

On Generalized Network Design Polyhedra

Corinne Feremans* Martine Labbé† Adam N. Letchford‡
Juan-José Salazar-González§

January 2009, Revised March 2010, Corrected March 2011

Abstract

In recent years, there has been an increased literature on so-called *Generalized Network Design Problems* (GNDPs), such as the Generalized Minimum Spanning Tree Problem and the Generalized Traveling Salesman Problem. In a GNDP, the node set of a graph is partitioned into *clusters*, and the feasible solutions must contain one node from each cluster. Up to now, the polyhedra associated with different GNDPs have been studied independently. The purpose of this paper is to show that it is possible, to a certain extent, to derive polyhedral results for all GNDPs simultaneously. Along the way, we point out some interesting connections to other polyhedra, such as the quadratic semi-assignment polytope and the boolean quadric polytope.

Keywords: Generalized minimum spanning tree problem, Generalized traveling salesman problem, Polyhedral combinatorics.

1 Introduction

In recent years, there has been an increased literature on so-called *Generalized Network Design Problems* (GNDPs). In such problems, one is given an undirected graph $G = (V, E)$, and the vertex set V is partitioned into m clusters V_1, \dots, V_m . The problem to be solved is then analogous to a

*Department of Technology, Telemark University College, Porsgrunn, Norway. E-mail: corinne.feremans@hit.no

†Department of Computer Science, Université Libre de Bruxelles, Belgium. E-mail: mllabbe@ulb.ac.be

‡Department of Management Science, Lancaster University, United Kingdom. E-mail: a.n.letchford@lancaster.ac.uk

§DEIOC, University of La Laguna, Tenerife, Spain. E-mail: jjsalaza@ull.es




Figure 1: A graph with 8 nodes, 4 clusters and 10 edges, with a GS in bold.

standard network optimization problem (such as the Traveling Salesman Problem, Minimum Spanning Tree Problem or Shortest Path Problem), but with the added requirement that the solution contain exactly (or at most, or at least) one node from each cluster.

The first GNDP, the *Generalized Traveling Salesman Problem* or GTSP, was introduced (apparently independently) by Henry-Labordere [10], Saskena [19] and Srivastava *et al.* [21]. Later, the *Generalized Minimum Spanning Tree Problem* or GMSTP was introduced by Myung *et al.* [13]. Since then, several other GNDPs have been introduced. We refer the reader to Feremans *et al.* [7] for a survey on GNDPs and their applications.

Up to now, the polyhedra associated with GNDPs have been studied more or less independently. For example, Fischetti *et al.* [9] focussed on GTSP polyhedra, whereas Feremans *et al.* [8] focussed on GMSTP polyhedra. Our goal in this paper is to show that it is possible, to a certain extent, to derive polyhedral results for all GNDPs simultaneously.

Our key concept is that of a *generalized subgraph* (GS). A GS is simply a subgraph of G containing exactly one node from each cluster. (We impose no other structure on the subgraph. In particular, it need not be connected, and it may have isolated nodes.)

Figure 1 shows a graph and a GS. The numbered dots represent the nodes, the ovals represent the clusters, and the lines represent the edges. The lines and dots in bold represent the GS.

Clearly, every feasible solution to a GNDP is a GS. Thus, if a linear inequality is valid for every incidence vector of a GS, it is also valid for all GNDPs. This observation led us to study the GS polyhedra themselves.

The structure of the paper is as follows. In Section 2, we formally define generalized subgraphs and the associated polyhedra, and provide some additional motives for studying them. In Section 3, we determine some fundamental properties of the polyhedra. In Section 4, we present various valid and facet-defining inequalities. Then, in Section 5, we consider the polyhedra obtained when, instead of requiring *exactly* one node per cluster, one requires either *at most* or *at least* one node per cluster. Finally, some conclusions and suggestions for future research are given in Section 6.

2 Definitions and Motivation

In this section, we define GS polyhedra formally (Subsection 2.1). We also point out some connections between GS polyhedra and some other polyhedra in the literature (Subsection 2.2), which provides further motivation for our study. We assume throughout that the reader is familiar with the basics of polyhedral theory (see, e.g., Nemhauser & Wolsey [14]).

2.1 Definitions

In this subsection, we define GS polyhedra formally. Along the way, we present some useful notation.

We define a binary variable x_v for each node $v \in V$, taking the value 1 if and only if node v is in the GS. We also define a binary variable y_e for each edge $e \in E$, taking the value 1 if and only if edge e is in the GS. For any set $S \subseteq V$, we let $x(S)$ denote $\sum_{v \in S} x_v$. Similarly, for any set $F \subset E$, we let $y(F)$ denote $\sum_{e \in F} y_e$. For any node set S , we let $E(S)$ denote the set of edges in E with both end-nodes in S . For any disjoint node sets S, T , we let $E(S : T)$ denote the set of edges in E with one node in S and the other in T . Finally, we let K denote $\{1, \dots, m\}$.

Given the above notation, a vector $(x, y) \in \{0, 1\}^{|V|+|E|}$ is the incidence vector of a GS if and only if the following inequalities are satisfied:

$$x(V_k) = 1 \quad (k \in K) \tag{1}$$

$$y_{uv} \leq x_u, y_{uv} \leq x_v \quad (\{u, v\} \in E). \tag{2}$$

We refer to the constraints (1) and (2) as *cluster constraints* and *variable upper bounds* (VUBs), respectively.

The above considerations lead naturally to a family of zero-one polytopes (bounded polyhedra). Specifically, we define:

$$P(G) = \text{conv} \left\{ (x, y) \in \{0, 1\}^{|V|+|E|} : (1), (2) \text{ hold} \right\}.$$

Here, as usual, ‘conv’ denotes the the convex hull of a set of points.

Clearly, given any specific GNDP that requires exactly one node per cluster, the convex hull of feasible solutions will be contained in $P(G)$ for some graph G . As a result, valid inequalities for GS polytopes can be used as cutting planes for GNDPs of that kind.

We remark that the VUBs (2) can be replaced by the following stronger inequalities:

$$y(E(\{v\} : V_k)) \leq x_v \quad (k \in K, v \in V \setminus V_k). \quad (3)$$

We call these *strengthened variable upper bounds* (SVUBs). The validity of SVUBs was noted in [9] for the GTSP and in [8] for the GMSTP.

2.2 Motivation

Although we are interested in GS polytopes primarily due to their application to GNDPs, it turns out that they are also closely related to some other important combinatorial optimization problems and associated polytopes.

First, we consider the *Quadratic Semi-Assignment Problem* (QSAP), an \mathcal{NP} -hard problem with many applications (see, e.g., Burkard & Çela [3]). In the QSAP, one has binary variables x_{ik} for $i = 1, \dots, n$ and $k = 1, \dots, m$. The task is to minimize a quadratic function of the form

$$\sum_{i=1}^n \sum_{k=1}^m c_{ik} x_{ik} + \sum_{i=1}^n \sum_{j=1}^n \sum_{1 \leq k < \ell \leq m} q_{ijk\ell} x_{ik} x_{j\ell}, \quad (4)$$

subject to constraints of the form:

$$\sum_{k=1}^m x_{ik} = 1 \quad (i = 1, \dots, n).$$

We have the following result:

Proposition 1 *The QSAP can be reduced to the problem of maximizing a*

non-negative linear function over $P(G)$.

