
Box Facets and Cut Facets of Lifted Multicut Polytopes

Lucas Fabian Naumann [1] **Jannik Irmai** [1] **Shengxian Zhao** [1 2] **Bjoern Andres** [1 2]

Abstract

The lifted multicut problem has diverse applications in the field of computer vision. Exact algorithms based on linear programming require an understanding of lifted multicut polytopes. Despite recent progress, two fundamental questions about these polytopes have remained open: Which lower box inequalities define facets, and which cut inequalities define facets? In this article, we answer the first question by establishing conditions that are necessary, sufficient and efficiently decidable. Toward the second question, we show that deciding facet-definingness of cut inequalities is NP-hard. This completes the analysis of canonical facets of lifted multicut polytopes.

1. Introduction

The *lifted multicut problem* (Keuper et al., 2015) is a combinatorial optimization problem whose feasible solutions relate one-to-one to the clusterings of a graph. A *clustering* or *decomposition* of a graph $G = (V, E)$ is a partition Π of the node set V such that for every $U \in \Pi$ the subgraph of G induced by U is connected. Horňáková et al. (2017) cast the lifted multicut problem in the form of a binary linear program in which costs are associated to binary variables x_{uw} that indicate for pairs of distinct nodes $u, w \in V$ whether these nodes are in the same cluster, $x_{uw} = 0$, or in distinct clusters, $x_{uw} = 1$.

Such variables are introduced for neighboring nodes (elements of E), but also for some non-neighboring nodes (elements of F, formally defined later) which allows to reward or penalize nodes being in the same cluster without changing the set of feasible clusterings. This property is employed, for example, in multiple object tracking, where nodes V represent object occurrences at different time steps, edges E only occur between occurrences in consecutive time steps, a cluster represents the time-continuous track of an object moving through time, and variables for non-neighboring nodes are used to reward similar occurrences apart in time being in the same cluster, i.e. being connected by a path of edges E whose variables get assigned 0 (Tang et al., 2017). For the special case of no variables corresponding to non-neighboring nodes ($F = \emptyset$), the lifted multicut problem specializes to the multicut problem (Deza et al., 1992; Chopra & Rao, 1993) and the correlation clustering problem (Bansal et al., 2004; Demaine et al., 2006).

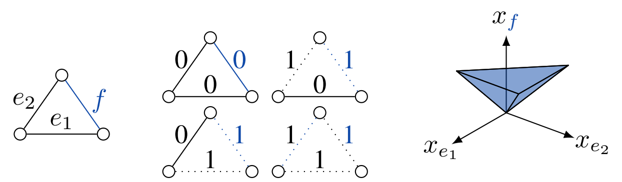

Figure 1. Depicted on the left is a graph $G = (V, E)$ with $E = \{e_1, e_2\}$ and an augmentation $\widehat{G} = (V, E \cup F)$ of G with $F = \{f\}$. Depicted in the middle are the four feasible solutions to the lifted multicut problem with respect to G and $\widehat{G}$. Depicted on the right is the lifted multicut polytope $\Xi_{G\widehat{G}}$. The figure is adopted from Andres et al. (2023).

In the following, we formally introduce the described binary linear program formulation of the lifted multicut problem:

Definition 1.1. (Horňáková et al., 2017, Def. 9) For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$, and any $c \in \mathbb{R}^{E \cup F}$, the instance of the *(minimum cost) lifted multicut problem* has the form

$$\min \quad \left\{ \sum_{e \in E \cup F} c_e \, x_e \;\middle|\; x \in X_{G\widehat{G}} \right\} \tag{1}$$

with $X_{G\widehat{G}}$ the set of all $x \in \{0,1\}^{E \cup F}$ that satisfy the following linear inequalities that we discuss in Section 3:

$$\forall C \in \text{cycles}(G)\ \forall e \in E_C\colon\ x_e \le \sum_{e' \in E_C \setminus \{e\}} x_{e'} \tag{2}$$

$$\forall uw \in F\ \forall P \in uw\text{-paths}(G)\colon\ x_{uw} \le \sum_{e \in E_P} x_e \tag{3}$$

$$\forall uw \in F\ \forall \delta \in uw\text{-cuts}(G)\colon\ 1 - x_{uw} \le \sum_{e \in \delta} (1 - x_e)\,. \tag{4}$$

[1]Faculty of Computer Science, TU Dresden [2]Center for Scalable Data Analytics and AI, Dresden/Leipzig. Correspondence to: Bjoern Andres <bjoern.andres@tu-dresden.de>.

Proceedings of the 41^{st} International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024.

We analyze the convex hull $\Xi_{G\widehat{G}} := \operatorname{conv} X_{G\widehat{G}}$ of the feasible set $X_{G\widehat{G}}$ in the real affine space $\mathbb{R}^{E\cup F}$, complementing properties established by Horňáková et al. (2017) and Andres et al. (2023) who call $\Xi_{G\widehat{G}}$ the *lifted multicut polytope* with respect to G and $\widehat{G}$ (see Figure 1 for an example). More specifically, we establish necessary, sufficient and efficiently decidable conditions for an inequality $0 \leq x_e$ with $e \in E \cup F$ to define a facet of $\Xi_{G\widehat{G}}$. Our proof involves an application of Menger's theorem (Menger, 1927). In addition, we show: Deciding whether a cut inequality (4) defines a facet of $\Xi_{G\widehat{G}}$ is NP-hard. In our proof, we first give a necessary and sufficient condition for the special case $|F| = 1$ and then show that deciding even this condition is NP-hard.

2. Related Work

The lifted multicut problem was introduced in the context of image and mesh segmentation by Keuper et al. (2015) and is discussed in further detail by Horňáková et al. (2017) and Andres et al. (2023). It has diverse applications, notably to the tasks of image segmentation (Keuper et al., 2015; Beier et al., 2016; 2017), video segmentation (Keuper, 2017; Keuper et al., 2020), and multiple object tracking (Tang et al., 2017; Nguyen et al., 2022; Kostyukhin et al., 2023). For these applications, local search algorithms are defined, implemented and compared empirically by Keuper et al. (2015); Levinkov et al. (2017). Two branch-and-cut algorithms for the lifted multicut problem are defined, implemented and compared empirically by Horňáková et al. (2017).

In order to significantly reduce the runtime of their branch-and-cut algorithm, Horňáková et al. (2017) are also the first to establish properties of lifted multicut polytopes, including its dimension $\dim \Xi_{G\widehat{G}} = |E \cup F|$ and a characterization of facets induced by cycle inequalities (2), path inequalities (3), upper box inequalities $x_e \leq 1$ for $e \in E \cup F$, and lower box inequalities $0 \leq x_e$ for $e \in E$. Moreover, they establish necessary conditions on facets of lifted multicut polytopes induced by cut inequalities (4) and lower box inequalities $0 \leq x_e$ for $e \in F$. Andres et al. (2023) describe an additional class of facets induced by so-called half-chorded odd cycle inequalities and show that these are facets also of a polytope isomorphic to the clique partitioning polytope (Deza et al., 1992; Grötschel & Wakabayashi, 1990; Deza & Laurent, 1997; Sørensen, 2002). Additionally, they establish the class of facets induced by so-called intersection inequalities, which is discovered based on a necessary condition for facets induced by cut inequalities. However, they do not make progress toward characterizing the facets of lifted multicut polytopes induced by cut inequalities themselves or lower box inequalities $0 \leq x_e$ for $e \in F$, which motivates the work we show in this article.

3. Preliminaries

For clarity, we adopt elementary terms and notation: Let $G = (V, E)$ be a graph. For any subset $A \subseteq E$, we write $\mathbb{1}_A \in \{0,1\}^E$ for the characteristic vector of the set A, i.e. $(\mathbb{1}_A)_e = 1 \Leftrightarrow e \in A$ for all $e \in E$. For any distinct $u, w \in V$, we write uw and wu as an abbreviation of the set $\{u, w\}$. We further call a path $P = (V_P, E_P)$ in G a *uw-path* in G if and only if its end-nodes are u and w. We call a set of edges $\delta \subseteq E$ a *uw-cut* of G if and only if every uw-path in G contains an edge of δ and the same does not hold for any proper subset $\delta' \subset \delta$.

Note that we interpret edges as two-elementary node sets, and will use them interchangeably with these sets. This means especially that for an edge $f = \{u, w\} = uw$ we use, e.g. f-path as an abbreviation of uw-path.

