Discrete Applied Mathematics 194 (2015) 65–80

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

The precedence constrained knapsack problem: Separating maximally violated inequalities

Daniel Espinoza[a,*], Marcos Goycoolea[b], Eduardo Moreno[c]

[a] Department of Industrial Engineering, Universidad de Chile, Chile
[b] School of Business, Universidad Adolfo Ibañez, Chile
[c] Faculty of Engineering and Sciences, Universidad Adolfo Ibañez, Chile

ARTICLE INFO

Article history:
Received 10 January 2013
Received in revised form 28 April 2015
Accepted 13 May 2015
Available online 4 June 2015

Keywords:
Lifting
Shrinking
Precedence-constrained knapsack problem
Induced cover inequality
Induced clique inequality
Separation problem

ABSTRACT

We consider the problem of separating maximally violated inequalities for the precedence constrained knapsack problem. Though we consider maximally violated constraints in a very general way, special emphasis is placed on induced cover inequalities and induced clique inequalities. Our contributions include a new partial characterization of maximally violated inequalities, a new safe shrinking technique, and new insights on strengthening and lifting. This work follows on the work of Boyd (1993), Park and Park (1997), van de Leensel et al. (1999) and Boland et al. (2011). Computational experiments show that our new techniques and insights can be used to significantly improve the performance of cutting plane algorithms for this problem.

1. Introduction

Given a directed graph $G = (V, A)$, vectors $a \in \mathbb{Z}_+^V$, $c \in \mathbb{Z}^V$, and $b \in \mathbb{Z}_+$, the precedence-constrained knapsack problem (PCKP) consists in solving a problem of the form,

$$\begin{array}{ll} \max & cx \\ s.t. & x \in P(G, a, b) \\ & x \in \{0, 1\}^V \end{array}$$

where

$$P(G, a, b) = \left\{x \in [0, 1]^V : ax \le b,\ x_i \le x_j\ \forall (i, j) \in A\right\}.$$

Aside from being an interesting problem in itself, PCKP is an important substructure of many common, more complex, integer programming problems. Important examples arise in the context of production scheduling problems, where a number of jobs must be scheduled for processing subject to limited resources, and where precedence relationships dictate that in order for some jobs to be processed other jobs must be processed as well. An example of such a problem that has received much attention in recent years is the open-pit mine production scheduling problem. In this problem, jobs represent discretized units of rock that must be extracted, and precedence relationships establish that units must be extracted from the surface on downwards. Integer programming formulations for open pit mine production scheduling problems having

* Corresponding author.
E-mail addresses: daespino@dii.uchile.cl (D. Espinoza), marcos.goycoolea@uai.cl (M. Goycoolea), eduardo.moreno@uai.cl (E. Moreno).

http://dx.doi.org/10.1016/j.dam.2015.05.020

PCKP as a substructure appear as early as in the 1960s and 1970s in the works of Johnson [16], Dagdelen and Johnson [13] and others. Since then, a number of articles have addressed PCKP specifically. These articles can broadly be subdivided in two groups: Those articles which develop algorithms to solve PCKP as an optimization problem, and those articles that analyze the polyhedral structure of $P(G, a, b) \cap \{0, 1\}^n$.

Practical open pit mining problems are very large, easily having tens or hundreds of millions of variables. As shown by Johnson and Niemi [17], however, PCKP is strongly $\mathcal{NP}$-complete. Thus, efforts to solve this problem to provable optimality gaps have either been based on approximation algorithms [11,21] or on enumerative methods such as branch-and-bound. An important component of any branch-and-bound solver is the LP relaxation solver. Much progress has been made in recent years solving the LP relaxations of PCKP generalizations with customized algorithms. Important examples include the work of Chicoisne et al. [8] and Bienstock and Zuckerberg [3]. Heuristics for these problems that use PCKP as a subproblem, or that use the PCKP linear programming relaxation as a subproblem, are described in Amaya et al. [1], Bley et al. [4], Chicoisne et al. [8] and Cullenbine et al. [12].

Polyhedral analyses, on the other hand, have focused on developing useful cutting planes. Boyd [7], Park and Park [20] and van de Leensel et al. [22] describe characterizations, separation algorithms, and strengthening techniques for an important class of cutting planes known as induced minimal cover inequalities. Boland et al. [5] extend previous results using clique inequalities. Bley et al. [4] test many of these ideas on open pit mine production scheduling problems. Despite important work in this problem, cutting plane techniques to date are still limited in terms of the instance sizes that can effectively be tackled computationally. This is a problem, because, if there is any hope of being able to solve large integer programming formulations of PCKP generalizations, such as those that appear in the context of open pit mining, cutting planes are likely to play an important role.