Proof. Note that any feasible QSAP solution satisfies $\sum_{i=1}^n \sum_{k=1}^m x_{ik} = n$ and

$$\sum_{i=1}^n \sum_{j=1}^n \sum_{1 \leq k < \ell \leq m} x_{ik} x_{j\ell} = \binom{m}{2}.$$

Therefore, minimizing (4) is equivalent to maximizing

$$\sum_{i=1}^n \sum_{k=1}^m (M - c_{ik}) x_{ik} + \sum_{i=1}^n \sum_{j=1}^n \sum_{1 \leq k < \ell \leq m} (M - q_{ijkl}) x_{ik} x_{j\ell}, \quad (5)$$

where M is a positive integer chosen to make all of the terms in (5) non-negative. Now, construct a complete m -partite graph G with m clusters, each containing n nodes. Assign a weight of $M - c_{ik}$ to the i th node in the k th cluster, and a weight of $M - q_{ijkl}$ to the edge connecting the i th node in the k th cluster to the j th node in the ℓ th cluster. Then, maximizing (5) is equivalent to finding a GS of maximum weight in G . \square

In fact, the polytopes introduced by Saito *et al.* [18] for the QSAP turn out to be nothing but the faces of $P(G)$ obtained by changing the SVUBs (3) to equations. Thus, the inequalities that we derive for $P(G)$ may also be useful for solving the QSAP.

Second, we mention that a variety of \mathcal{NP} -hard *Constraint Satisfaction Problems* (CSPs) can also be reduced, in a similar way, to the problem of maximizing a non-negative linear function over $P(G)$. This includes, for example, the problem **MAX CAPACITY REPRESENTATIVES** of Bellare [2], the problem **UNIQUE GAMES** of Khot [11], and the problems **MAX 2-CONJ** and **MAX 2-CSP** of Serna *et al.* [20]. (We omit details of the reductions, for the sake of brevity.) Thus, the inequalities derived in this paper may also be useful for solving CSPs.

We remark that some other well-known \mathcal{NP} -hard combinatorial optimization problems, such as **MAX-CUT**, **MAX-DI-CUT** and **MAX-2-SAT**, are special cases of **MAX 2-CSP**. Such problems can be reduced to the problem of maximizing a non-negative linear function over $P(G)$ for graphs G that have only two nodes per cluster.

Finally, we also remark that the ‘partial constraint satisfaction polytope’ studied by Koster *et al.* [12] is essentially identical to the QSAP polytope

studied by Saito *et al.* [18]. The inequalities presented in the two papers are, however, different.

3 Fundamental Properties of the Polyhedra

In this section, we determine some of the fundamental properties of $P(G)$. For ease of notation, we assume in this section and the next that G is a complete m -partite graph. (Note that an ‘intra-cluster’ edge, i.e., an edge with both end-nodes in the same cluster, can never be part of a GS.)

3.1 Dimension

First, we determine the dimension of the polytope:

Theorem 1 $P(G)$ is of dimension $|V| + |E| - m$.

Proof. Since the equations (1) are linearly independent, the dimension is certainly no larger than $|V| + |E| - m$. Now, for each $k \in K$, let $v(k)$ be an arbitrary node in V_k , and let V^* be the union of these selected nodes. Consider the following extreme points of $P(G)$:

- the point obtained by setting $x_v = 1$ for all $v \in V^*$, $x_v = 0$ for all $v \in V \setminus V^*$, and $y_e = 0$ for all $e \in E$;
- for all $k \in K$, and for all $v \in V_k \setminus \{v(k)\}$, the point obtained by taking the first point, changing $x_{v(k)}$ from 1 to 0, and changing x_v from 0 to 1;
- for all $e = \{u, v\} \in E$, an arbitrary point such that $y_e = 1$ and $y_f = 0$ for all $f \in E \setminus \{e\}$.

These $|V| + |E| - m + 1$ points are affinely independent. □

3.2 Canonical form

The fact that $P(G)$ is not full-dimensional complicates matters, since each facet is defined by an infinite number of linear inequalities. The following theorem enables us to associate one ‘canonical’ inequality with each (non-trivial) facet:

Theorem 2 (Canonical Form) *Let F be a facet of $P(G)$, and suppose that F is not defined by a lower bound of the form $y_e \geq 0$. Then there is a unique inequality of the form*

$$\sum_{e \in E} \beta_e y_e \leq \sum_{v \in V} \alpha_v x_v + \gamma, \quad (6)$$

that defines the face F and satisfies the following conditions:

- α , β and γ are all non-negative;
- in each cluster, there exists at least one node u for which $\alpha_u = 0$.

Proof. Clearly, we can assume that the inequality is in ‘ \leq ’ form. Moreover, by adding or subtracting suitable multiples of the cluster constraints (1), we can ensure that $\alpha \geq 0$ and that the second condition is satisfied.

Now, let $e \in E$ be an arbitrary edge. If the facet is not defined by the inequality $y_e \geq 0$, then there exists at least one extreme point on the facet satisfying $y_e = 1$. From this, we can obtain another extreme point of $P(G)$ by setting $y_e = 0$. Therefore $\beta_e \geq 0$.

Finally, for each $k \in K$, let $v(k)$ be a node in V_k such that $\alpha_{v(k)} = 0$. By setting $x_{v(k)}$ to 1 for all k , and all other variables to zero, we obtain an extreme point of $P(G)$ for which the slack of the inequality (6) is equal to γ . Thus $\gamma \geq 0$. \square

Note that the SVUBs (3) are in canonical form.

3.3 Cluster addition and node cloning

Next, we present two simple results that enable new facet-defining inequalities to be derived from known ones.

Theorem 3 (Cluster Addition) *Suppose that an inequality of the form (6) is valid for $P(G)$. Let $G' = (V', E')$ be a graph obtained from G by adding another cluster V_{m+1} , consisting of a single node u , together with the $|V|$ additional edges needed to make G complete $(m + 1)$ -partite. Then the inequality (6) is valid for $P(G')$. Moreover, the inequality defines a facet of $P(G')$ if and only if it defines a facet of $P(G)$.*

Proof. See the appendix. \square

Theorem 4 (Node Cloning) *Suppose that an inequality of the form (6) is valid for $P(G)$. Let $u \in V$ be a specified node, and let G' be the graph obtained from G by ‘cloning’ node u . That is, a new node u' is added to V , in the same cluster as u , and the edge $\{u', v\}$ is added if and only if $\{u, v\} \in E$. Then the ‘cloned’ inequality*

$$\sum_{e \in E} \beta_e y_e + \sum_{v: \{u, v\} \in E} \alpha_{uv} y_{u'v} \leq \sum_{v \in V} \alpha_v x_v + \alpha_u x_{u'} + \gamma$$

is valid for $P(G')$. Moreover, if there exists a node $w \neq u$ such that $\alpha_w > 0$, then the cloned inequality defines a facet of $P(G')$ if and only if the original inequality defines a facet of $P(G)$.

Proof. See the appendix. □

3.4 A connection with the boolean quadric polytope

Finally, we consider the special case of $P(G)$ that arises when each cluster contains exactly two nodes. It turns out that, in this case, there is a connection between $P(G)$ and the so-called *boolean quadric polytope*. The boolean quadric polytope of order m , denoted by BQP_m , is defined as follows (Padberg [15]):

$$BQP_m = \text{conv} \left\{ (z, Z) \in \{0, 1\}^{m + \binom{m}{2}} : Z_{ij} = z_i z_j \ (1 \leq i < j \leq m) \right\}.$$

The boolean quadric polytope is a fundamental problem in quadratic 0 – 1 optimization, and it is also an affine image of the so-called *cut polytope*, which is the polytope associated with the well-known max-cut problem [1, 5]. The boolean quadric and cut polytopes have been studied in great depth (see Deza & Laurent [6] for an extensive survey).

We have the following theorem:

Theorem 5 *Let m be fixed, and as usual let $K = \{1, \dots, m\}$. Let $G_m = (V_m, E_m)$ be a complete m -partite graph with m clusters and 2 nodes per cluster. For $k \in K$, let $u(k)$ denote one of the two nodes in the cluster V_k , and let $v(k)$ denote the other node. We define the following affine mapping, that maps each point $(z^*, Z^*) \in \mathbb{R}^{m + \binom{m}{2}}$ to a point $(x^*, y^*) \in \mathbb{R}^{|V_m| + |E_m|}$:*

- $x_{u(k)}^* = z_k^*$ and $x_{v(k)}^* = 1 - z_k^*$ for all $k \in K$;

- $y_{u(k),u(\ell)}^* = Z_{k\ell}^*$ and $y_{v(k),v(\ell)}^* = 1 - z_k^* - z_\ell^* + Z_{k\ell}^*$ for all $\{k, \ell\} \subset K$;
- $y_{u(k),v(\ell)}^* = z_k^* - Z_{k\ell}^*$ for all $\{k, \ell\} \subset K$.