Properties of Feasible Solutions. We discuss briefly (2)–(4) in Definition 1.1; for details, we refer to Andres et al. (2023, Proposition 3). The inequalities (2) state that no cycle in G intersects with the set $\{e \in E \mid x_e = 1\}$ in precisely one edge. This property is equivalent to the existence of a clustering Π of G such that for any $uw \in E$: $x_{uw} = 0$ if and only if there exists a cluster $U \in \Pi$ such that $uw \subseteq U$. The inequalities (3) and (4) together state for any $uw \in F$ that $x_{uw} = 0$ if and only if there exists a uw-path (V_P, E_P) in G with all edges $e \in E_P$ such that $x_e = 0$, i.e. if and only if there exists a cluster $U \in \Pi$ such that $uw \subseteq U$.

One consequence of these properties that we apply in this article is that each set of clusters A has a vector $x^A \in X_{G\widehat{G}}$ such that the clustering induced by this vector contains exactly the clusters in A and otherwise singleton clusters (see Figure 2 for examples):

Definition 3.1. For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$ and any disjoint node sets $A \subseteq 2^V$ such that for all $U \in A$ the subgraph $G[U]$ of G induced by U is connected, we denote by $x^A \in \{0,1\}^{E\cup F}$ the unique vector for which $x^A_{uw} = 0 \Leftrightarrow \exists U \in A\colon uw \subseteq U$.

Lemma 3.2. *For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$ and any disjoint node sets $A \subseteq 2^V$ such that for all $U \in A$ the subgraph $G[U]$ of G induced by U is connected, $x^A \in X_{G\widehat{G}}$.*

Proof. Firstly, $\Pi = (A \setminus \{\emptyset\}) \cup \{\{v\} \mid v \in V \setminus \cup_{U\in A} U\}$ is a clustering of G. Secondly, x^A is such that for any $uw \in E \cup F$ we have $x^A_{uw} = 0$ if and only if there is a $U \in \Pi$ such that $uw \subseteq U$. Thus, $x^A \in X_{G\widehat{G}}$. □

Geometry of Convex Polytopes. We recall terms and facts about the geometry of convex polytopes that we will apply: An inequality is said to be *valid* for a polytope P if and only if it is satisfied by all $x \in P$. For an inequality $a^T x \leq$

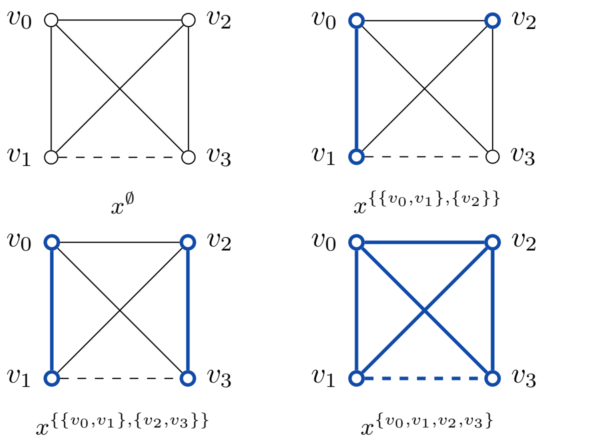

Figure 2. Depicted above are four graphs G (solid edges) with corresponding augmentations $\widehat{G}$ (dashed edges), which are marked to illustrate a vector x^A as given in Definition 3.1, that is stated below them. Nodes depicted in blue are contained in a cluster of A and edges depicted blue get assigned 0 by x^A.

α that is valid for a convex polytope P, the set $P_{a\alpha} := \{x \in P \mid a^T x = \alpha\}$ of those points in P that satisfy the inequality at equality is a maximal extremal face, or *facet*, of P if and only if $1 + \dim \operatorname{aff} P_{a\alpha} = \dim \operatorname{aff} P$, i.e. if and only if the dimension of the affine span of $P_{a\alpha}$ is (only) one less than the dimension of the affine span of P. For example, the inequality $x_f \leq 1$ defines a facet of the polytope depicted in Figure 1 because it is valid and the intersection of the hyperplane defined by $x_f = 1$ with the 3-dimensional polytope is 2-dimensional. In contrast, the inequality $x_{e_1} \leq 1$ is valid but not facet-defining because the intersection of the hyperplane defined by $x_{e_1} = 1$ and the 3-dimensional polytope is only 1-dimensional.

In order to prove that an inequality $a^T x \leq \alpha$ that is valid for a convex polytope P defines a facet of P, it is sufficient to construct $\dim \operatorname{aff} P - 1$ many linearly independent points in the difference space $P_{a\alpha} - P_{a\alpha} = \{x - y \mid x, y \in P_{a\alpha}\}$ because $\dim \operatorname{aff} P_{a\alpha} = \dim \operatorname{lin}(P_{a\alpha} - P_{a\alpha})$. In order to prove that an inequality $a^T x \leq \alpha$ that is valid for a convex polytope $P \subseteq \mathbb{R}^n$ with $\dim \operatorname{aff} P = n$ does not define a facet of P, it is sufficient to show that all points in $P_{a\alpha}$ satisfy another, orthogonal equality, for this implies $2 + \dim \operatorname{aff} P_{a\alpha} \leq \dim \operatorname{aff} P$.

Separators and Cut Nodes. For any graph $G = (V, E)$, any distinct $u, w \in V$ and any $S \subseteq V$, we call S a *uw-separator* of G and say that u and w are *separated* by S in G if and only if every uw-path in G contains a node of S. We call S *proper* if and only if $u \notin S$ and $w \notin S$.

For any graph $G = (V, E)$ and any $u, v, w \in V$ such that $S = \{v\}$ is a uw-separator of G, we call v a *uw-cut-node* of G. We call it *proper* if and only if S is *proper*. We let $C_{uw}(G)$ denote the set of all proper uw-cut-nodes of G.

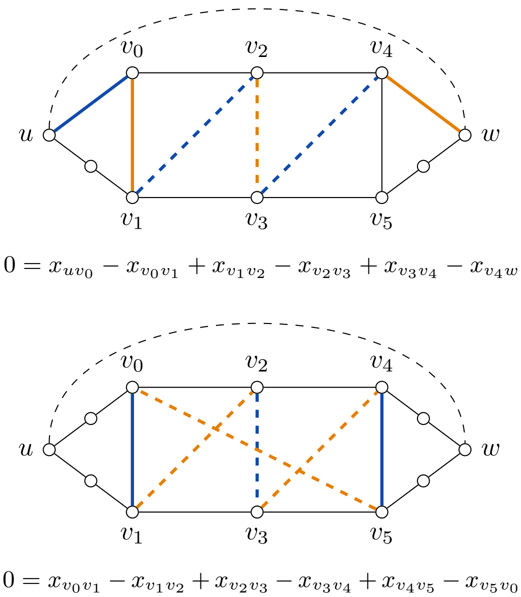

Figure 3. Depicted above are two examples of a graph G (solid edges) and augmentation $\widehat{G}$ (dashed edges) such that a condition of Theorem 4.1 is violated for the inequality $0 \leq x_{uw}$. In the example at the top, the path with the edge set $\{uv_0, v_0v_1, v_1v_2, v_2v_3, v_3v_4, v_4w\}$ violates (i). In the example at the bottom, the cycle with the edge set $\{v_0v_1, v_1v_2, v_2v_3, v_3v_4, v_4v_5, v_5v_0\}$ violates (ii). For both cases, Equation (5) from the proof of Theorem 4.1 is stated. Edges depicted in blue occur with a positive sign in this equation, and edges depicted in orange occur with a negative sign.

4. Lower Box Facets

In this section, we establish necessary, sufficient and efficiently decidable conditions for a lower box inequality $0 \leq x_{uw}$ with $uw \in E \cup F$ to define a facet of a lifted multicut polytope $\Xi_{G\widehat{G}}$. Examples for necessity of these conditions are shown in Figure 3.

Theorem 4.1. *For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$ and any $uw \in E \cup F$, the lower box inequality $0 \leq x_{uw}$ is facet-defining for $\Xi_{G\widehat{G}}$ if and only if the following two conditions hold:*

(i) There exists no simple path in $\widehat{G}$ of length at least one, besides $(\{u, w\}, \{uw\})$, whose end-nodes are uw-cut-nodes of G and whose edges are uw-separators of G.

(ii) There exists no simple cycle in $\widehat{G}$ whose edges are uw-separators of G.

In the remainder of this section, we prove a structural lemma and then apply this lemma in order to prove Theorem 4.1.

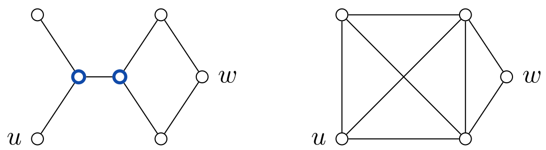

Figure 4. Depicted on the left is a graph G, and depicted on the right is the corresponding auxiliary graph G' whose construction is described in the proof of Lemma 4.2. The nodes depicted in blue are proper uw-cut-nodes of G and get removed in the construction of G'.