In this article, given a fractional point $x^* \in P(G, a, b)$, we are interested in efficiently finding an inequality $\alpha x \leq \beta$ that is valid for all x in $P(G, a, b) \cap \{0, 1\}^n$ and maximally violated by x^*. In the context of this paper, we will say that a valid inequality $\alpha x \leq \beta$ is maximally violated by x^* if it maximizes $\frac{(\alpha x - \beta)}{\|\alpha\|_1}$, where $\|\cdot\|_1$ represents the one-norm in $\mathbb{R}^n$. Specifically, we are interested in extending the work of Boyd [7], Park and Park [20], van de Leensel et al. [22] and Boland et al. [5] so as to tackle significantly larger instances of PCKP and its generalizations.

For this we introduce new shrinking techniques that can be used to reduce the separation problem in any given instance of PCKP to an equivalent separation problem in a smaller instance. This shrinking procedure is safe in the sense that it guarantees that the most violated cuts in the original problem can be mapped to equally violated cuts in the shrunken problem, and vice-versa. Moreover, within this shrunken graph, we identify a very small set of nodes in which the support of maximally violated constraint coefficients must be contained. Finally, we introduce a new way of strengthening general valid inequalities for PCKP, and remark how the lifting techniques of Park and Park [20], originally proposed for minimal induced cover inequalities, can be generalized to broader classes of inequalities.

This article is organized as follows. In Section 2 we review important results from the literature and introduce the notation we will use throughout the paper. In Section 3 we characterize maximally violated inequalities and introduce the concept of *break-points*, which will be used throughout the paper. In Section 4 we show how to shrink the original graph in order to find maximally violated inequalities in a smaller problem. Moreover, we show that, if this shrinking operations is obtained using break-points, then we can map maximally violated inequalities obtained for the shrunken problem to maximally violated inequalities for the original one. In Section 5 we show that it is possible to obtain even further reductions. In Section 6 we show how to obtain strengthened inequalities by using lifting procedures. Finally, in Section 7, we present computational results that show the usefulness of the proposed methodologies.

2. Definitions, assumptions, and background material

In this section we establish the assumptions and notation that will be used throughout the article. More important, we survey some important prior results concerning two classes of valid inequalities for PCKP, the minimal induced cover inequalities and the clique inequalities. These results, which can be found in previous literature, will be the starting point for our developments in later sections.

Definition 1. Consider a directed graph $G = (V, A)$ with no directed cycles. We say that $C \subseteq V$ is a *closure* in G if $i \in C$ implies $j \in C$ for all $(i, j) \in A$. That is, if set C is closed under the precedence relationships defined by graph G. Given any set $S \subseteq V$, we define the smallest closure containing S in graph $G = (V, A)$ as follows:

$$cl(G, S) = S \cup \{j \in V : \text{ there is a path in } G \text{ from some } i \in S \text{ to } j\}.$$

To simplify notation, when graph G is clear from context we will write $cl(S)$ instead of $cl(G, S)$. For $i \in V$, we will write $cl\ (i)$ instead of $cl\ (\{i\})$.

Definition 2. Consider a directed graph $G = (V, A)$ with no directed cycles. We say that $R \subseteq V$ is a *reverse closure* in G if $j \in R$ implies $i \in R$ for all $(i, j) \in A$. That is, if set R is closed under the reverse of precedence relationships defined by graph G. Given any set $S \subseteq V$, we define the smallest reverse closure containing S in graph $G = (V, A)$ as follows:

$$rcl(G, S) = S \cup \{i \in V : \text{there is a path in } G \text{ from } i \text{ to some } j \in S\}.$$

To simplify notation, when graph G is clear from context we will write $rcl(S)$ instead of $rcl(G, S)$. For $i \in V$, we will write rcl (i) instead of rcl $(\{i\})$.

Definition 3. We say that (G, a, b, c) defines an instance of PCKP if,

- $G = (V, A)$ is a directed graph,
- $a \in \mathbb{Z}^V$ represents non-negative node-weights,
- $b \in \mathbb{Z}$ is a non-negative scalar and
- $c \in \mathbb{Z}^V$ represents the cost of each $v \in V$.

