On imposing connectivity constraints in integer programs[*]

Yiming Wang   Austin Buchanan   Sergiy Butenko

February 5, 2017

Abstract

In many network applications, one searches for a *connected* subset of vertices that exhibits other desirable properties. To this end, this paper studies the *connected subgraph polytope* of a graph, which is the convex hull of subsets of vertices that induce a connected subgraph. Much of our work is devoted to the study of two nontrivial classes of valid inequalities. The first are the a, b-separator inequalities, which have been successfully used to enforce connectivity in previous computational studies. The second are the indegree inequalities, which have previously been shown to induce all nontrivial facets for trees. We determine the precise conditions under which these inequalities induce facets and when each class fully describes the connected subgraph polytope. Both classes of inequalities can be separated in polynomial time and admit compact extended formulations. However, while the a, b-separator inequalities can be lifted in linear time, it is NP-hard to lift the indegree inequalities.

Keywords: connected subgraph polytope; maximum-weight connected subgraph; prize-collecting Steiner tree; contiguity;

1 Introduction

In many clustering and network analysis applications, one searches for a *connected* subset of vertices that exhibits other desirable properties. To this end, this paper studies the *connected subgraph polytope* of a graph, which is the convex hull of subsets of vertices that induce a connected subgraph. This serves as a basic model of connectivity which can provide insights into more constrained problems.

In this paper, we consider a simple graph $G = (V, E)$ and usually let $n = |V|$ and $m = |E|$. The neighborhood of a vertex v is $N(v) := \{w \in V \mid \{v, w\} \in E\}$. We consider 0-vertex and 1-vertex graphs to be connected.

Definition 1. *The connected subgraph polytope of a graph $G = (V, E)$ on n vertices is*

$$\mathscr{P}(G) := \operatorname{conv} \left\{ x^S \in \{0, 1\}^n \mid G[S] \text{ is connected} \right\},$$

where x^S denotes the characteristic vector of $S \subseteq V$.

This requirement of *induced connectivity* is common in a variety of applications: the construction of virtual backbones for wireless sensor networks [6, 11], cluster detection in social network analysis [20, 27] and bioinformatics [3, 4, 10], forest harvesting [7], political districting [14], energy distribution [18], and computer vision [8, 26]. It also appears in resource scheduling problems when enforcing "contiguous ones" in a binary decision vector (here, G is a path graph). For an excellent introduction to modeling induced connectivity constraints in integer programs, consult [7] and the references therein.

It is important to distinguish these connectivity constraints from those encountered in other network design problems. Many times, one is tasked with choosing *edges* to meet some vertex- or edge-connectivity requirements among a set of vertices. For example, in the Steiner tree problem one is tasked with choosing a minimum-cost subset of edges to connect a *subset* of specified terminals. In several other network design problems, edges are chosen to meet some edge-connectivity or vertex-connectivity constraints for the *entire* vertex set. For many of these problems, there are thorough polyhedral studies [23].

However, the connected subgraph polytope is not nearly as well-studied. One can find integer programming formulations for induced connectivity [1, 2] and for the related prize-collecting Steiner tree problem [18, 19], but these papers work in a higher-dimensional space and the work is focused more on developing effective branch-and-cut algorithms.

In contrast, this paper is devoted to developing a thorough understanding of the facial structure of the connected subgraph polytope in the original space of variables. This is motivated, in part, by recent computational successes [6, 7, 13] that rely on vertex variables and a, b-separator inequalities to impose connectivity constraints.

$$(a, b\text{-separator inequality}) \qquad x_a + x_b - \sum_{i \in C} x_i \leq 1.$$

Here, a and b are nonadjacent vertices and $C \subseteq V \setminus \{a, b\}$ is an a, b-separator, i.e., there is no path from a to b in $G - C$.

This focus on vertex variables has proven useful in a variety of contexts. Carvajal et al. [7] handle forest planning instances two to three times as large as previous studies. Buchanan et al. [6] solve, in ten seconds, more instances of the minimum connected dominating set problem than any previous approach could solve in an hour. Fischetti et al. [13] solve, in seconds, benchmark instances of Steiner tree problems that were never solved by previous approaches "even after days of computation." Despite these significant computational improvements, fundamental questions about induced connectivity polyhedra remain unanswered. For example, when do the a, b-separator inequalities induce facets? We answer this and many other questions about the a, b-separator inequalities.

We also consider a class of valid inequalities called indegree inequalities. These inequalities are interesting because they have been shown to induce all nontrivial facets of $\mathscr{P}(G)$ when G is a tree [17]. As far as we know, no one has studied them for arbitrary graphs. They are defined as follows. A vector $d \in \mathbb{R}^n$ is said to be an indegree vector if for some orientation of G the indegree of each vertex v is d_v. For each indegree vector d of G, there is a corresponding indegree inequality.

$$(\text{indegree inequality}) \qquad \sum_{i \in V} (1 - d_i) x_i \leq 1.$$

It turns out that these inequalities are valid for arbitrary graphs and can induce facets even when G is not a tree.

There can be exponentially many a, b-separator and indegree inequalities. Often, this leads researchers to develop branch-and-cut algorithms that rely on (hopefully) efficient separation algorithms. For this reason, it is important to study the separation problems associated with these two classes of inequalities. The separation problem for the a, b-separator inequalities can be solved in polynomial time by a standard reduction to the maximum flow problem (as noted by, e.g., [13]). We show that the separation problem for the indegree inequalities can be solved in linear time.

Another popular approach to handle a large number of inequalities in the original space of variables is to search for small extended formulations. We provide positive results on this front, showing that the feasible regions defined by the separator and indegree inequalities admit polynomial-size extended formulations. Indeed, in the case of indegree inequalities, there is a linear-size extended formulation. Thus, for any tree G, there is a linear-size extended formulation for $\mathscr{P}(G)$, even though $\mathscr{P}(G)$ has $2^{n-1} + n$ facets [17].

Lifting is important when studying the connected subgraph polytope $\mathscr{P}(G)$ if one wants to generate facets of $\mathscr{P}(G)$ from subgraphs of G. We show that the lifting problem associated with the a, b-separator inequalities can be solved in linear time. However, we show that this is not expected to be the case for the indegree inequalities. In fact, it is strongly NP-hard even in very restricted classes of graphs.

Another natural question is—When do these classes of inequalities fully describe $\mathscr{P}(G)$? As mentioned earlier, the indegree inequalities have been shown to induce all nontrivial facets for trees [17]. We generalize this result, showing that the indegree inequalities induce all facets of $\mathscr{P}(G)$ (aside from nonnegativity bounds) if and only if G is a forest. We also show that the a, b-separator inequalities induce all nontrivial facets of $\mathscr{P}(G)$ if and only if G has no independent set of three vertices. This generalizes work of Biha et al. [5], who (essentially) prove similar results for *line* graphs. These analytical results suggest that indegree inequalities may be useful for sparse, treelike graphs, and separator inequalities may be useful for very dense graphs. Our numerical experiments corroborate this intuition.

1.1 Our contributions

In Section 2, we study some fundamental properties of $\mathscr{P}(G)$, including its dimension and when the trivial 0-1 bounds induce facets. We also show that all facets of $\mathscr{P}(G)$ can be derived from its components' polytopes and vice versa. In many cases, this simplifies our analysis, as we can suppose that G is connected. Many of our proofs rely on lifting arguments, so we also provide some background information about lifting. We also show that any facet-defining inequality of $\mathscr{P}(G)$ has at most $\alpha(G)$ positive coefficients, where $\alpha(G)$ is the independence number of G. This will prove useful later.

In Section 3, we study the a, b-separator inequalities. We show that, assuming the input graph is connected, an a, b-separator inequality induces a facet if and only if the separator is a minimal a, b-separator. We then show that one can lift a vertex into a separator inequality in linear time. While there are exponentially many (facet-defining) separator inequalities, we provide a compact extended formulation for a separator-based relaxation $Q(G)$. Generalizing a result of Biha et al. [5], we show that $Q(G)$ and $\mathscr{P}(G)$ coincide precisely when the graph has no independent set of three vertices, i.e., $\alpha(G) \leq 2$.

3

In Section 4, we study the indegree inequalities. We show that they are valid for arbitrary graphs. Moreover, they are facet-defining if and only if the oriented graph D satisfies the property that if there is a directed s-t walk, then it is unique. (In particular, D should be acyclic.) We show that lifting a vertex into a given indegree inequality is NP-hard. Then we provide a linear-time algorithm to separate indegree inequalities and a linear-size extended formulation for an indegree-based relaxation $Q'(G)$. Generalizing results of [17], we show that $Q'(G)$ and $\mathscr{P}(G)$ coincide precisely when G is a forest.

In Section 5, we provide a brief set of numerical experiments. As expected, the indegree inequalities provide strong relaxations when the graph is tree-like, and give poor LP bounds when the graph is sufficiently dense. Similarly, the a, b-separator inequalities are very weak for tree-like graphs, but the LP bound becomes stronger for denser graphs.

1.2 Preliminaries and related work

The connected subgraph polytope $\mathscr{P}(G)$ of a graph G has close connections with the maximum-weight connected subgraph (MWCS) problem. Indeed, $\mathscr{P}(G)$ is essentially the feasible region for the MWCS problem.

Problem: Maximum-Weight Connected Subgraph (MWCS).
Input: a graph $G = (V, E)$, a weight w_i (possibly negative) for each $i \in V$.
Output: a maximum-weight subset $S \subseteq V$ such that $G[S]$ is connected.

Here, the *weight* of a subset $S \subseteq V$ of vertices is $w(S) = \sum_{i \in S} w_i$.

The MWCS problem is NP-hard, even in planar graphs of maximum degree three with all weights either $+1$ or -1 [16]. This suggests that $\mathscr{P}(G)$ must be "complicated" even in very restricted classes of graphs.

We use the following lemma throughout the paper. It is rather simple, but since we use it so often it is stated explicitly.

Lemma 1 (folklore). *Let $ax \leq b$ and $cx \leq d$ be valid inequalities for a* full-dimensional *polyhedron P such that (a, b) and (c, d) are not scalar multiples of each other. Then, the aggregated inequality $(a + c)x \leq (b + d)$ cannot induce a facet of P.*

For any other notation or necessary background knowledge on polyhedral theory and graph theory, consult [9, 21, 25].

In related work, Biha et al. [5] study the *edge-induced* connected subgraph polytope of a graph G, which we will denote by $\mathscr{P}_E(G)$. It is defined as the convex hull of *edge* subsets that induce a connected subgraph. They show, among other things, that, for connected graphs G, $\mathscr{P}_E(G)$ is fully described by the 0-1 bounds and edge-based analogues of the a, b-separator inequalities if and only if G has no matching of 3 edges. This is implied by our results through the use of the *line graph*. Recall that the line graph $L(G)$ of a graph $G = (V, E)$ has vertex set E and two vertices of $L(G)$ are adjacent if the corresponding edges in G share an endpoint. A graph H is said to be line if there exists a graph G such that $H = L(G)$.