Let BQP'_m be the image of BQP_m under this mapping. Then $P(G_m)$ is the downward monotonization of BQP'_m with respect to the y variables. That is,

$$P(G_m) = \left\{ (x, y) \in [0, 1]^{|V_m|+|E_m|} : \exists (x, y') \in BQP'_m : y \leq y' \right\}.$$

Proof. Let (z^*, Z^*) be an extreme point of BQP_m , and let (x^*, y^*) be the corresponding extreme point of BQP'_m . It follows immediately from the definitions that:

- $(x^*, y^*) \in \{0, 1\}^{|V_m|+|E_m|}$;
- $x^*(V_k) = 1$ for $k \in K$;
- $y_{ij}^* = 1$ if and only if $x_i^* = x_j^* = 1$, for all $\{i, j\} \in E_m$.

Thus, (x^*, y^*) is an extreme point of $P(G_m)$. Moreover, any vector obtained from (x^*, y^*) by changing a y variable from 1 to 0 is also an extreme point of $P(G_m)$, since deleting an edge from a GS yields another GS. Thus, every extreme point of the downward monotonization of BQP'_m is an extreme point of $P(G_m)$. In a similar way, one can show that every extreme point of $P(G_m)$ is an extreme point of the downward monotonization of BQP'_m . \square

In Subsection 4.2, we will use Theorem 5 to derive valid inequalities for $P(G)$.

4 Valid Inequalities and Facets

In this section, we present some specific valid inequalities, and show that they define facets of $P(G)$ under certain conditions. For the sake of brevity, we omit the proofs of some of the results. In all such cases, the omitted details are easy, but tedious.

4.1 Trivial facets

The following two propositions describe some trivial facets of $P(G)$:

Proposition 2 For each $e \in E$, the lower bound $y_e \geq 0$ defines a facet of $P(G)$.

Proof. Simply note that all of the affinely independent points listed in the proof of Theorem 1 satisfy the lower bound at equality, apart from one. \square

Proposition 3 The SVUBs (3) define facets of $P(G)$, unless $|V_k| > 1$ and $\{v\}$ is itself a cluster.

Proof. If $|V_k| > 1$ and $\{v\}$ is a cluster, then the SVUB is dominated by the cluster constraints $x_v = 1$ and $x(V_k) = 1$, together with the VUBs $y_{uv} \leq x_u$ for all $u \in V_k$. Now, Proposition 7 of [8] states that, in all other cases, the SVUB inequality defines a facet of the GMTSP polytope. The GMTSP polytope is contained in $P(G)$, and has dimension $|V| + |E| - m - 1$, because it satisfies the equation $y(E) = m - 1$ in addition to the equations satisfied by $P(G)$. Therefore, there exist $|V| + |E| - m$ affinely independent extreme points of $P(G)$ that satisfy the SVUB at equality and also satisfy $y(E) = m - 1$. To complete the proof, we need an extreme point satisfying the SVUB at equality, but not satisfying $y(E) = m - 1$. For this, we can take any extreme point satisfying $y(E) = 0$ and $x_v = 0$. \square

Note that the bounds $0 \leq x_v \leq 1$ for all v , and $y_e \leq 1$ for all e , do not define facets in general, since they are implied by the inequalities in the above propositions, together with the cluster constraints (1).

We remark that, when $|V_k| = 1$ for all $k \in K$, $P(G)$ is completely described by the cluster constraints $x_v = 1$ for all v and the bounds $0 \leq y_e \leq 1$ for all e . On the other hand, there is no hope of obtaining a complete linear description of $P(G)$ when $|V_k| = 2$ for all k , since optimizing a linear function over $P(G)$ is \mathcal{NP} -hard even in this special case (see the penultimate paragraph of Subsection 2.2).

4.2 Inequalities from the boolean quadric polytope

Now, we will show how to derive valid inequalities for $P(G)$ from valid inequalities for the boolean quadric polytope. We use exactly the same notation as in Subsection 3.4.

We will find the following three lemmas helpful:

Lemma 1 *If an inequality of the form $\alpha^T z + \beta^T Z \leq \gamma$ is valid for BQP_m , then the inequality*

$$\sum_{k=1}^m \alpha_k x_{u(k)} + \sum_{1 \leq k < \ell \leq m} \beta_{k\ell} y_{u(k),u(\ell)} \leq \gamma$$

is valid for BQP'_m .

Proof. This follows trivially from the way BQP'_m is constructed. \square

Lemma 2 *If an inequality is valid for BQP'_m , there is at least one valid inequality for BQP'_m that defines the same face and has canonical form.*

Proof. It follows from the way BQP'_m is constructed that it satisfies the following equations:

- $x(V_k) = 1$ for all $k \in K$;
- $y_{u(k),v(\ell)} = x_{u(k)} - y_{u(k),u(\ell)}$ for all $\{k, \ell\} \subset K$.
- $y_{v(k),v(\ell)} = 1 - x_{u(k)} - x_{u(\ell)} + y_{u(k),u(\ell)}$ for all $\{k, \ell\} \subset K$.

By adding or subtracting suitable multiples of these equations, one can always bring any valid inequality into the desired form. \square

Lemma 3 *If an inequality is valid for BQP'_m , and has canonical form, then it is valid also for $P(G_m)$.*

Proof. This follows from the fact that $P(G_m)$ is the downward monotonicization of BQP'_m with respect to the y variables. \square

The above three lemmas enable one to take any valid inequality for BQP_m , and convert it into a canonical valid inequality for $P(G_m)$. Moreover, by applying cluster addition and node cloning (Subsection 3.3), one can then derive canonical valid inequalities for $P(G)$, for more general graphs G .

We have been able to derive many interesting canonical inequalities for $P(G)$ in this way. For the sake of brevity, we give just two examples, based on inequalities presented in Padberg [15].

Proposition 4 (Odd Ring Inequalities) *Let V_1, \dots, V_c be clusters, where $c \geq 3$. For $k = 1, \dots, c$, let cluster V_k be partitioned into two subsets, called V_k^1 and V_k^2 . Moreover, let C denote $\{1, \dots, c\}$ and let $D \subseteq C$ be such that $|D|$ is odd. Then the following ‘odd ring’ inequality is valid for $P(G)$, where indices are taken modulo c (see Figure 2 for an illustration):*

$$\sum_{k \in C \setminus D} y(E(V_k^1 : V_{k+1}^1)) + \sum_{k \in D} y(E(V_k^1 : V_{k+1}^2)) \leq \sum_{k \in C \setminus D} x(V_{k+1}^1) + \lfloor |D|/2 \rfloor. \quad (7)$$

Proof. Padberg ([15], p. 154) showed that the following ‘odd cycle’ inequalities are valid for BQP_m , for any $c \geq 3$ and any $D \subset C$ such that $|D|$ is odd:

$$\sum_{k \in C \setminus D} Z_{k,k+1} + \sum_{k \in S_0} z_k \leq \sum_{k \in D} Z_{k,k+1} + \sum_{k \in S_2} z_k + \lfloor |D|/2 \rfloor,$$

where $S_0 = \{k \in D : k - 1 \in D\}$ and $S_2 = \{k \in C \setminus D : k - 1 \in C \setminus D\}$. This, together with Lemma 1, implies that the (non-canonical) inequality

$$\sum_{k \in C \setminus D} y_{u(k),u(k+1)} + \sum_{k \in S_0} x_{u(k)} \leq \sum_{k \in D} y_{u(k),u(k+1)} + \sum_{k \in S_2} x_{u(k)} + \lfloor |D|/2 \rfloor \quad (8)$$

is valid for BQP'_m . Now, from the proof of Lemma 2, all points in BQP'_m satisfy the following equations:

$$y_{u(k),u(k+1)} = x_{u(k)} - y_{u(k),v(k+1)} \quad (\forall k \in D).$$

Together with (8), this implies that the canonical inequality

$$\sum_{k \in C \setminus D} y_{u(k),u(k+1)} + \sum_{k \in D} y_{u(k),v(k+1)} \leq \sum_{k \in C \setminus D} x_{u(k+1)} + \lfloor |D|/2 \rfloor$$

is valid for BQP'_m . By Lemma 3, this inequality is valid also for $P(G_m)$. Validity for $P(G)$ then follows from Theorem 4. \square