Lemma 4.2. *Let $G = (V, E)$ be a graph and let $u, w \in V$. Any simple cycle $C = (V_C, E_C)$ with $V_C \subseteq V$ and $E_C \subseteq \binom{V}{2}$ such that no $v \in V_C$ is a uw-cut-node of G and every $e \in E_C$ is a uw-separator of G is even.*

Proof of Lemma 4.2. In a first step, we construct for any $u, w \in V$ and any cycle $C = (V_C, E_C)$ as defined in the lemma an auxiliary graph $G' = (V', E')$ by removing from G the set $C_{uw}(G)$ of all proper uw-cut-nodes and connecting remaining nodes, which are connected in G by a path of only proper uw-cut-nodes, by additional edges, i.e. by setting

$$\begin{aligned} V' &= V \setminus C_{uw}(G) \\ E' &= \Big\{ st \in \tbinom{V'}{2} \;\Big|\; \exists (V_P, E_P) \in st\text{-paths}(G) \colon \\ &\qquad V_P \setminus st \subseteq C_{uw}(G) \Big\} \cup \Big(E \cap \tbinom{V'}{2} \Big) . \end{aligned}$$

An example of this construction is shown in Figure 4.

In a second step, we now show that G' has the following properties:

(i) $V_C \cup \{u, w\} \subseteq V'$ and $E_C \subseteq \binom{V'}{2}$;

(ii) there exist no proper uw-cut-nodes of G';

(iii) all $e \in E_C$ are proper uw-separators of G'.

Property (i) follows directly from the construction of the auxiliary graph G'.

Assume (ii) does not hold. Then there exists a $v \in C_{uw}(G')$. It follows that $v \notin \{u, w\}$ and, by construction of G', that $v \notin C_{uw}(G)$. Thus, v is no uw-cut-node of G and there exists a uw-path (V_P, E_P) in G such that $v \notin V_P$. By construction of G', we can create a uw-path $(V_{P'}, E_{P'})$ in G' with $v \notin V_{P'}$ by replacing all subpaths of (V_P, E_P) whose internal nodes are in $C_{uw}(G)$ with edges in $E' \setminus E$. The existence of such a uw-path $(V_{P'}, E_{P'})$ contradicts $v \in C_{uw}(G')$.

Assume (iii) does not hold. Then there exists an $e \in E_C$ that is not a proper uw-separator of G'. As $e \cap \{u, w\} = \emptyset$ by assumption, e is also no uw-separator of G'. Thus, there exists a uw-path $(V_{P'}, E_{P'})$ in G' with $e \cap V_{P'} = \emptyset$. By construction of G', we can define a uw-path (V_P, E_P) in G from $(V_{P'}, E_{P'})$ by replacing all edges in $E_{P'} \setminus E$ with paths in G whose internal nodes are in $C_{uw}(G)$. For this path, $e \cap V_P = \emptyset$ because $e \cap V_{P'} = \emptyset$ (see above) and $e \cap C_{uw}(G) = \emptyset$ (by assumption). The existence of such a uw-path (V_P, E_P) contradicts e being a uw-separator of G.

In a third step, we now prove that C is even: Menger's theorem (Menger, 1927) states that for two distinct non-adjacent nodes $a, b \in V'$, the number of internally node-disjoint ab-paths in G' is equal to the minimal size of proper ab-separators of G'. By (i), u and w are in V'. Furthermore, they are distinct and non-adjacent in G', as otherwise every uw-separator of G' would contain u or w, in contradiction to the elements of E_C being proper uw-separators of G' by (iii). As $C_{uw}(G') = \emptyset$ by (ii), and all edges in E_C are proper uw-separators of G' by (iii), the minimal size of proper uw-separators of G' is two. Thus, there exist precisely two internally node-disjoint uw-paths $P_1 = (V_{P_1}, E_{P_1})$ and $P_2 = (V_{P_2}, E_{P_2})$ in G', by Menger's theorem.

W.l.o.g., we enumerate the nodes in the cycle (V_C, E_C): For $n := |E_C|$, let $v \colon \mathbb{Z}_n \to V_C$ such that $E_C = \{v_j v_{j+1} \mid j \in \mathbb{Z}_n\}$. As $v_0 v_1$ is a uw-separator of G' by (iii), the paths P_1 and P_2 each contain v_0 or v_1. Moreover, as these paths are internally node-disjoint, precisely one of them contains v_0, the other v_1. Assume w.l.o.g. that $v_0 \in V_{P_1}$ and $v_1 \in V_{P_2}$. By (iii), any $v_j v_{j+1} \in E_C$ with $j \in \{1, \dots, n-2\}$ is a uw-separator of G'. Thus:

$$\begin{aligned} V_{P_1} \cap V_C &= \big\{ v_{2j} \mid j \in \{0, \dots, \lfloor \tfrac{n-1}{2} \rfloor\} \big\} \\ V_{P_2} \cap V_C &= \big\{ v_{2j+1} \mid j \in \{0, \dots, \lfloor \tfrac{n-2}{2} \rfloor\} \big\} . \end{aligned}$$

If C were odd, n would be odd. Thus, $n-1$ would be even. Consequently, it would follow that $v_{n-1} v_0 \cap V_{P_2} = \emptyset$, in contradiction to $v_{n-1} v_0$ being a uw-separator of G' by (iii). Thus, C must be even. □

Proof of Theorem 4.1. Assume there exists a path or cycle $H = (V_H, E_H)$ of uw-separators of G as defined in the theorem. W.l.o.g., fix enumerations of the nodes and edges of H as follows: Let $n := |E_H|$. If H is a path, let $v \colon \{0, \dots, n\} \to V_H$ and $e \colon \{0, \dots, n-1\} \to E_H$ such that $\forall j \in \{0, \dots, n-1\} \colon e_j = v_j v_{j+1}$ and $E_H = \{e_j \mid j \in \{0, \dots, n-1\}\}$. If H is a cycle, let $v \colon \mathbb{Z}_n \to V_H$ and $e \colon \mathbb{Z}_n \to E_H$ such that $\forall j \in \mathbb{Z}_n \colon e_j = v_j v_{j+1}$ and $E_H = \{e_j \mid j \in \mathbb{Z}_n\}$. If H is a cycle containing uw-cut-nodes of G, assume further and w.l.o.g. that $v_0 = v_n$ is such a uw-cut-node. Finally, consider the partition $\{E_0, E_1\}$ of E_H into even and odd edges, i.e.

$$\begin{aligned} E_0 &= \big\{ e_{2j} \in E_H \;\big|\; j \in \{0, \dots, \lfloor \tfrac{n-1}{2} \rfloor\} \big\} \\ E_1 &= \big\{ e_{2j+1} \in E_H \;\big|\; j \in \{0, \dots, \lfloor \tfrac{n-2}{2} \rfloor\} \big\} . \end{aligned}$$

We will prove that $\sigma = \{x \in X_{G\widehat{G}} \mid x_{uw} = 0\}$ is not a facet of $\Xi_{G\widehat{G}}$ by showing that all $x \in \sigma$ satisfy the additional orthogonal equality

$$0 = \sum_{j \in \{0,\dots,n-1\}} (-1)^j \, x_{e_j} \,. \tag{5}$$

This then implies $2 + \dim \operatorname{aff} \sigma \leq |E \cup F| = \dim \Xi_{G\widehat{G}}$ where the last equality stems from the full-dimensionality of the lifted multicut polytope by Theorem 7 of Horňáková et al. (2017).

More specifically, we will prove for every $x \in \sigma$ the existence of a bijection

$$\vartheta_x \colon \; E_0 \cap x^{-1}(1) \;\to\; E_1 \cap x^{-1}(1) \,.$$

Using these bijections, we conclude for every $x \in \sigma$ that the number of elements in the sum of (5) taking the value $+1$ is equal to the number of elements taking the value -1, and thus that the equality holds.

We now show that these bijections exist. Let $x \in \sigma$. As $x_{uw} = 0$, the clustering of G induced by x has a cluster containing both u and w. Let V_{uw} be the node set of that cluster. If $n = 1$, then H is a path $(\{v_0, v_1\}, \{e_0\})$. Thus, $E_1 \cap x^{-1}(1) = \emptyset$ because $E_1 = \emptyset$. Moreover, $E_0 \cap x^{-1}(1) = \emptyset$ as v_0 and v_1 are uw-cut-nodes of G and thus elements of V_{uw}, which implies $x_{e_0} = 0$. In this case, $\vartheta_x = \emptyset$ and (5) specializes to $x_{e_0} = 0$, which is satisfied.