The result of Biha et al. can thus be stated in terms of our problem as follows. Supposing G is a line graph, $\mathscr{P}(G)$ is fully described by the 0-1 bounds and a, b-separator inequalities if and only if G has no independent set of three vertices. We note that their result does not imply ours. For example, our result shows that the separator-based LP relaxation is

not integral for the claw graph, but this graph is not a line graph, and thus their theorem does not apply. Similarly, our results imply that the (vertex-induced) connected subgraph polytope of the 6-vertex wheel graph is completely described by the 0-1 bounds and separator inequalities, but again we cannot apply the theorem of Biha et al. since the 6-vertex wheel is not line.

2 Fundamental Properties of the Connected Subgraph Polytope

In this section, we describe fundamental properties of $\mathscr{P}(G)$, including when the 0-1 bounds induce facets. Lifting arguments are a primary tool used in this paper's proofs, so we also provide some background information about lifting.

2.1 Trivial facets

Proposition 1 (dimension; 0-1 facets). *The connected subgraph polytope $\mathscr{P}(G)$ of a graph $G = (V, E)$ is full-dimensional. Moreover, for each $i \in V$,*

1. *$x_i \geq 0$ induces a facet, and*

2. *$x_i \leq 1$ induces a facet if and only if G is connected.*

Proof. The usual $n + 1$ affinely independent points $\mathbf{0}$ and e_i, $i = 1, \ldots, n$ suffice to show full dimension. The points $\mathbf{0}$ and e_j, $j \neq i$ show that $x_i \geq 0$ induces a facet. When G is connected, consider the vertices $i = v_1, v_2, \ldots, v_n$ in a preordering (i.e., when first visited in a depth-first search) starting from i. Then the n affinely independent points $\sum_{j=1}^{k} e_{v_j}$ for $k = 1, \ldots, n$ show that $x_i \leq 1$ induces a facet. When G is not connected, then consider a vertex j that belongs to a different component of G than i. Then the valid inequalities $x_i + x_j \leq 1$ and $-x_j \leq 0$ imply $x_i \leq 1$, meaning that, by Lemma 1, $x_i \leq 1$ cannot induce a facet. $\square$

Lemma 2. *Consider a facet-defining inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ of $\mathscr{P}(G)$. Then $\pi_0 \geq 0$. Further, the inequality is (a scalar multiple of) some nonnegativity bound $-x_j \leq 0$ if and only if $\pi_0 = 0$.*

Proof. As the empty set is assumed to induce a connected subgraph, $\pi_0 \geq 0$. The 'only if' direction is trivial, so suppose $\pi_0 = 0$. Then $\pi_i \leq 0$ for each vertex $i \in V$ (since the trivial graphs are assumed to be connected). Further suppose that at least two coefficients are negative, say π_u and π_v. Then $\sum_{i \in V} \pi_i x_i \leq 0$ is implied by the valid inequalities $\pi_u x_u \leq 0$ and $\sum_{i \in V \setminus \{u\}} \pi_i x_i \leq 0$ so (by Lemma 1) $\sum_{i \in V} \pi_i x_i \leq 0$ cannot be facet-defining. $\square$

2.2 Generating facets from components

Proposition 1 shows that the facets of $\mathscr{P}(G)$ depend on whether G is connected. We expound upon this in the following theorem, showing that $\mathscr{P}(G)$ is determined by its components' connected subgraph polytopes.

Theorem 1. *Let $\{G_j = (V_j, E_j)\}_j$ be the (connected) components of a graph $G = (V, E)$ and consider $\pi \in \mathbb{R}^n$ and $\pi_0 > 0$. The following are equivalent.*

1. *For each G_j, the inequality $\sum_{i \in V_j} \pi_i x_i \le \pi_0$ induces a facet of $\mathscr{P}(G_j)$.*

2. *The inequality $\sum_{i \in V} \pi_i x_i \le \pi_0$ induces a facet of $\mathscr{P}(G)$.*

Proof. The proof is relatively straightforward and is omitted. $\qquad\square$

The inequalities below are the so-called *clique inequalities* discussed by [7] in the context of induced connectivity.

Corollary 1. *If $U \subseteq V$ contains exactly one vertex from each connected component of G, then the inequality $\sum_{i \in U} x_i \le 1$ induces a facet of $\mathscr{P}(G)$.*

Proof. Let the components of G be $\{G_j = (V_j, E_j)\}_j$. By Proposition 1, for any $i \in U \cap V_j$, the inequality $x_i \le 1$ induces a facet of $\mathscr{P}(G_j)$. Then, by Theorem 1, $\sum_{i \in U} x_i \le 1$ induces a facet of $\mathscr{P}(G)$. $\qquad\square$

Corollary 2. *For a graph $G = (V, E)$ and an independent set $S \subseteq V$, the inequality $\sum_{i \in S} x_i \le 1$ induces a facet of $\mathscr{P}(G[S])$.*

2.3 Basics of lifting

Corollary 2 shows that we can easily generate the facet-defining inequality $\sum_{i \in S} x_i \le 1$ for $\mathscr{P}(G[S])$, where S is an independent set. However, we want facet-defining inequalities for $\mathscr{P}(G)$, and this inequality is perhaps not even valid for $\mathscr{P}(G)$. Loosely speaking, *lifting* is the procedure whereby this or other *seed* inequalities are altered so that they induce facets of $\mathscr{P}(G)$.

Lemma 3 (Lifting zero-valued variables for $\mathscr{P}(G)$)**.** *If the inequality $\sum_{i \in V \setminus \{v\}} \pi_i x_i \le \pi_0$ induces a facet of $\mathscr{P}(G - v)$, then the inequality*

$$(\pi_0 - \zeta)x_v + \sum_{i \in V \setminus \{v\}} \pi_i x_i \le \pi_0$$

induces a facet of $\mathscr{P}(G)$, where

$$\zeta = \max_{S \subseteq V} \left\{ \sum_{i \in V \setminus \{v\}} \pi_i x_i^S \ \middle| \ x_v^S = 1 \text{ and } G[S] \text{ is connected} \right\}.$$

Proof. This follows by well-known results (cf. Prop. 1.1 on p. 261 of [21]). $\qquad\square$

This lifting principle provides a way to generate facets for $\mathscr{P}(G)$ from facets of its subgraphs' polytopes. A key idea is that this can be applied sequentially based on some lifting order. This machinery is vital for our proofs.

Lemma 4 (Bounds on lifting)**.** *Suppose $\sum_{i \in V \setminus \{v\}} \pi_i x_i \le \pi_0$ induces a facet of $\mathscr{P}(G - v)$ and that $\pi_0 > 0$. Then, when lifting in v, the objective ζ of the lifting problem satisfies:*

1. *if v is isolated, then $\zeta = 0$;*

2. *if v is not isolated, then $\pi_0 \le \zeta \le |N(v)|\pi_0$.*

Proof. If vertex v is isolated in the graph $G = (V, E)$, then the only feasible solution is $\{v\}$, in which case $\zeta = 0$. So, we will now suppose that $N(v) \neq \emptyset$.

Consider an optimal solution $D \subseteq V$ to the lifting problem. Here, $v \in D$ and $G[D]$ is connected. Suppose $N(v) = \{u_1, \ldots, u_s\}$. Partition $D' := D \setminus \{v\}$ into s (some possibly empty) subsets as follows. Let D_1 denote the set of vertices in D' connected to u_1 by some path of $G[D']$. Then for $p = 2, \ldots, s$, let D_p denote the vertices of $D' \setminus (D_1 \cup \cdots \cup D_{p-1})$ that are connected to u_p by some path in $G[D']$. Each $G[D_p]$ is a connected subgraph of $G - v$, so by the validity of the seed inequality,

$$\sum_{i \in V \setminus \{v\}} \pi_i x_i^{D_p} = \sum_{i \in D_p} \pi_i \leq \pi_0,$$

implying that

$$\begin{aligned}
\zeta &= \sum_{i \in V \setminus \{v\}} \pi_i x_i^D \\
&= \sum_{p=1}^{s} \left(\sum_{i \in D_p} \pi_i \right) + \sum_{i \in V \setminus D} \pi_i x_i^D \\
&\leq s\pi_0 + 0 \\
&= |N(v)|\pi_0.
\end{aligned}$$

Now we show that $\zeta \geq \pi_0$ when $N(v) \neq \emptyset$. Pick $u \in N(v)$. Let $G' = (V', E')$ be the connected component of $G - v$ that includes u. Then, by Theorem 1, $\sum_{i \in V'} \pi_i x_i \leq \pi_0$ induces a facet of $\mathscr{P}(G')$. Moreover, there must be at least one connected vertex subset $D \subseteq V'$ containing u for which $\sum_{i \in V'} \pi_i x_i^D = \pi_0$, since otherwise the inequality could not induce a facet. Then, $G[D \cup \{v\}]$ is connected and D has weight π_0, so $\zeta \geq \pi_0$. $\qquad\square$

2.4 The number of positive coefficients in a facet-defining inequality

Lemma 5. *Suppose that $\sum_{i \in V} \pi_i x_i \leq \pi_0$ is valid for $\mathscr{P}(G)$. If vertices u and v are adjacent and $\pi_v \geq 0$, then the following inequality is also valid.*

$$(\pi_u + \pi_v)x_u + 0x_v + \sum_{i \in V \setminus \{u,v\}} \pi_i x_i \leq \pi_0.$$

Proof. Suppose that $G[S]$ is connected, and consider the following two cases.

- If $u \in S$, then $S' = S \cup \{v\}$ is also connected, so

$$\begin{aligned}
&(\pi_u + \pi_v)x_u^S + 0x_v^S + \sum_{i \in V \setminus \{u,v\}} \pi_i x_i^S \\
&= (\pi_u + \pi_v)x_u^{S'} + 0x_v^{S'} + \sum_{i \in V \setminus \{u,v\}} \pi_i x_i^{S'} = \sum_{i \in V} \pi_i x_i^{S'} \leq \pi_0.
\end{aligned}$$

- If $u \notin S$, then since $x_u^S = 0$ and $0x_v^S \leq \pi_v x_v^S$, we have

$$(\pi_u + \pi_v)x_u^S + 0x_v^S + \sum_{i \in V \setminus \{u,v\}} \pi_i x_i^S \leq \sum_{i \in V} \pi_i x_i^S \leq \pi_0.$$

7

Thus, the inequality is valid in both cases, and is valid in general. □

Lemma 6. *In a facet-defining inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ of $\mathscr{P}(G)$, no pair of adjacent vertices can have positive coefficients.*

Proof. Suppose that adjacent vertices u and v have positive coefficients. Then, by Lemma 5, the following inequalities are valid.

$$(\pi_u + \pi_v)x_u + 0x_v + \sum_{i \in V \setminus \{u,v\}} \pi_i x_i \leq \pi_0;$$

$$0x_u + (\pi_u + \pi_v)x_v + \sum_{i \in V \setminus \{u,v\}} \pi_i x_i \leq \pi_0.$$

Multiplying the first inequality by $\beta := \pi_u/(\pi_u + \pi_v)$ and the second by $1 - \beta$ shows, by Lemma 1, that $\sum_{i \in V} \pi_i x_i \leq \pi_0$ cannot induce a facet. □

Recall that the independence number $\alpha(G)$ of a graph G is the size of its largest independent set.