Given a set $C \subseteq V$ and a vector $y \in \mathbb{R}^V$ we will convene that $y(C) = \sum_{i \in C} y_i$. Since we are only interested in the feasible set of an instance of PCKP, we note (G, a, b) throughout the document, omitting the c parameter. Also, in order to simplify notation and proofs, we will make the following working assumptions regarding instances of PCKP.

Definition 4. We say that an instance of PCKP (G, a, b) satisfies our working assumptions if,

- $G = (V, A)$ has no directed cycles,
- For every $i, j, k \in V$, $(i, j), (j, k) \in A$ implies $(i, k) \in A$,
- There are no arcs of the form (i, i) with $i \in V$,
- $a(cl(i)) \le b$ for all $i \in V$ and
- $a(V) > b$.

Note that there is no loss of generality from our working assumptions. In fact, if G were not acyclic, we could iteratively collapse all of the variables associated with nodes in any directed cycle into a single variable. In this way we would obtain an equivalent instance of PCKP. Adding (or removing) arcs to satisfy the second and third conditions does not change the solution space of PCKP, nor does it change the set $P(G, a, b)$. If $a(cl(k)) > b$ for $k \in V$, then $x_k = 0$ for all $x \in P(G, a, b) \cap \{0, 1\}^n$. In this case variable x_k may as well be eliminated. If $a(V) \le b$, then the convex hull of $P(G, a, b) \cap \{0, 1\}^n$ is fully described by inequalities $x_i \le x_j$ for $(i, j) \in A$ and $x_i \ge 0$ for $i \in V$. This follows from the total unimodularity of the precedence constraints matrix. Thus, if $a(V) \le b$ there are no interesting valid inequalities to study. Finally observe that under our working assumptions, sets $cl(S)$ and $rcl(S)$ can be defined as follows,

$$cl(S) = S \cup \{j \in V : (i, j) \in A \text{ for some } i \in S\} \text{ and}$$
$$rcl(S) = S \cup \{i \in V : (i, j) \in A \text{ for some } j \in S\}.$$

Theorem 5 (Boyd, [7]). *If (G, a, b) defines an instance of PCKP satisfying our working assumptions, then $P(G, a, b)$ is full dimensional.*

Given an instance of PCKP (G, a, b) satisfying our working assumptions, we focus on finding strong inequalities $\alpha x \le \beta$ that are valid for $P(G, a, b) \cap \{0, 1\}^V$. In what follows, we formalize our notion of strong valid inequalities.

Definition 6. Let (G, a, b) define an instance of PCKP satisfying our working assumptions. Consider two inequalities $\alpha x \le \beta$, and $\alpha' x \le \beta'$, both valid for $P(G, a, b) \cap \{0, 1\}^V$. We say that inequality $\alpha' x \le \beta'$ is *stronger* than inequality $\alpha x \le \beta$ if (a) for every $\hat{x} \in P(G, a, b)$ we have that $\alpha' \hat{x} \le \beta'$ implies $\alpha \hat{x} \le \beta$, and (b) there exists $\hat{x} \in P(G, a, b)$ such that $(\alpha' \hat{x} - \beta')/\|(\alpha', \beta')\| < (\alpha \hat{x} - \beta)/\|(\alpha, \beta)\|$.[1]

From a polyhedral analysis point of view it is only natural to try and generalize known families of valid inequalities for the knapsack problem (KP) to PCKP. The first to do so was Boyd [7], who, extending the cover inequalities of Balas and Jeroslow [2] and the $(1, k)$-configuration inequalities of Padberg [18], first introduced the notion of induced cover inequalities.

Definition 7. Let (G, a, b) define an instance of PCKP satisfying our working assumptions. We say that $C \subseteq V$ is an *induced cover* of (G, a, b) if $a(cl(C)) > b$. If C is an induced cover, then the valid inequality

$$x(C) \le |C| - 1 \tag{1}$$

is known as the *induced cover inequality* associated with C. In this document we define an induced cover C to be *minimal* if $a\,(cl(C \setminus \{i\})) \le b$ for all $i \in C$. If C is a minimal induced cover, we say that inequality (1) is a *minimally induced cover inequality*, or MIC inequality for short.

Note that the definition of MIC inequalities that we use coincides with that used by Park and Park [20]. However, Boyd [7] and van de Leensel et al. [22], require that an induced cover $C \subseteq V$ satisfies $a\,(cl(C) \setminus \{i\}) \le b$ for all $i \in C$ in order to be minimal.