We now consider $n \geq 2$. For every $e_j = v_j v_{j+1} \in E_0 \cap x^{-1}(1)$, we define:

$$\vartheta_x(e_j) = \begin{cases} e_{j-1} & \text{if } v_j \notin V_{uw} \\ e_{j+1} & \text{if } v_{j+1} \notin V_{uw} \end{cases} . \tag{6}$$

We show that ϑ_x is well-defined: Let $e_j \in E_0 \cap x^{-1}(1)$. In general, at least one of v_j and v_{j+1} is not in V_{uw} because $x_{e_j} = 1$, and at most one of v_j and v_{j+1} is not in V_{uw} because e_j is a uw-separator of G. Thus, ϑ_x assigns e_j a unique element. It remains to show that this element is in $E_1 \cap x^{-1}(1)$.

Firstly, we show $\vartheta_x(e_j) \in E_1$. Clearly, it holds for $j \in \{1, \dots, n-2\}$, that $\vartheta_x(e_j) \in E_1$. We regard the remaining cases of $j \in \{0, n-1\}$. Let first $j = 0$. For H a path or cycle with uw-cut-node, $\vartheta_x(e_0) = e_1 \in E_1$ because $v_0 \in V_{uw}$ as v_0 is a uw-cut-node of G. For H a cycle without uw-cut-node, we distinguish $v_0 \in V_{uw}$ and $v_0 \notin V_{uw}$. If $v_0 \in V_{uw}$, then $\vartheta_x(e_0) = e_1 \in E_1$. If $v_0 \notin V_{uw}$, then $\vartheta_x(e_0) = e_{n-1} \in E_1$ because $n-1$ is odd by Lemma 4.2. Let now $j = n-1$. For H a path or cycle with uw-cut-node, $\vartheta_x(e_{n-1}) = e_{n-2} \in E_1$ because $v_n \in V_{uw}$ as v_n is a uw-cut-node of G. For H a cycle without uw-cut-node, $e_{n-1} \notin E_0 \cap x^{-1}(1)$ because $n-1$ is odd by Lemma 4.2.

Secondly, we show $x_{\vartheta_x(e_j)} = 1$. By definition of ϑ_x, e_j and $\vartheta_x(e_j)$ share a node $v \notin V_{uw}$. As $\vartheta_x(e_j)$ is a uw-separator of G, the other node of $\vartheta_x(e_j)$ is in V_{uw} and therefore $x_{\vartheta_x(e_j)} = 1$. Thus, $\vartheta_x(e_j) \in E_1 \cap x^{-1}(1)$, and ϑ_x is well-defined.

We show that ϑ_x is surjective: Let $e_j \in E_1 \cap x^{-1}(1)$. As $x_{e_j} = 1$, either $v_j \notin V_{uw}$ or $v_{j+1} \notin V_{uw}$. If $v_j \notin V_{uw}$, then $e_{j-1} \in E_0 \cap x^{-1}(1)$ and $\vartheta_x(e_{j-1}) = e_j$. If $v_{j+1} \notin V_{uw}$, then $e_{j+1} \in E_0 \cap x^{-1}(1)$ and $\vartheta_x(e_{j+1}) = e_j$. Thus, ϑ_x is surjective.

We show that ϑ_x is injective: Assume ϑ_x is not injective. Then there exists a $j \in \{0, \dots, n-1\}$ such that $e_j \in E_1 \cap x^{-1}(1)$ and $e_{j-1}, e_{j+1} \in E_0 \cap x^{-1}(1)$ such that $\vartheta_x(e_{j-1}) = e_j = \vartheta_x(e_{j+1})$, by definition of ϑ_x. This implies $v_j, v_{j+1} \notin V_{uw}$, which contradicts e_j being a uw-separator. By this contradiction, ϑ_x is injective.

Altogether, we have shown that ϑ_x is well-defined, surjective and injective, and thus a bijection. This concludes the proof of necessity.

Assume now that (i) and (ii) are satisfied. We prove that $0 \leq x_{uw}$ is facet-defining by constructing $|E \cup F| - 1$ linearly independent vectors in $\operatorname{lin}(\sigma - \sigma)$, implying $\dim \operatorname{aff} \sigma = \dim \operatorname{lin}(\sigma - \sigma) = |E \cup F| - 1$ and thus that $\operatorname{aff} \sigma$ is a facet of $\Xi_{G\widehat{G}}$. In particular, we construct the characteristic vectors of all $st \in E \cup F \setminus \{uw\}$. For this construction, we distinguish the following cases:

1. st is not a uw-separator of G;
2. st is a uw-separator of G and neither s nor t is a uw-cut-node;
3. precisely one node of st is a uw-cut-node.

Note that no $st \in E \cup F \setminus \{uw\}$ is such that both s and t are uw-cut-nodes, as otherwise the path $(\{s,t\}, \{st\})$ would violate (i). Thus, this distinction of cases is complete.

For the first case, let $st \in E \cup F \setminus \{uw\}$ such that st is not a uw-separator of G. By this property, there exists a uw-path $(V_{P_{uw}}, E_{P_{uw}})$ in G that contains neither s nor t. Let further $(V_{P_{st}}, E_{P_{st}})$ be an st-path in G. If $G[V_{P_{uw}} \cup V_{P_{st}}]$ is not connected, we define:

$$\begin{aligned} V_1 &= \{V_{P_{uw}}, V_{P_{st}}\} & V_2 &= \{V_{P_{uw}}, V_{P_{st}} \setminus \{s,t\}\} \\ V_3 &= \{V_{P_{uw}}, V_{P_{st}} \setminus \{s\}\} & V_4 &= \{V_{P_{uw}}, V_{P_{st}} \setminus \{t\}\} \,. \end{aligned}$$

Otherwise, we define:

$$\begin{aligned} V_1 &= \{V_{P_{uw}} \cup V_{P_{st}}\} & V_2 &= \{V_{P_{uw}} \cup V_{P_{st}} \setminus \{s,t\}\} \\ V_3 &= \{V_{P_{uw}} \cup V_{P_{st}} \setminus \{s\}\} & V_4 &= \{V_{P_{uw}} \cup V_{P_{st}} \setminus \{t\}\} \,. \end{aligned}$$

In both cases, it is easy to see for $i \in \{1, \dots, 4\}$ that $G[U]$ is connected for all $U \in V_i$ and thus $x^{V_i} \in X_{G\widehat{G}}$, by Lemma 3.2. It further holds, $\mathbb{1}_{\{st\}} = -x^{V_1} - x^{V_2} + x^{V_3} + x^{V_4}$ and $x^{V_i}_{uw} = 0$, as for all $pq \in E \cup F$:

- $x^{V_i}_{pq} = 1$ for $i = 1, \dots, 4$ if $\nexists U \in V_1\colon \{p,q\} \subseteq U$
- $x^{V_i}_{pq} = 0$ for $i = 1, \dots, 4$ if $\exists U \in V_2\colon \{p,q\} \subseteq U$
- $x^{V_i}_{pq} = 0$ for $i = 1, 3$ and $x^{V_i}_{pq} = 1$ for $i = 2, 4$ if $s \in \{p,q\}, t \notin \{p,q\}$ and $\exists U \in V_1\colon \{p,q\} \subseteq U$
- $x^{V_i}_{pq} = 0$ for $i = 1, 4$ and $x^{V_i}_{pq} = 1$ for $i = 2, 3$ if $t \in \{p,q\}, s \notin \{p,q\}$ and $\exists U \in V_1\colon \{p,q\} \subseteq U$
- $x^{V_1}_{pq} = 0$ and $x^{V_i}_{pq} = 1$ for $i = 2, 3, 4$ if $\{p,q\} = \{s,t\}$.

It follows from $x^{V_i}_{uw} = 0$ that $x^{V_i} \in \sigma$. Thus, $\mathbb{1}_{\{st\}} = -x^{V_1} - x^{V_2} + x^{V_3} + x^{V_4} \in \mathrm{lin}(\sigma - \sigma)$, which concludes the first case.