Proposition 2. *There is a facet-defining inequality of $\mathscr{P}(G)$ with $\alpha(G)$ positive coefficients. None have more.*

Proof. Let S be a maximum independent set of G. The first claim follows by lifting the inequality $\sum_{i \in S} x_i \leq 1$, which induces a facet of $\mathscr{P}(G[S])$ by Corollary 2. The second claim follows by Lemma 6. □

We show that the 0-1 bounds, on their own, provide an $\alpha(G)$ polyhedral approximation of $\mathscr{P}(G)$.

Proposition 3. *The inclusions $\mathscr{P}(G) \subseteq [0,1]^n \subseteq \alpha(G)\mathscr{P}(G)$ hold and are sharp.*

Proof. The first inclusion is trivial. Consider $x^* \in [0,1]^n$ and a facet-defining inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ of $\mathscr{P}(G)$. Let $S = \{i \in V \mid \pi_i > 0\}$. If $|S| = 0$, then $\pi_0 = 0$, since otherwise no feasible point could satisfy it at equality. By Lemma 2, the inequality must be a nonnegativity bound $\pi_j x_j \leq 0$, so

$$\sum_{i \in V} \pi_i x_i^* = \pi_j x_j^* \leq 0 = \alpha(G)\pi_0.$$

Now suppose $|S| \geq 1$, implying $\pi_0 > 0$. Then, by Lemma 6, S must be an independent set. For each vertex $i \in S$, we have $\pi_i \leq \pi_0$ and $x_i^* \leq 1$, so

$$\sum_{i \in V} \pi_i x_i^* \leq \sum_{i \in S} \pi_i x_i^* \leq |S|\pi_0 \leq \alpha(G)\pi_0.$$

Thus $x^* \in \alpha(G)\mathscr{P}(G)$. The inclusions are sharp for complete graphs K_n. □

3   Separator Inequalities

In this section, we study the separator inequalities. We show that, assuming the graph is connected, the a, b-separator inequality

$$(a, b\text{-separator inequality}) \qquad x_a + x_b - \sum_{i \in C} x_i \leq 1$$

induces a facet of $\mathscr{P}(G)$ if and only if C is a minimal a, b-separator. The problem of lifting a vertex into an a, b-separator inequality is shown to be linear-time solvable. Also, the separation problem for these inequalities is polynomial-time solvable, meaning that we can optimize over the linear programming relaxation $Q(G)$ for $\mathscr{P}(G)$ in polynomial time via the ellipsoid method [15].

$$Q(G) := \{x \in [0,1]^n \mid x \text{ satisfies all separator inequalities}\}.$$

We also provide a compact extended formulation for $Q(G)$ via multicommodity flows, so we need not rely on the ellipsoid method to optimize over $Q(G)$.

A natural question to ask is—when is formulation $Q(G)$ tight, i.e., $\mathscr{P}(G) = Q(G)$? We show that this is the case precisely when the graph has no independent set of three vertices, i.e., $\alpha(G) \leq 2$. This result is interesting, in part, because there can be exponentially many inequalities defining $Q(G)$ even when $\alpha(G) = 2$. An example is shown in Figure 1, where the vertices within each rectangle form a clique. A minimal a, b-separator can be created by choosing, for each i, one vertex from $\{c_i, d_i\}$. The number of such separators is $2^{n/2-1}$.


Figure 1: A graph G with $\alpha(G) = 2$, but many minimal a, b-separators. Vertices within a rectangle form a clique.

3.1   Separator facets

We provide a good characterization for when an a, b-separator inequality induces a facet. Recall that an a, b-separator is said to be *minimal* if no proper subset of it is an a, b-separator.

Theorem 2 (a, b-separator facets)**.** *Consider a connected graph $G = (V, E)$; distinct, non-adjacent vertices a and b; and a vertex subset $C \subseteq V \setminus \{a, b\}$. Then, the inequality*

$$x_a + x_b - \sum_{i \in C} x_i \leq 1 \tag{1}$$

induces a facet of $\mathscr{P}(G)$ if and only if C is a minimal a, b-separator.

Proof. ($\implies$) Suppose that C is not an a, b-separator. Then there exists a path from a to b in $G - C$. Let P be the set of vertices in the path (including a and b). Then $G[P]$ is connected, but

$$x_a^P + x_b^P - \sum_{i \in C} x_i^P = x_a^P + x_b^P = 2 > 1,$$

so $x_a + x_b - \sum_{i \in C} x_i \leq 1$ is not valid. This shows that C is an a, b-separator. Now suppose C is not a *minimal a, b-separator*. Then there exists $c \in C$ such that $C \setminus \{c\}$ is an a, b-separator. Then, by Lemma 1, the two valid inequalities $-x_c \leq 0$ and $x_a + x_b - \sum_{i \in C \setminus \{c\}} x_i \leq 1$ show that inequality (1) cannot induce a facet. This shows that C is a minimal a, b-separator.

($\impliedby$) Suppose that C is a minimal a, b-separator and define

$$A := \{v \in V \mid v \text{ and } a \text{ belong to the same component of } G - C\}$$
$$B := \{v \in V \mid v \text{ and } b \text{ belong to the same component of } G - C\}$$
$$D := V \setminus (A \cup B \cup C).$$

The proof continues by sequential lifting; we provide only a brief sketch. First, $x_a + x_b \leq 1$ induces a facet of $\mathscr{P}(G[A \cup B])$ by Corollary 1. Then, $x_a + x_b - \sum_{i \in C} x_i \leq 1$ induces a facet of $\mathscr{P}(G[A \cup B \cup C])$, by lifting in the vertices of C arbitrarily. Finally, $x_a + x_b - \sum_{i \in C} x_i \leq 1$ induces a facet of $P(G) = P(G[A \cup B \cup C \cup D])$, by lifting in the vertices of D based on their distance from (a vertex of) C; specifically, lift closer vertices first. $\qquad \square$

3.2   Lifting separator inequalities in linear time

Suppose that we have a facet-defining a, b-separator inequality $x_a + x_b - \sum_{i \in C} x_i \leq 1$ for $\mathscr{P}(G - v)$, and that we want to lift in vertex v so that the inequality induces a facet of $\mathscr{P}(G)$. Consider the following algorithm, recalling that the coefficient for variable x_v will be $\pi_v = \pi_0 - \zeta = 1 - \zeta$.

1. let $A := \{i \in V \mid i \text{ and } a \text{ belong to the same component of } G - C\}$;

2. let $B := \{i \in V \mid i \text{ and } b \text{ belong to the same component of } G - C\}$;

3. if $N(v) = \emptyset$, then return $\zeta = 0$;

4. if $A \neq B$, then return $\zeta = 1$;

5. if $A = B$, then return $\zeta = 2$.

Proposition 4. *The above algorithm (optimally) lifts vertex v into a given a, b-separator inequality in linear time.*

Proof. Since we can construct sets A and B in linear time via breadth-first search, the algorithm runs in linear time. If $N(v) = \emptyset$, then $\zeta = 0$ by Lemma 4, so the algorithm is correct. So suppose $N(v) \neq \emptyset$, and consider two cases.

In the first case, suppose that $v \in A$ and $v \in B$. Thus $A = B$. Further, A is connected with weight 2 and contains v, so $\zeta \geq 2$. The reverse inequality is easy to see, so $\zeta = 2$, and the algorithm is correct.

In the second case, suppose that $v \notin A$ or $v \notin B$. Thus C is an a,b-separator for G, so $A \neq B$ and $\zeta \leq 1$. The inequality $\zeta \geq 1$ follows by Lemma 4, so $\zeta = 1$, and the algorithm is correct. $\qquad\square$

3.3 A compact extended formulation for the separator-based relaxation

Again, consider a simple graph $G = (V, E)$. Denote by $\mathcal{E}$ the set of all directed edges (u, v) and (v, u) whose undirected counterparts $\{u, v\}$ belong to E. Order the vertices of G arbitrarily, and denote by $\overline{\mathcal{E}}$ the set of ordered pairs (a, b) of nonadjacent vertices for which a appears before b in the vertex ordering. The polytope $F(G)$ is the set of all (x, f) satisfying the following constraints.

$$-x_i + \sum_{j \in N(i)} f_{ij}^{ab} \leq 0, \ \forall i \in V, \ \forall (a, b) \in \overline{\mathcal{E}} \tag{2}$$

$$x_a + x_b - \left(\sum_{j \in N(a)} f_{aj}^{ab} - \sum_{j \in N(a)} f_{ja}^{ab} \right) \leq 1, \ \forall (a, b) \in \overline{\mathcal{E}} \tag{3}$$

$$\sum_{j \in N(i)} f_{ji}^{ab} - \sum_{j \in N(i)} f_{ij}^{ab} = 0, \ \forall i \in V \setminus \{a, b\}, \ \forall (a, b) \in \overline{\mathcal{E}} \tag{4}$$

$$0 \leq f_{ij}^{ab} \leq 1, \ \forall (i, j) \in \mathcal{E}, \ \forall (a, b) \in \overline{\mathcal{E}} \tag{5}$$

$$0 \leq x_i \leq 1, \ \forall i \in V. \tag{6}$$

Proposition 5. *The separator-based relaxation $Q(G)$ for $\mathscr{P}(G)$ admits an extended formulation of size $O(n^2(m + n))$. Indeed, $\mathrm{proj}_x(F(G)) = Q(G)$.*

Proof. The equality $\mathrm{proj}_x(F(G)) = Q(G)$ holds by max-flow/min-cut ideas, and the size is easy to see. $\qquad\square$

3.4 When the separator-based relaxation is tight

This subsection is devoted to proving the following theorem, which generalizes a result of Biha et al. [5].

Theorem 3. *The equality $\mathscr{P}(G) = Q(G)$ holds if and only if $\alpha(G) \leq 2$.*

As a consequence of Theorem 3 and the ability to optimize over $Q(G)$ in polynomial time, we have the following corollary.

Corollary 3. *If $\alpha(G) \leq 2$, then the MWCS problem is polynomial-time solvable.*

One direction of the proof of Theorem 3 is easier and is shown first.

Lemma 7. *If $\mathscr{P}(G) = Q(G)$, then $\alpha(G) \leq 2$.*

Proof. By the contrapositive. Suppose that G has an independent set S of three vertices. By Proposition 2, there is a facet-defining inequality of $\mathscr{P}(G)$ that has three positive coefficients, but this is not true for $Q(G)$. Each of $\mathscr{P}(G)$ and $Q(G)$ has a unique half-space representation (up to scalar multiples) since they are full-dimensional, but we have seen that the facets of $Q(G)$ and $\mathscr{P}(G)$ are different, so $\mathscr{P}(G) \neq Q(G)$. $\qquad\square$

The converse is more complicated and requires two lemmata.

Lemma 8. *Suppose that $\sum_{i \in V} \pi_i x_i \leq \pi_0$ induces a facet of $\mathscr{P}(G)$. If π_u and π_v are its only positive coefficients, then $\pi_u = \pi_v = \pi_0$.*

Proof. Since $G[\{u\}]$ and $G[\{v\}]$ are connected, this implies that $\pi_u \leq \pi_0$ and $\pi_v \leq \pi_0$. If $\pi_u + \pi_v \leq \pi_0$, then any 0-1 solution $x^* \in \mathscr{P}(G)$ satisfying the inequality at equality must have $x_u^* = x_v^* = 1$, implying that the face of $\mathscr{P}(G)$ where $\sum_{i \in V} \pi_i x_i = \pi_0$ has dimension at most $n - 2$, meaning that the inequality cannot induce a facet. Thus, we will assume that $\pi_u + \pi_v > \pi_0$.

