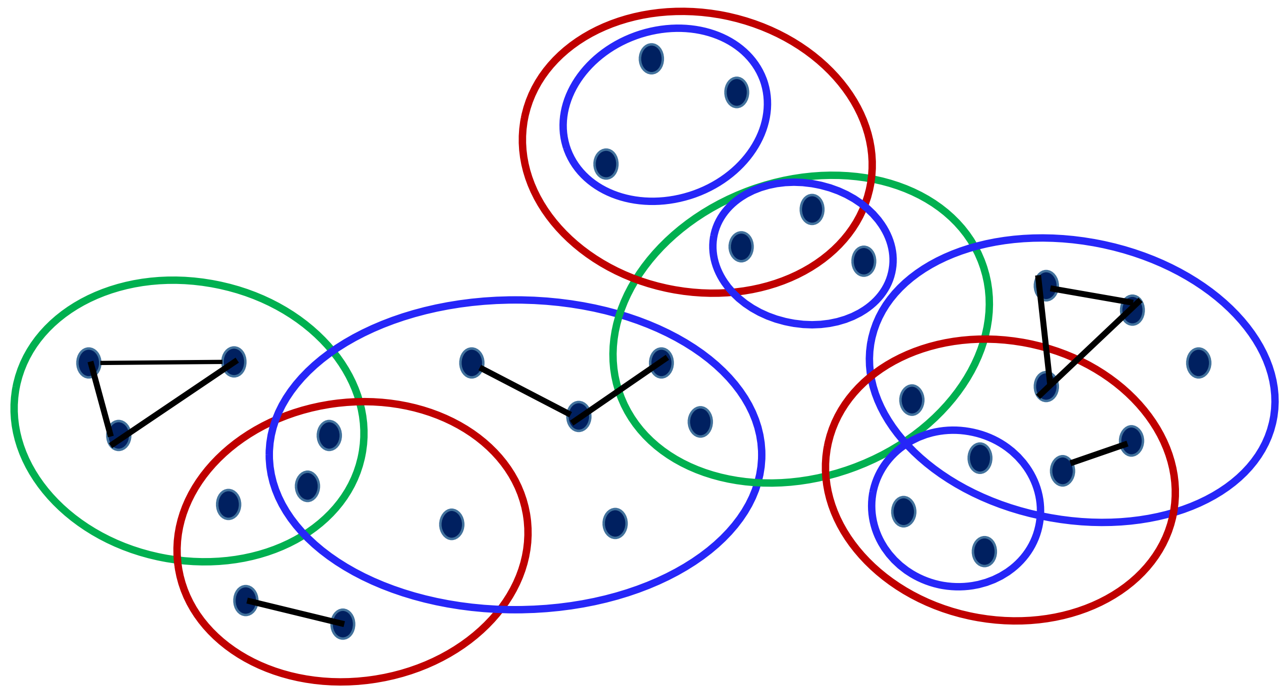

α-acyclic hypergraphs

- Theorem: The multilinear polytope of an α-acyclic hypergraph of rank r has an extended formulation with at most $O(2^r|V|)$ variables and inequalities.

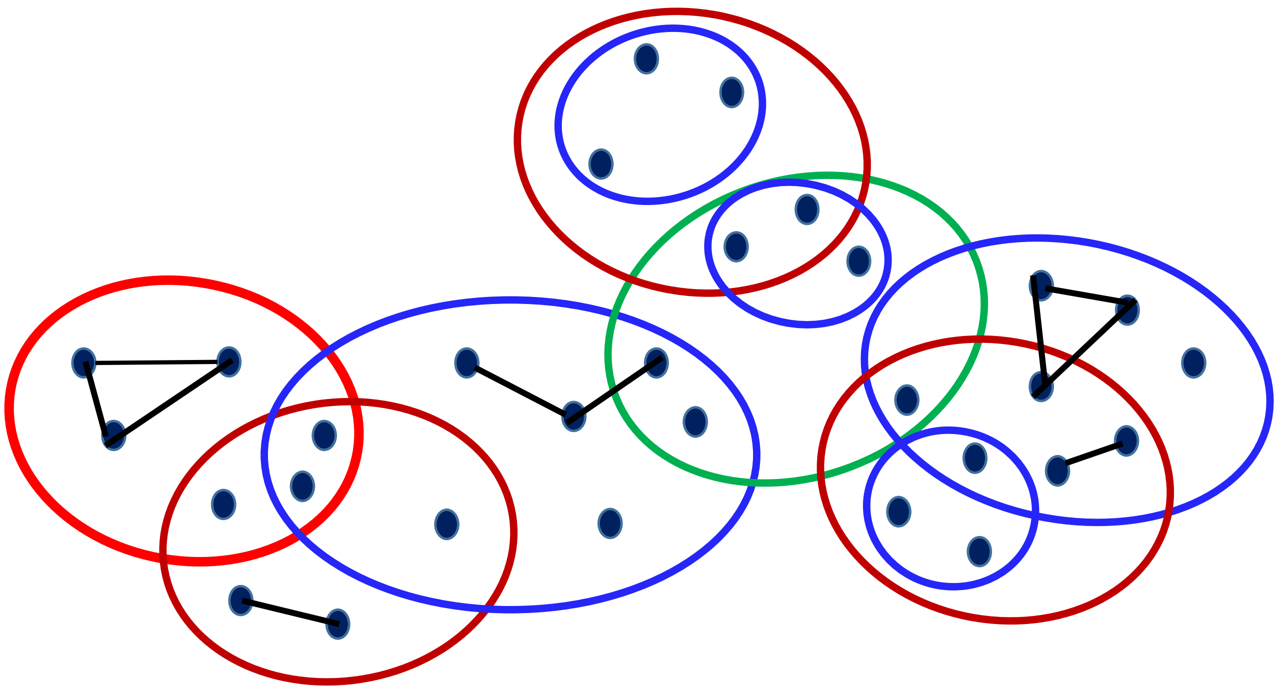

α-acyclic hypergraphs

- Theorem: The multilinear polytope of an α-acyclic hypergraph of rank r has an extended formulation with at most $O(2^r|V|)$ variables and inequalities.

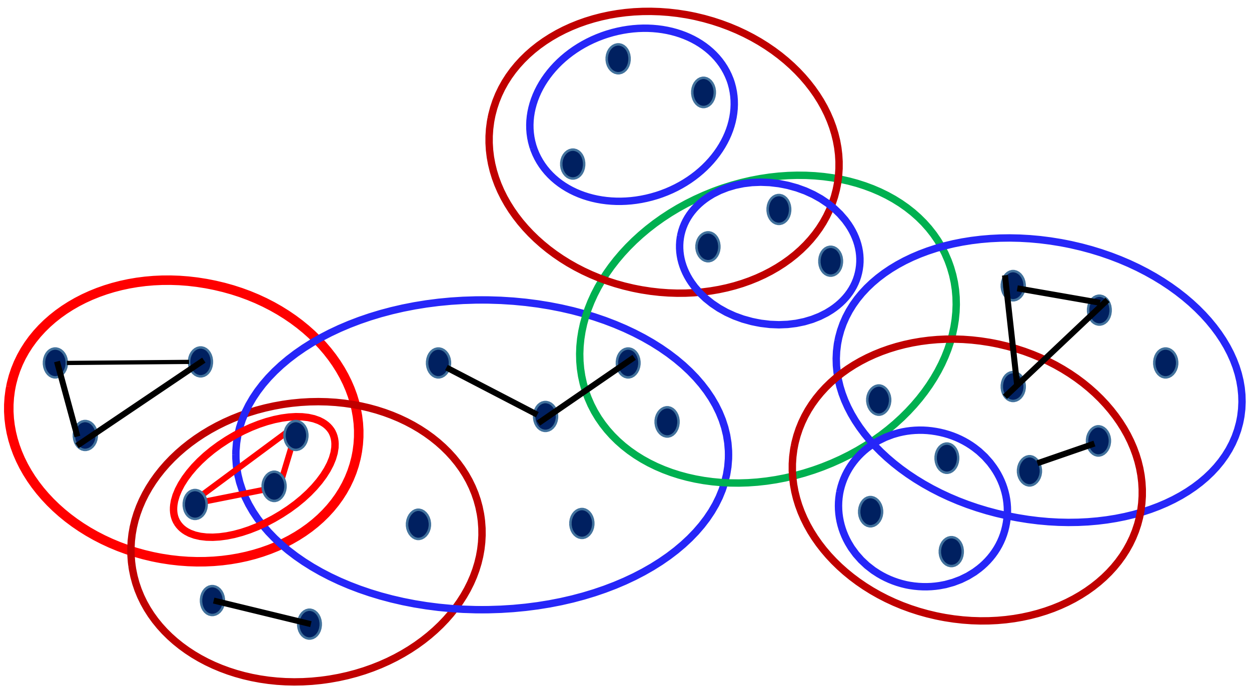

α-acyclic hypergraphs

- Theorem: The multilinear polytope of an α-acyclic hypergraph of rank r has an extended formulation with at most $O(2^r|V|)$ variables and inequalities.

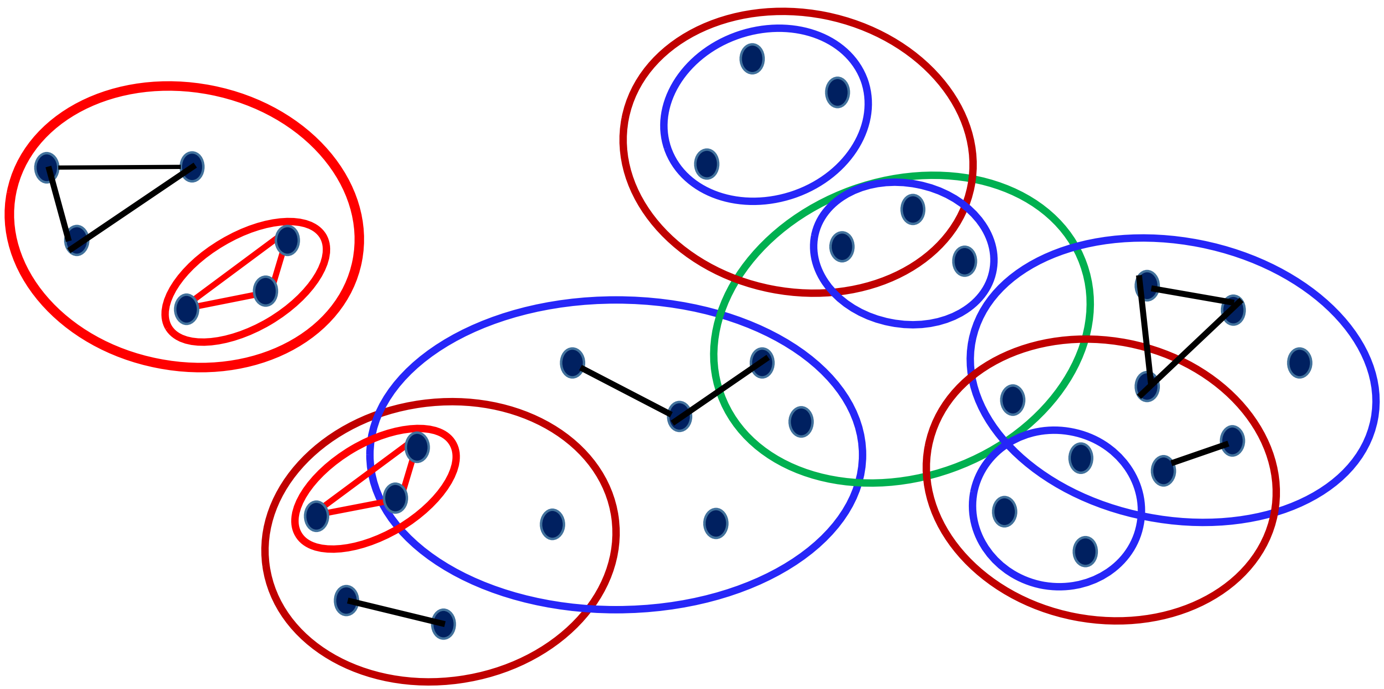

The multilinear polytope of β-acyclic hypergraphs

- A β-cycle of length q for $q \geq 3$ in G is a sequence $v_1, e_1, v_2, e_2, \ldots, v_q, e_q, v_1$ such that $v_1, v_2, \ldots, v_q$ are distinct nodes, $e_1, e_2, \ldots, e_q$ are distinct edges, and v_i belongs to e_{i-1}, e_i and no other edges for all $i = 1, \ldots, q$.

- **Theorem:** The multilinear polytope of β-acyclic hypergraphs has a polynomial-size extended formulation with at most $(r-1)|V| + |E|$ variables ($(r-2)|V|$ additional variables) and at most $(3r-4)|V| + 4|E|$ inequalities.