Proposition 5 (Odd Clique Inequalities) *Let V_1, \dots, V_c be clusters, and assume $c \geq 3$ and odd. For $k = 1, \dots, c$, let cluster V_k be partitioned into two subsets, called V_k^1 and V_k^2 . For any integer $1 \leq r \leq \lfloor c/2 \rfloor$, the following*

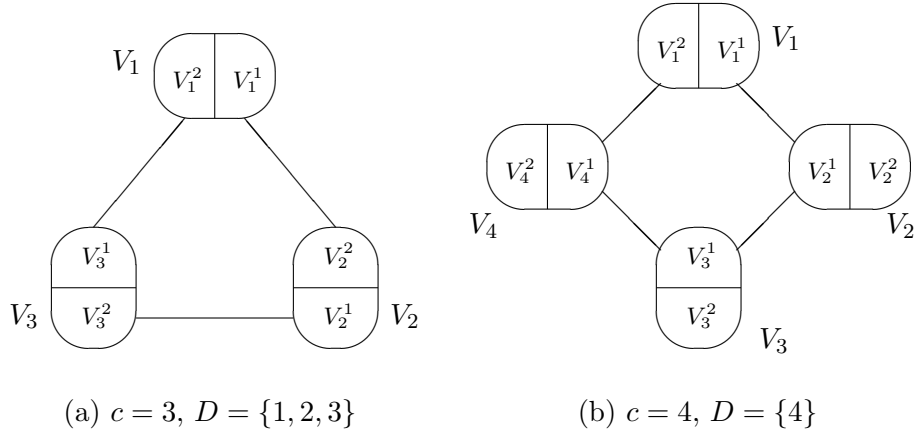


Figure 2: Representation of two odd ring inequalities. Edges represent y variables with a coefficient of 1.

‘odd clique’ inequality is valid for $P(G)$, where indices are taken modulo c (see Figure 3 for an illustration):

$$\sum_{k=1}^c \sum_{\ell=1}^{\lfloor c/2 \rfloor} y(E(V_k^1 : V_{k+\ell}^2)) \leq (\lfloor c/2 \rfloor - r) \sum_{k=1}^c x(V_k^1) + r(r+1)/2. \quad (9)$$

Proof. Padberg ([15], p. 149) showed that the following ‘clique’ inequalities define facets of BQP_m , for any $c \geq 3$ and for any $1 \leq r \leq c-2$:

$$r \sum_{k=1}^c z_k \leq \sum_{1 \leq k < \ell \leq c} Z_{k\ell} + r(r+1)/2. \quad (10)$$

Taking c odd and $1 \leq r \leq \lfloor c/2 \rfloor$, and applying Lemma 1, we have the following (non-canonical) valid inequality for BQP'_m :

$$r \sum_{k=1}^c x_{u(k)} \leq \sum_{k=1}^c \sum_{\ell=1}^{\lfloor c/2 \rfloor} y_{u(k), u(k+\ell)} + r(r+1)/2. \quad (11)$$

Now, from the proof of Lemma 2, all points in BQP'_m satisfy the following equations:

$$y_{u(k), u(k+\ell)} = x_{u(k)} - y_{u(k), v(k+\ell)} \quad (\forall k, \ell).$$




Figure 3: Representation of an odd clique inequality for $c = 5$. Edges represent y variables with a coefficient of 1.

Together with (11), this implies that the canonical inequality

$$\sum_{k=1}^c \sum_{\ell=1}^{\lfloor c/2 \rfloor} y_{u(k),v(k+\ell)} \leq (\lfloor c/2 \rfloor - r) \sum_{k=1}^c x_{u(k)} + r(r+1)/2$$

is valid for BQP'_m . By Lemma 3, this inequality is valid also for $P(G_m)$. Validity for $P(G)$ then follows from Theorem 4. \square

It turns out that the odd ring and odd clique inequalities are always facet-defining:

Theorem 6 *Odd ring inequalities define facets of $P(G)$.*

Proof. See the appendix. \square

Theorem 7 *Odd clique inequalities define facets of $P(G)$.*

Proof. See the appendix. \square

We remark that the ‘odd cycle’ inequalities for the GMTSP, presented by Feremans *et al.* [8], are the special case of the odd ring inequalities obtained when $D = C$. Moreover, it can be shown that the ‘cycle’ inequalities for the partial constraint satisfaction polytope, presented in Koster *et al.* [12], are equivalent to odd ring inequalities, in the sense that the cycle and odd ring inequalities define the same facets of that polytope. The odd clique inequalities, on the other hand, appear to be entirely new.

4.3 Odd circulant inequalities

The valid inequalities in the previous subsection involved partitions of clusters into two components. In this subsection, we introduce a class of canonical inequalities for $P(G)$ that involves partitions of clusters into more than two components.

Proposition 6 (Odd Circulant Inequalities) *Let $c \geq 3$ be an odd integer, let V_1, \dots, V_c be clusters, and let d be an integer between 1 and $\lfloor c/2 \rfloor$. Moreover, let each cluster V_k , for $k = 1, \dots, c$, be partitioned into $2d$ sets, called V_k^ℓ for $\ell = 1, \dots, 2d$. The following ‘odd circulant’ inequality is valid for $P(G)$, where indices are taken modulo c (see Figure 4 for an illustration):*

$$\sum_{k=1}^c \sum_{\ell=1}^d y(E(V_k^\ell : V_{k+\ell}^{2d+1-\ell})) \leq \lfloor c/2 \rfloor. \quad (12)$$

Proof. For a given k , if we sum together the cluster constraint $x(V_k) = 1$ and $|V_k|$ suitable SVUBs, we obtain:

$$\sum_{\ell=1}^d y(E(V_k^\ell : V_{k+\ell}^{2d+1-\ell})) + \sum_{\ell=1}^d y(E(V_k^{d+\ell} : V_{k-\ell}^{d+1-\ell})) \leq 1.$$

Summing these constraints over all k and simplifying, we obtain:

$$2 \sum_{k=1}^c \sum_{\ell=1}^d y(E(V_k^\ell : V_{k+\ell}^{2d+1-\ell})) \leq c.$$

Dividing by two and rounding down the right-hand side, we obtain the odd circulant inequality. \square

It turns out that the odd circulant inequalities are always facet-defining:

Theorem 8 *Odd circulant inequalities define facets of $P(G)$.*

Proof. See the appendix. \square

We remark that the odd circulant inequalities contain some other known inequalities as special cases. When $d = \lfloor c/2 \rfloor$, they reduce to the *odd clique matching* inequalities, introduced by Feremans *et al.* [8] in the context of




Figure 4: Representation of an odd circulant inequality for $c = 5$ and $d = 2$.

the GMSTP. On the other hand, when $d = 1$, they reduce to the *odd cycle* inequalities, also due to Feremans *et al.* [8], that we mentioned in the previous subsection.

5 Two Related Polytopes: $P^{\leq}(G)$ and $P^{\geq}(G)$

In this section, we consider the polytopes obtained when, instead of requiring *exactly* one node per cluster, one requires either *at most* or *at least* one node per cluster. We deal with the ‘at most’ variant in Subsection 5.1 and the ‘at least’ variant in Subsection 5.2. We will see that the ‘at least’ variant is considerably more complicated than the other two variants.