We claim that $S := \{i \in V \mid \pi_i < 0\}$ is a u, v-separator. Suppose not, then there exists a path from u to v in $G[V \setminus S]$. Let P be the set vertices in the path. This implies that $\sum_{i \in V} \pi_i x_i^P = \pi_u + \pi_v > \pi_0$, which contradicts the validity of $\sum_{i \in V} \pi_i x_i \leq \pi_0$.

For contradiction purposes, suppose that at least one of π_u and π_v is less than π_0. Without loss of generality, suppose that $\pi_u < \pi_0$. Now, let $S' \subseteq S$ be a minimal u, v-separator, and define

$$\pi_{max} := \max\{\pi_i \mid i \in S'\}$$

$$\epsilon := \frac{1}{2} \min\{-\pi_{max}, \pi_0 - \pi_u\}.$$

Note that $\pi_{max} < 0$ and $\pi_0 - \pi_u > 0$, so $\epsilon > 0$. Also, $\pi_u + \epsilon < \pi_0$, and for every $i \in S'$, we have $\pi_i + \epsilon < 0$. Further, let

$$R = V \setminus (S' \cup \{u, v\}).$$

Then consider the following inequalities.

$$(\pi_u + \epsilon)x_u + \pi_v x_v + \sum_{i \in S'}(\pi_i - \epsilon)x_i + \sum_{i \in R} \pi_i x_i \leq \pi_0 \tag{7}$$

$$(\pi_u - \epsilon)x_u + \pi_v x_v + \sum_{i \in S'}(\pi_i + \epsilon)x_i + \sum_{i \in R} \pi_i x_i \leq \pi_0. \tag{8}$$

If these inequalities were valid, then they would imply $\sum_{i \in V} \pi_i x_i \leq \pi_0$, thus showing (by Lemma 1) that $\sum_{i \in V} \pi_i x_i \leq \pi_0$ cannot induce a facet, a contradiction. The rest of the proof is devoted to showing that inequalities (7) and (8) are indeed valid when $\pi_u < \pi_0$.

Consider $D \subseteq V$ such that $G[D]$ is connected. There are two cases. In the first case, $|D \cap \{u, v\}| \leq 1$. Then, since $\pi_i \leq 0$ for any $i \in R \subseteq V \setminus \{u, v\}$ and $\pi_i - \epsilon < \pi_i + \epsilon < 0$ for any $i \in S'$,

$$(\pi_u + \epsilon)x_u^D + \pi_v x_v^D + \sum_{i \in S'}(\pi_i - \epsilon)x_i^D + \sum_{i \in R} \pi_i x_i^D$$

$$\leq (\pi_u + \epsilon)x_u^D + \pi_v x_v^D$$

$$\leq \max\{\pi_u + \epsilon, \pi_v\} \leq \pi_0.$$

12

The same logic shows that inequality (8) is valid when $|D \cap \{u, v\}| \leq 1$.

In the second case, $|D \cap \{u, v\}| = 2$. Since S' is a u, v-separator and both u and v belong to D, there exists $w \in D \cap S'$. Then, since $\pi_i \leq 0$ for any $i \in R \subseteq V \setminus \{u, v\}$, we have

$$
(\pi_u + \epsilon)x_u^D + \pi_v x_v^D + \sum_{i \in S'}(\pi_i - \epsilon)x_i^D + \sum_{i \in R}\pi_i x_i^D
$$

$$
\leq (\pi_u + \epsilon)x_u^D + \pi_v x_v^D + (\pi_w - \epsilon)x_w^D + \sum_{i \in S' \setminus \{w\}}\pi_i x_i^D + \sum_{i \in R}\pi_i x_i^D
$$

$$
= \pi_u x_u^D + \pi_v x_v^D + \sum_{i \in V \setminus \{u,v\}}\pi_i x_i^D
$$

$$
= \sum_{i \in V}\pi_i x_i^D \leq \pi_0.
$$

Thus, inequality (7) is valid when $|D \cap \{u, v\}| = 2$.

Finally, we show that inequality (8) is valid when $|D \cap \{u, v\}| = 2$. Since u and v belong to D and $G[D]$ is connected, there is a path from u to v in $G[D]$. Moreover, at least one of these u-v paths crosses only one vertex, say w, from $S' \cap D$. This holds by minimality of S'. Let P be the set of vertices in this particular u-v path. Then, since $\pi_i \leq 0$ for any $i \in R \subseteq V \setminus \{u, v\}$, and $\pi_i + \epsilon < 0$ for any $i \in S'$, we have

$$
(\pi_u - \epsilon)x_u^D + \pi_v x_v^D + \sum_{i \in S'}(\pi_i + \epsilon)x_i^D + \sum_{i \in R}\pi_i x_i^D
$$

$$
= (\pi_u - \epsilon)x_u^D + \pi_v x_v^D + (\pi_w + \epsilon)x_w^D + \sum_{i \in S' \setminus \{w\}}(\pi_i + \epsilon)x_i^D + \sum_{i \in R}\pi_i x_i^D
$$

$$
\leq (\pi_u - \epsilon)x_u^D + \pi_v x_v^D + (\pi_w + \epsilon)x_w^D + \sum_{i \in R \cap P}\pi_i x_i^D
$$

$$
= \pi_u x_u^P + \pi_v x_v^P + \pi_w x_w^P + \sum_{i \in R \cap P}\pi_i x_i^P
$$

$$
= \sum_{i \in V}\pi_i x_i^P \leq \pi_0.
$$

$\square$

Lemma 9. *If a facet-defining inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ of $\mathscr{P}(G)$ has exactly two positive coefficients, then it is a separator inequality.*

Proof. Let the positive coefficients be π_a and π_b. By Lemma 8, $\pi_a = \pi_b = \pi_0$. Define

$$
C = \{i \in V \mid \pi_i = -\pi_0\}
$$
$$
S = \{i \in V \mid -\pi_0 < \pi_i < 0\}
$$
$$
R = \{i \in V \mid \pi_i < -\pi_0\}.
$$

We claim that $R = \emptyset$. If not, there is a vertex $v \in R$, and no subset S of vertices that contains v can have weight π_0. Thus all points on the face where the inequality holds at equality also belong to the face where $-x_v \leq 0$, but this cannot happen.

Thus, we can write the facet-defining inequality as

$$\pi_0 x_a + \pi_0 x_b - \sum_{i \in C} \pi_0 x_i + \sum_{i \in S} \pi_i x_i \leq \pi_0. \tag{9}$$

Now see that $C \cup S$ must be an a, b-separator. If not, then there is a path P from a to b in $G[V \setminus (C \cup S)]$, yielding the contradiction that

$$2\pi_0 = \pi_a + \pi_b = \sum_{i \in V} \pi_i x_i^P \leq \pi_0.$$

If $S = \emptyset$, then inequality (9) is an a, b-separator inequality, as desired. So suppose that $S \neq \emptyset$ and consider the following subsets of vertices.

$$A = \{v \in V \mid v \text{ and } a \text{ belong to the same component of } G[V \setminus (C \cup S)]\}$$
$$B = \{v \in V \mid v \text{ and } b \text{ belong to the same component of } G[V \setminus (C \cup S)]\}$$
$$S_A = \{s \in S \mid N(s) \cap A \neq \emptyset\}$$
$$S_B = \{s \in S \mid N(s) \cap B \neq \emptyset\}.$$

We argue that $S_A \cap S_B = \emptyset$. Otherwise, for any vertex $v \in S_A \cap S_B$, the set $D := A \cup B \cup \{v\}$ is connected, so

$$2\pi_0 + \pi_v = \sum_{i \in V} \pi_i x_i^D \leq \pi_0.$$

This implies that $\pi_v \leq -\pi_0$, which contradicts that $v \in S$. Thus, the three sets S_A, S_B, and $S \setminus (S_A \cup S_B)$ partition S.

We claim that $S_A \cup S_B \neq \emptyset$. For contradiction purposes, suppose that $S_A = S_B = \emptyset$. Then C is an a, b-separator, so $\pi_0 x_a + \pi_0 x_b - \sum_{i \in C} \pi_0 x_i \leq \pi_0$ is valid, and, for $i \in S$, the inequality $\pi_i x_i \leq 0$ is valid. Then, by Lemma 1, inequality (9) cannot induce a facet. Thus $S_A \cup S_B \neq \emptyset$.

Now, choose an $\epsilon > 0$ such that $\pi_i + \epsilon \leq 0$ for each $i \in S_A \cup S_B$. We show that inequality (10) below is valid; the proof for inequality (11) is similar.

$$\sum_{i \in V \setminus (S_A \cup S_B)} \pi_i x_i + \sum_{i \in S_A} (\pi_i + \epsilon) x_i + \sum_{i \in S_B} (\pi_i - \epsilon) x_i \leq \pi_0 \tag{10}$$

$$\sum_{i \in V \setminus (S_A \cup S_B)} \pi_i x_i + \sum_{i \in S_A} (\pi_i - \epsilon) x_i + \sum_{i \in S_B} (\pi_i + \epsilon) x_i \leq \pi_0. \tag{11}$$

Suppose that $D \subseteq V$ induces a connected subgraph. If $a \notin D$ or $b \notin D$, then inequality (10) holds, so suppose $a, b \in D$. Now, if $D \cap C \neq \emptyset$, then

$$\sum_{i \in V \setminus (S_A \cup S_B)} \pi_i x_i^D + \sum_{i \in S_A} (\pi_i + \epsilon) x_i^D + \sum_{i \in S_B} (\pi_i - \epsilon) x_i^D$$
$$\leq \pi_0 x_a^D + \pi_0 x_b^D - \sum_{i \in C} \pi_0 x_i^D \leq \pi_0.$$

14

Now suppose $D \cap C = \emptyset$. Consider a shortest path from a to b in $G[D]$ measured in terms of the number of vertices used from $S \cap D$. Let P be the vertices along this path. Note that $|P \cap S_A| = |P \cap S_B| = 1$, so

$$\sum_{i \in V \setminus (S_A \cup S_B)} \pi_i x_i^D + \sum_{i \in S_A} (\pi_i + \epsilon) x_i^D + \sum_{i \in S_B} (\pi_i - \epsilon) x_i^D$$

$$\leq \sum_{i \in V \setminus (S_A \cup S_B)} \pi_i x_i^P + \sum_{i \in S_A} (\pi_i + \epsilon) x_i^P + \sum_{i \in S_B} (\pi_i - \epsilon) x_i^P$$

$$= \sum_{i \in V} \pi_i x_i^P \leq \pi_0.$$

So, in both cases, inequality (10) is valid.

Thus inequalities (10) and (11) are valid. But, by Lemma 1, this contradicts that inequality (9) induces a facet. So, $S = \emptyset$, and inequality (9) is (a scalar multiple of) an a, b-separator inequality. □

Lemma 10. *If $\alpha(G) \leq 2$, then $\mathscr{P}(G) = Q(G)$.*

Proof. Consider an arbitrary facet-defining inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ of $\mathscr{P}(G)$. Let $S = \{i \in V \mid \pi_i > 0\}$. By Proposition 2, $|S| \leq \alpha(G) \leq 2$. Consider the following three cases. In each case, we show that the inequality (or a scalar multiple thereof) is already in the description of $Q(G)$.