- The defining inequalities are very sparse: at most four variables with non-zero coefficients. All coefficients are ± 1 and all right-hand sides are $0/1$.

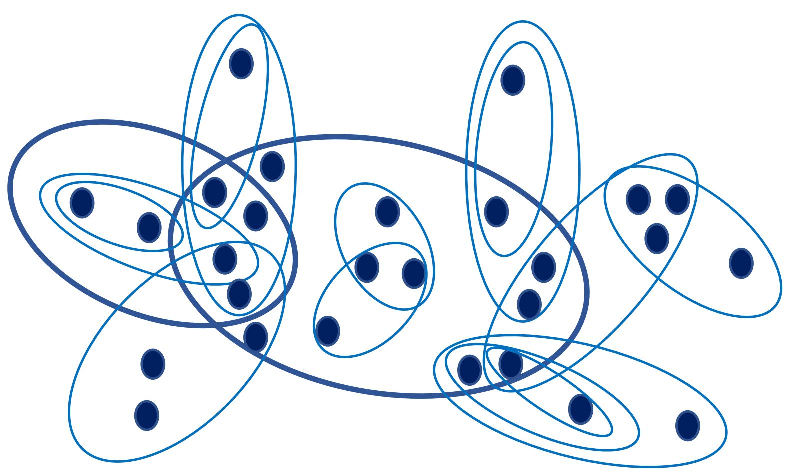

The multilinear polytope of β-acyclic hypergraphs

- **Theorem:** The multilinear polytope of β-acyclic hypergraphs has a polynomial-size extended formulation with at most $(r-1)|V|+|E|$ variables ($(r-2)|V|$ additional variables) and at most $(3r-4)|V|+4|E|$ inequalities (ADP and AK, 2023).

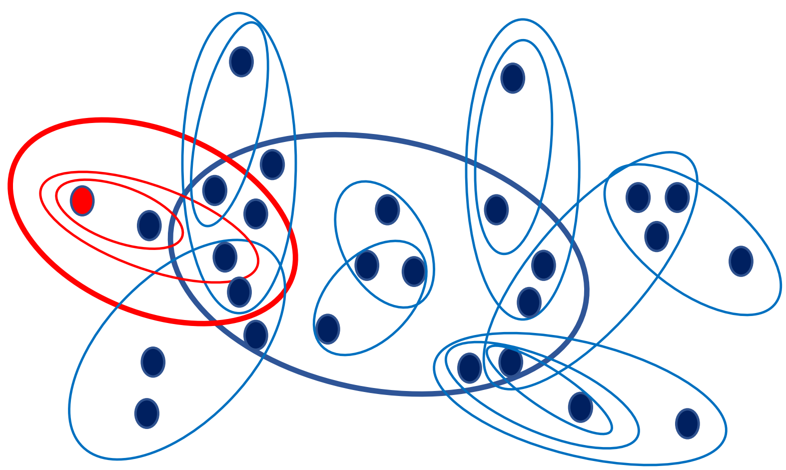

- A node is a nest point (β-leaf) if the edges containing it are totally ordered.

- A hypergraph is β-acyclic iff we can recursively remove nest points till obtaining an empty set.

The multilinear polytope of β-acyclic hypergraphs

- **Theorem:** The multilinear polytope of β-acyclic hypergraphs has a polynomial-size extended formulation with at most $(r-1)|V|+|E|$ variables ($(r-2)|V|$ additional variables) and at most $(3r-4)|V|+4|E|$ inequalities (ADP and AK, 2023).

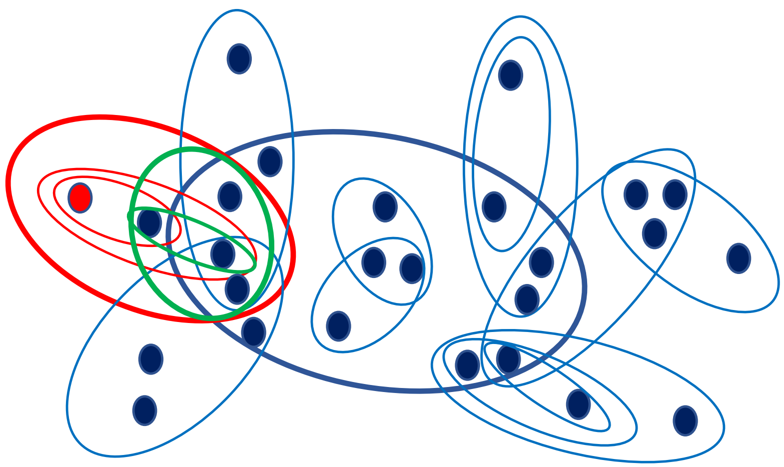

The multilinear polytope of β-acyclic hypergraphs

- **Theorem:** The multilinear polytope of β-acyclic hypergraphs has a polynomial-size extended formulation with at most $(r-1)|V|+|E|$ variables ($(r-2)|V|$ additional variables) and at most $(3r-4)|V|+4|E|$ inequalities (ADP and AK, 2023).

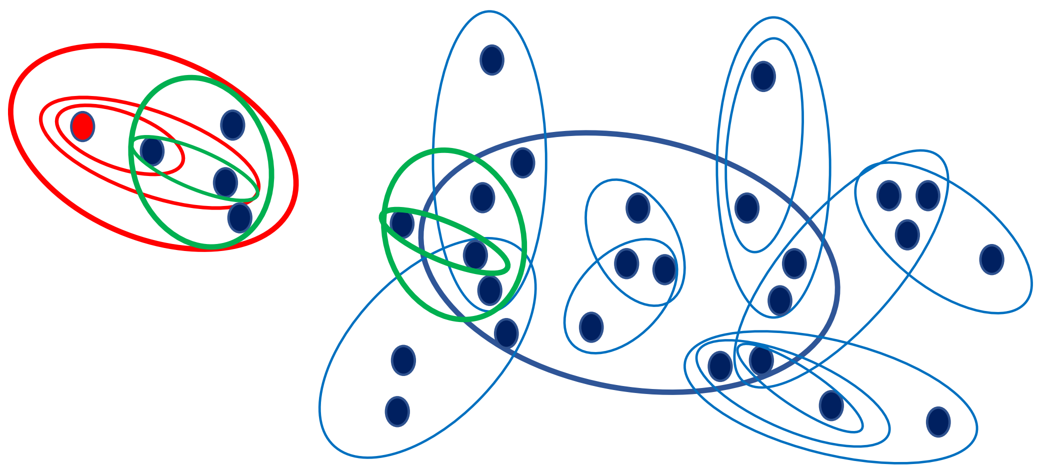

- The multilinear polytope of a pointed hypergraph $G=(V,E)$ consists of $5|V|+2$ inequalities.

The original space

- Convex hull characterizations and valid inequalities for the multilinear polytope in the original space
- Strong LP relaxations for general mixed-integer polynomial optimization problems

Standard linearization of multilinear sets

- Replace each multilinear term $z_e = \prod_{v \in e} z_v$, by its convex hull over the unit hypercube and use $\bigcap_i \mathrm{conv}(\mathcal{S}_i) \supseteq \mathrm{conv}(\bigcap_i \mathcal{S}_i)$ to obtain the standard linearization ${\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G)$ of $\mathcal{S}(G)$:

$$
\begin{aligned}
{\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G) \quad &= \Big\{ z: \quad z_v \leq 1,\ \forall v \in V, z_e \geq 0,\ z_e \geq \sum_{v \in e} z_v - |e| + 1,\ \forall e \in E, \\
& \qquad z_e \leq z_v, \forall v \in e,\ \forall e \in E \Big\}.
\end{aligned}
$$

- Recall that $\mathsf{QP}(G) = \mathsf{QP}^{\mathsf{LP}}(G)$ iff G is an acyclic graph (Padberg 89).

- Theorem: ${\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G) = \mathcal{P}_{\mathrm{m}}(G)$ if and only if G is a Berge-acyclic hypergraph.

The standard linearization vs. the convex hull relaxation

- Theorem: ${\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G) = \mathcal{P}_{\mathrm{m}}(G)$ if and only if G is a Berge-acyclic hypergraph.

- Proof sketch:

  – If G has a Berge-cycle of length two; i.e., $E(C) = \{e_1, e_2\}$ with $|e_1 \cap e_2| \geq 2$, the following is valid for $\mathcal{S}_G$:

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

  Consider $\tilde{z}_v = 1$ for all $v \in e_2 \setminus e_1$, $\tilde{z}_v = 1/2$ for all $v \in e_1$, $\tilde{z}_v = 0$ for the remaining nodes in G, $\tilde{z}_{e_1} = 1/2$, $\tilde{z}_{e_2} = 0$, $\tilde{z}_e = 1$ for all $e \subseteq e_2 \setminus e_1$, $\tilde{z}_e = 0$ for all $e \not\subseteq e_1 \cup e_2$ and $\tilde{z}_e = 1/2$ for all remaining edges in G. $\tilde{z} \in {\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G)$. Substituting $\tilde{z}$ in the above inequality yields $|e_2 \setminus e_1| + 1/2 - 0 \not\leq |e_2 \setminus e_1|$.

  – Let C be a Berge-cycle of minimum length t, where $t \geq 3$. Since $|e_i \cap e_j| \leq 1$ for all $e_i, e_j \in E$, the subhypergraph $G_{V(C)}$ is a graph consisting of a chordless cycle. To show $\mathcal{P}_{\mathrm{m}}(G) \subset {\mathcal{P}_{\mathrm{m}}}^{LP}(G)$ is suffices to show that $\mathcal{P}_{\mathrm{m}}(G_{V(C)}) \subset {\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G_{V(C)})$. The polytope $\mathcal{P}_{\mathrm{m}}(G_{V(C)})$ is integral while ${\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G_{V(C)})$ is not integral.

$\Rightarrow$ if G contains a Berge-cycle, we have $\mathcal{P}_{\mathrm{m}}(G) \subset {\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G)$

The standard linearization vs. the convex hull relaxation

- Suppose that G is a Berge-acyclic hypergraph. Then there exists an edge $\tilde{e}$ of G such that $\tilde{e} \cap \{v : \exists e \in E(G) \setminus \tilde{e}, v \in e\} = \{\tilde{v}\}$, for some $\tilde{v} \in V(G)$.