5.1 The polytope $P^{\leq}(G)$

To handle the case of at most one node per cluster, we must change the cluster constraints (1) from equations to inequalities as follows:

$$\sum_{v \in V_k} x_v \leq 1 \quad (k \in K). \quad (13)$$

Accordingly, we define the following polytope:

$$P^{\leq}(G) = \text{conv} \left\{ (x, y) \in \{0, 1\}^{|V|+|E|} : (2), (13) \text{ hold} \right\}.$$

As before, we assume for simplicity that G is a complete m -partite graph.

Note that, by definition, $P(G)$ is the face of $P^{\leq}(G)$ defined by the equations (1). Using this fact, together with the results given in Section 3, one can easily prove the following results:

Proposition 7 $P^{\leq}(G)$ is full-dimensional, i.e., of dimension $|V| + |E|$.

Proposition 8 If an inequality of the form $\beta y \leq \alpha x + \gamma$ defines a facet of $P^{\leq}(G)$, and is not equivalent to a bound of the form $y_e \geq 0$ or a cluster constraint of the form (13), then α , β and γ are non-negative.

Proposition 9 Cluster addition applies to all inequalities that define facets of $P^{\leq}(G)$. That is, if an inequality of the form $\beta y \leq \alpha x + \gamma$ defines a facet of $P^{\leq}(G)$, and G' is defined as in Theorem 3, then the inequality also defines a facet of $P^{\leq}(G')$.

Proposition 10 Node cloning applies to all inequalities that define facets of $P^{\leq}(G)$, apart from bounds of the form $y_e \geq 0$. That is, if an inequality defines a facet of $P^{\leq}(G)$, and is not a bound of the given form, then the cloned inequality defines a facet of $P^{\leq}(G')$, where the cloned inequality and G' are defined as in Theorem 4.

Proposition 11 For each $e \in E$, the lower bound $y_e \geq 0$ defines a facet of $P^{\leq}(G)$.

Proposition 12 For each $k \in K$, the cluster constraint $x(V_k) \leq 1$ defines a facet of $P^{\leq}(G)$.

The situation with ‘canonical form’ (Theorem 2) is, however, more complicated. In the first place, we have the following result:

Proposition 13 There exist non-trivial inequalities that define facets of $P^{\leq}(G)$, yet do not have canonical form.

Proof. Let $m = 3$, $|V| = 4$, $V_1 = \{1, 2\}$, $V_2 = \{3\}$ and $V_3 = \{4\}$, and assume that G is complete 3-partite. One can check (either by hand or with the aid of a computer) that the following inequality defines a facet of $P^{\leq}(G)$:

$$y_{13} + y_{24} + y_{34} \leq x_3 + x_4. \quad (14)$$

This inequality is not canonical, because V_2 does not contain a node u whose variable x_u has a coefficient of zero in the inequality (and the same applies to V_3). \square

We remark that the inequality (14) can be viewed as a ‘degenerate’ odd ring inequality, in which $c = 3$, $D = \{3\}$, and V_2^2 and V_3^2 are empty.

On the other hand, we have the following positive result:

Theorem 9 *If an inequality defines a facet of $P(G)$, and it has canonical form, then it also defines a facet of $P^\leq(G)$.*

Proof. Since the inequality defines a facet of $P(G)$, there exist $|V| + |E| - m$ affinely independent extreme points of $P(G)$ that satisfy the inequality at equality. These are all extreme points of $P^\leq(G)$ as well. So, to complete the proof, we need m additional affinely independent extreme points.

Now consider an arbitrary cluster V_k . Since the inequality is in canonical form, there exists a node $u \in V_k$ such that $\alpha_u = 0$. Moreover, there must exist at least one extreme point in our collection of $|V| + |E| - m$ extreme points satisfying $x_u = 1$ (since otherwise all of the extreme points would satisfy the equation $x_u = 0$, and the inequality would not define a facet of $P(G)$). Taking one such extreme point, and setting x_u to 0, yields an extreme point of $P^\leq(G)$ that is not an extreme point of $P(G)$, yet satisfies the inequality at equality. Moreover, this extreme point is affinely independent of the previous ones, since it does not satisfy the cluster constraint $x(V_k) = 1$. Repeating this for all clusters k , we obtain the desired m additional extreme points. \square

This immediately yields the following corollary:

Corollary 1 *SVUB, odd ring, odd clique and odd circulant inequalities define facets of $P^\leq(G)$.*

Thus, all of those inequalities can be used to tackle GNDPs in which at most one node may be selected from each cluster.

We close this subsection with the following theorem, which may also be of interest:

Theorem 10 *Suppose that $|V_k| = 1$ for all $k \in K$. Then:*

- $P^\leq(G)$ is completely described by the VUBs (2), the upper bounds $x_v \leq 1$ for all v , and the lower bounds $y_e \geq 0$ for all e .
- One can optimize a linear function over $P^\leq(G)$ in polynomial time, by solving a max-flow/min-cut problem in a graph with $m + 2$ nodes.

Proof. See the appendix. \square

5.2 The polytope $P^{\geq}(G)$

To handle the case of at least one node per cluster, we must change the cluster constraints (1) from equations to inequalities as follows:

$$\sum_{v \in V_k} x_v \geq 1 \quad (k \in K). \quad (15)$$

Accordingly, we define the following polytope:

$$P^{\geq}(G) = \text{conv} \left\{ (x, y) \in \{0, 1\}^{|V|+|E|} : (2), (15) \text{ hold} \right\}.$$

Although $P^{\geq}(G)$ may seem at first sight to be a minor variant of the other two polytopes, it is in fact significantly more complex. Consider, for example, the following three points:

- We can no longer assume that G is m -partite, since it is now possible for a GS to include one or more ‘intra-cluster’ edges (edges with both end-nodes in the same cluster).
- The SVUBs (3) are no longer valid, since the left-hand side can now be greater than 1 (unless only a single node in V_k is adjacent to v , in which case the inequality reduces to a standard VUB).
- Node cloning is no longer guaranteed to work. For example, if we take a VUB of the form $y_{uv} \leq x_u$, and clone node v , we obtain the SVUB $y_{uv} + y_{uv'} \leq x_u$, which is not valid.

In the light of the first point, it makes sense to assume that G is a complete graph. The following two propositions can then be proven easily:

Proposition 14 $P^{\geq}(G)$ is full-dimensional, i.e., of dimension $|V| + |E|$, unless there is a cluster V_k such that $|V_k| = 1$, in which case the equation $x(V_k) = 1$ is valid.

Proposition 15 If an inequality of the form $\beta y \leq \alpha x + \gamma$ defines a facet of $P^{\geq}(G)$, and is not equivalent to a bound of the form $x_v \leq 1$ or $y_e \geq 0$, then α and β are non-negative.

Moreover, although node cloning does not apply to valid inequalities for $P^{\geq}(G)$, a general form of cluster addition applies, as expressed in the following theorem:

Theorem 11 *Suppose that an inequality of the form (6) is valid for $P^\geq(G)$, and let k be an arbitrary positive integer. Let $G' = (V', E')$ be a graph obtained from G by adding another cluster V_{m+1} , containing k nodes, together with the $k|V| + \binom{k}{2}$ additional edges needed to make G complete. Then the inequality (6) is valid for $P^\geq(G')$. Moreover, it defines a facet of $P^\geq(G')$ if and only if it defines a facet of $P^\geq(G)$.*

Proof. See the appendix. □

Armed with Theorem 11, one can easily prove the following results:

Proposition 16 *For each $v \in V$, the upper bound $x_v \leq 1$ defines a facet of $P^\geq(G)$, unless $\{v\}$ is itself a cluster, in which case the equation $x_v = 1$ is valid.*

Proposition 17 *For each $e \in E$, the lower bound $y_e \geq 0$ defines a facet of $P^\geq(G)$.*

Proposition 18 *For each $\{u, v\} \in E$, the VUBs $y_{uv} \leq x_u$ and $y_{uv} \leq x_v$ define facets of $P^\geq(G)$, unless $\{u, v\}$ is itself a cluster, in which case the stronger inequality $y_{uv} \leq x_u + x_v - 1$ is valid and facet-defining.*

Now, consider again the SVUBs (3). Although they are not valid for $P^\geq(G)$, they can be ‘lifted’ to make them valid, by adding a suitable multiple of the cluster constraint (15). This leads to the following result:

Proposition 19 *The following ‘lifted SVUBs’ define facets of $P^\geq(G)$:*

$$y(E(\{v\} : V_k)) \leq x_v + x(V_k) - 1 \quad (k \in K, v \in V \setminus V_k).$$

Proof. Suppose that a feasible solution includes r of the edges in $E(\{v\} : V_k)$. If $r = 0$, the inequality is implied by the cluster constraint (15). If, on the other hand, $r > 0$, the node v must also be included in the feasible solution, along with at least r nodes from the cluster V_k . This shows that the inequality is valid. The fact that it also defines a facet can be easily proved for $m = 2$, by exhibiting a suitable collection of affinely independent extreme points satisfying the inequality at equality. The result for general m then follows by Theorem 11. □

It may be that other facet-defining inequalities for $P(G)$ can be lifted in a similar way to obtain facets of $P^{\geq}(G)$.