In the first case, suppose $|S| = 0$. Recall that $\pi_0 \geq 0$ by Lemma 2. Since no variable has a positive coefficient, π_0 cannot be positive, since otherwise no point in $\mathscr{P}(G)$ could satisfy the inequality at equality. Thus $\pi_0 = 0$. Then, by Lemma 2, the inequality is a (scalar multiple of a) nonnegativity bound.

In the second case, $|S| = 1$, and suppose $S = \{j\}$. Then, $\pi_0 \geq \pi_j > 0$, since $G[\{j\}]$ is connected. Further, $\pi_0 = \pi_j$, since otherwise no point in $\mathscr{P}(G)$ satisfies the inequality at equality. Now, the inequality $\pi_j x_j \leq \pi_0$ is valid, and $0 x_j + \sum_{i \in V \setminus \{j\}} \pi_i x_i \leq 0$ is valid since $\pi_i \leq 0$ for every $i \in V \setminus \{j\}$. If $\pi_i = 0$ for every $i \in V \setminus \{j\}$, then $\sum_{i \in V} \pi_i x_i \leq \pi_0$ is a scalar multiple of $x_j \leq 1$, as desired. Otherwise, there is vertex $k \in V \setminus \{j\}$ with $\pi_k < 0$. Then the inequality $0 x_j + \sum_{i \in V \setminus \{j\}} \pi_i x_i \leq 0$ discussed previously is not the $0 x \leq 0$ inequality, and it, along with $\pi_j x_j \leq \pi_j$ implies $\sum_{i \in V} \pi_i x_i \leq \pi_0$, so by Lemma 1, the inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ cannot induce a facet, a contradiction.

In the third and final case, $|S| = 2$. Then, by Lemma 9, the facet-defining inequality is a separator inequality.

Thus, in every case, the facet-defining inequality $\sum_{i \in V} \pi_i x_i \leq \pi_0$ of $\mathscr{P}(G)$ is already a part of the description of $Q(G)$. Thus $Q(G) \subseteq \mathscr{P}(G)$. The reverse inclusion holds since $Q(G)$ is a relaxation for $\mathscr{P}(G)$, so $\mathscr{P}(G) = Q(G)$. □

4   Indegree Inequalities

In this section, we study the indegree inequalities. For a graph $G = (V, E)$, a vector $d \in \mathbb{R}^n$ is said to be an *indegree vector* if for some orientation $D = (V, A)$ of G the indegree of each

vertex v is d_v. For each indegree vector d of G, there is a corresponding *indegree inequality*.

$$\text{(indegree inequality)} \qquad \sum_{i \in V}(1 - d_i)x_i \leq 1. \qquad (12)$$

The indegree inequalities are interesting because of the following theorem.

Theorem 4 (Theorem 3.6 on p. 168 of [17])**.** *If $G = (V, E)$ is a tree, then the following equality holds.*

$$\mathscr{P}(G) = \left\{ x \in \mathbb{R}^n_+ \mid x \text{ satisfies all indegree inequalities} \right\}. \qquad (13)$$

Moreover, each of the indegree inequalities induces a facet when G is a tree.

Lemma 11. *The indegree inequalities are valid for $\mathscr{P}(G)$ for arbitrary G.*

Proof. The proof for arbitrary G is the same as for trees [17]. Suppose that $S \subseteq V$ induces a connected subgraph. This implies that the number of edges with both endpoints in S is at least $|S| - 1$. Hence, for any indegree vector d, we have $\sum_{i \in S} d_i \geq |S| - 1$, implying that

$$\sum_{i \in V}(1 - d_i)x_i^S = |S| - \sum_{i \in S} d_i \leq |S| + (1 - |S|) = 1.$$

$\square$

We note that the indegree inequalities can be facet-defining, even when the graph G has undirected cycles. For example, consider the 4-vertex cycle graph

$$C_4 = ([4], \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{1, 4\}\}) \,.$$

If we orient the edges away from vertices 1 and 3, we get the facet-defining indegree inequality $x_1 - x_2 + x_3 - x_4 \leq 1$. However, if we orient the edges into a directed cycle, then we get the inequality $0x_1 + 0x_2 + 0x_3 + 0x_4 \leq 1$ which is not facet-defining. If we orient the edges away from vertex 1 and towards vertex 3, then we get the inequality $x_1 + 0x_2 - x_3 + 0x_4 \leq 1$, which is also not facet-defining. This is illustrated in Figure 2. Later we will provide the


Figure 2: Orientations of C_4. The leftmost orientation leads to a facet; the others do not.

exact conditions for an indegree inequality to induce a facet. We also show that the separation problem for these inequalities is polynomial-time solvable. In fact, we can find a most-violated inequality in linear time.

Consider the polyhedron $Q'(G)$ defined by the indegree inequalities and nonnegativity bounds. It is a relaxation for $\mathscr{P}(G)$ by Lemma 11.

$$Q'(G) := \left\{ x \in \mathbb{R}^n_+ \mid x \text{ satisfies all indegree inequalities} \right\}$$

We provide a linear-size extended formulation $F'(G)$ for $Q'(G)$, so we need not rely on the ellipsoid method to optimize over it. We also show that the equality $\mathscr{P}(G) = Q'(G)$ holds if and only if G is a forest.

16

4.1 Indegree facets

In the following, a subset of vertices is said to be *tight* for an inequality if its characteristic vector satisfies it at equality.

Lemma 12. *Suppose that $S \subseteq V$ induces a connected subgraph of G. Then, S is tight for an indegree inequality if and only if*

1. *S induces a tree in G; and*

2. *each edge of E having exactly one endpoint in S is oriented out of S.*

Proof. ($\Longleftarrow$) Since S induces a tree in G, the number of edges with both endpoints in S is $|S|-1$. Hence these edges contribute $|S|-1$ to the sum $\sum_{i \in S} d_i$. The edges with one endpoint in S do not contribute to this sum, as they are oriented away from S. Thus $\sum_{i \in S} d_i = |S|-1$, meaning that

$$\sum_{i \in V}(1-d_i)x_i^S = \sum_{i \in S}(1-d_i) = |S| - (|S|-1) = 1.$$

($\Longrightarrow$) It is straightforward to show that if $G[S]$ is not a tree, or if one of the edges is directed towards S, then the quantity $\sum_{i \in V}(1-d_i)x_i^S$ is at most zero, in which case S cannot be tight. $\qquad\square$

Lemma 13. *If there are two directed s-t walks in the orientation $D = (V, A)$ of G, then the corresponding indegree inequality does not induce a facet of $\mathscr{P}(G)$.*

Proof. For the indegree inequality to induce a facet, there must be a tight set S, containing t, such that $G[S]$ is connected. We argue that the vertices along the two *s-t* walks must belong to S. Suppose not, then not all edges with one endpoint in S point away from S, so, by Lemma 12, S is not a tight set, a contradiction. Thus all vertices along these two *s-t* walks belong to S. However, this shows that S does not induce a tree in G, which, by Lemma 12, contradicts that S is a tight set. Thus no such S exists, and the indegree inequality cannot induce a facet. $\qquad\square$

As a consequence of Lemma 13, if D has a directed cycle, then the corresponding indegree inequality cannot induce a facet.

Theorem 5. *The indegree inequality corresponding to an orientation $D = (V, A)$ of G induces a facet of $\mathscr{P}(G)$ if and only if for every $u, v \in V$ there is at most one directed u-v walk in D.*

Proof. The 'only if' direction of the proof follows by Lemma 13. To show the 'if' direction, consider an orientation $D = (V, A)$ of G satisfying the assumptions. In this case, D is acyclic, so there exists a topological ordering of its vertices. We prove that the inequality is facet-defining by lifting in the vertices according to the topological ordering. We will start with the seed inequality $x_w \leq 1$, where w is the first vertex in the topological ordering. We claim that each time a vertex v is lifted into the inequality, the lifting problem objective ζ_v is d_v, meaning that $\pi_v = \pi_0 - d_v = 1 - d_v$, in which case the resulting facet-defining inequality is the desired indegree inequality. For each vertex v, let

$$D_v := \{v\} \cup \{u \in V \mid \text{there is a directed } u\text{-}v \text{ path in } D\}.$$

We use induction on the position of the vertex in the topological ordering. If v is first in the topological ordering, we have our seed inequality $x_v \leq 1$. Now suppose that v is not first in the topological ordering. Note that each $u \in D_v \setminus \{v\}$ is earlier than v in the topological ordering and $G[D_v]$ is connected, so D_v is feasible when lifting in v. We argue that $G[D_v]$ is a tree. If it is not the case, there is an undirected cycle subgraph (V', E') of $G[D_v]$. Let a and b be the first and the last vertices of V', respectively, in the topological ordering. Then there are two directed a-b paths in the orientation of (V', E'), hence there are at least two directed a-b paths in D, which contradicts the assumption. Now, since $G[D_v]$ is a tree and by the induction assumption that $\pi_u = \pi_0 - \zeta_u = 1 - d_u$ for vertices u prior to v in the topological ordering,

$$\sum_{u \in D_v \setminus \{v\}} \pi_u = \sum_{u \in D_v \setminus \{v\}} (1 - d_u) = (|D_v| - 1) - (|E(G[D_v])| - d_v) = d_v.$$

So $\zeta_v \geq d_v$. Meanwhile, by Lemma 4, $\zeta_v \leq d_v$, so $\zeta_v = d_v$ as desired. $\square$

It is natural to ask—Given a graph, does there exist a facet-defining indegree inequality? In other words, does the graph admit an orientation satisfying the conditions of Theorem 5? We will call this MULTITREE ORIENTATION.

Problem: MULTITREE ORIENTATION.
Input: a simple graph $G = (V, E)$.
Question: Is there an orientation of G in which, for every $s, t \in V$, there is at most one (directed) s-t walk?

We have been unable to find a "good characterization" for when there is a facet-defining indegree inequality. The following theorem explains why.

Theorem 6 (Eppstein [12]). MULTITREE ORIENTATION *is NP-complete.*

Corollary 4. *Given a graph $G = (V, E)$, it is NP-complete to determine whether there is a facet-defining indegree inequality of $\mathscr{P}(G)$.*

Proof. Follows immediately from Theorems 5 and 6. $\square$

In contrast, the same problem for the separator inequalities is easy.

4.2 Lifting indegree inequalities is NP-hard

One may wonder how difficult it is to generate a facet-defining inequality for $\mathscr{P}(G)$ via a specified lifting order. We show that this problem is hard even when the seed inequality is a facet-defining indegree inequality.

Theorem 7. *Lifting a vertex into an indegree inequality is strongly NP-hard. This holds even when the graph is bipartite and 2-degenerate.*

Proof. The reduction is from 3OCC-3SAT, a special case of 3SAT in which each variable appears at most three times and each literal appears at most twice. This is NP-complete (cf. Theorem 16.5 of [22]). Let the instance $\Phi = \bigwedge_{j=1}^{m}(c_j^1 \vee c_j^2 \vee c_j^3)$ of 3OCC-3SAT be defined over

18

Figure 3: Variable gadget (left) and clause gadget (right)

variables $x_1, \ldots, x_n$. We construct a graph G and a (facet-defining) indegree inequality for $\mathscr{P}(G-v)$ for which the lifting problem for v has objective $2n+m$ if and only if Φ is satisfiable.