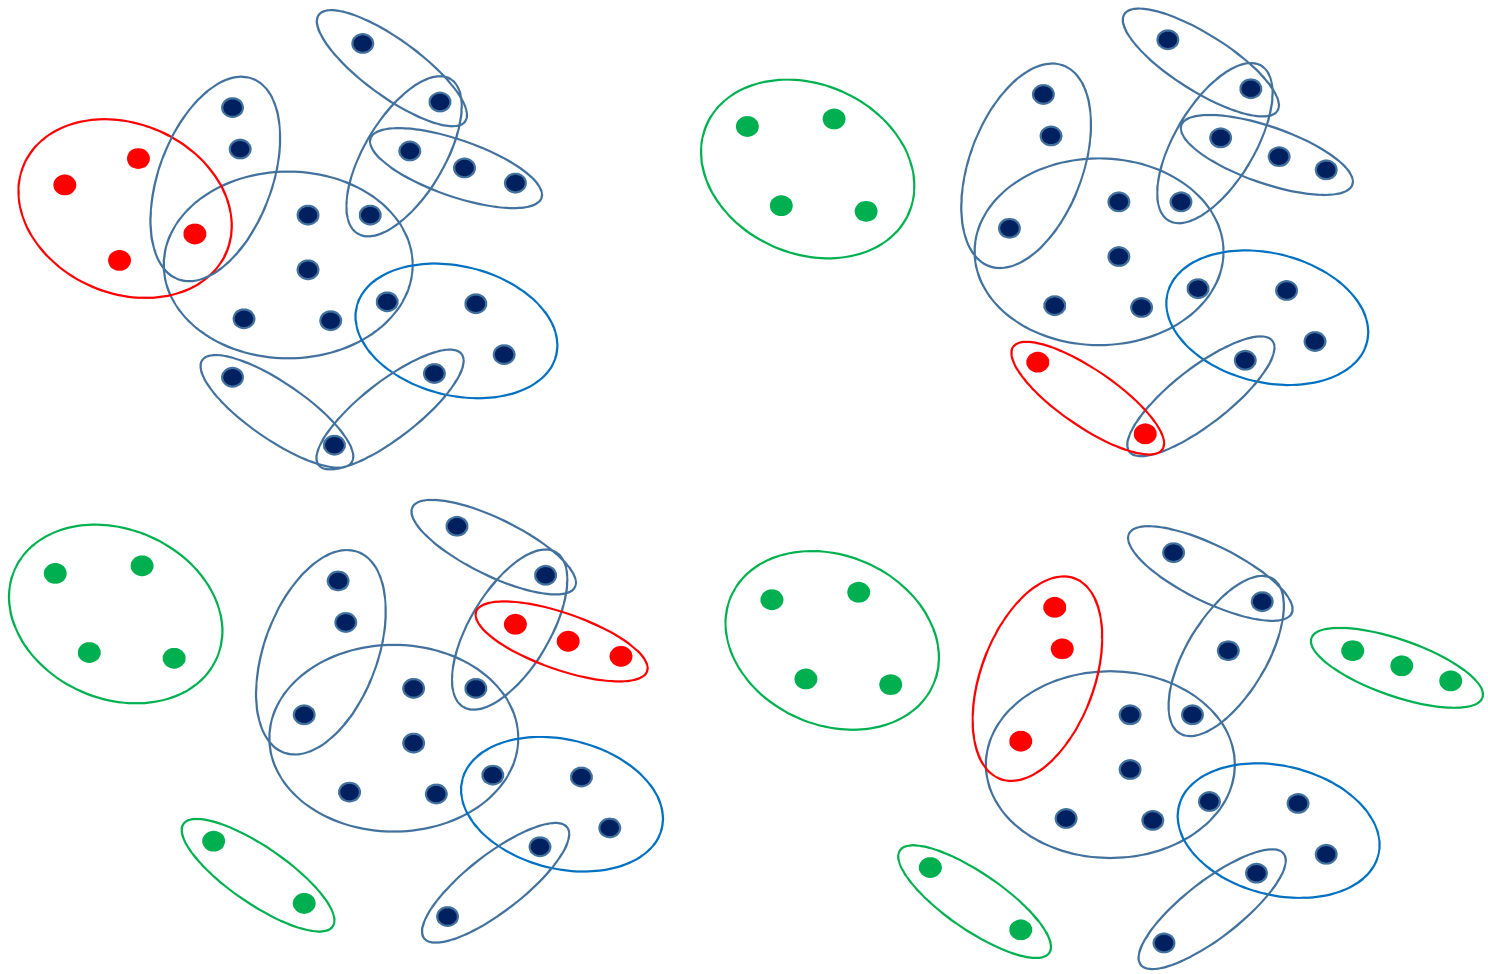

$\Rightarrow$ if G is Berge-acyclic, we have $\mathcal{P}_{\rm m}(G) = {\mathcal{P}_{\rm m}}^{\rm LP}(G)$

Flower inequalities

- Let $e_0 \in E$ and let e_k, $k \in K$, be the set of all edges adjacent to e_0. Let $T \subseteq K$ such that

$$\left|(e_0 \cap e_i) \setminus \bigcup_{j \in T \setminus \{i\}} (e_0 \cap e_j)\right| \geq 2, \quad \forall i \in T.$$

- The flower inequality centered at e_0 with neighbors e_k, $k \in T$ is:

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

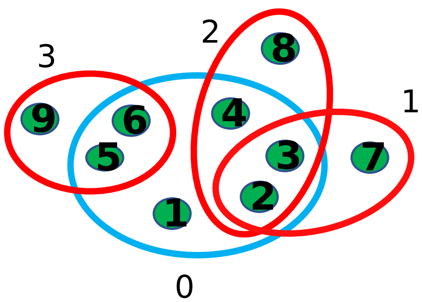

$$z_1 + z_4 + z_5 + z_6 + z_{e_1} - z_{e_0} \leq 4$$

$$z_1 + z_5 + z_6 + z_{e_2} - z_{e_0} \leq 3,$$

$$z_1 + z_2 + z_3 + z_4 + z_{e_3} - z_{e_0} \leq 4,$$

$$z_1 + z_4 + z_{e_1} + z_{e_3} - z_{e_0} \leq 3, \; z_1 + z_{e_2} + z_{e_3} - z_{e_0} \leq 2$$

- The flower relaxation $\mathcal{P}_{\mathrm{m}}{}^F(G)$ is obtained by adding the system of flower inequalities centered at each edge of G to $\mathcal{P}_{\mathrm{m}}{}^{\mathrm{LP}}(G)$.

The flower relaxation vs. the convex hull relaxation

- Theorem: ${\mathcal{P}_{\mathrm{m}}}^{F}(G) = \mathcal{P}_{\mathrm{m}}(G)$ if and only if G is a γ-acyclic hypergraph.

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

- The separation problem for flower inequalities over general hypergraphs is NP-hard (reduction from 3D matching).

- The separation problem for flower inequalities and running intersection inequalities for fixed-rank hypergraphs can be solved in polynomial-time. More precisely, in $O(|E|(r2^r|E| + 2^{r^2}r^3))$ operations.

Flower relaxation vs. recursive McCormick relaxations

- Recursive McCormick (RMC) relaxations are among the most popular convexification techniques for binary polynomial optimization.

- Write each multilinear term $z_e = \prod_{v \in e} z_v$ in a lifted space of variables as a collection of bilinear equations $z_{I \cup J} = z_I z_J$.

- Example: consider $z_{1234} = z_1 z_2 z_3 z_4$, then we have $z_{12} = z_1 z_2$, $z_{34} = z_3 z_4$, $z_{1234} = z_{12} z_{34}$ or $z_{1234} = z_1 z_{234}$, $z_{234} = z_2 z_{34}$ and $z_{34} = z_3 z_4$.

- The quality and the size of these relaxations depend on the recursive sequence and finding an optimal sequence amounts to solving a difficult combinatorial optimization problem.

- Theorem: Let G be a hypergraph and let $\mathcal{P}_{\mathrm{m}}{}^{\mathrm{RMC}}(G)$ denote an RMC relaxation of $\mathcal{S}(G)$. Then $\mathcal{P}_{\mathrm{m}}{}^{\mathrm{F}}(G) \subseteq \mathcal{P}_{\mathrm{m}}{}^{\mathrm{RMC}}(G)$.

Running intersection inequalities

- A multiset F of subsets of a finite set V has the running intersection property if there exists an ordering $p_1, p_2, \ldots, p_m$ of the sets in F such that

$$\forall k \in \{2, \ldots, m\}, \exists j < k : \; N(p_k) := p_k \cap \Big(\bigcup_{i<k} p_i \Big) \subseteq p_j.$$

  We refer to such an ordering as a running intersection ordering of F.

- Let e_0 and e_k, $k \in K$, be a collection of edges adjacent to e_0 such that $\tilde{E} := \{e_0 \cap e_k : k \in K\}$ has the running intersection property. Consider a running intersection ordering of $\tilde{E}$. For each $k \in K$, let $w_k \subseteq N(e_0 \cap e_k)$ such that $w_k \in \emptyset \cup V \cup E$. We define a running intersection inequality centered at e_0 with neighbours e_k, $k \in K$ as

$$-\sum_{k \in K} z_{w_k} + \sum_{v \in e_0 \setminus \bigcup_{k \in K} e_k} z_v + \sum_{k \in K} z_{e_k} - z_{e_0} \leq \omega - 1,$$

  where we define $z_\emptyset = 0$, and $\omega = \Big|\{k \in K : w_k = \emptyset\}\Big| + \Big|\{e_0 \setminus \bigcup_{k \in K} e_k\}\Big|$.

Running intersection inequalities

- Letting $w_k = \emptyset$ for all $k \in K$, running intersection inequalities simplify to flower inequalities.

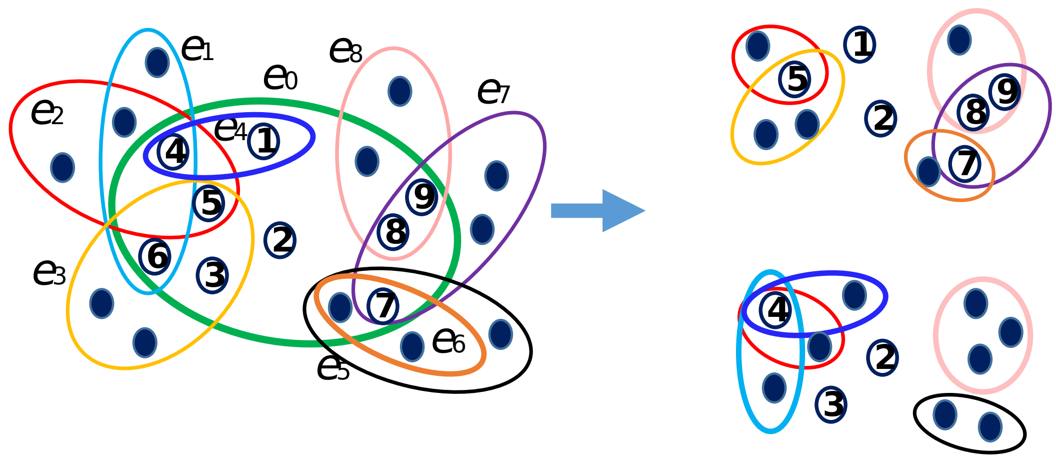

$$-z_{v_5} - z_{v_7} - z_{v_8} + z_{v_1} + z_{v_2} + z_{e_2} + z_{e_3} + z_{e_6} + z_{e_7} + z_{e_8} - z_{e_0} \leq 3,$$
$$-z_{v_5} - z_{v_7} - z_{v_9} + z_{v_1} + z_{v_2} + z_{e_2} + z_{e_3} + z_{e_6} + z_{e_7} + z_{e_8} - z_{e_0} \leq 3,$$
$$-2z_{v_4} + z_{v_2} + z_{v_3} + z_{e_1} + z_{e_2} + z_{e_4} + z_{e_5} + z_{e_8} - z_{e_0} \leq 4$$

- Any running intersection ordering of $\tilde{E}$ leads to the same system of running intersection inequalities centered at e_0 with neighbors e_k, $k \in K$.

The running intersection relaxation

- The running intersection relaxation ${\mathcal{P}_{\mathrm{m}}}^{\mathrm{RI}}(G)$ is the polytope obtained by adding to ${\mathcal{P}_{\mathrm{m}}}^{\mathsf{LP}}(G)$ all possible running intersection inequalities for $\mathcal{S}(G)$.

- If $\mathcal{P}_{\mathrm{m}}(G)$ is not β-acyclic, then $\mathcal{P}_{\mathrm{m}}(G) \subset {\mathcal{P}_{\mathrm{m}}}^{\mathrm{RI}}(G)$.

- Let G be a β-acyclic hypergraph. Suppose that there exist no three edges $e_0, e_1, e_2 \in E$ such that $|e_0 \cap e_1 \cap e_2| \geq 2$, $(e_0 \cap e_1) \setminus e_2 \neq \emptyset$, and $(e_0 \cap e_2) \setminus e_1 \neq \emptyset$. Then $\mathcal{P}_{\mathrm{m}}(G) = {\mathcal{P}_{\mathrm{m}}}^{\mathrm{RI}}(G)$.

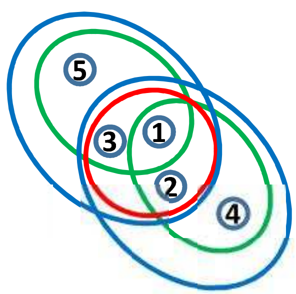

$$-z_1 + z_{123} + z_{124} + z_{135} - z_{1234} - z_{1235} \leq 0$$

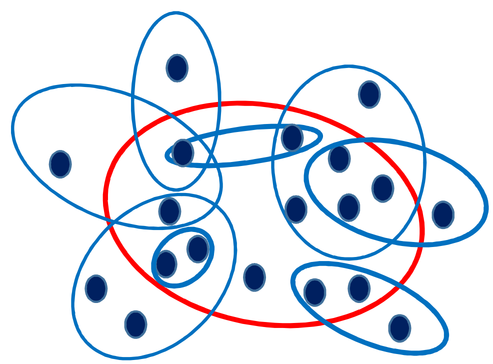

What about the multilinear polytope of β-acyclic hypergraphs?

- From a computational perspective, sparsity is key to the effectiveness of cutting planes in a branch-and-cut framework.

- For a rank r hypergraph, flower inequalities contain at most $\frac{r}{2}$ nonzero coefficients, and running intersection inequalities contain at most $2(r-1)$ nonzero coefficients.

- For β-acyclic hypergraphs, $\mathcal{P}_{\mathrm{m}G}$ may contain dense facet-defining inequalities with $\theta(|E|)$ nonzero coefficients.

- In practice, we almost always have $r \ll |E|$.