For the second case, consider the set H of all $st \in E \cup F \setminus \{uw\}$ such that st is a uw-separator of G and neither s nor t is a uw-cut-node of G. Let $st \in H$ and let $v \in st$. As v is no uw-cut-node, there exists a uw-path $(V_{P_{uw}}, E_{P_{uw}})$ in G that does not contain v. Let further $(V_{P_{st}}, E_{P_{st}})$ be an st-path in G and let $P = (V_P, E_P) = (V_{P_{uw}} \cup V_{P_{st}}, E_{P_{uw}} \cup E_{P_{st}})$. With $E_{\widehat{G}}(P, v) = \{vv' \in (E \cup F) \cap \binom{V_P}{2}\}$ denoting the set of edges of $\widehat{G}$ containing v whose nodes are in V_P, we first show that $\mathbb{1}_{E_{\widehat{G}}(P,v)} \in \mathrm{lin}(\sigma - \sigma)$. If $G[V_{P_{uw}} \cup V_{P_{st}}]$ is not connected, we define:

$$V_1 = \{V_{P_{uw}}, V_{P_{st}}\} \qquad V_4 = \{V_{P_{uw}}, V_{P_{st}} \setminus \{v\}\}\,.$$

Otherwise, we define:

$$V_1 = \{V_{P_{uw}} \cup V_{P_{st}}\} \qquad V_4 = \{V_{P_{uw}} \cup V_{P_{st}} \setminus \{v\}\}\,.$$

Analogously to the previous case, we get $\mathbb{1}_{E_{\widehat{G}}(P,v)} = -x^{V_1} + x^{V_4} \in \mathrm{lin}(\sigma - \sigma)$. Denoting by $E_{\widehat{G}}(v) = \{vv' \in (E \cup F)\}$ the set of edges of $\widehat{G}$ containing v and noting that

$$\mathbb{1}_{\{st\}} = \mathbb{1}_{E_{\widehat{G}}(P,v)} - \sum_{e \in E_{\widehat{G}}(P,v) \setminus \{st\}} \mathbb{1}_{\{e\}} \tag{7}$$

and $E_{\widehat{G}}(P,v) \subseteq E_{\widehat{G}}(v)$, we see that it is sufficient for proving $\mathbb{1}_{\{st\}} \in \mathrm{lin}(\sigma - \sigma)$ to show that there exists a node $v \in st$ such that $\mathbb{1}_{\{e\}} \in \mathrm{lin}(\sigma - \sigma)$ for all $e \in E_{\widehat{G}}(v) \setminus \{st\}$.

Next, we define a sequence $\{H_j\}_{j \in \mathbb{N}_0}$ of subsets of H (for an example see Figure 5) and show iteratively that the characteristic vectors of their elements are in $\mathrm{lin}(\sigma - \sigma)$ using (7). For any $j \in \mathbb{N}_0$, we define:

$$H_j = \Big\{st \in H \setminus \cup_{k<j} H_k \mid \exists v \in st\, \forall e \in E_{\widehat{G}}(v) \setminus \{st\}\colon \\ e \text{ is no } uw\text{-separator of } G \vee\; e \in \cup_{k<j} H_k \Big\}\,. \tag{8}$$

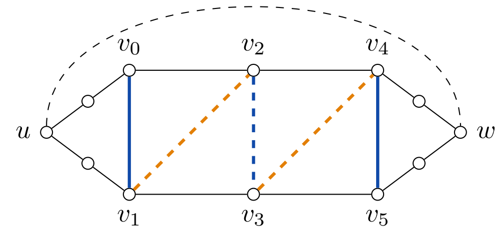

Figure 5. Depicted above is an example of a graph G (solid edges) and augmentation $\widehat{G}$ (dashed edges) that fulfills the conditions of Theorem 4.1 for $0 \leq x_{uw}$. Essential for the sufficiency proof of this theorem is that the introduced edge sets H and H_j for $j \in \mathbb{N}_0$ have the property $H \subseteq \cup_{j \geq 0} H_j$. In the given example, $H = \{v_0v_1, v_1v_2, v_2v_3, v_3v_4, v_4v_5\}$, $H_0 = \{v_0v_1, v_4v_5\}$, $H_1 = \{v_1v_2, v_3v_4\}$, $H_2 = \{v_2v_3\}$ and $H_j = \emptyset$ for $j \geq 3$. Thus, $H \subseteq \cup_{j \geq 0} H_j$.

By this definition, for any $st \in H_0$ there exists a $v \in st$ such that all $e \in E_{\widehat{G}}(v) \setminus \{st\}$ are no uw-separators of G. Thus, it follows from the previous case that $\mathbb{1}_{\{e\}} \in \mathrm{lin}(\sigma - \sigma)$ for all $e \in E_{\widehat{G}}(v) \setminus \{st\}$. Consequently, $\mathbb{1}_{\{st\}} \in \mathrm{lin}(\sigma - \sigma)$ by (7). Let now $j > 0$ and assume that the characteristic vectors of all elements in $\cup_{k<j} H_k$ are in $\mathrm{lin}(\sigma - \sigma)$. By definition, for any $st \in H_j$ there exists a $v \in st$ such that any $e \in E_{\widehat{G}}(v) \setminus \{st\}$ is either no uw-separator of G and thus $\mathbb{1}_{\{e\}} \in \mathrm{lin}(\sigma - \sigma)$ by the previous case, or is in $\cup_{k<j} H_k$ and thus $1_{\{e\}} \in \mathrm{lin}(\sigma - \sigma)$ by assumption. Consequently, $\mathbb{1}_{\{st\}} \in \mathrm{lin}(\sigma - \sigma)$ by (7).

For completing the second case, it remains to show that we have constructed the characteristic vectors of all elements in H by this, i.e. that $H \subseteq \cup_{j \geq 0} H_j$. This follows directly from Claim 4.3, which is proven in Appendix A.

Claim 4.3. *If (i) and (ii) are satisfied, the set $\{H_j \mid j \in \mathbb{N}_0 \wedge H_j \neq \emptyset\}$ is a partition of H.*

For the last case, let $st \in E \cup F \setminus \{uw\}$ such that precisely one node of st, say t, is a uw-cut-node. We construct P and show $\mathbb{1}_{E_{\widehat{G}}(P,s)} \in \mathrm{lin}(\sigma - \sigma)$ analogously to the previous case and again have

$$\mathbb{1}_{\{st\}} = \mathbb{1}_{E_{\widehat{G}}(P,s)} - \sum_{e \in E_{\widehat{G}}(P,s) \setminus \{st\}} \mathbb{1}_{\{e\}}\,. \tag{9}$$

For any $e = s's \in E_{\widehat{G}}(P,s) \setminus \{st\}$, s' is no uw-cut-node of G, as otherwise the path $(\{s', s, t\}, \{s's, st\})$ would violate (i). Consequently, $\mathbb{1}_{\{e\}} \in \mathrm{lin}(\sigma - \sigma)$ by the previous two cases. It follows from (9) that $\mathbb{1}_{\{st\}} \in \mathrm{lin}(\sigma - \sigma)$, which concludes the third case. Altogether, we have constructed $|E \cup F| - 1$ linearly independent vectors in $\mathrm{lin}(\sigma - \sigma)$ and have thus established sufficiency of the specified conditions. □

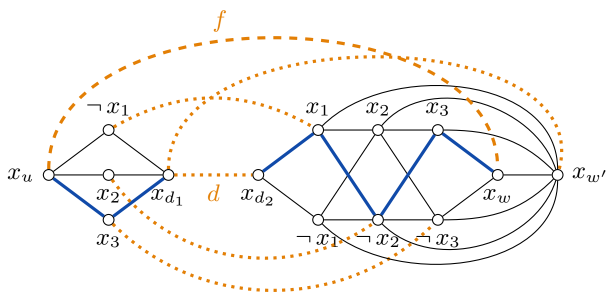

Figure 6. Depicted above is an example of the reduction from 3-SAT used in the proof of Lemma 5.3. Graphs G and $\widehat{G}$ are constructed from the instance of the 3-SAT problem given by $\neg x_1 \vee x_2 \vee x_3$. The additional edge f as well as the edges in the f-cut δ are depicted in orange. The f_d-path with respect δ, given by the blue edges and d, corresponds to the solution of the 3-SAT problem instance: $\varphi(x_1) =$ FALSE, $\varphi(x_2) =$ FALSE and $\varphi(x_3) =$ TRUE.

5. NP-Hardness of Deciding Cut Facets

In this section, we prove that it is NP-hard to decide facet-definingness of cut inequalities (4) for lifted multicut polytopes. We do so in two steps: Firstly, we establish a necessary and sufficient condition for facet-definingness of cut inequalities for lifted multicut polytopes in the special case $|F| = 1$ (Lemma 5.2). Secondly, we show that deciding this condition for these specific polytopes is NP-hard (Lemma 5.3). Together, this implies that facet-definingness is NP-hard to decide for cut inequalities of general lifted multicut polytopes (Theorem 5.4).