For each variable x_i and for each clause c_j in the 3OCC-3SAT instance, construct a gadget, as shown in Figure 3. Connect the gadgets as follows. Connect each literal $x_i(\overline{x}_i)$ from a clause gadget (denoted by, say, c_j^1 in Figure 3) to either a literal x_i^1 or to x_i^2 ($\overline{x}_i^1$ or $\overline{x}_i^2$) from the corresponding variable gadget. Because each literal appears in at most two clauses, we can suppose that no pair of clause vertices are connected to the same variable gadget literal. This is illustrated in Figure 4. Finally, add a new vertex v and connect it to every clause vertex of the type c_j^k and to all vertices of the type y_i^1, y_i^2, $\overline{y}_i^1$, and $\overline{y}_i^2$. Since the number of vertices of G is only $12n + 6m + 1$ the reduction is polynomial.

First see that G is bipartite, with partitions A and B:

$$A = \{v\} \cup \left(\bigcup_{i=1}^{n} \{l_i, r_i, b_i, t_i, x_i^1, x_i^2, \overline{x}_i^1, \overline{x}_i^2\} \right) \cup \left(\bigcup_{j=1}^{m} \{d_j^1, d_j^2, d_j^3\} \right);$$

$$B = \left(\bigcup_{i=1}^{n} \{y_i^1, y_i^2, \overline{y}_i^1, \overline{y}_i^2\} \right) \cup \left(\bigcup_{j=1}^{m} \{c_j^1, c_j^2, c_j^3\} \right).$$

Now we show G is 2-degenerate. Suppose not; then there is a subgraph H of G in which all vertices have degree at least three. Then H cannot contain a vertex of the type d_j^k, l_i, r_i, t_i, or b_i, since these vertices have degree at most two in G. Now, if those vertices do not belong to H, then H cannot contain a vertex of the type $y_i^1, y_i^2, \overline{y}_i^1, \overline{y}_i^2$, or c_j^k. This implies that $V(H) \subseteq \{v\} \cup \left(\bigcup_{i=1}^{n} \{x_i^1, x_i^2, \overline{x}_i^1, \overline{x}_i^2\} \right)$, meaning that $V(H)$ is independent, but this contradicts that H has minimum degree at least three.

Consider the indegree inequality (14) for $G - v$ that is obtained by orienting the edges from $A \setminus \{v\}$ to B. It induces a facet of $\mathscr{P}(G - v)$ by Theorem 5.

$$\sum_{i \in A \setminus \{v\}} x_i - \sum_{i \in B} 2x_i \leq 1. \tag{14}$$

19

Figure 4: The construction of the graph $G - v$ when given the 3OCC-3SAT instance $\Phi = (x_1 \vee x_2 \vee x_3) \wedge (\overline{x}_1 \vee \overline{x}_2 \vee x_4) \wedge (x_1 \vee x_2 \vee \overline{x}_4)$.

Now, consider the problem of lifting v into inequality (14), i.e., solving for

$$
\zeta := \max_{S \subseteq V} \left\{ \sum_{i \in A \setminus \{v\}} x_i^S - \sum_{i \in B} 2x_i^S \ \middle| \ x_v^S = 1 \text{ and } G[S] \text{ is connected} \right\}.
$$

Claim 1: There is an optimal solution $D \subseteq V$ to the lifting problem so that

- for each i, either $\{y_i^1, y_i^2\} \subseteq D$ or $\{\overline{y}_i^1, \overline{y}_i^2\} \subseteq D$, but not both; and

- for each j, exactly one of c_j^1, c_j^2, and c_j^3 belongs to D.

If an optimal solution $D \subseteq V$ to the lifting problem does not fit these criteria, we modify it so that it does. Recognize that $v \in D$ and consider these six cases.

1. Three or more of $\{y_i^1, y_i^2, \overline{y}_i^1, \overline{y}_i^2\}$ belong to D. Without loss of generality, suppose that $\{y_i^2, \overline{y}_i^1, \overline{y}_i^2\} \subseteq D$, thus we can assume that $x_i^2 \in D$. Then $D' = D \setminus \{y_i^2, x_i^2\}$ is connected, contains v, and has a larger weight than D, a contradiction.

2. Two of $\{y_i^1, y_i^2, \overline{y}_i^1, \overline{y}_i^2\}$ belong to D. If either $\{y_i^1, y_i^2\} \subseteq D$ or $\{\overline{y}_i^1, \overline{y}_i^2\} \subseteq D$, then Claim 1 is satisfied. Otherwise, without loss of generality, suppose that y_i^1 and $\overline{y}_i^1$ belong to D. Then t_i cannot belong to D by connectivity. We can assume that x_i^1 belongs to D. Now, $D' = D \cup \{\overline{y}_i^2, t_i\} \setminus \{x_i^1, y_i^1\}$ is connected, contains v, and has the same weight.

3. One of $\{y_i^1, y_i^2, \overline{y}_i^1, \overline{y}_i^2\}$ belongs to D. Without loss of generality, suppose that y_i^1 belongs to D. Then $D' = D \cup \{y_i^2, t_i, r_i\}$ is connected, contains v, and has the same weight.

20

4. None of $\{y_i^1, y_i^2, \bar{y}_i^1, \bar{y}_i^2\}$ belong to D. Then $D' = D \cup \{y_i^1, y_i^2, t_i, b_i, l_i, r_i\}$ is connected, contains v, and has the same weight.

5. Two or more of $\{c_j^1, c_j^2, c_j^3\}$ belong to D. Without loss of generality, suppose that $c_j^1, c_j^2 \in D$. We can assume that $d_j^1, d_j^2, d_j^3 \in D$, and that c_j^1 has a neighbor, say $w \in A$, from a variable gadget and w also belongs to D. Then $D' = D \setminus \{c_j^1, d_j^2, w\}$ is connected, contains v, and has the same weight.

6. None of $\{c_j^1, c_j^2, c_j^3\}$ belong to D. Then $D' = D \cup \{c_j^1, d_j^1, d_j^2\}$ is connected, contains v, and has the same weight.

These steps can be applied repeatedly until D satisfies the claim.

Claim 2: $\zeta \leq 2n + m$. Consider an optimal solution $D \subseteq V$ to the lifting problem that satisfies Claim 1. See that any weight $+1$ vertex in D must have a weight -2 neighbor in D. There are $2n+m$ vertices of weight -2 in D and each has three weight $+1$ neighbors in G. So,

$$
\begin{aligned}
\zeta &= \sum_{i \in A \setminus \{v\}} x_i^D - \sum_{i \in B} 2x_i^D \\
&= \sum_{i \in A \setminus \{v\}} x_i^D - 2(2n + m) \\
&\leq 3(2n + m) - 2(2n + m) = 2n + m.
\end{aligned}
$$

Claim 3: If Φ is satisfiable, then $\zeta \geq 2n + m$. Given a satisfying assignment x^* for Φ, construct a solution D to the lifting problem as follows.

- For each i: if $x_i^* = 1$, choose $\bar{y}_i^1$ and $\bar{y}_i^2$; otherwise, select y_i^1 and y_i^2. Note that this is, in a sense, the opposite of the satisfying assignment.

- For each j: the satisfying assignment makes clause j evaluate to true by some literal, say c_j^k; choose vertex c_j^k and also the neighboring vertex from the variable gadget.

- Add v and all positive-weight vertices that neighbor a previously chosen vertex.

This solution D is feasible, since all negative-weight vertices are adjacent to v, and their positive-weight neighbors were chosen. One negative-weight vertex was chosen from each clause gadget, and two negative-weight vertices were selected from each variable gadget. So, there are $2n + m$ vertices of negative weight in D. Each of these negative-weight vertices has three positive-weight neighbors. All that remains is to demonstrate that no two negative-weight vertices of D share a neighbor of positive weight. The proof of this is straightforward but tedious, so we omit it. Thus D has weight $(2n + m)(-2 + 3(1)) = 2n + m$.

Claim 4: If $\zeta \geq 2n + m$, then Φ is satisfiable. Consider an optimal solution $D \subseteq V$ that satisfies Claim 1 and has weight at least $2n + m$. Then,

- for each i, either $\{y_i^1, y_i^2\} \subseteq D$ or $\{\bar{y}_i^1, \bar{y}_i^2\} \subseteq D$, but not both; and
- for each j, exactly one of c_j^1, c_j^2, and c_j^3 belongs to D.

The following assignment x^* will be shown to satisfy Φ. For each i: if $\{y_i^1, y_i^2\} \subseteq D$, then set $x_i^* = 0$; otherwise, set $x_i^* = 1$. Then $\zeta = 2n + m$ by Claim 2, and this equality holds if and only if no two negative-weight vertices in D have a common neighbor (of positive weight).

We argue that, for each j, x^* makes clause j evaluate to true. Let c_j^k be the vertex from clause j that belongs to D. Suppose that the neighbor of c_j^k from the variable gadget is

- x_i^δ for some $\delta \in \{1, 2\}$. Then, y_i^δ does not belong to D, so $\bar{y}_i^1$ and $\bar{y}_i^2$ belong to D and thus $x_i^* = 1$, which satisfies clause j.

- $\bar{x}_i^\delta$ for some $\delta \in \{1, 2\}$. Then, $\bar{y}_i^\delta$ does not belong to D, so y_i^1 and y_i^2 belong to D and thus $x_i^* = 0$, which satisfies clause j.

So x^* is a satisfying assignment.

By Claims 2, 3, and 4, the lifting problem for v has objective $2n + m$ if and only Φ is satisfiable. Then, since 3OCC-3SAT is NP-hard and since the reduction is polynomial, the problem of lifting vertex v into the indegree inequality (14) is NP-hard. $\qquad \square$

4.3   Separating indegree inequalities in linear time

Given $x^* \in \mathbb{R}^n$, consider the following separation algorithm for the indegree inequalities.

1. For each edge $\{u, v\} \in E$ do

    - If $x_u^* > x_v^*$ then orient edge $\{u, v\}$ as (u, v).
    - Else orient it as (v, u).