Example

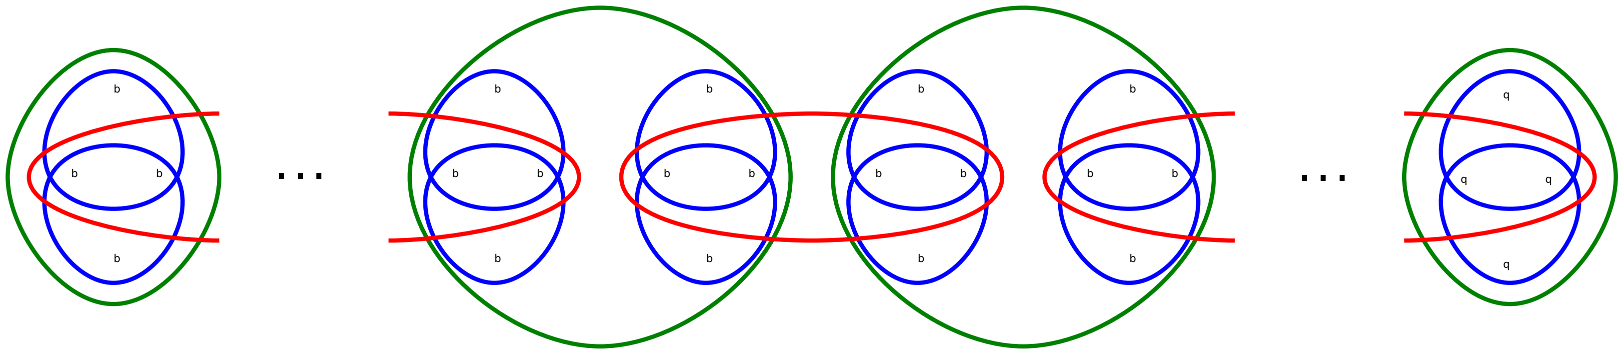

- Let $n \geq 2$ and consider the β-acyclic hypergraph $G = (V, E)$ with $V = \bigcup_{i \in [n]} V^i$, $E = H \cup \bigcup_{i \in [n]} E^i$, where $V^1 = \{v_3^1, v_4^1, v_7^1, v_8^1\}$, $V^i = \{v_1^i, \cdots, v_8^i\}$ for all $i \in [n-1] \setminus \{1\}$, $V^n = \{v_1^n, v_2^n, v_5^n, v_6^n\}$,

$$H = \Big\{\{v_3^i, v_4^i, v_1^{i+1}, v_2^{i+1}\},\ i \in [n-1]\Big\}$$

$$E^1 = \Big\{\{v_3^1, v_4^1, v_7^1\}, \{v_3^1, v_4^1, v_8^1\}, V^1\Big\}$$

$$E^i = \Big\{\{v_1^i, v_2^i, v_5^i\}, \{v_1^i, v_2^i, v_6^i\}, \{v_3^i, v_4^i, v_7^i\}, \{v_3^i, v_4^i, v_8^i\}, V^i\Big\}, \quad \forall i \in [n-1] \setminus \{1\}$$

$$E^n = \Big\{\{v_1^n, v_2^n, v_5^n\}, \{v_1^n, v_2^n, v_6^n\}, V^n\Big\}.$$

- The following inequality containing $|E|$ nonzero coefficients defines a facet of $\mathcal{P}_{\mathrm{m}}(G)$:

$$-\sum_{i \in [n]} z_{V^i} - \sum_{e \in H} z_e + \sum_{i \in [n]} \sum_{e \in E^i \setminus \{V^i\}} z_e \leq 2n - 3.$$

The complete edge relaxation

- Let $G = (V, E)$ be a hypergraph and let $\bar{E} \subseteq E$ the set of maximal edges of G. For each $\bar{e} \in \bar{E}$, let $G^{\bar{e}}$ be the complete hypergraph with the node set $\bar{e}$. We define the complete edge relaxation of $\mathcal{P}_{\mathrm{m}}(G)$, denoted by $\mathcal{P}_{\mathrm{m}}{}^{\mathrm{CE}}(G)$, as the polytope obtained by putting together the descriptions of $\mathcal{P}_{\mathrm{m}}(G^{\bar{e}})$ for all $\bar{e} \in \bar{E}$.

- The edge complete relaxation is stronger than the running intersection relaxation which in turn is stronger than the flower relaxation.

- Theorem: $\mathcal{P}_{\mathrm{m}}(G) = \mathcal{P}_{\mathrm{m}}{}^{\mathrm{CE}}(G)$ if and only if G is an α-acyclic hypergraph.

Numerical Experiments

- We characterize each problem by its degree (d), number of variables (n), number of constraints (q), and density (ν).
- Polynomial problems of degree 3 with

$$(n, \nu) \in \{(10, 0.75), (15, 0.25), (15, 0.15), (20, 0.1), (20, 0.05)\},$$

  and multilinear problems of degree 3 with

$$(n, \nu) \in \{(10, 1.0), (15, 0.5), (20, 0.15), (20, 0.1), (25, 0.05), (30, 0.02)\}.$$

- Polynomial problems of degree 4 with

$$(n, \nu) \in \{(10, 0.25), (10, 0.15), (15, 0.05), (15, 0.02), (20, 0.01)\},$$

  and multilinear problems of degree 4 with

$$(n, \nu) \in \{(10, 1.0), (15, 0.15), (20, 0.02), (20, 0.01), , (25, 0.01), (25, 0.005)\}.$$

- In both sets, we let $q \in \{0, n/5, n/2, n\}$. For each combination, 5 random instance are generated.
- Relative/absolute optimality tolerance $= 10^{-6}$ and time limit $= 500s$.

220 polynomial optimization problems of degree three

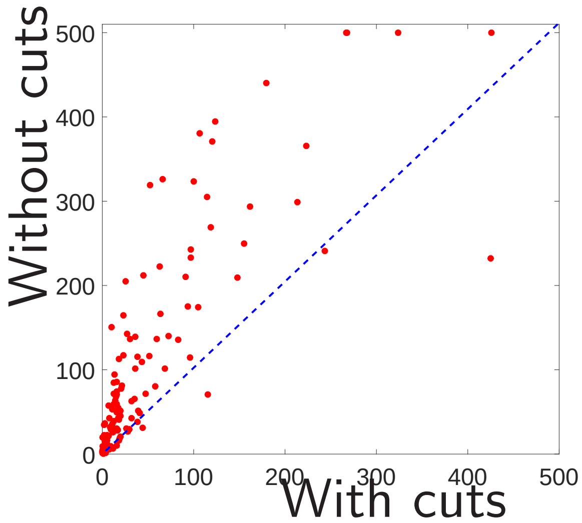

(a) CPU time (s)

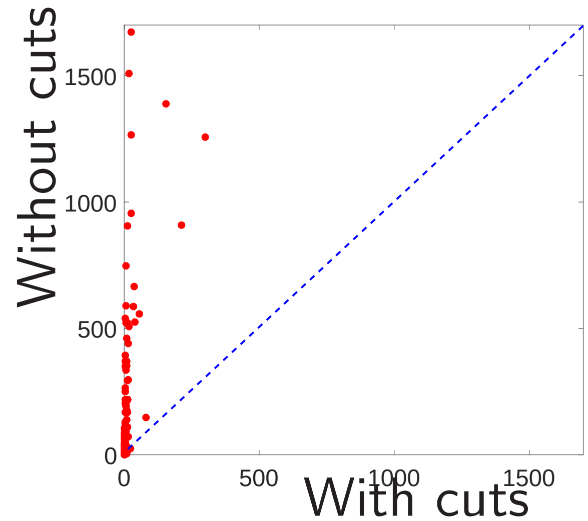

(b) Iterations

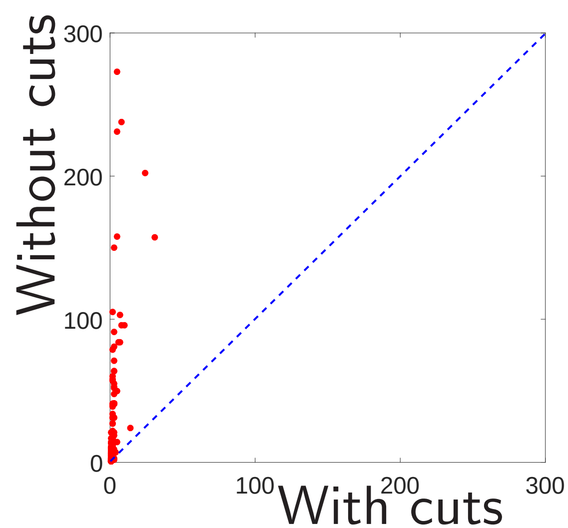

(c) Memory

- **Average reductions of 60% in CPU time, 78% in number of nodes, and 70% in maximum number of nodes in memory.**

220 polynomial optimization problems of degree four

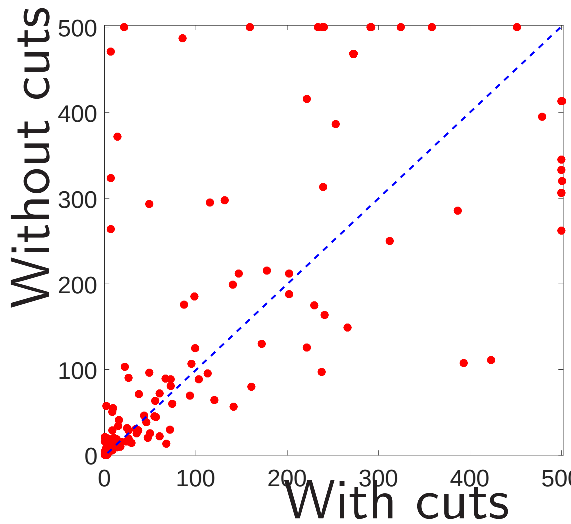

(d) CPU time (s)

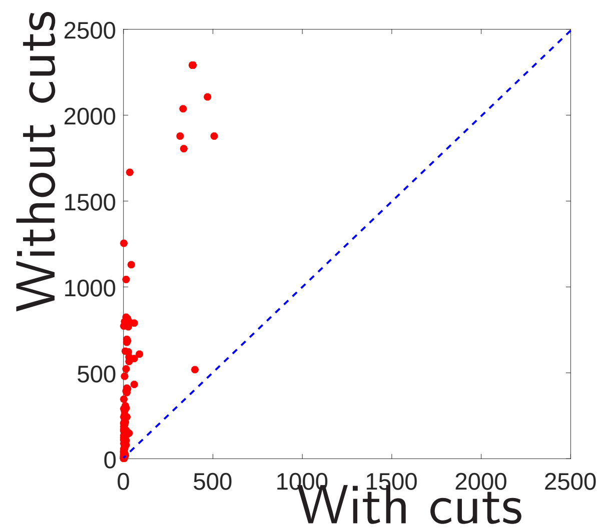

(e) Iterations

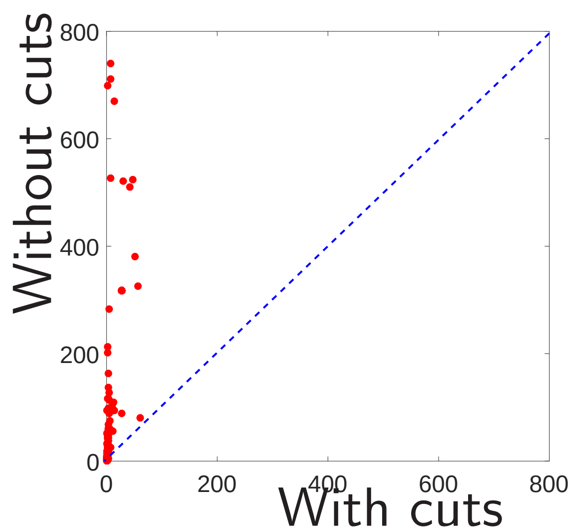

(f) Memory

- **Average reductions of 43% in CPU time, 76% in number of nodes, and 72% in maximum number of nodes in memory.**

Numerical Experiments – computer vision instances

- The purpose of image restoration is to estimate the original image from the degraded data. An image is modeled as a $l \times h$ matrix where each binary element x_{ij} represents a pixel.
- The image restoration problem is defined as the objective function $f(x) = H(x) + L(x)$ to be minimized, where $H(x)$ is linear and models similarity between the input blurred image and the output, $L(x)$ is a multilinear function of degree four and models smoothness.
- Test set taken from [CramaRodrigez16] with images sizes $\{10 \times 10\}$, $\{10 \times 15\}$, $\{15 \times 15\}$.

Effect of adding cuts	CPU time	Iterations	Nodes
Better by a factor at least 2	17 (38%)	10 (23%)	10 (23%)
Between 30% and 100% better	13 (30%)	0 (0%)	0 (0%)
Difference smaller than 30%	14 (32%)	34 (77%)	34 (77%)
Between 30% and 100% worse	0 (0%)	0 (0%)	0 (0%)
Worse by a factor of at least 2	0 (0%)	0 (0%)	0 (0%)

- **Average reductions of 63% in CPU time, 42% in number of iterations, and 30% in maximum number of nodes in memory.**

Summary of the first lecture

- We define the multilinear set as the feasible region of a linearized binary polynomial optimization problem and we define the multilinear polytope as the convex hull of the multilinear set.

- We represent multilinear sets/polytopes by hypergraphs.