We begin by introducing a structure fundamental to this discussion, paths crossing a cut in precisely one edge that have no other edge of the cut as chord:

Definition 5.1. For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$, any $f \in F$, any f-cut δ of G and any $d \in \delta$, we call an f-path (V_P, E_P) in G an f_d-path in G with respect to δ if and only if it holds for all $d' \in \delta \setminus \{d\}$ that $d' \not\subseteq V_P$.

We proceed by stating the two lemmata and the theorem in terms of f_d-paths.

Lemma 5.2. *For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$, any $f \in F$ and any f-cut δ of G, it is necessary for the cut inequality $1 - x_f \leq \sum_{e \in \delta}(1 - x_e)$ to be facet-defining for $\Xi_{G\widehat{G}}$ that an f_d-path in G with respect to δ exists for all $d \in \delta$. For the special case of $F = \{f\}$, this condition is also sufficient.*

Lemma 5.3. *For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$, any $f \in F$ and any f-cut δ of G, it is NP-hard to decide if an f_d-path in G with respect to δ exists for all $d \in \delta$, even for the special case of $F = \{f\}$.*

Theorem 5.4. *For any connected graph $G = (V, E)$, any augmentation $\widehat{G} = (V, E \cup F)$ with $E \cap F = \emptyset$, any $f \in F$ and any f-cut δ of G, it is NP-hard to decide if the cut inequality $1 - x_f \leq \sum_{e \in \delta}(1 - x_e)$ is facet-defining for $\Xi_{G\widehat{G}}$, even for the special case of $F = \{f\}$.*

In the remainder of this section, we prove first Theorem 5.4 and then Lemma 5.2 and Lemma 5.3.

Proof of Theorem 5.4. In case $F = \{f\}$, a cut inequality is facet defining if and only if there exists an f_d-path in G with respect to δ for all $d \in \delta$, by Lemma 5.2. Deciding if such paths exist is NP-hard, by Lemma 5.3. Together, this implies NP-hardness of deciding facet-definingness, even for the special case of $F = \{f\}$. □

Proof of Lemma 5.2. Necessity of an equivalent statement was already proven as Condition $C1$ of Theorem 5 of Andres et al. (2023).

We now show sufficiency. For this, let $F = \{f\} = \{uw\}$, let $\sigma = \{x \in X_{G\widehat{G}} \mid 1 - x_f = \sum_{d \in \delta}(1 - x_d)\}$ and assume that there exists an f_d-path in G with respect to δ for all $d \in \delta$. We prove that the cut inequality with respect to f and δ is facet-defining under the specified conditions by explicitly constructing $|E \cup F| - 1 = |E|$ linearly independent vectors in $\operatorname{lin}(\sigma - \sigma)$, implying $\dim \operatorname{aff} \sigma = \dim \operatorname{lin}(\sigma - \sigma) = |E|$ and thus that $\operatorname{aff} \sigma$ is a facet of $\Xi_{G\widehat{G}}$. In particular, we first construct the characteristic vectors of the elements in $E \setminus \delta$ and then $\mathbb{1}_{\{d,f\}}$ for all $d \in \delta$.

For any $e \in E \setminus \delta$, define:

$$V_1 = \{e\} \qquad V_2 = \emptyset\,.$$

As $G[U]$ is connected for any $U \in V_1$ and $U \in V_2$, we have $x^{V_1}, x^{V_2} \in X_{G\widehat{G}}$, by Lemma 3.2. It further holds $\mathbb{1}_{\{e\}} = -x^{V_1} + x^{V_2}$ and, for $j \in \{0, 1\}$, that $x^{V_j}_{uw} = 1$ and $x^{V_j}_{d} = 1$ for all $d \in \delta$, as for all $pq \in E \cup F$:

- $x^{V_1}_{pq} = 1$ and $x^{V_2}_{pq} = 1$ if $\nexists U \in V_1 \colon \{p, q\} \subseteq U$
- $x^{V_1}_{pq} = 0$ and $x^{V_2}_{pq} = 1$ if $\exists U \in V_1 \colon \{p, q\} \subseteq U\,.$

It follows from $x^{V_j}_{uw} = 1$ and $x^{V_j}_{d} = 1$ for all $d \in \delta$ that $x^{V_j} \in \sigma$. Thus, $\mathbb{1}_{\{e\}} = -x^{V_1} + x^{V_2} \in \operatorname{lin}(\sigma - \sigma)$.

For any $d \in \delta$, there exists an f_d-path $P = (V_P, E_P)$ in G with respect to δ according to our assumptions. We assume w.l.o.g. that this path is chordless and define:

$$V_1 = \{V_P\} \qquad V_2 = \emptyset\,.$$

Analogously to the previous case, we get $x^{V_1} \in X_{G\widehat{G}}$, $x^{V_2} \in \sigma$ and $\mathbb{1}_{E_P \cup \{f\}} = -x^{V_1} + x^{V_2}$. Using the same distinction of cases as before, we further get $x^{V_1}_{uw} = 0$ and, as P is an f_d path, $x_d^{V_1} = 0$ and $x_{d'}^{V_1} = 1$ for all $d' \in \delta \setminus \{d\}$, implying $x^{V_1} \in \sigma$. Consequently, $\mathbb{1}_{E_P \cup \{f\}} = -x^{V_1} + x^{V_2} \in \mathrm{lin}(\sigma - \sigma)$. We now note that the characteristic vector associated with f and d can be written as

$$\mathbb{1}_{\{f,d\}} = \mathbb{1}_{E_P \cup \{f\}} - \sum_{e \in E_P \setminus \{d\}} \mathbb{1}_{\{e\}} . \quad (10)$$

As $\mathbb{1}_{\{e\}} \in \mathrm{lin}(\sigma - \sigma)$ for all $e \in E_P \setminus \{d\}$ by the previous case, this implies $\mathbb{1}_{\{f,d\}} \in \mathrm{lin}(\sigma - \sigma)$. Altogether, we have constructed $|E|$ linearly independent vectors in $\mathrm{lin}(\sigma - \sigma)$ and have thus established sufficiency of the specified condition. □

Proof of Lemma 5.3. For showing NP-hardness, we use a reduction from the NP-hard 3-SAT problem with exactly three literals per clause and no duplicating literals within clauses (Schaefer, 1978). For any instance of this 3-SAT problem, with variables $x_1, x_2, \ldots, x_n$ and clauses $C_1, C_2, \ldots, C_m$, we construct in polynomial time an instance of our decision problem and show that it has a solution if and only if the instance of the 3-SAT problem has a solution. An example of this construction is depicted in Figure 6. We begin by defining two graphs, G_1 and G_2, which will be the components of G induced by the f-cut δ of our original decision problem.

In the first graph $G_1 = (V_1, E_1)$, there are $3m + 2$ nodes which are organized in $m + 2$ fully-connected layers. For $j \in \{0, 1, \ldots, m + 1\}$, we denote the set of nodes in the j-th layer by V_{1j}. The 0-th layer contains a single node u and the $m + 1$-th layer a single node d_1. The remaining m layers correspond to the m clauses $C_1, C_2, \ldots, C_m$ and contain three nodes each. The edges between consecutive layers are the only edges in E_1. For $j \in \{1, 2, \ldots, m\}$, we label each node in the j-th layer by a different literal in C_j. For completeness, we label u (respectively d_1) by a unique auxiliary propositional variable x_u (respectively x_{d_1}). For any $v \in V_1$, we let $l(v)$ denote the label of that node.

The second graph $G_2 = (V_2, E_2)$ is such that $V_1 \cap V_2 = \emptyset$ and $E_1 \cap E_2 = \emptyset$. It consists of $2n + 3$ nodes which are organized in $n + 3$ fully-connected layers. For $k \in \{0, 1, \ldots, n + 2\}$, we denote the set of nodes in the k-th layer by V_{2k}. The 0-th layer contains a single node d_2, the $n + 1$-th layer a single node w and the $n + 2$-th layer a single node w', which is connected to all other nodes of G_2, besides d_2, by a set of edges $E_2' \subseteq E_2$. The remaining n layers correspond to the n variables $x_1, x_2, \ldots, x_n$ and contain two nodes each. The edges between consecutive layers and the edges in E_2' are the only edges in E_2. For $k \in \{1, 2, \ldots, n\}$, we label one node in the k-th layer by x_k and the other by $\neg x_k$. Again, we label w (respectively d_2 and w') by a unique auxiliary propositional variable x_w (respectively x_{d_2} and $x_{w'}$) and denote the label of any $v \in V_2$ by $l(v)$.