2. Let d be the indegree vector obtained from orientation in step 1.

3. If $\sum_{i \in V}(1 - d_i)x_i^* > 1$ then return the inequality $\sum_{i \in V}(1 - d_i)x_i \leq 1$.

4. Else certify that x^* satisfies all indegree inequalities.

Theorem 8. *The above algorithm solves the separation problem for the indegree inequalities in linear time. In fact, it finds a most-violated inequality.*

Proof. Clearly the algorithm runs in time $O(n + m)$, so we must only prove correctness. Let $\mathbf{d}$ denote the set of all indegree vectors. Consider the set of all orientations of G satisfying the property that if $\{u, v\} \in E$ and $x_u^* > x_v^*$ then edge $\{u, v\}$ is oriented as (u, v). Let $\mathbf{d}^*$ denote the set of corresponding indegree vectors. The indegree vector returned by the algorithm belongs to $\mathbf{d}^*$, and we argue that any indegree vector from $\mathbf{d}^*$ corresponds to a most-violated inequality (if any exist). For an indegree vector $d \in \mathbf{d}$, define

$$f(d) := \sum_{i \in V}(1 - d_i)x_i^*.$$

The problem of finding a most-violated inequality is that of finding $d \in \mathbf{d}$ that maximizes $f(d)$. Consider $d' \in \mathbf{d} \setminus \mathbf{d}^*$. By definition of $\mathbf{d}^*$ this means the orientation corresponding to d' has an oriented edge (v, u) with $x_u^* > x_v^*$. Consider the indegree vector d'' obtained by flipping this edge's orientation to (u, v). Then, $d_v'' = d_v' + 1$ and $d_u'' = d_u' - 1$, so

$$f(d') = -x_u^* + x_v^* + f(d'') < f(d'').$$

So, d' is suboptimal. It is easy to see that any two indegree vectors $d^1, d^2 \in \mathbf{d}^*$ satisfy $f(d^1) = f(d^2)$. Thus, any indegree vector from $\mathbf{d}^*$ is optimal. $\qquad \square$

4.4 A linear-size extended formulation for the indegree-based relaxation

We propose an extended formulation for the indegree-based relaxation. Denote by $F'(G)$ the set of $(x, y) \in \mathbb{R}^{n+m}$ satisfying

$$y_e - x_v \leq 0 \quad \text{and} \quad y_e - x_u \leq 0, \ \forall e = \{u, v\} \in E \tag{15}$$

$$\sum_{i \in V} x_i - \sum_{e \in E} y_e \leq 1 \tag{16}$$

$$x_i \geq 0, \ \forall i \in V. \tag{17}$$

Theorem 9. *The polyhedron $Q'(G)$ defined by the indegree inequalities and nonnegativity bounds admits a size $O(m + n)$ extended formulation.*

Proof. The polyhedron $F'(G)$ clearly has size $O(m+n)$, so we must only show $\text{proj}_x(F'(G)) = Q'(G)$.

To show $Q'(G) \subseteq \text{proj}_x(F'(G))$, consider $x \in Q'(G)$. For each edge $e = \{u, v\} \in E$, let $y_e = \min\{x_u, x_v\}$. We claim that $(x, y) \in F'(G)$. Clearly, (x, y) satisfies the constraints (15) and (17). Consider the indegree vector d obtained by orienting the edges as in the separation algorithm. In this case, $\sum_{\{u,v\} \in E} \min\{x_u, x_v\} = \sum_{i \in V} d_i x_i$. So,

$$\sum_{e \in E} y_e = \sum_{\{u,v\} \in E} \min\{x_u, x_v\} = \sum_{i \in V} d_i x_i \geq \sum_{i \in V} x_i - 1,$$

where the last inequality holds since x satisfies all indegree inequalities. This shows that (x, y) satisfies constraint (16), and thus $(x, y) \in F'(G)$.

To show $\text{proj}_x(F'(G)) \subseteq Q'(G)$, consider $(x, y) \in F'(G)$. Construct an alternative point (x, y') where, for each edge $e = \{u, v\}$, we have $y'_e = \min\{x_u, x_v\}$. Clearly $(x, y') \in F'(G)$ as well. To show $x \in Q'(G)$, consider an arbitrary indegree vector d, and see that

$$\sum_{i \in V} (1 - d_i) x_i = \sum_{i \in V} x_i - \sum_{i \in V} d_i x_i$$

$$\leq \sum_{i \in V} x_i - \sum_{\{u,v\} \in E} \min\{x_u, x_v\}$$

$$= \sum_{i \in V} x_i - \sum_{e \in E} y'_e \leq 1.$$

The inequality holds because $\sum_{\{u,v\} \in E} \min\{x_u, x_v\} \leq \sum_{i \in V} d_i x_i$ for any indegree vector d. Thus, x satisfies all indegree inequalities and is nonnegative by (17), so $x \in Q'(G)$. $\square$

4.5 When the indegree-based relaxation is tight

Theorem 10. *The equality $\mathscr{P}(G) = Q'(G)$ holds if and only if G is a forest.*

Proof. The 'if' direction follows from Theorems 1 and 4. To prove the 'only if' direction, suppose that G is not a forest. Then G has an undirected cycle subgraph (V', E'). Let $v \in V'$ be one of the cycle's vertices. Pick a subset $U \subset V$ of vertices, containing v, that contains exactly one vertex from each component of G. Then, by Corollary 1, the inequality

$\sum_{i \in U} x_i \leq 1$ induces a facet of $\mathscr{P}(G)$. However, we argue that this inequality is not an indegree inequality. If it were, then there is an indegree vector d satisfying $d_v = 0$ and $d_i = 1$ for each $i \in V' \setminus \{v\}$, hence

$$\sum_{i \in V'} d_i = d_v + \sum_{i \in V' \setminus \{v\}} d_i = 0 + (|V'| - 1).$$

However, the oriented edges from E' contribute $|E'|$ indegrees among V', i.e., $\sum_{i \in V'} d_i \geq |E'| = |V'|$. Each of $\mathscr{P}(G)$ and $Q'(G)$ has a unique half-space representation (up to scalar multiples) since they are full-dimensional, but the facets of $Q'(G)$ and $\mathscr{P}(G)$ are different, so $\mathscr{P}(G) \neq Q'(G)$. $\qquad\square$

5 The Relative Strength of Indegree and Separator Inequalities

In this section, we study the relative strength of the LP relaxations given by the indegree and/or separator inequalities. We provide some theoretical insights, as well as some numerical experiments.

We remark that the indegree-based and separator-based LP relaxations are incomparable in general. For example, the indegree inequalities and nonnegativity bounds fully describe the connected subgraph polytope of the claw graph $K_{1,3}$, while the separator-based LP relaxation does not. On the other hand, the separator inequalities and 0-1 bounds fully describe the connected subgraph polytope of the diamond graph $K_4 - e$, but the indegree-based LP relaxation does not.

We note that the indegree inequalities can be useful, even for dense graphs. For an extreme example, consider the class of (dense) MWCS instances where the input is the complete bipartite graph $K_{q,q}$ and the "left" vertices A are weighted 1 and the "right" vertices B are weighted $1 - q$. The separator-based LP relaxation gives a bound no less than $q/2$, as evidenced by the vector $x^* \in Q(G)$ that has $x_i^* = \frac{1}{2}$ for $i \in A$, and $x_i^* = 0$ for $i \in B$. However, the addition of the single indegree inequality $\sum_{i \in A} x_i + \sum_{i \in B}(1 - q)x_i \leq 1$ proves that 1 is the optimal MWCS objective. (This indegree inequality is obtained by orienting the edges from A to B.)

To better understand the relative strength of these inequalities with respect to the graph's density, we perform some numerical experiments. Specifically, we compare their LP objectives on random instances of the maximum-weight connected subgraph (MWCS) problem. In both cases, we include the trivial 0-1 bounds. For our tests, we consider $G(n, p)$ random graphs with uniform random weights, as well as a class of graphs similar to the $K_{q,q}$ example from above. An anonymous referee remarked that this second class exhibits a problem structure similar to the quasi-bipartite instances studied for Steiner tree [24].

For ease of implementation, we optimize over the associated extended formulations. Thus, due to the size of the separator-based extended formulation, we are limited to relatively small instances. Although we do not report solution times here, the time to solve the separator-based extended formulation was typically three orders of magnitude longer than the time to solve the indegree-based extended formulation. Of course, using the extended formulation is not the quickest way to optimize over the separator-based LP relaxation in practice.

5.1 Experiments for $G(n, p)$ random graphs

The first set of test instances are random graphs on $n = 50$ vertices in which each edge exists independently with probability $p \in \{0.01, 0.02, \ldots, 0.25\}$ and the vertex weights are (independently) uniformly distributed integers between -50 and $+50$. For each value of p, we construct five instances.

The mean LP objectives over the five instances are provided in Table 1. For comparison purposes, we also provide the LP objective obtained when optimizing over just the 0-1 bounds, the (integer) MWCS objective, and the LP bound obtained by using both indegree inequalities and separator inequalities. A standard preprocessing procedure for the MWCS problem is to contract an edge if its incident vertices both have nonnegative weight. Table 2 provides the same information as Table 1, but after this preprocessing.

As can be seen in Tables 1 and 2, the separator inequalities and indegree inequalities together ("Both") provide strong LP bounds for our test instances. In fact, before preprocessing, the gap is zero for 122 of the 125 instances (and 123 of 125 after preprocessing). As our theoretical results would suggest, the indegree inequalities give strong bounds for very sparse graphs, but weaken as the graphs become denser. Also consistent with our theoretical results, the separator inequalities give poor LP bounds for very sparse graphs, and improve for denser graphs. However, once the graphs become sufficiently dense, the LP relaxation using only the 0-1 bounds is enough. By comparing Tables 1 and 2, one sees that the preprocessing step often makes the graph sparser, thus improving the resulting indegree-based LP bound; however, the separator-based LP bound does not change.

5.2 Experiments for instances similar to the $K_{q,q}$ example

The second set of test instances are also random graphs on $n = 50 = 2q$ vertices, but they are constructed to be similar to the $K_{q,q}$ example from above. They have a large independent set with positive weights, and all other vertices have large negative weights. The instances are constructed by taking half of the vertices and designating them as an independent set; all of the other possible edges exist independently with probability $p \in \{0.1, 0.2, \ldots, 1\}$. As in the $K_{q,q}$ example, the independent set vertices are each given weight 1, and the other nodes are given weight $1 - q = -24$. By this construction, the optimal MWCS objective is 1. Table 3 reports the LP bounds for each value of p (averaged over five instances).

As can be seen in Table 3, the indegree inequalities give optimal LP bounds for very dense instances—even at $p = 0.6$, while the separator-based LP relaxation is always weak. The strength of the indegree-based LP relaxation deteriorates quickly around $p = 0.8$, but it is still stronger than the separator-based bound at $p = 1$. Indeed, it can be proven that this dominance holds over all instances from this class. The separator-based LP objective is at least $\frac{1}{2}q$ regardless of p, while the indegree-based objective is at most $\frac{1}{3}q + \frac{2}{3}$, and this is sharp when $p = 1$.

6 Conclusion

In this paper, we provide foundational knowledge about the connected subgraph polytope $\mathscr{P}(G)$. We study two classes of valid inequalities called separator inequalities and indegree inequalities, which have been previously shown to be computationally useful and theoretically

interesting. We determine when these inequalities induce facets, when they fully describe the connected subgraph polytope, how to separate them in polynomial time, and how to craft polynomial-size extended formulations. We show that it is NP-hard to lift the indegree inequalities, but lifting is easy for separator inequalities.

As we have seen, the separator-based relaxation $Q(G)$ coincides with $\mathscr{P}(G)$ when the graph has no independent set of three vertices, a class of very dense graphs. On the other hand, the indegree inequalities give a perfect description for forests—a class of sparse graphs. This suggests that the indegree inequalities should give strong LP relaxations for very sparse graphs, and the separator inequalities should provide strong LP relaxations for sufficiently dense graphs. These intuitions are corroborated by numerical experiments on random instances of the maximum-weight connected subgraph (MWCS) problem. Moreover, these inequalities, when used together, give optimal LP bounds for 122 out of 125 MWCS test instances on $G(n, p)$ random graphs. This paper focuses on the theoretical aspects of the separator and indegree inequalities. An interesting topic for future research is to determine how these inequalities should best be used for computational purposes.