- Acyclic hypergraphs in increasing degree of generality:

$$\text{Berge}-\text{acyclic} \subset \gamma\ -\text{acyclic} \subset \beta\ -\text{acyclic} \subset \alpha\ -\text{acyclic}$$

- The complexity of facial structure of the multilinear polytope is closely related to the acyclicity degree of the corresponding hypergraph

- Optimizing a linear function over the multilinear polytope of α-acyclic hypergraphs is NP-hard in general; however if the rank is bounded, then the multilinear polytope admits a polynomial-size extended formulation with at most $2^r|V|$ variables and inequalities.

- The multilinear polytope of β-acyclic hypergraphs admits a polynomial-size extended formulation with at most $r|V| + |E|$ variables and inequalities.

Summary of the first lecture

- In the original space:
  - The standard linearization is the multilinear polytope if and only if the hypergraph is Berg-acyclic.
  - The flower relaxation is the multilinear polytope if and only if the hypergraph is γ-acyclic. May have exponentially many facets.
  - We do not have a characterizaton of the multilinear polytope of β-acyclic hypergraphs in the original space.

- The complete edge relaxation is the multilinear polytope if and only if the hypergraph is α-acyclic. Has at most $2^r|V|$ variables and inequalities.

Beyond hypergraph acyclicity

- We present a new framework that
  - unifies all prior results on the existence of polynomial-size extended formulations, and
  - provides polynomial-size extended formulations for the multilinear polytope of hypergraphs with β-cycles

Binary polynomial optimization

- Binary polynomial optimization is the problem of maximizing a multivariate polynomial function over the set of binary points.

- Based on the encoding of the polynomial function, we obtain two popular optimization problems: multilinear optimization and pseudo-Boolean optimization.

- With any $G = (V, E)$, and $c \in \mathbb{R}^{V \cup E}$, we associate the multilinear optimization problem:

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

Signed hypergraphs

- 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 parallel if $e = e'$, and they are identical if $e = e'$ and $\eta_s = \eta_{s'}$.

- We consider signed hypergraphs with no identical signed edges but often with parallel signed edges.

Pseudo-Boolean optimization

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

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

where

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

- Define $z_s := \prod_{v \in s} \sigma_s(z_v)$ for all $s \in S$:

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

Pseudo-Boolean sets and polytopes

- We define the pseudo-Boolean set of the signed hypergraph $H = (V, S)$, as:

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

  and we refer to its convex hull as the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$.

- If $\eta_s(v) = +1$ for all $v \in s$ and all $s \in S$, then the pseudo-Boolean set/polytope coincides with the multilinear set/polytope.

- Unlike the multilinear polytope, the pseudo-Boolean polytope is NOT full dimensional.

- Let $s_1 = \{1^-, 2^+, 3^+\}$, $s_2 = \{1^+, 2^+, 3^+\}$, $s_3 = \{2^+, 3^+\}$, then

$$z_{s_3} = z_{s_1} + z_{s_2}.$$

- Note that $z_{s_1} = (1 - z_1) z_2 z_3$, $z_{s_2} = z_1 z_2 z_3$ and $z_{s_3} = z_2 z_3$.

The underlying hypergraph vs the multilinear hypergraph

- The underlying hypergraph of a signed hypergraph H is the hypergraph obtained from H by ignoring signs and dropping parallel edges.
- The pseudo-Boolean optimization problem over a signed hypergraph $H = (V, S)$ can be reformulated as a multilinear optimization problem over a hypergraph, which we call the multilinear hypergraph $\mathrm{mh}(H)$ of H.
- Let the underlying hypergraph of H be β-acyclic; then the multilinear hypergraph of H may contain many β-cycles.

$$z_s = (1 - z_1)(1 - z_2)(1 - z_3)$$

The pseudo-Boolean polytope versus the multilinear polytope

- Recall that the multilinear polytope is a special case of the pseudo-Boolean polytope.

- Let $H = (V, S)$ be a signed hypergraph. If we have the description of the multilinear polytope of the multilinear hypergraph $\mathrm{mh}(H)$, then we can obtain a description of the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$ using:

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

- Then an extended formulation for $\mathcal{P}_{\mathrm{pB}}(H)$ is given by the description of $\mathcal{P}_{\mathrm{m}}(G)$, where $G := \mathrm{mh}(H)$ together with the following equations:

$$z_s = \sum_{e \in E} d_e z_e + d_0, \quad \forall s \in S.$$

- However, $\mathcal{P}_{\mathrm{m}}(G)$ with $G = \mathrm{mh}(H)$ may contain exponentially many more variables than $\mathcal{P}_{\mathrm{pB}}(H)$.

The recursive inflate and decompose framework

- Main ingredients:

1. A sufficient condition for decomposability of pseudo-Boolean polytopes.

2. A polynomial-size extended formulation for the pseudo-Boolean polytope of pointed signed hypergraphs, which appears as a result of applying the decomposition technique.

3. The inflation operation that we use to transform a large class of signed hypergraphs to those for which our results of Parts 1 and 2 are applicable.

Decomposability of pseudo-Boolean polytopes

- 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 $H_2 := (V_2, S_2)$.

- We say $\mathcal{P}_{\mathrm{pB}}(H)$ is decomposable into $\mathcal{P}_{\mathrm{pB}}(H_1)$ and $\mathcal{P}_{\mathrm{pB}}(H_2)$, if the system comprised of a description of $\mathcal{P}_{\mathrm{pB}}(H_1)$ and a description of $\mathcal{P}_{\mathrm{pB}}(H_2)$, is a description of $\mathcal{P}_{\mathrm{pB}}(H)$.

- Theorem: Assume the underlying hypergraph of H has a nest point v. Let $s_1 \subseteq s_2 \subseteq \cdots \subseteq s_k$ be the signed edges of H containing v, and assume S contains $s_i - v$ for all $i \in [k]$. Then $\mathcal{P}_{\mathrm{pB}}(H)$ is decomposable into $\mathcal{P}_{\mathrm{pB}}(H_1)$ and $\mathcal{P}_{\mathrm{pB}}(H_2)$, where $H_1 := (V_1, S_v \cup P_v)$, V_1 is the underlying edge of s_k, $S_v := \{s_1, \ldots, s_k\}$, $P_v := \{s_i - v : |s_i - v| \geq 2,\ i \in [k]\}$, and $H_2 := H - v$.

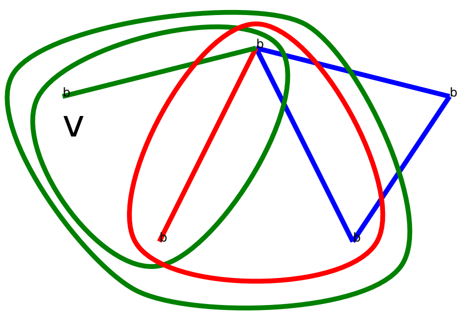

G

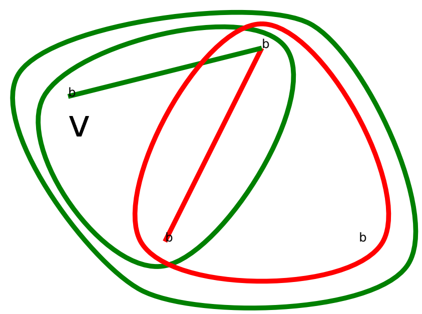

G_1

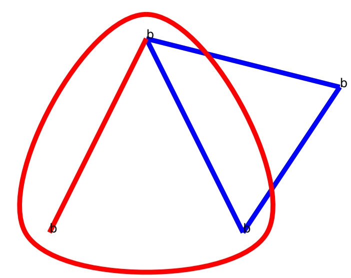

G_2

The pseudo-Boolean polytope of pointed signed hypergraphs

- Consider a signed hypergraph $H = (V, S)$ and let $v \in V$ be a nest point 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 s a pointed signed hypergraph if V coincides with the underlying edge of the signed edge of maximum cardinality in S_v and $S = S_v \cup P_v$.

- Theorem: Let $H = (V, S)$ be a pointed signed hypergraph. Then $\mathcal{P}_{\mathrm{pB}}(H)$ 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 $\mathcal{P}_{\mathrm{pB}}(H)$ are $0, \pm 1$.

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is β-acyclic. Then the pseudo-Boolean polytope has a polynomial-size extended formulation with at most $O(r|S||V|)$ variables and inequalities.

Inflation of signed edges

- Let $H = (V, S)$ be a signed hypergraph, let $s \in S$, and let $e \subseteq V$ such that $s \subset e$. let $I(s, e)$ be the set of all possible signed edges s' parallel to e such that $\eta_s(v) = \eta_{s'}(v)$ for every $v \in s$. Then $H' = (V, S')$ is obtained from H by inflating s to e if $S' = S \cup I(s, e) \setminus \{s\}$.

- Theorem: Let $H' = (V, S')$ be obtained from H by inflating s to e. Then an extended formulation of $\mathcal{P}_{\mathrm{pB}}(H)$ can be obtained by an extended formulation of $\mathcal{P}_{\mathrm{pB}}(H')$ and

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

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

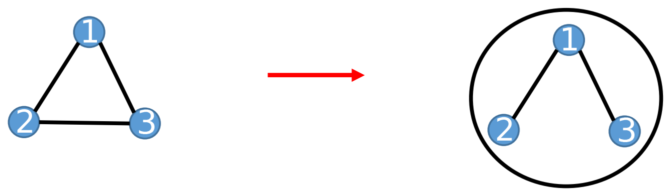

$s_1 = \{v_1^+, v_2^+\}$, $s_2 = \{v_1^+, v_3^+\}$, $s_3 = \{v_3^+, v_2^+\}$

$s_4 = \{v_1^-, v_2^+, v_3^+\}$, $s_5 = \{v_1^+, v_2^+, v_3^+\}$

$z_{s_3} = z_{s_4} + z_{s_5}$

Inflation of signed edges

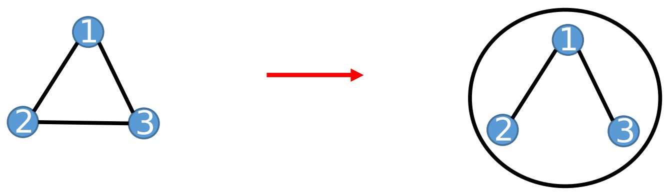

$s_1 = \{v_1^+, v_2^+\}, s_2 = \{v_1^+, v_3^+\}, s_3 = \{v_3^+, v_2^+\}$

$s_4 = \{v_1^-, v_2^+, v_3^+\}, s_5 = \{v_1^+, v_2^+, v_3^+\}$

$z_{s_3} = z_{s_4} + z_{s_5}$

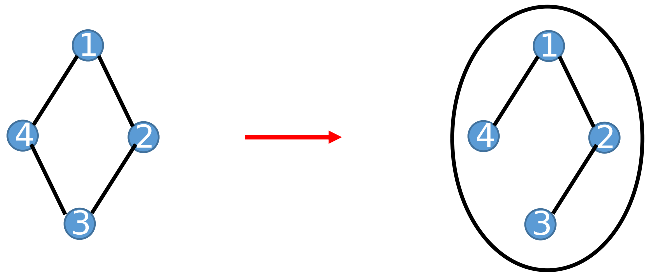

$s_1 = \{v_1^+, v_2^+\}, s_2 = \{v_2^+, v_4^+\}, s_3 = \{v_3^+, v_4^+\}, s_4 = \{v_1^+, v_4^+\}$

$s_5 = \{v_1^+, v_2^+, v_3^+, v_4^+\}, s_6 = \{v_1^-, v_2^+, v_3^+, v_4^+\}$

$s_7 = \{v_1^+, v_2^-, v_3^+, v_4^+\}, s_8 = \{v_1^-, v_2^-, v_3^+, v_4^+\}$

$z_{s_3} = z_{s_5} + z_{s_6} + z_{s_7} + z_{s_8}$

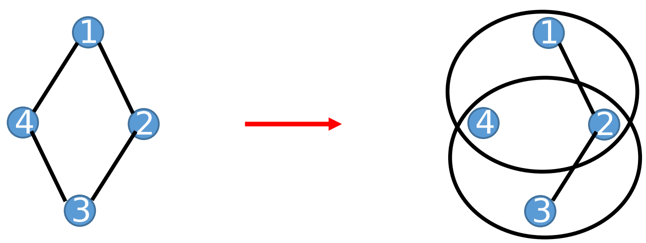

$s_1 = \{v_1^+, v_2^+\}, s_2 = \{v_2^+, v_4^+\}, s_3 = \{v_3^+, v_4^+\}, s_4 = \{v_1^+, v_4^+\}$

$s_5 = \{v_2^+, v_3^+, v_4^+\}, s_6 = \{v_2^-, v_3^+, v_4^+\}$

$s_7 = \{v_1^+, v_2^+, v_4^+\}, s_8 = \{v_1^+, v_2^-, v_4^+\}$

$z_{s_3} = z_{s_5} + z_{s_6}, z_{s_4} = z_{s_7} + s_{s_8}$

Applications of inflation

- Consider a signed hypergraph $H = (V, S)$. Suppose that each $s \in S$ contains at least $|V| - k$ nodes. Then the pseudo-Boolean polytope has an extended formulation with $O(2^k|V||S|)$ variables and inequalities.