[1] An example of this situation is when the face induced in $P(G, a, b)$ by (α', β') strictly contains the face induced in $P(G, a, b)$ by (α, β); in this case, there exists $\hat{x} \in P(G, a, b)$ satisfying $\alpha' \hat{x} = \beta'$ and $\alpha \hat{x} < \beta$.

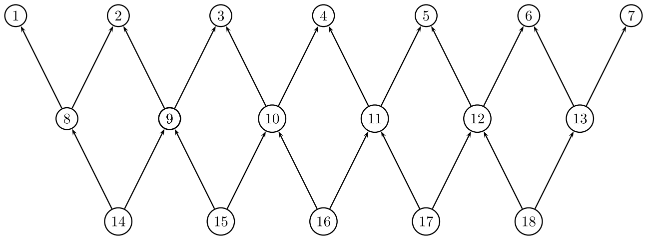

Fig. 1. Example PCKP: $a_i = 1,\ i = 1, \dots, 18,\ b = 10$, note that $C = \{15, 16, 17\}$ is a MIC, but not a facet for the problem. In particular, x_3, x_4, x_5 must be one whenever $x_{15} + x_{16} + x_{17} = 2$.

Remark. Let (G, a, b) represent an instance of PCKP satisfying our working assumptions. If C is a minimal induced cover, then $i, j \in C$ implies $(i, j) \notin A$. In fact, if $i, j \in C$ and $(i, j) \in A$, then $cl(C) = cl(C - \{j\})$, thus contradicting the minimality of C.

As observed by Bley et al. [4], given $x^* \in P(G, a, b)$ it is possible to find a maximally violated MIC inequality by solving a simple integer programming problem.

Given any induced cover C, there always exists a minimal induced cover $C' \subseteq C$. Further, if C' is an induced cover and $C' \subsetneq C$, then the induced cover inequality associated with C' is stronger than that associated with C. As observed by Boyd, this does not imply that minimal induced cover inequalities are facet-defining. To see this, consider the example illustrated in Fig. 1 and the MIC inequality $x_{15} + x_{16} + x_{17} \le 2$. All feasible points satisfying this inequality at equality also satisfy $x_{10} = x_{11} = 1$. Since we know that the feasible region is full dimensional (Theorem 5), it follows that the inequality cannot be facet defining. Boyd [7] and later Park and Park [20] describe conditions under which MIC inequalities are facet-defining.

In the case that a MIC inequality is not facet-defining it can be strengthened through lifting techniques to obtain a constraint of the form

$$\sum_{i\in C} x_i + \sum_{i\in cl(C)} \gamma_i(1 - x_i) + \sum_{j\in V\setminus cl(C)} \eta_j x_j \le |C| - 1, \tag{2}$$

where η, γ are non-negative vectors in $\mathbb{R}^V$. Going back to the example illustrated in Fig. 1, we have that, if we strengthen the MIC inequality $x_{15}+x_{16}+x_{17} \le 2$ on x_{10} and then x_{11}, we obtain the facet-defining inequality $x_{15}+x_{16}+x_{17}-x_{10}-x_{11} \le 0$.

Park and Park [20] showed that, given a minimally induced cover C, it is possible to lift the corresponding MIC inequality on variables x_i with $i \in cl(C)$ in polynomial time. Moreover, they showed that this was possible in such a way as to guarantee that the resulting inequality was maximally violated. Furthermore, van de Leensel et al. [22] showed that it was possible to speed-up the algorithm of Park and Park to run in $\mathcal{O}(|V| \log^* |V|)$. Park and Park [20] also introduced a heuristic for lifting x_i with $i \in rcl(C)$. Van de Leensel et al. [22] studied the optimal lifting problem associated with these variables as well as the conditions required to ensure that the resulting inequality would be facet-defining. Van de Leensel et al. [22] also proposed a pseudo-polynomial time algorithm for the case in which G is a tree.

Definition 8. Let (G, a, b) define an instance of PCKP satisfying our working assumptions. We say that $C \subseteq V$ is an *induced clique* of (G, a, b) if $a(cl(\{i, j\})) > b$ for all $i, j \in C$ such that $i \ne j$. If C is an induced clique, then the valid inequality

$$x(C) \le 1, \tag{3}$$

is known as the *induced clique inequality* associated with C.