We construct a third graph $G = (V, E)$ such that $V = V_1 \cup V_2$ and $E = E_1 \cup E_2 \cup \delta$ with

$$\delta = \{d_1 d_2, d_1 w'\} \cup \{st \subseteq V_1 \cup V_2 \mid s \in V_1 \wedge t \in V_2 \wedge l(s) = \neg\, l(t)\} .$$

Finally, we define a fourth graph $\widehat{G} = (V, E \cup F)$ such that $F = \{f\} = \{uw\}$. Note that G is connected and that δ is an f-cut of G, partitioning it into V_1 and V_2. Note also that $|F| = 1$, covering the part of the lemma claiming NP-hardness also for this special case.

For brevity, we introduce the symbol $d := d_1 d_2$, as this edge of the cut and its f_d-paths will be of particular importance in the remainder of the proof. For the same reason, we henceforth mean by an f_d-path an f_d-path in G with respect to δ and establish properties of such paths in the following claim, which is proven in Appendix A.

Claim 5.5. *The graph G has the following properties:*

(i) For any clause C_j and any f_d-path (V_P, E_P), there exists a literal in C_j that is labeled by a node from $V_P \cap V_1$.

(ii) Any f_d-path that contains a node in V_1 labeled $\neg\, x_k$ (respectively x_k) does not contain a node labeled x_k (respectively $\neg\, x_k$).

Using Claim 5.5, we show that the 3-SAT formula is satisfiable if and only if there exists an $f_{d'}$-path for every $d' \in \delta$, which finishes the reduction. We do so in two steps: Firstly, we show that the 3-SAT formula is satisfiable if and only if there exists an f_d-path for the specific edge $d \in \delta$. Secondly, we show that there always exists an $f_{d'}$-path for every other edge $d' \in \delta \setminus \{d\}$. This second statement is thereby necessary, as otherwise, even when the 3-SAT problem instance has a solution, the corresponding cut inequality might still not be facet-defining if there exists a $d' \in \delta \setminus \{d\}$ for which no $f_{d'}$-path exists.

Let $P = (V_P, E_P)$ be an f_d-path. We construct an assignment of truth values φ to the variables $x_1, x_2, \ldots, x_n$ satisfying the corresponding 3-SAT problem instance by setting $\varphi(x_k) = $ TRUE for all $k \in \{1, \ldots, n\}$ if and only if there exists a node $v \in V_P \cap V_1$ such that $l(v) = x_k$. Assume this assignment would not satisfy the 3-SAT problem instance. Then there exists a clause C_j assigning FALSE to all of its labels. By (i), there exists a node $v \in V_P \cap V_1$ that is labeled by a literal in C_j. If $l(v) = x_k$ for some variable x_k, then $\varphi(x_k) = $ TRUE, leading C_j to be true. If $l(v) = \neg\, x_k$, then $\varphi(x_k) = $ FALSE by (ii), leading C_j to be true as well. Consequently, such a clause C_j where all literals get assigned FALSE cannot exist and φ is a solution to the given 3-SAT problem instance.

Let now φ be an assignment of truth values to the variables $x_1, x_2, \dots, x_n$ that satisfies the corresponding instance of the 3-SAT problem. In the following, we will show that an f_d-path $P = (V_P, E_P)$ in G is given by

$$V_P = \{u, u_1, \dots, u_m, d_1, d_2, w_1, \dots, w_n, w\}$$
$$E_P = \{uu_1, u_1u_2, \dots, u_md_1, d_1d_2, d_2w_1, w_1w_2, \dots w_nw\},$$

where $u_j \in V_{1j}$ (respectively $w_k \in V_{2k}$) has a label that gets assigned TRUE by φ for all $j \in \{1, \dots, m\}$ (respectively $k \in \{1, \dots, n\}$). It is easy to see that P is an f-path in G. It remains to show that it is an f_d-path, i.e. that there exist no $d^* = d_1^*d_2^* \in \delta \setminus \{d\}$ such that $d^* \subseteq V_P$. Assume there exists such a d^*. As $w' \notin V_P$, it holds then $d^* \in \delta \setminus \{d, d_1w'\}$. By construction of δ, it follows $l(d_1^*) = \neg\, l(d_2^*)$. As both $l(d_1^*)$ and $\neg\, l(d_2^*)$ need to get assigned TRUE by φ according to the construction of P, this is a contradiction. Thus, there exists no such d^*, and P is an f_d-path. For an example of this correspondence between f_d-paths and solutions of the given 3-SAT problem instance, see again Figure 6.

Next, we regard the other edges of the cut. Let $d' = d_1'd_2' \in \delta \setminus \{d\}$ be an edge in the cut except d. We assume w.l.o.g. that $d_1' \in V_{1i}$ for some $i \in \{1, \dots, m+1\}$ and regard the path $P = (V_P, E_P)$ given by

$$V_P = \{u, u_1, u_2, \dots, u_{i-1}, d_1', d_2', w', w\}$$
$$E_P = \{uu_1, u_1u_2, \dots, u_{i-1}d_1', d_1'd_2', d_2'w', w'w\} \setminus \{w'\},$$

where u_j is an arbitrary node in V_{1j} such that $l(u_j) \neq l(d_1')$ for all $j \in \{1, \dots i-1\}$, and taking the set difference with $\{w'\}$ in the definition of E_P is necessary for P being a path in case $d_2' = w'$. Note that such u_j are guaranteed to exist as we consider the 3-SAT problem with exactly three literals per clause and no duplicated literals within clauses. Again, it is easy to see that P is an f-path, and it remains to show that there exists no $d^* = d_1^*d_2^* \in \delta \setminus \{d'\}$ such that $d^* \subseteq V_P$. Assume there exists such a d^*. Then one if its nodes, say d_1^*, must be in V_1 and its other node must be in V_2. We make a case distinction on whether $d_1^* \in V_1 \setminus \{d_1\}$ or $d_1^* = d_1$. If $d_1^* \in V_1 \setminus \{d_1\}$, then $d^* \in \delta \setminus \{d, d_1w'\}$. By construction of δ, it follows $l(d_1^*) = \neg\, l(d_2^*)$. As $d_2^* \in V_2 \cap V_P = \{d_2', w', w\}$ and $l(d_1^*) \neq l(d_1')$ according to the construction of P, this is a contradiction. On the other hand, if $d_1^* = d_1$, it holds by construction of P that $i = m+1$. As $d_1' \in V_{1i} = V_{1m+1} = \{d_1\}$, it follows $d_1' = d_1 = d_1^*$. Furthermore, as d_1d_2 and d_1w' are the only edges in δ containing d_1 and $d' \neq d$, we get $d_2' = w'$. Thus, we especially have $d_2 \notin V_2 \cap V_P = \{w', w\}$, leading to $d_2^* = w'$ when using the same argument as before. Consequently, $d^* = d'$ which contradicts $d^* \in \delta \setminus \{d'\}$. As both cases lead to a contradiction, there does not exist such a d^* and P is an f_d-path. This finishes the reduction from the 3-SAT problem and the proof of the lemma. □

6. Conclusion

We characterize in terms of efficiently decidable conditions the facets of lifted multicut polytopes induced by lower box inequalities. In addition, we show that deciding facet-defininingness of cut inequalities for lifted multicut polytopes is NP-hard, even for the special case of $|F| = 1$. Toward the design of cutting plane algorithms for the lifted multicut problem, our hardness result does not rule out the existence of inequalities strengthening the cut inequalities for which facet-defininingness and possibly also the separation problem can be solved efficiently. The search for such inequalities is one direction of future work. In our proof, we identify a structure (paths crossing the cut that have an edge of the cut as a chord) that complicates the characterization of cut inequalities. This structure exists for cuts (edge subsets, discussed in this article) but does not exist for separators (node subsets, not discussed in this article). This observation motivates the study of non-local connectedness with respect to separators instead of cuts.

Acknowledgements

This work is partly supported by the Federal Ministry of Education and Research of Germany through DAAD Project 57616814 (SECAI).

Impact Statement

This theoretical paper presents work whose goal is to advance the field of machine learning, more specifically clustering. As for all advances in this field, there are many potential societal consequences of our work, regarding the application of clustering algorithms for video surveillance, also some with negative impact. However, we do not feel that the implications of this paper differ from those of other contributions to that field and must be specifically highlighted here.

References

Andres, B., Di Gregorio, S., Irmai, J., and Lange, J.-H. A polyhedral study of lifted multicuts. *Discrete Optimization*, 47:100757, 2023. doi: 10.1016/j.disopt.2022.100757.

Bansal, N., Blum, A., and Chawla, S. Correlation clustering. *Machine Learning*, 56(1–3):89–113, 2004. doi: 10.1023/B:MACH.0000033116.57574.95.

Beier, T., Andres, B., Köthe, U., and Hamprecht, F. A. An efficient fusion move algorithm for the minimum cost lifted multicut problem. In *European Conference on Computer Vision (ECCV)*, 2016. doi: 10.1007/978-3-319-46475-6_44.