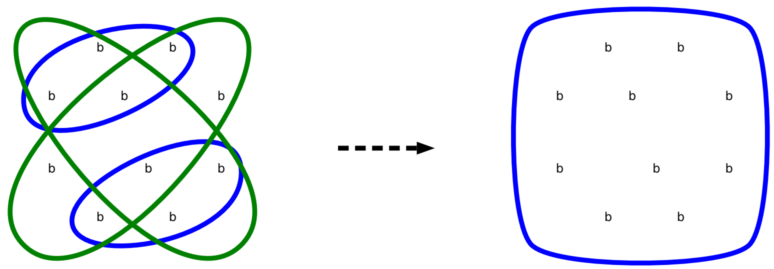

- Consider a signed hypergraph $H = (V, S)$ of rank r. For each $s \in S$, among all maximal signed edges of H containing s, denote by f_s one with minimum cardinality. Let k be such that $|f_s| - |s| \leq k$ for all $s \in S$. Let $\bar{S}$ denote the set of maximal signed edges of H. If the underlying hypergraph of $(V, \bar{S})$ is β-acyclic, then the pseudo-Boolean polytope has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities.

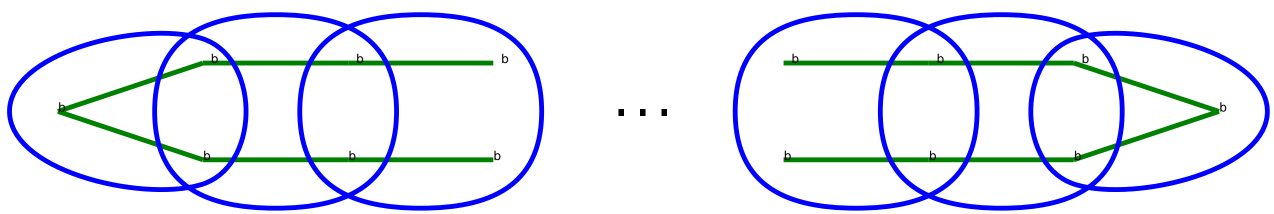

The Recursive inflate-and-decompose (RID) framework

- Input: A signed hypergraph $H = (V, S)$, Output: An extended formulation for $\mathcal{P}_{\rm pB}(H)$.

- 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 $\mathcal{P}_{\rm pB}(\bar{H}^{(i)})$ is available, then we are done. Otherwise, go to Step 2.

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

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

The Recursive inflate-and-decompose (RID) framework

- RID provides a polynomial-size extended formulation for $\mathcal{P}_{\mathrm{pB}}(H)$ if the following conditions are satisfied:
  - In Step 1, the algorithm should terminate, only if a polynomial-size extended formulation for $\mathcal{P}_{\mathrm{pB}}(\bar{H}^{(i)})$ is available.
  - 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 $|V|, |S|$.

- A simple way to obtain a nest point in Step 2 is to inflate each signed edge containing $\bar{v}$ to the union of all signed edges containing $\bar{v}$.

α-acyclic hypergraphs with log-poly ranks

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is α-acyclic. Then $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with at most $O(3^r|V|)$ variables and inequalities.

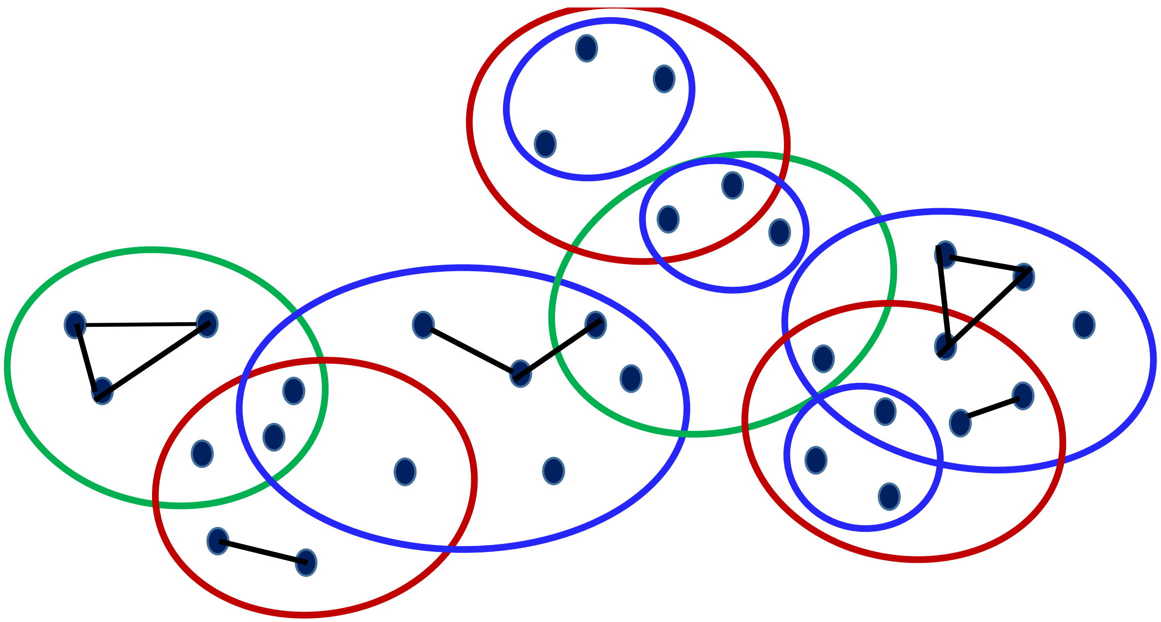

- A node $v \in V$ is an α-leaf if the set of edges containing v has a maximal element for inclusion.

- A hypergraph is α-acyclic iff we can recursively remove α-leaves till obtaining an empty set.

α-acyclic hypergraphs with log-poly ranks

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is α-acyclic. Then $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with at most $O(3^r|V|)$ variables and inequalities.

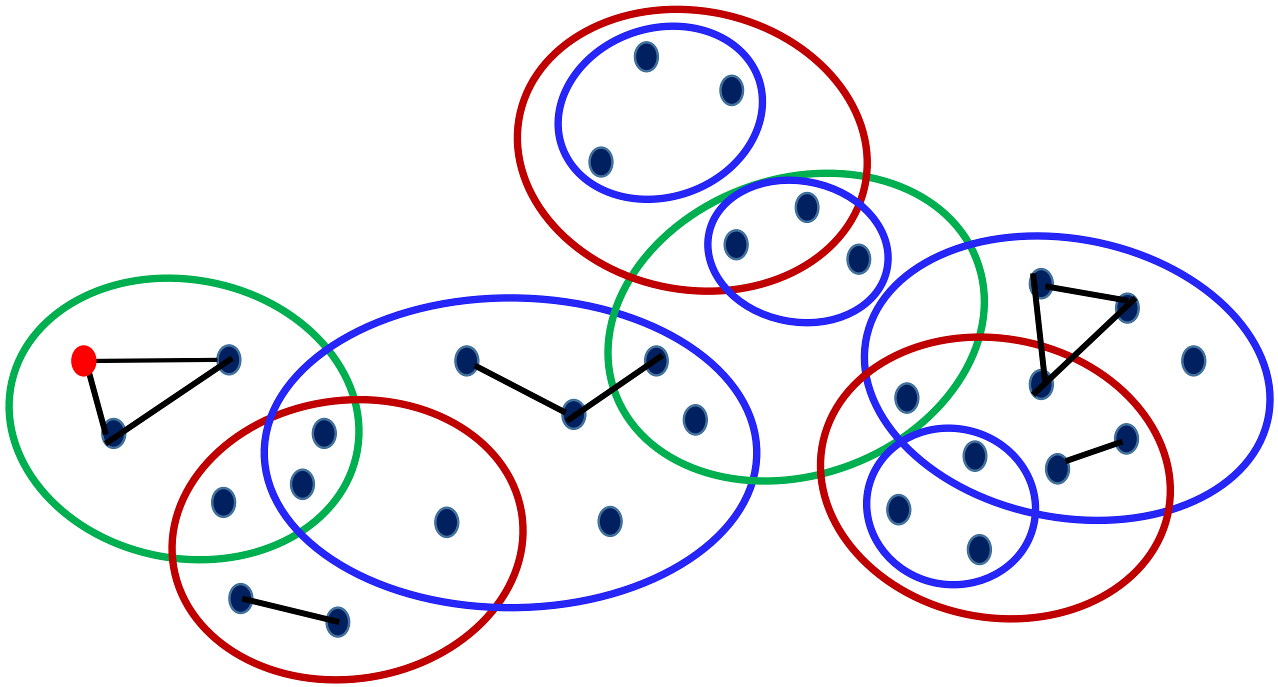

α-acyclic hypergraphs with log-poly ranks

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is α-acyclic. Then $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with at most $O(3^r|V|)$ variables and inequalities.

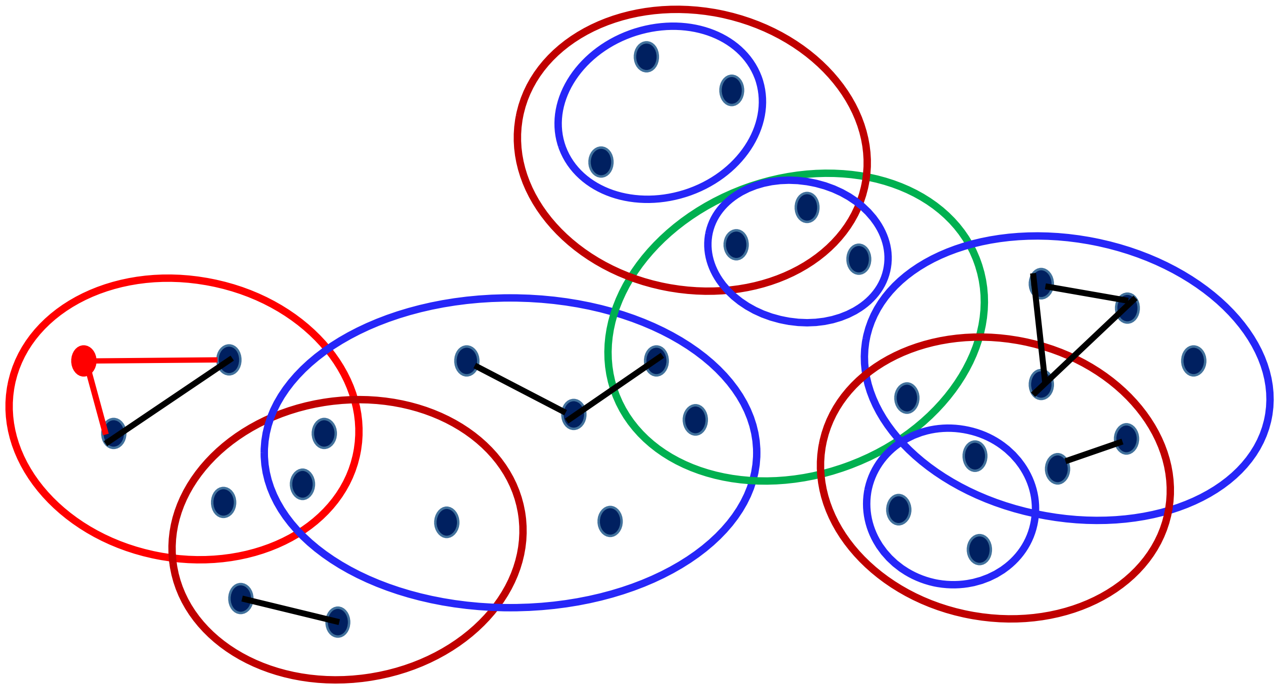

α-acyclic hypergraphs with log-poly ranks

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph is α-acyclic. Then $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with at most $O(3^r|V|)$ variables and inequalities.