Acknowledgements

References

[1] E. Álvarez-Miranda, I. Ljubić, and P. Mutzel. The maximum weight connected subgraph problem. In M. Jünger and G. Reinelt, editors, *Facets of Combinatorial Optimization*, pages 245–270. Springer, Berlin, 2013.

[2] E. Álvarez-Miranda, I. Ljubić, and P. Mutzel. The rooted maximum node-weight connected subgraph problem. In C. Gomes and M. Sellmann, editors, *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, volume 7874 of *Lecture Notes in Computer Science*, pages 300–315, Berlin, 2013. Springer.

[3] C. Backes, A. Rurainski, G.W. Klau, O. Müller, D. Stöckel, A. Gerasch, J. Küntzer, D. Maisel, N. Ludwig, M. Hein, et al. An integer linear programming approach for finding deregulated subgraphs in regulatory networks. *Nucleic acids research*, 40(6):e43, 2012.

[4] M. Bailly-Bechet, C. Borgs, A. Braunstein, J. Chayes, A. Dagkessamanskaia, J.-M. François, and R. Zecchina. Finding undetected protein associations in cell signaling by belief propagation. *Proceedings of the National Academy of Sciences*, 108(2):882–887, 2011.

[5] M.D. Biha, H.L.M. Kerivin, and P.H. Ng. Polyhedral study of the connected subgraph problem. *Discrete Mathematics*, 338(1):80–92, 2015.

[6] A. Buchanan, J.S. Sung, S. Butenko, and E.L. Pasiliao. An integer programming approach for fault-tolerant connected dominating sets. *INFORMS Journal on Computing*, 27(1):178–188, 2015.

[7] R. Carvajal, M. Constantino, M. Goycoolea, J.P. Vielma, and A. Weintraub. Imposing connectivity constraints in forest planning models. *Operations Research*, 61(4):824–836, 2013.

[8] C.-Y. Chen and K. Grauman. Efficient activity detection with max-subgraph search. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1274–1281. IEEE, 2012.

[9] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Heidelberg, fourth edition, 2010.

[10] M.T. Dittrich, G.W. Klau, A. Rosenwald, T. Dandekar, and T. Müller. Identifying functional modules in protein–protein interaction networks: an integrated exact approach. *Bioinformatics*, 24(13):i223–i231, 2008.

[11] D.-Z. Du and P.J. Wan. *Connected Dominating Set: Theory and Applications*, volume 77 of *Springer Optimization and Its Applications*. Springer, New York, 2013.

[12] D. Eppstein. When does a graph admit an orientation in which there is at most one s-t walk? Theoretical Computer Science StackExchange. URL: http://cstheory.stackexchange.com/questions/31412/when-does-a-graph-admit-an-orientation-in-which-there-is-at-most-one-s-t-walk (version: 2015-05-07).

[13] M. Fischetti, M. Leitner, I. Ljubić, M. Luipersbeck, M. Monaci, M. Resch, D. Salvagnin, and M. Sinnl. Thinning out Steiner trees: a node-based model for uniform edge costs. *Mathematical Programming Computation*, 2016. To Appear.

[14] R.S. Garfinkel and G.L. Nemhauser. Optimal political districting by implicit enumeration techniques. *Management Science*, 16(8):B–495, 1970.

[15] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, second edition, 1993.

[16] D.S. Johnson. The NP-completeness column: an ongoing guide. *Journal of Algorithms*, 6(1):145–159, 1985.

[17] B. Korte, L. Lovász, and R. Schrader. *Greedoids*, volume 4 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 1991.

[18] I. Ljubić, R. Weiskircher, U. Pferschy, G.W. Klau, P. Mutzel, and M. Fischetti. An algorithmic framework for the exact solution of the prize-collecting Steiner tree problem. *Mathematical Programming*, 105(2-3):427–449, 2006.

[19] A. Lucena and M.G.C. Resende. Strong lower bounds for the prize collecting Steiner problem in graphs. *Discrete Applied Mathematics*, 141(1):277–294, 2004.

[20] J. Moody and D.R. White. Structural cohesion and embeddedness: A hierarchical concept of social groups. *American Sociological Review*, pages 103–127, 2003.

[21] G.L. Nemhauser and L.A. Wolsey. *Integer and Combinatorial Optimization*, volume 18 of *Discrete Mathematics and Optimization*. Wiley, New York, 1988.

[22] C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Dover Publications, Mineola, 1998.

[23] S. Raghavan and T.L. Magnanti. Network connectivity. In M. Dell'Amico, F. Maffioli, and S. Martello, editors, *Annotated Bibliographies in Combinatorial Optimization*, pages 335–354. John Wiley and Sons, New York, 1997.

[24] S. Rajagopalan and V.V. Vazirani. On the bidirected cut relaxation for the metric Steiner tree problem. In *SODA*, volume 99, pages 742–751, 1999.

[25] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, Berlin, 2003.

[26] S. Vijayanarasimhan and K. Grauman. Efficient region search for object detection. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1401–1408. IEEE, 2011.

[27] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*, volume 8 of *Strutural Analysis in the Social Sciences*. Cambridge University Press, New York, 1994.

Table 1: This table reports LP bounds for MWCS instances on $G(n, p)$ graphs using: the 0-1 bounds (0-1), the indegree inequalities and 0-1 bounds (Ind), the separator inequalities and 0-1 bounds (Sep), and both the indegree and separator inequalities and 0-1 bounds (Both). The numbers in the table are averages taken over five instances. If Ind or Sep is stronger than the other (on average), it is bolded. If Both is stronger than Ind and Sep, then it is bolded. If IP is not equal to Both, then it is bolded. If 0-1 equals IP, it is bolded.

p	n	m	0-1	Ind	Sep	Both	IP
0.01	50	17.0	631.0	**107.0**	315.5	107.0	107.0
0.02	50	25.0	666.4	**166.6**	333.2	**157.1**	**149.0**
0.03	50	37.4	546.4	**168.1**	273.2	**142.8**	**138.2**
0.04	50	46.8	602.8	436.3	**371.2**	**361.8**	361.8
0.05	50	66.4	685.0	612.7	**502.0**	502.0	502.0
0.06	50	73.4	609.2	579.7	**519.2**	519.2	**518.8**
0.07	50	89.2	591.6	582.4	**543.0**	543.0	543.0
0.08	50	94.8	557.4	550.8	**504.2**	504.2	504.2
0.09	50	109.8	739.0	737.6	**688.2**	688.2	688.2
0.10	50	120.6	596.4	593.3	**564.6**	564.6	564.6
0.11	50	129.6	530.6	529.4	**506.2**	506.2	506.2
0.12	50	145.0	654.0	653.6	**643.6**	643.6	643.6
0.13	50	163.8	675.0	675.0	**673.6**	673.6	673.6
0.14	50	168.8	648.4	648.4	**648.2**	648.2	648.2
0.15	50	185.8	622.4	622.4	**607.6**	607.6	607.6
0.16	50	199.8	557.6	557.6	**556.4**	556.4	556.4
0.17	50	200.8	726.6	726.6	**723.0**	723.0	723.0
0.18	50	216.6	628.6	628.6	**624.4**	624.4	624.4
0.19	50	235.4	581.6	581.6	**578.6**	578.6	578.6
0.20	50	231.2	527.6	527.6	**522.0**	522.0	522.0
0.21	50	256.4	606.8	606.8	**605.6**	605.6	605.6
0.22	50	268.2	605.8	605.8	**603.2**	603.2	603.2
0.23	50	282.4	**658.2**	658.2	658.2	658.2	658.2
0.24	50	292.0	**632.4**	632.4	632.4	632.4	632.4
0.25	50	298.8	**608.6**	608.6	608.6	608.6	608.6

Table 2: This table provides the same information as Table 1, but after performing the preprocessing step of contracting edges whose incident vertices have nonnegative weight.

			After Contracting Edges				
p	n	m	0-1	Ind	Sep	Both	IP
0.01	45.2	12.2	631.0	**107.0**	315.5	107.0	107.0
0.02	41.4	16.2	666.4	**160.8**	333.2	**156.2**	**149.0**
0.03	46.0	33.0	546.4	**154.7**	273.2	**138.2**	138.2
0.04	37.2	31.0	602.8	391.8	**371.2**	**361.8**	361.8
0.05	35.6	44.8	685.0	574.8	**502.0**	502.0	502.0
0.06	31.0	44.2	609.2	562.7	**519.2**	519.2	**518.8**
0.07	30.6	56.2	591.6	564.6	**543.0**	543.0	543.0
0.08	32.4	59.2	557.4	525.7	**504.2**	504.2	504.2
0.09	30.6	64.4	739.0	712.8	**688.2**	688.2	688.2
0.10	30.0	67.0	596.4	582.0	**564.6**	564.6	564.6
0.11	31.4	75.0	530.6	516.4	**506.2**	506.2	506.2
0.12	26.6	64.0	654.0	645.6	**643.6**	643.6	643.6
0.13	25.6	61.8	675.0	673.8	**673.6**	673.6	673.6
0.14	24.6	58.0	648.4	648.2	648.2	648.2	648.2
0.15	26.4	73.4	622.4	615.9	**607.6**	607.6	607.6
0.16	28.4	83.0	557.6	556.7	**556.4**	556.4	556.4
0.17	23.6	60.4	726.6	724.0	**723.0**	723.0	723.0
0.18	26.2	82.4	628.6	626.9	**624.4**	624.4	624.4
0.19	27.2	90.6	581.6	580.6	**578.6**	578.6	578.6
0.20	30.2	107.8	527.6	524.2	**522.0**	522.0	522.0
0.21	26.0	92.4	606.8	605.6	605.6	605.6	605.6
0.22	26.0	95.6	605.8	604.5	**603.2**	603.2	603.2
0.23	25.4	93.2	**658.2**	658.2	658.2	658.2	658.2
0.24	25.2	95.6	**632.4**	632.4	632.4	632.4	632.4
0.25	25.2	93.2	**608.6**	608.6	608.6	608.6	608.6

Table 3: This table reports LP bounds for MWCS instances that are similar to the $K_{q,q}$ example. One of the five instances at $p = 0.1$ is disconnected; all others are connected.

p	n	m	0-1	Ind	Sep	Both	IP
0.1	50	89.6	25	1.0	12.5	1.0	1.0
0.2	50	187.2	25	1.0	12.5	1.0	1.0
0.3	50	276.0	25	1.0	12.5	1.0	1.0
0.4	50	356.0	25	1.0	12.5	1.0	1.0
0.5	50	455.8	25	1.0	12.5	1.0	1.0
0.6	50	553.4	25	1.0	12.5	1.0	1.0
0.7	50	636.8	25	1.5	12.5	1.3	1.0
0.8	50	740.0	25	4.9	12.5	3.7	1.0
0.9	50	828.8	25	7.1	12.5	5.4	1.0
1.0	50	925.0	25	9.0	12.5	6.9	1.0