Proposition 9. *(Boland et al. [5]) Let (G, a, b) define an instance of PCKP satisfying our working assumptions. If C is an induced clique in (G, a, b), then $i, j \in C$ implies $(i, j) \notin A$. Further, $i, j \in C$ implies that there is no $k \in V$ such that $(k, i) \in A$ and $(k, j) \in A$.*

Proof. First, suppose $i, j \in C$ and $(i, j) \in A$. Since $i, j \in C$ we know that $a(cl(\{i, j\})) > b$. However, $(i, j) \in A$ implies $cl(j) \subseteq cl(i)$, thus $cl(\{i, j\}) \subseteq cl(\{i\})$. This in turn implies $a(cl(\{i\})) > b$ which contradicts our working assumptions. Second, suppose $i, j \in C$ and that there exists $k \in V$ such that $(k, i) \in A$ and $(k, j) \in A$. In this case $cl\{i, j\} \subseteq cl(k)$. However, in this case, $a(cl(\{i, j\})) > b$ implies $a(cl(\{k\})) > b$, which contradicts our working assumptions. □

Boland et al. [5] observe that it is possible to obtain facet-defining clique inequalities by sequentially up-lifting variables from covers of size 2. Based on this observation they propose a polynomial-time running heuristic that can generate strengthened induced clique inequalities.

3. Maximally violated inequalities

In this section we introduce some tools for characterizing maximally violated inequalities. As we will see, these tools also allow us to characterize maximally violated inequalities belonging to specific classes of constraints, such as MIC and Clique Inequalities, and their lifted variants.

Definition 10. Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. We say that an inequality $\alpha x \leq \beta$ valid for $P(G, a, b) \cap \{0, 1\}^V$ is maximally violated by x^* if (α, β) maximizes $\frac{(\alpha x^* - \beta)}{\|\alpha\|_1}$.

Before further discussing maximally violated inequalities we introduce the notion of *break-points* and *weight-balanced inequalities*.

Definition 11. Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. Define

$$B[x^*] = \{j \in V : x_i^* < x_j^*,\ \forall (i, j) \in A\}.$$

We say that $B[x^*]$ is the set of precedence *break-points* associated with G and x^*. If $j \in V$ is such that there is no arc $(i, j) \in A$, then we assume $j \in B[x^*]$. For each $i \in V$ define,

$$B[x^*, i] = B[x^*] \bigcap rcl(i).$$

Proposition 12. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$.*

- *For every $i \notin B[x^*]$, there exists $j \in B[x^*, i]$ such that $x_j^* = x_i^*$.*
- *If $i, j \in V$ are such that $B[x^*, i] = B[x^*, j]$ then $x_i^* = x_j^*$.*

Proof. We only need to prove that $x_i^* = \max_{j \in B[x^*, i]} x_j^*$. Since $B[x^*, i] \subseteq rcl(i)$, then $x_i^* \geq x_j^*$ for all $j \in B[x^*, i]$. If $i \in B[x^*]$ then the result holds trivially. Suppose that $i \notin B[x^*]$ then there exists $j \in rcl(i)$ such that $x_i^* = x_j^*$. If $j \notin B[x^*]$ then applying this argument recursively we find a $j' \in B[x^*, i]$ such that $x_i^* = x_{j'}^*$, concluding the proof. □

Definition 13. Let (G, a, b) define an instance of PCKP satisfying our working assumptions. We say that an inequality $\alpha x \leq \beta$ is *weight-balanced* with respect to G and x^* if $\alpha_i \leq 0$ for all $i \notin B[x^*]$.

The following proposition describes an important characterization of violated inequalities that we will later see is very helpful for separating cliques and MICs.

Proposition 14. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions. Consider $x^* \in P(G, a, b)$ and an inequality $\alpha x \leq \beta$ valid for $P(G, a, b) \cap \{0, 1\}^V$. There exists an inequality $\bar{\alpha} x \leq \beta$, valid for $P(G, a, b) \cap \{0, 1\}^V$, such that (1) $\bar{\alpha} x \leq \beta$ is weight-balanced with respect to x^*, (2) $\alpha x^* = \bar{\alpha} x^*$, and (3) $\|\bar{\alpha}\|_1 \leq \|\alpha\|_1$.*

Proof. Suppose $\alpha x \leq \beta$ defines a valid inequality for $P(G, a, b)$ that is not weight-balanced with respect to x^*. By definition, because $\alpha x \leq \beta$ is not weight-balanced, there must exist $i \notin B[x^*]$ such that $\alpha_i > 0$. From Proposition 12 we know there exists $j \in B[x^*, i]$ such that $x_i^* = x_j^*$.