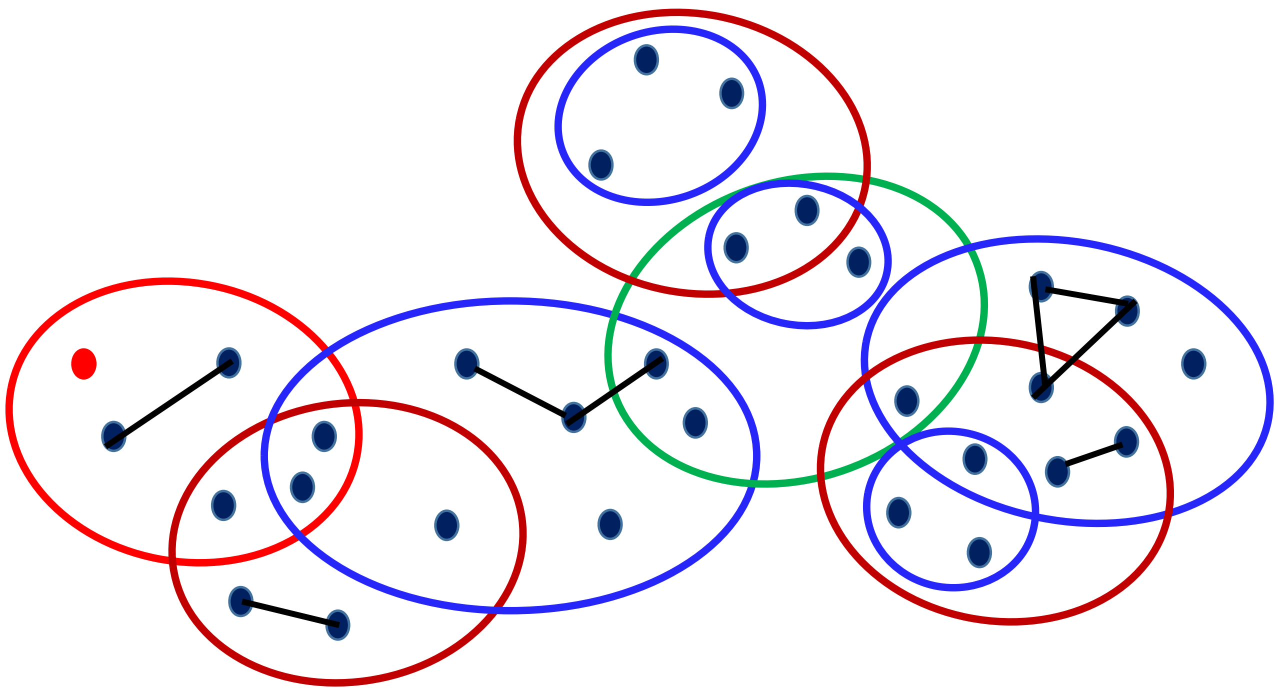

Nest-sets

- Let $G = (V, E)$ be a hypergraph and let $N \subseteq V$. We say that N is a nest-set of G, if the set
$$\{e \setminus N : e \in E, e \cap N \neq \emptyset\},$$
is totally ordered with respect to inclusion. If $|N| = 1$, then N contains a nest point of G (Lanzinger 2023).

- Let $N_1 \cdots N_t$ for some $t \geq 1$ be pairwise disjoint subsets of V such that $\cup_{i\in[t]} N_i = V$. We say that $\mathcal{N} = N_1, \cdots, N_t$ is a nest set elimination order of G, if N_1 is a nest set of G, N_2 is a nest-set of $G - N_1$, and so on.

- Given a nest set elimination order $\mathcal{N}$ of G, the nest-set width of this elimination order $\mathrm{nsw}_{\mathcal{N}}(G)$, is the maximum cardinality of any element in $\mathcal{N}$.

- The nest-set width of G $\mathrm{nsw}(G)$, is the minimum value of $\mathrm{nsw}_{\mathcal{N}}(G)$ over all nest set elimination orders $\mathcal{N}$ of G.

- $\mathrm{nsw}(G) = 1$, if and only if G is a β-acyclic hypergraph.

Nest-sets, nest-set width, and nest-set gap

- Let $G = (V, E)$, and let $V' \subseteq V$; we define the gap of G induced by V' as

$$\mathrm{gap}(G, V') := \max\left\{|V'| - |e \cap V'| : e \in E,\ e \cap V' \neq \emptyset\right\}.$$

- Given a nest-set elimination order $\mathcal{N}$ of G, we define nest-set gap of $\mathcal{N}$ as

$$\mathrm{nsg}_{\mathcal{N}}(G) := \max\left\{\mathrm{gap}(G - N_1 - \cdots - N_{i-1}, N_i) : i \in [t]\right\},$$

- We then have

$$\mathrm{nsg}_{\mathcal{N}}(G) \leq \mathrm{nsw}_{\mathcal{N}}(G) - 1.$$

- The nest-set gap of G $\mathrm{nsg}(G)$, as the minimum value of $\mathrm{nsg}_{\mathcal{N}}(G)$ over all nest-set elimination orders $\mathcal{N}$ of G.

- Example: Consider a hypergraph $G = (V, E)$ whose edge set E consists of all subsets of V of cardinality $|V| - 1$. Letting $\mathcal{N} = V \setminus \{\bar{v}\}, \{\bar{v}\}$ for some $\bar{v} \in V$, we have $\mathrm{nsw}(G) = |V| - 1$, while $\mathrm{nsg}(G) = 1$.

Nest-set width and nest-set gap

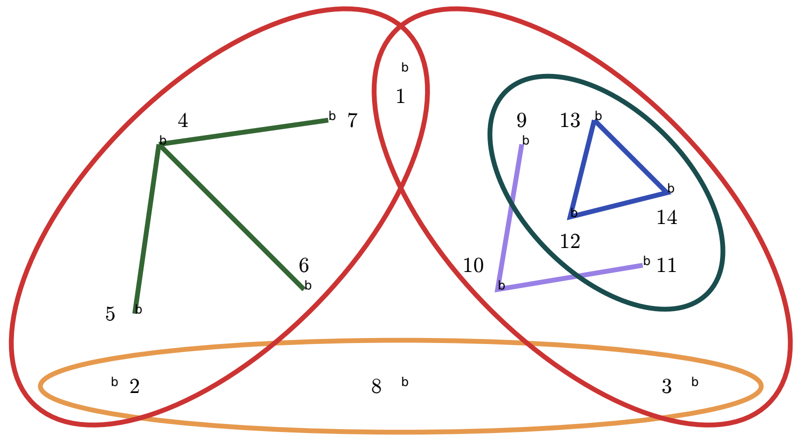

A hypergraph G with $\mathrm{nsw}(G) = 2$ and $\mathrm{nsg}(G) = 1$.
$\mathcal{N} = \{8\}, \{7\}, \{6\}, \{5\}, \{4\}, \{12, 13\}, \{14\}, \{9, 10\}, \{11\}, \{2, 3\}, \{1\}$. G contains β-cycles of length three.

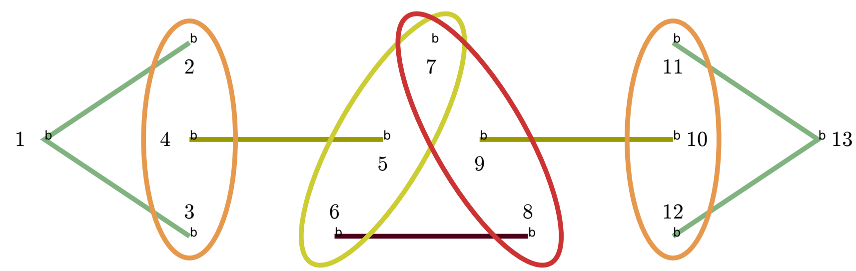

A hypergraph G with $\mathrm{nsw}(G) = 2$ and $\mathrm{nsg}(G) = 1$.
$\mathcal{N} = \{1, 2\}, \{3\}, \{4\}, \{5\}, \{6, 7\}, \{8\}, \{9\}, \{10\}, \{11, 12\}, \{13\}$. G contains β-cycles of length three.

Nest-set width and nest-set gap

- If G has a β-cycle of length ℓ, then $\text{nsw}(G) \geq \ell - 1$.
- However, the converse does not hold.

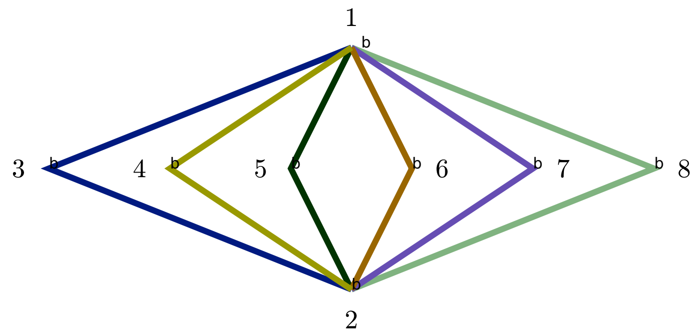

This graph contains cycles of length four only, however, its nest-set width is $|V| - 1$.

- Question: Is it possible to give a characterization of nest-set width in terms of the size of the β-cycles in the hypergraph?

Hypergraphs with small nest-set gaps

- Deciding if $\mathrm{nsw}(G) \leq k$ for any integer k is NP-complete. However, when parameterized by k, this problem is fixed-parameter tractable (Lanzinger 2023):

- There exists a $2^{O(k^2)}\mathrm{poly}(|V|, |E|)$ time algorithm that takes as input hypergraph $G = (V, E)$ and integer $k \geq 1$ and returns a nest set elimination $\mathcal{N}$ with $\mathrm{nsw}_{\mathcal{N}}(G) = k$ if one exists, or rejects otherwise.

- Question: What is the complexity of checking whether the nest-set gap of a hypergraph is bounded?

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph $G = (V, E)$ satisfies $\mathrm{nsg}(G) \leq k$. Then the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities. In particular, if $k \in O(\log \mathrm{poly}(|V|, |S|))$, then $\mathcal{P}_{\mathrm{pB}}(H)$ has a polynomial-size extended formulation. Moreover, all coefficients and right-hand side constants in the system defining $\mathcal{P}_{\mathrm{pB}}(H)$ are $0, \pm 1$.

Hypergraphs with small nest-set widths

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph $G = (V, E)$ satisfies $\mathrm{nsg}(G) \leq k$. Then the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities.

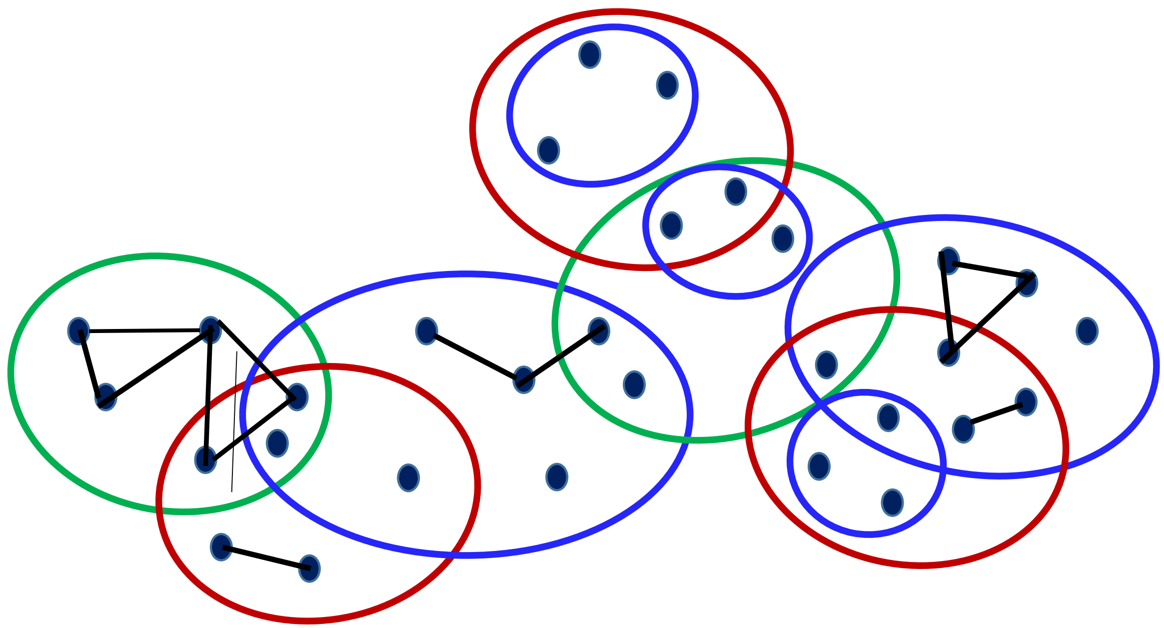

Hypergraphs with small nest-set widths

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph $G = (V, E)$ satisfies $\mathrm{nsw}(G) \leq k$. Then the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities.