Beier, T., Pape, C., Rahaman, N., Prange, T., Berg, S., Bock, D. D., Cardona, A., Knott, G. W., Plaza, S. M., Scheffer, L. K., Köthe, U., Kreshuk, A., and Hamprecht, F. A. Multicut brings automated neurite segmentation closer to human performance. *Nature Methods*, 14:101–102, 2017. doi: 10.1038/nmeth.4151.

Chopra, S. and Rao, M. The partition problem. *Mathematical Programming*, 59(1):87–115, 1993. doi: 10.1007/BF01581239.

Demaine, E. D., Emanuel, D., Fiat, A., and Immorlica, N. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2–3):172–187, 2006. doi: 10.1016/j.tcs.2006.05.008.

Deza, M. M. and Laurent, M. *Geometry of Cuts and Metrics*. Springer, 1997. doi: 10.1007/978-3-642-04295-9.

Deza, M. M., Grötschel, M., and Laurent, M. Clique-web facets for multicut polytopes. *Mathematics of Operations Research*, 17(4):981–1000, 1992. doi: 10.1287/moor.17.4.981.

Grötschel, M. and Wakabayashi, Y. Facets of the clique partitioning polytope. *Mathematical Programming*, 47 (1):367–387, 1990. doi: 10.1007/BF01580870.

Horňáková, A., Lange, J.-H., and Andres, B. Analysis and optimization of graph decompositions by lifted multicuts. In *International Conference on Machine Learning (ICML)*, 2017. doi: 10.5555/3305381.3305540.

Keuper, M. Higher-order minimum cost lifted multicuts for motion segmentation. In *International Conference on Computer Vision (ICCV)*, 2017. doi: 10.1109/ICCV.2017.455.

Keuper, M., Levinkov, E., Bonneel, N., Lavoué, G., Brox, T., and Andres, B. Efficient decomposition of image and mesh graphs by lifted multicuts. In *International Conference on Computer Vision (ICCV)*, 2015. doi: 10.1109/ICCV.2015.204.

Keuper, M., Tang, S., Andres, B., Brox, T., and Schiele, B. Motion segmentation and multiple object tracking by correlation co-clustering. *Transactions on Pattern Analysis and Machine Intelligence*, 42(1):140–153, 2020. doi: 10.1109/TPAMI.2018.2876253.

Kostyukhin, V., Keuper, M., Ibragimov, I., Owtscharenko, N., and Cristinziani, M. Improving primary-vertex reconstruction with a minimum-cost lifted multicut graph partitioning algorithm. *Journal of Instrumentation*, 18(07): P07013, 2023. doi: 10.1088/1748-0221/18/07/P07013.

Levinkov, E., Uhrig, J., Tang, S., Omran, M., Insafutdinov, E., Kirillov, A., Rother, C., Brox, T., Schiele, B., and Andres, B. Joint graph decomposition and node labeling: Problem, algorithms, applications. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. doi: 10.1109/CVPR.2017.206.

Menger, K. Zur allgemeinen Kurventheorie. *Fundamenta Mathematicae*, 10:96–115, 1927.

Nguyen, D. M. H., Henschel, R., Rosenhahn, B., Sonntag, D., and Swoboda, P. LMGP: lifted multicut meets geometry projections for multi-camera multi-object tracking. In *Computer Vision and Pattern Recognition (CVPR)*, 2022. doi: 10.1109/CVPR52688.2022.00866.

Schaefer, T. J. The complexity of satisfiability problems. In *Symposium on Theory of Computing (STOC)*, 1978. doi: 10.1145/800133.804350.

Sørensen, M. M. A note on clique-web facets for multicut polytopes. *Mathematics of Operations Research*, 27(4): 740–742, 2002. doi: 10.1287/moor.27.4.740.301.

Tang, S., Andriluka, M., Andres, B., and Schiele, B. Multiple people tracking by lifted multicut and person reidentification. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. doi: 10.1109/CVPR.2017.394.

A. Additional Proofs

Proof of Claim 4.3. It follows directly from (8) that $H_j \cap H_k = \emptyset$ for any distinct $j, k \in \mathbb{N}_0$ and that $\cup_{j\geq 0} H_j \subseteq H$. Let $H_\infty = H \setminus \cup_{j\geq 0} H_j$, it remains to show that $H_\infty = \emptyset$. Assume this does not hold, then there exists an $st \in H_\infty$, and thus especially a simple st-path $(\{s,t\},\{st\})$ in $\widehat{G}$ whose edges are all in H_∞. We show that such a path cannot exist given (i) and (ii).

Assume there exist simple paths in $\widehat{G}$ whose edges are all in H_∞. Let $P = (V_P, E_P)$ with $E_P \subseteq H_\infty$ be one of those simple paths with maximum length, let $p, q \in V$ be its end-nodes and let $e_p, e_q \in E_P$ be the unique edges in E_P containing p and q, respectively. Recall that, by definition of H, all edges in E_P are uw-separators of G and no node in V_P is a uw-cut-node of G. By (8), there exists a $qq' \in E_{\widehat{G}}(q) \setminus \{e_q\}$ such that qq' is a uw-separator and $qq' \notin \cup_{j\geq 0} H_j$. By definition of H, this is equivalent to q' being either a uw-cut-node of G or $qq' \in H_\infty$. It is not possible that $qq' \in H_\infty$, as either $q' \in V_P$ and there exists a cycle in $(V_P, E_P \cup \{qq'\})$ that violates (ii), or $q' \notin V_P$ and $(V_P \cup \{q'\}, E_P \cup \{qq'\})$ is a simple path in G whose edges are all in H_∞, contradicting P to be the longest such path. Thus, q' must be a uw-cut-node. Further, it holds $q' \notin V_P$ as no node in V_P is a uw-cut-node of G. Analogously, there must exist a uw-cut-node $p' \in V \setminus V_P$ such that $p'p \in E_{\widehat{G}}(p) \setminus \{e_p\}$.

If $p' = q'$, the simple cycle $(V_P \cup \{p'\}, E_P \cup \{p'p, qq'\})$ violates (ii). If $p' \neq q'$, the simple $p'q'$-path $(V_P \cup \{p', q'\}, E_P \cup \{p'p, qq'\})$ violates (i). As both cases lead to a contradiction, there cannot exist simple paths in $\widehat{G}$ whose edges are all in H_∞, and thus especially no $st \in H_\infty$. □

Proof of Claim 5.5. For proving (i) and (ii), we first show that any f_d-path contains one node from each layer of G besides $V_{2n+2} = \{w'\}$. Assume this does not hold. Then there exists an f_d-path $P = (V_P, E_P)$ and a layer V_{1j} with $j \in \{0, \dots, m+1\}$ or V_{2k} with $k \in \{0, \dots, n+1\}$ such that no node in this layer is contained in P. As P is an f_d-path, it holds that $d_1 \in V_P$ and $\binom{V_P}{2} \cap \delta = \{d\}$. As $d_1 w' \in \delta$, this especially implies that $w' \notin V_P$ and thus $E_P \cap E_2' = \emptyset$. Hence, E_P must be a subset of the remaining edges $E_1 \cup E_2 \cup \{d\} \setminus E_2'$. As these edges only exist between consecutive layers and P contains $u \in V_{1,0}$ and $w \in V_{2,n+1}$, having a layer in-between for which P does not contain a node would imply P not being connected and thus results in a contradiction.

Assume (i) does not hold. Then there exists a clause C_j and an f_d-path P such that no node in $V_P \cap V_1$ is labeled by a literal in C_j. By construction of the labels, this would imply that there exists no node in P that is in V_{1j}, contradicting the discussion of the previous paragraph.

Assume (ii) does not hold. Then there exists an f_d-path $P = (V_P, E_P)$ containing an $s \in V_1 \cap V_P$ with $l(s) = \neg x_k$ (respectively x_k) and a $t \in V_P$ with $l(t) = x_k$ (respectively $\neg x_k$). We make a case distinction depending on whether t is in V_1 or V_2. Suppose $t \in V_1$. By the discussion of the first paragraph, P contains a $v \in V_{2k} \cap V_P$ with either $l(v) = x_k$ or $l(v) = \neg x_k$. By construction of δ, it follows either $sv \in \delta \setminus \{d\}$ or $tv \in \delta \setminus \{d\}$, contradicting P to be an f_d-path. Suppose $t \in V_2$. In this case, $st \in \delta \setminus \{d\}$ by construction of δ, contradicting P to be an f_d-path. As both cases lead a contradiction, such nodes s and t cannot exist. □