Finally, we consider the cluster constraints (15). Surprisingly, it turns out that they *never* define facets of $P^{\geq}(G)$. This is an immediate consequence of the following result:

Proposition 20 (Tree Inequalities) *Let $k \in K$ be such that $|V_k| \geq 2$. Let $E(V_k)$ be the set of edges with both end-nodes in V_k and let $T \subset E(V_k)$ be the edge set of a tree spanning the nodes in V_k . Then the ‘tree’ inequality*

$$y(T) \leq x(V_k) - 1 \tag{16}$$

defines a facet of $P^{\geq}(G)$.

Proof. If a feasible solution includes none of the edges in T , we have $y(T) = 0$, and the inequality is implied by the cluster constraint (15). If, on the other hand, the solution includes r of the edges in T , with $r > 0$, it must also include at least $r + 1$ of the nodes in V_k (since the r edges define a forest). This shows that the inequality is valid. The fact that it also defines a facet can be easily proved for $m = 1$, by exhibiting a suitable collection of affinely independent extreme points satisfying the inequality at equality. The result for general m then follows by Theorem 11. \square

Observe that the inequality $y_{uv} \leq x_u + x_v - 1$, mentioned in the proof of Proposition 18, is a tree inequality.

We conjecture that, when $m = 1$, $P^{\geq}(G)$ is completely described by the tree inequalities, the VUBs, the upper bounds $x_v \leq 1$ for all $v \in V$ and the lower bounds $y_e \geq 0$ for all $e \in E$. In general, however, it seems that $P^{\geq}(G)$ is remarkably complex. Using the software PORTA [4], we have found facet-defining inequalities that we have been unable to classify, even for $m = 2$. Thus, $P^{\geq}(G)$ may deserve further study.

6 Concluding Remarks

Up to now, the polyhedra associated with Generalized Network Design Problems (GNDPs) have been studied largely independently. This paper is the first to treat them all in a unified way. Moreover, to our knowledge, we are

the first to study the polyhedra associated with GNDPs in which at least one node per cluster is required.

Of course, the inequalities that we have presented in this paper are not guaranteed to define facets of the polyhedra associated with every GNDP. Even if they do not, however, they may still be useful as cutting planes. Moreover, it should be borne in mind that our inequalities can also be applied to the quadratic semi-assignment problem and various constraint satisfaction problems.

An important topic for future research is the construction of exact or heuristic *separation algorithms* for the odd ring, odd clique and odd circulant inequalities mentioned in Section 4, and the tree inequalities mentioned in Subsection 5.2. Such algorithms would be essential if one wished to use these inequalities within a cutting-plane or branch-and-cut algorithm for GNDPs and related problems.

Acknowledgements: The second author was partially supported by the Communauté française de Belgique under Actions de Recherche Concertées (ARC). The third author was supported by the Engineering and Physical Sciences Research Council under grant EP/D072662/1. The fourth author was supported by the Ministerio de Ciencia e Innovación under grant MTM2009-14039-C06-01.

References

- [1] F. Barahona and A.R. Mahjoub, On the cut polytope, *Math Program*, 36 (1986), 157–173.
- [2] M. Bellare, Interactive proofs and approximation: Reductions from two provers in one round, *Proceedings of 2nd Israel Symposium on Theory and Computing Systems*, IEEE Press, Silver Spring, MD, 1993, pp. 266–274.
- [3] R.E. Burkard and E. Çela, “Quadratic and three-dimensional assignments”, *Annotated Bibliographies in Combinatorial Optimization*, M. Dell’Amico, F. Maffioli & S. Martello (Editors), Wiley, New York, 1997, pp. 373–391.

- [4] T. Christof and A. Loebel, PORTA (polyhedron representation transformation algorithm), Software package, available for download at <http://www.zib.de/Optimization/Software/Porta/>
- [5] C. De Simone, The cut polytope and the boolean quadric polytope, *Discr Math* 79 (1989), 71–75.
- [6] M.M. Deza and M. Laurent, *Geometry of cuts and metrics*, Springer, Berlin, 1997.
- [7] C. Feremans, M. Labbé, and G. Laporte, Generalized network design problems, *Eur J Opl Res* 148 (2003), 1–13.
- [8] C. Feremans, M. Labbé, and G. Laporte, The generalized minimum spanning tree problem: Polyhedral analysis and branch-and-cut algorithm, *Networks* 43 (2004), 71–86.
- [9] M. Fischetti, J.J. Salazar-González, and P. Toth, The symmetric generalized traveling salesman polytope, *Networks* 26 (1995), 113–123.
- [10] A.L. Henry-Labordere, The record balancing problem: A dynamic programming solution of a generalized traveling salesman problem, *Revue Francaise d’Informatique de Recherche Operationnelle B-2* (1969), 43–49.
- [11] S. Khot, On the power of unique 2-prover 1-round games, *Proceedings of 34th Annual ACM Symposium on Theory of Computing*, ACM Press, 2002, pp. 767–775.
- [12] A.M.C.A. Koster, S.P.M. Van Hoesel, and A.W.J. Kolen, The partial constraint satisfaction problem: Facets and lifting theorems, *Oper Res Lett* 23 (1998), 89–97.
- [13] Y.S. Myung, C.H. Lee, and D.W. Tcha, On the generalized minimum spanning tree problem, *Networks* 26 (1995), 231–241.
- [14] G.L. Nemhauser and L.A. Wolsey, *Integer and combinatorial optimization*, Wiley, New York, 1988.
- [15] M.W. Padberg, The boolean quadric polytope: Some characteristics, facets and relatives, *Math Program* 45 (1989), 139–172.

- [16] J.-C. Picard and H.D. Ratliff, Minimum cuts and related problems, *Networks* 5 (1975), 357–370.
- [17] W.R. Pulleyblank and J. Edmonds, “Facets of 1-matching polyhedra”, *Hypergraph Seminar*, C. Berge & D. Ray-Chaudhuri (Editors), Springer, Berlin, 1974, pp. 214-242.
- [18] H. Saito, T. Fujie, T. Matsui, and S. Matuura, A study of the quadratic semi-assignment polytope, *Discr Opt* 6 (2009), 37–50.
- [19] J.P. Saskaena, Mathematical model of scheduling clients through welfare agencies, *J Canadian Oper Res Soc* 8 (1970), 185–200.
- [20] M. Serna, L. Trevisan, and F. Xhafa, The approximability of non-Boolean satisfiability problems and restricted integer programming, *Theor Comput Sci* 332 (2005), 123–139.
- [21] S.S. Srivastava, S. Kumar, R.C. Garg, and P. Sen, Generalized traveling salesman problem through n sets of nodes, *J Canadian Oper Res Soc* 7 (1969), 97–101.

Appendix

Proof of Theorem 3: Validity is trivial. So, assume that the inequality defines a facet of $P(G)$. Then, there exist $|V| + |E| - m$ affinely independent extreme points of $P(G)$ satisfying the inequality at equality. From these it is trivial to construct $|V| + |E| - m$ affinely independent extreme points of $P(G')$ that also satisfy the inequality at equality: just set x_u to 1, and y_e to 0 for all edges incident on u .