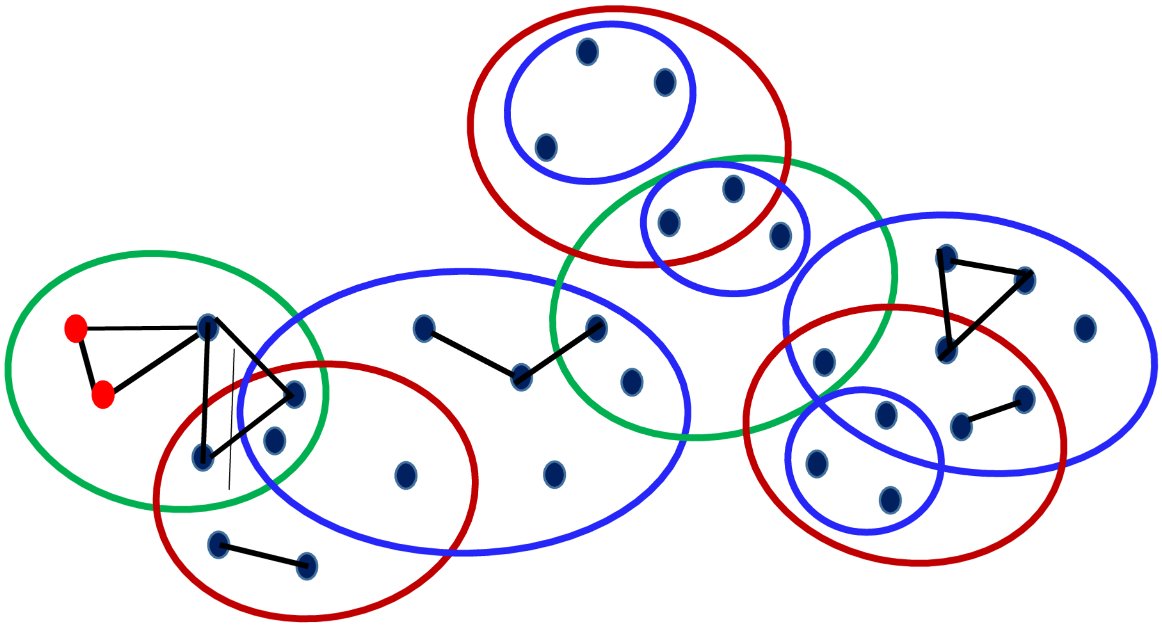

Hypergraphs with small nest-set widths

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph $G = (V, E)$ satisfies $\mathrm{nsw}(G) \leq k$. Then the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with $O(r2^k|V||S|)$ variables and inequalities.

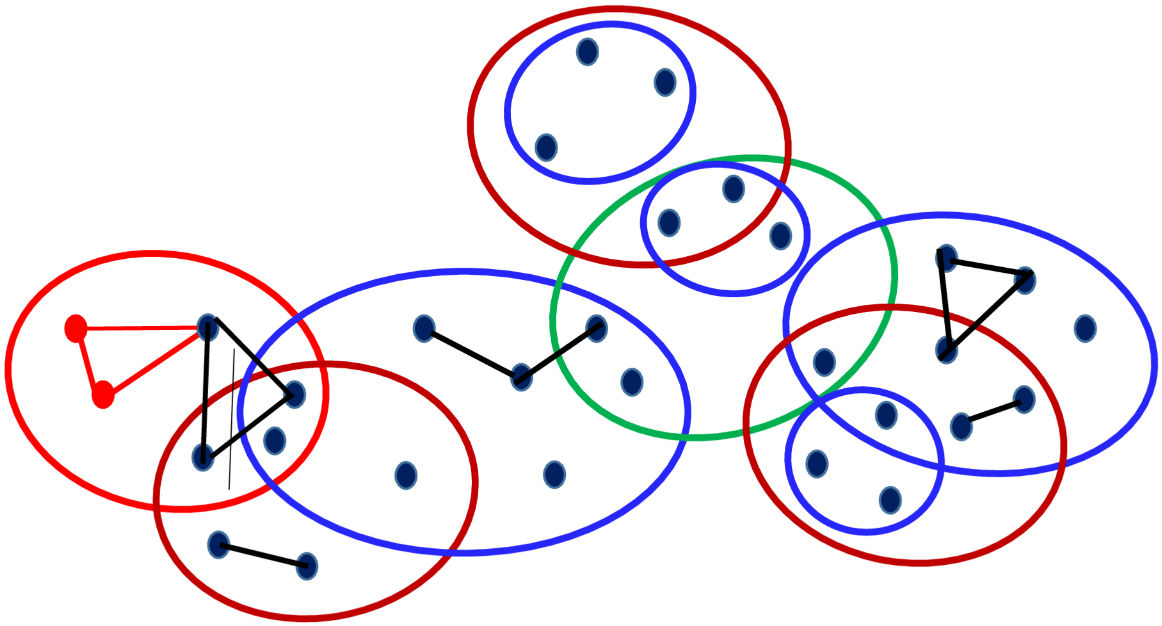

Hypergraphs with small nest-set widths

- Theorem: Let $H = (V, S)$ be a signed hypergraph of rank r whose underlying hypergraph $G = (V, E)$ satisfies $\mathrm{nsw}(G) \leq k$. Then the pseudo-Boolean polytope $\mathcal{P}_{\mathrm{pB}}(H)$ has an extended formulation with $O(r2^k|V|^2|S|)$ variables and inequalities.

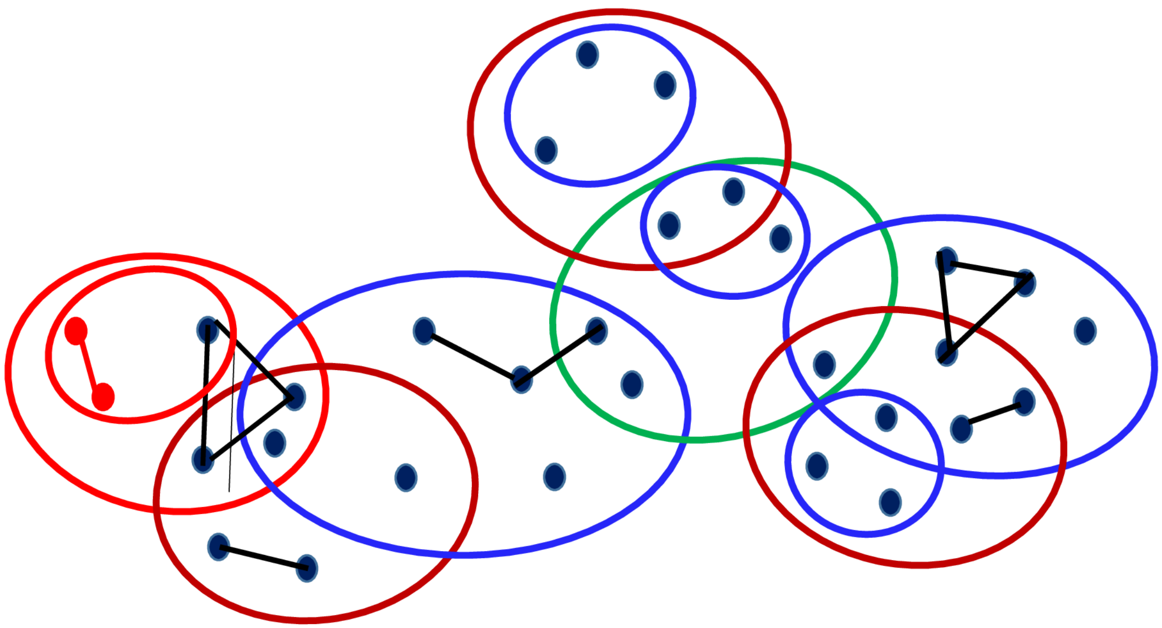

The extension complexity of the Boolean quadric polytope

- The extension complexity of a polytope is defined as the minimum number of inequalities and equalities in an extended formulation of the polytope.

- Theorem: There exists a universal constant $\delta > \frac{1}{20}$ such that, for every graph G with n nodes, we have $\mathrm{xc}(\mathsf{QP}(G)) \geq 2^{\Omega((\mathrm{tw}(G))^{\delta}+\log n)}$ (Fiorini et al, 2019).

- Corollary: For any hypergraph $G = (V, E)$, denote by G' the graph (V, E'), where E' contains all edges in E of cardinality two. Then there exists a universal constant $\delta > \frac{1}{20}$ such that for every hypergraph G with n nodes, we have $\mathrm{xc}(\mathcal{P}_{\mathrm{m}}(G)) \geq 2^{\Omega((\mathrm{tw}(G'))^{\delta}+\log n)}$.

- This lower bound implies an exponential extension complexity for hypergraphs with unbounded treewidth in which all subsets of cardinality two of each edge are also edges of the hypergraph; i.e., $\mathrm{tw}(G) = \mathrm{tw}(G')$. However, this lower bound can be arbitrarily weak, because $\mathrm{tw}(G') \ll \mathrm{tw}(G)$, in general.

The extension complexity of the pseudo-Boolean polytope

- Theorem: There exists constant $\delta > \frac{1}{20}$ such that, for every hypergraph G with n nodes with $r \in O(\log \mathrm{poly}(\mathrm{tw}(G)))$, there exists a signed hypergraph H with the underlying hypergraph G for which $\mathrm{xc}(\mathcal{P}_{\mathrm{pB}}(H)) \geq 2^{\Omega((\mathrm{tw}(G))^{\delta} + \log n)}$.

- Proof Sketch:
  - Let $G = (V, E)$, and let $G' = (V, E')$ be the intersection graph of G; $\mathrm{tw}(G') = \mathrm{tw}(G)$.
  - For each edge $e \in E'$, there exists an edge $g \in E$ such that $e \subseteq g$; denote by $g(e)$ one such edge of E.
  - Define the signed hypergraph $\bar{H} = (V, \bar{S})$ obtained from G' by inflating each edge $e \in E'$ to $g(e)$.
  - Denote by r the rank of G; we have $|g(e)| \leq r$, therefore $|\bar{S}| \leq 2^r \cdot |E'|$.
  - Since $|E'| \leq |V|^2$ and r is a log-poly function in $\mathrm{tw}(G)$, $|\bar{S}|$ is upper bounded by a polynomial in $\mathrm{tw}(G)$.
  - The underlying hypergraph of $\bar{H}$ is a partial hypergraph of G; let $H = (V, S)$ be a signed hypergraph with the underlying hypergraph G such that $S \supseteq \bar{S}$; $\mathrm{xc}(\mathcal{P}_{\mathrm{pB}}(H)) \geq \mathrm{xc}(\mathcal{P}_{\mathrm{pB}}(\bar{H}))$.
  - An extended formulation for $\mathrm{QP}(G')$ is given by an extended formulation for $\mathcal{P}_{\mathrm{pB}}(\bar{H})$ together with $|E'|$ equalities containing $|\bar{S}|$ variables. Since $\mathrm{xc}(\mathrm{QP}(G')) \geq 2^{\Omega((\mathrm{tw}(G))^{\delta} + \log n)}$, we get $\mathrm{xc}(\mathcal{P}_{\mathrm{pB}}(\bar{H})) \geq 2^{\Omega((\mathrm{tw}(G))^{\delta} + \log n)}$.

Open questions

- Statement 1: There exists a universal positive constant δ such that, for every signed hypergraph H with n nodes whose rank is upper bounded by a log-poly function in $\mathrm{tw}(H)$, we have $\mathrm{xc}(\mathcal{P}_{\mathrm{pB}}(H)) \geq 2^{\Omega((\mathrm{tw}(H))^{\delta}+\log n)}$.

- Statement 2: There exists a universal positive constant δ such that, for every hypergraph G with n nodes whose rank is upper bounded by a log-poly function in $\mathrm{tw}(\mathrm{G})$, we have $\mathrm{xc}(\mathcal{P}_{\mathrm{m}}(G)) \geq 2^{\Omega((\mathrm{tw}(G))^{\delta}+\log n)}$.

- If Statement 1 is true, then Statement 2 is true.

- These remain open even if we let the rank to be a constant.

References

- A. Del Pia and A. Khajavirad, The complete edge relaxation for binary polynomial optimization, working paper, 2025.
- A. Del Pia and A. Khajavirad, Beyond hypergraph acyclicity: limits of tractability for pseudo-Boolean optimization, arXiv:2410.23045, 2025.
- A. Del Pia and A. Khajavirad, The pseudo-Boolean polytope and polynomial-size extended formulations for binary polynomial optimization, Mathematical Programming Series A, 2024.
- A. Del Pia and A. Khajavirad, A polynomial-size extended formulation for the multilinear polytope of beta-acyclic hypergraphs, Mathematical Programming Series A, 2023.
- A. Del Pia and A. Khajavirad, The running intersection relaxation of the multilinear polytope, Mathematics of Operations Research, 2021.
- A. Del Pia, A. Khajavirad, and N. V. Sahinidis, On the impact of running-intersection inequalities for globally solving polynomial optimization problems, Mathematical Programming Computation, 2020.
- A. Del Pia and A. Khajavirad, The multilinear polytope for acyclic hypergraphs, SIAM Journal on Optimization, 2018.
- A. Del Pia and A. Khajavirad, On decomposability of multilinear sets, Mathematical Programming Series A, 2018.
- A. Del Pia and A. Khajavirad, A polyhedral study of binary polynomial programs, Mathematics of Operations Research, 2017.