One needs an additional $|V|$ affinely independent extreme points of $P(G')$ to prove that the inequality defines a facet of $P(G')$. Note that, for every $v \in V$, there is at least one point in our collection of extreme points of $P(G)$ satisfying $x_v = 1$. (If there were not, then the facet would be contained in the hyperplanes $x_v = 0$ and $y_e = 0$ for all e incident on v , a contradiction.) For each $v \in V$, we construct an additional affinely independent extreme point by taking such an extreme point of $P(G)$ and setting x_u and y_{uv} to 1, and y_e to 0 for all other edges incident on u .

Finally, suppose that the original inequality is valid for $P(G)$, but not facet-defining. Then it is dominated by two or more other valid inequalities, that are valid also for $P(G')$. Therefore, the original inequality cannot define a facet of $P(G')$. \square

Proof of Theorem 4: Validity follows easily from the fact that a GS in G' cannot contain the nodes u and u' simultaneously.

Now assume that the original inequality defines a facet of $P(G)$. Then, there exist $|V| + |E| - m$ affinely independent extreme points of $P(G)$ satisfying the original inequality at equality. From these it is trivial to construct $|V| + |E| - m$ affinely independent extreme points of $P(G')$ that satisfy the cloned inequality at equality: just set $x_{u'}$ to 0, and $y_{u'v}$ to 0 for all edges incident on u' . To show that the cloned inequality defines a facet of $P(G')$, one needs an additional $r + 1$ affinely independent extreme points of $P(G)$, where r is the number of edges incident on u' .

Now suppose that there exists a node $w \neq u$ such that $\alpha_w > 0$. Suppose we take our collection of affinely independent extreme points of $P(G)$, and construct a 0-1 matrix as follows: there are $r + 1$ columns, one representing node u and the other representing edges incident on u . The first column contains the values taken by x_u at each point, and the remaining columns

contain the values taken by the corresponding y variables at each point. Then, the rank of this matrix must be $r + 1$. (If the rank were not $r + 1$, then all of the extreme points would satisfy an equation of the form $\alpha'_u x_u + \sum_{v:\{u,v\} \in E} \beta'_v y_{uv} = \gamma'$, showing that the original inequality either did not define a facet, or had a positive α coefficient only for node u .)

So, let us take $r + 1$ extreme points of $P(G)$ corresponding to a non-singular submatrix of that matrix. For each such extreme point (x^*, y^*) , we can construct an additional affinely independent extreme point of $P'(G)$ by setting x_u to 0, y_e to 0 for all edges incident on u , $x_{u'}$ to x_u^* , and $y_{u'v}$ to y_{uv}^* for all nodes v incident on u' .

Finally, suppose that the original inequality is valid for $P(G)$, but not facet-defining. Then it is dominated by two or more other valid inequalities. The cloned version of the original inequality is dominated by the cloned versions of these other valid inequalities, and therefore it cannot define a facet of $P(G')$. \square

Proof of Theorem 6: From Theorems 3 and 4, we can assume that $V = V_1 \cup \dots \cup V_c$ and that $|V_k^1| = |V_k^2| = 1$ for $k = 1, \dots, c$. Then, we can simplify notation by assuming that $V_k^1 = \{v_k^1\}$ and $V_k^2 = \{v_k^2\}$ for $k = 1, \dots, c$. Also let \bar{x}_u denote $1 - x_u$ for any $u \in V$. The odd ring inequality can then be written in the following form:

$$\sum_{k \in C \setminus D} y(v_k^1, v_{k+1}^1) + \sum_{k \in D} y(v_k^1, v_{k+1}^2) + \sum_{k \in C \setminus D} \bar{x}(v_{k+1}^1) \leq \lfloor (2c - |D|)/2 \rfloor. \quad (17)$$

Note that there are $2c - |D|$ terms on the left-hand side of (17). Let G^* be a graph with one node for each of those terms, and an edge between two nodes if and only if the corresponding terms cannot both take the value 1 simultaneously. Then G^* is a chordless circuit of odd cardinality, i.e., an odd hole, and extreme points of $P(G)$ satisfying the inequality (17) at equality correspond to maximum cardinality stable sets in G^* .

Since G^* is an odd hole, there exist $2c - |D|$ maximum cardinality stable sets in G^* . For each such stable set, construct an extreme point of $P(G)$ by setting the corresponding terms in (17) to 1, setting x_u to 1 for any node incident on an edge with $y_e = 1$, setting the remaining x variables (if any) to arbitrary feasible values, and setting y_e to 0 for all edges not involved in the inequality. The resulting $2c - |D|$ extreme points of $P(G)$ satisfy the

inequality (17) at equality and are affinely independent.

Now note that, for any given $k \in D$, there exists exactly one extreme point in the above-mentioned collection of extreme points that has $y_e = 0$ for every edge with an end-node in V_k . Taking such an extreme point and exchanging the values of $x(v_k^1)$ and $x(v_k^2)$, we obtain an additional affinely independent extreme point satisfying the odd ring inequality at equality. There are $|D|$ such points.

Next, note that, if we take any two nodes i, j of an odd hole, there exists a maximum cardinality stable set including both i and j , another including i but not j , a third including j but not i , and a fourth including neither i nor j . This means that, given any two nodes u, w in G belonging to different clusters, there exists at least one extreme point of $P(G)$ that satisfies the inequality (17) at equality and has $x_u = x_w = 1$. Taking such an extreme point, setting y_{uw} to 1 and setting y_e to 0 for all other edges not involved in the inequality yields an additional affinely independent point. There are $c(2c - 3)$ such points.

The total number of points obtained is $c(2c - 1)$, which is equal to the dimension of $P(G)$ under the stated conditions. \square

Proof of Theorem 7: As in the proof of Theorem 6, we can assume that $V = V_1 \cup \dots \cup V_c$ and that $V_k^1 = \{v_k^1\}$ and $V_k^2 = \{v_k^2\}$ for $k = 1, \dots, c$.

Now, Padberg [15] showed that an extreme point of the boolean quadric polytope satisfies a clique inequality (10) at equality if and only if $\sum_{k=1}^c z_k \in \{r, r + 1\}$. From this it follows that an extreme point of $P(G)$ satisfies the odd clique inequality (9) at equality if and only if it has the following two properties:

- $\sum_{k=1}^c x(v_k^1) \in \{r, r + 1\}$
- $y(v_k^1 : v_\ell^1) = x(v_k^1)x(v_\ell^1)$ for all k and ℓ .

Such extreme points will be called *roots*.

Now, suppose that the odd clique inequality does not define a facet. Then there is a valid inequality $\beta y \leq \alpha x + \gamma$, not equivalent to the odd clique inequality, such that every root satisfies $\beta y = \alpha x + \gamma$. From Theorem 2, we can assume that $\alpha(v_k^2) = 0$ for $k = 1, \dots, c$. Moreover, we must have $\beta_e = 0$ for every edge e having a zero coefficient in the odd clique inequality.

(This is so because, for each such edge, there exists a root having $y_e = 1$. By changing the value of y_e from 1 to 0, we obtain another root.)

Now, for a given $k \in \{1, \dots, c\}$, let (x^*, y^*) be an arbitrary root such that $\sum_{\ell=1}^c x^*(v_\ell^1) = r$ and $x^*(v_k^1) = 0$, and let (x', y') be an arbitrary root with $x'(v_\ell^1) = x^*(v_\ell^1)$ for all $\ell \neq k$, but with $x'(v_k^1) = 1$. Define the two sets

$$S = \{\ell \in \{k+1, \dots, k + \lfloor c/2 \rfloor\} : x^*(v_\ell^2) = 1\}$$

and

$$T = \{\ell \in \{k - \lfloor c/2 \rfloor, \dots, k-1\} : x^*(v_\ell^1) = 0\}.$$

A comparison of the two roots shows that

$$\alpha(v_k^1) + \sum_{\ell \in S} \beta(v_k^1 : v_\ell^2) = \sum_{\ell \in T} \beta(v_k^2 : v_\ell^1). \quad (18)$$

Since this holds for all possible S and T , we must have:

$$\beta(v_k^1 : v_{k+s}^2) = \beta(v_k^2 : v_{k-t}^1) \quad (\forall s = 1, \dots, \lfloor c/2 \rfloor, t = 1, \dots, \lfloor c/2 \rfloor).$$

This implies, by rotational symmetry, that β_e is a constant for all edges having a non-zero coefficient in the odd clique inequality. Then, setting $|S| = 0$ and $|T| = \lfloor c/2 \rfloor - r$ in (18), we see that $\alpha(v_k^1)$ must equal $\lfloor c/2 \rfloor - r$ times that constant, for all k . The inequality $\beta y \leq \alpha x + \gamma$ is therefore equivalent to or dominated by the odd clique inequality, which is a contradiction. \square

Proof of Theorem 8: From Theorems 3 and 4, we can assume that $V = V_1 \cup \dots \cup V_c$ and that $V_k^\ell = \{v_k^\ell\}$ for $k = 1, \dots, c$ and $\ell = 1, \dots, 2d$. We define an auxiliary graph $G' = (V', E')$, which is the graph with node set $\{1, \dots, c\}$ and an edge between nodes i and j if and only if there is an edge in G between V_i and V_j having a positive coefficient in the odd circulant inequality. Note that each extreme point of $P(G)$ that satisfies the inequality at equality corresponds to a maximum cardinality matching in G' .

We will need the following four properties of G' :

1. For any $k \in V'$, there exist a maximum cardinality matching in G' that has no edge incident on node k .
2. For any pair of edges in E' having a common end-node, there exists a

maximum cardinality matching in G' containing one of them.

3. For any pair of node-disjoint edges in E' , there exists a maximum cardinality matching in G' containing both of them.
4. One can find $|E'|$ maximum cardinality matchings in G' whose incidence vectors are affinely independent.

The first property is obvious. The fourth property follows from the characterization of the facets of the matching polytope given in [17]. The other two properties can be easily proved either by induction or by brute-force enumeration of cases.

Now, suppose that the equation $\alpha x + \beta y = \gamma$ is satisfied by every extreme point of $P(G)$ satisfying the odd circulant inequality at equality. Using the cluster constraints (1), we can assume that $\alpha(v_k^1) = 0$ for $k = 1, \dots, c$. Now, the first property of G' implies that, for any $1 \leq k \leq c$, there exists an extreme point of $P(G)$ that satisfies the odd circulant inequality at equality, has $x(v_k^1) = 1$, and has $y_e = 0$ for every edge e with an end-node in V_k . Note that, for any $1 < \ell \leq 2d$ we can set $x(v_k^1)$ to 0 and set $x(v_k^\ell)$ to 1, while still satisfying the inequality at equality. This implies that $\alpha_v = 0$ for all $v \in V$.

Now consider an arbitrary edge e in G that has a coefficient of zero in the inequality, and let v_i^k and v_j^ℓ be its end-nodes. If we set y_e to 1, we force $x(v_i^k)$ and $x(v_j^\ell)$ to be 1. Moreover, there is a unique edge incident on v_i^k that has a coefficient of 1 in the inequality, and the same for v_j^ℓ . Let e' and e'' denote these edges. We now consider two cases:

- The edges e' and e'' have a common end-node. In this case, by the second property of G' , there exists an extreme point of $P(G)$ that has $y_e = y_{e'} = y_{e''} = 1$ and satisfies the inequality at equality. Since we can change y_e to 0 and still satisfy the inequality at equality, we must have $\beta_e = 0$.
- The edges e' and e'' are node disjoint. In this case, by the third property of G' , there exists an extreme point of $P(G)$ that has $y_e = y_{e'} + y_{e''} = 1$ and satisfies the inequality at equality. As in the first case, this implies that $\beta_e = 0$.

The equation $\alpha x + \beta y = \gamma$ therefore has a non-zero coefficient only for the edges that have a non-zero coefficient in the inequality. Now, by the

fourth property of G' , the inequality $\alpha x + \beta y = \gamma$ must be equivalent to the odd circulant inequality. \square

Proof of Theorem 10: Let M be the constraint matrix defined by the VUBs, upper bounds and lower bounds. The entries of M take the values 0, 1 and -1 , and no row of M contains two non-zero coefficients of the same sign. Therefore M is totally unimodular, and the constraints mentioned define an integral polytope.

Now, consider the problem of optimizing a linear function $\alpha x + \beta y$ over $P^{\leq}(G)$, and assume without loss of generality that one is maximizing. If $\beta_{uv} \leq 0$ for some edge $\{u, v\}$, then it is optimal to set y_{uv} to 0. If on the other hand $\beta_{uv} > 0$, then y_{uv} will equal $x_u x_v$ in any optimal solution. Therefore, maximizing $\alpha x + \beta y$ over $P^{\leq}(G)$ is equivalent to maximizing the following quadratic function of the x variables:

$$\sum_{v \in V} \alpha_v x_v + \sum_{\{u, v\} \in E'} \beta_{uv} x_u x_v,$$

where $E' = \{\{u, v\} \in E : \beta_{uv} > 0\}$. Now, maximizing a quadratic function in m binary variables, in the special case where all quadratic terms are non-negative, can be reduced to a max-flow/min-cut problem in a graph with $m + 2$ nodes (Picard & Ratliff [16]). \square

Proof of Theorem 11: This is similar to that of Theorem 3, but a more involved argument is needed to prove that the inequality defines a facet of $P^{\geq}(G')$ if it defines a facet of $P^{\geq}(G)$.

So, assume that the inequality defines a facet of $P^{\geq}(G)$. Then, there exist $|V| + |E|$ affinely independent extreme points of $P^{\geq}(G)$ satisfying the inequality at equality. We construct $|V| + |E|$ corresponding extreme points of $P^{\geq}(G')$ by setting x_u to 1 for some $u \in V_{m+1}$, setting $x_{u'}$ to 0 for all $u' \in V_{m+1} \setminus \{u\}$, and setting y_e to 0 for all edges incident on at least one node in V_{m+1} .

We now construct k additional extreme points, one for each node in V_{m+1} . To this end, let (x^*, y^*) be one of the $|V| + |E|$ extreme points of $P^{\geq}(G')$ defined above, and let u be the node in V_{m+1} for which $x_u^* = 1$. We construct $k - 1$ additional extreme points as follows: for each node $u' \in V_{m+1} \setminus \{u\}$, we take (x^*, y^*) and change the value of $x_{u'}$ from 0 to 1.

To construct one additional extreme point, we change x_u to 0 and change u' to 1 for some $u' \in V_{m+1} \setminus \{u\}$.

Next, we construct $k|V|$ additional extreme points, one for each edge in $E(V : V_{m+1})$. As in the proof of Theorem 3, for every $v \in V$, there is at least one extreme point of $P^{\geq}(G)$ in our collection that satisfies $x_v = 1$. We take such an extreme point and, for each $u \in V_{m+1}$, convert it to an extreme point of $P^{\geq}(G')$ by setting x_u and y_{uv} to 1, $x_{u'}$ to 0 for all $u' \in V_{m+1} \setminus \{u\}$, and y_e to 0 for all other edges incident on at least one node in V_{m+1} .

Finally, we construct $\binom{k}{2}$ additional extreme points, one for each edge in $E(V_{m+1})$. To this end, for any $\{u, u'\} \in E(V_{m+1})$, take one of the extreme points of $P^{\geq}(G)$ in our collection, set $x_u, x_{u'}$ and $y_{uu'}$ to 1, set $x_{\tilde{u}}$ to 0 for all $\tilde{u} \in V_{m+1} \setminus \{u, u'\}$, and set y_e to 0 for all other edges incident on at least one node in V_{m+1} . \square