First, since $x_j \leq x_i$ is valid for $P(G, a, b) \cap \{0, 1\}^V$, it follows that $\alpha x + \alpha_i(x_j - x_i) \leq \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$ as well. Second, note that $\alpha x^* = \alpha x^* + \alpha_i(x_j^* - x_i^*) = \bar{\alpha} x^*$. Finally, note that $\|\alpha'\|_1 \leq \|\alpha\|_1$. In fact,

$$\begin{aligned} \|\alpha'\|_1 = \sum_k |\alpha'_k| &= \left(\sum_{k \neq i,j} |\alpha'_k|\right) + |\alpha'_i| + |\alpha'_j| \\ &= \left(\sum_{k \neq i,j} |\alpha_k|\right) + 0 + |\alpha_i + \alpha_j| \leq \|\alpha\|_1. \end{aligned}$$

Now $\bar{\alpha}$ can be obtained from α by the following recursive procedure: If α is not weight-balanced with respect to x^*, let $i \notin B[x^*]$ be such that $\alpha_i > 0$ and let $j \in B[x^*, i]$ such that $x_i^* = x_j^*$. Define α' such that $\alpha' x = \alpha x + \alpha_i(x_j - x_i)$. If α' is not weight-balanced, redefine α as α' and repeat. Otherwise, set $\bar{\alpha}$ to α: the recursion is complete.

The recursion terminates because

$$|\{i \notin B[x^*] : \bar{\alpha}_i > 0\}| < |\{i \notin B[x^*] : \alpha_i > 0\}|. \quad \square$$

The following theorem follows directly from Proposition 14. As we will see, this is a very practical result for separating valid inequalities.

Theorem 15. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. There exists a maximally violated inequality with respect to x^* that is weight-balanced with respect to x^*.*

A point of concern regarding Theorem 15 is that it says nothing concerning specific classes of inequalities. For example, if $x(C) \leq |C| - 1$ is a maximally violated MIC inequality that is not weight-balanced, it is possible that after applying the weight-balancing procedure described in Proposition 14, we might obtain an inequality that is not a MIC inequality. The same concern applies to induced clique inequalities. Is it also true that within specific classes of inequalities such as cliques and MICs, there are maximally violated inequalities that are weight-balanced? This is addressed in the following propositions.

Proposition 16. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. Let C be a minimally induced cover. There exists a minimally induced cover $\bar{C} \subseteq B[x^*]$ such that (a) constraint $x(\bar{C}) \leq |\bar{C}| - 1$ is weight-balanced and (b) $x(\bar{C}) \leq |\bar{C}| - 1$ is at least as violated as $x(C) \leq |C| - 1$.*

Proof. Since C is a minimal induced cover we know $a(cl(C)) > b$. From Proposition 12 we know that for all $i \in C$ there exists $j(i) \in B[x^*, i]$ such that $x^*_i = x^*_{j(i)}$. Let $C' = \{j(i) : i \in C\}$. Observe that $|C'| \leq |C|$. Further, since either $j(i) = i$ or $(j(i), i) \in A$ for all $i \in C$ we know that $cl(C) \subseteq cl(C')$. Since $a \geq 0$ it follows that C' is an induced cover. Let $\bar{C} \subseteq C'$ be a minimal induced cover. It follows that $|\bar{C}| \leq |C|$, and, $\sum_{i\in\bar{C}}(1 - x^*_i) \leq \sum_{i\in C'}(1 - x^*_i) \leq \sum_{i\in C}(1 - x^*_{j(i)}) = \sum_{i\in C}(1 - x^*_i)$. Hence, $x^*(\bar{C}) - (|\bar{C}| - 1) \geq x^*(C) - (|C| - 1)$, and so, $\frac{x^*(\bar{C})-(|\bar{C}|-1)}{|\bar{C}|} \geq \frac{x^*(C)-(|C|-1)}{|C|}$. Further, since $\bar{C} \subseteq B[x^*]$ and all coefficients of inequality $x(\bar{C}) \leq |\bar{C}| - 1$ are positive, we conclude that it is weight-balanced. □

Proposition 17. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. Consider an induced clique $C \subseteq V$. There exists an induced clique $\bar{C} \subseteq B[x^*]$ such that (a) constraint $x(\bar{C}) \leq 1$ is weight-balanced and (b) constraint $x(\bar{C}) \leq 1$ is at least as violated as $x(C) \leq 1$.*

Proof. Since $x(C) \leq 1$ is a valid clique inequality we know that if $i, j \in C$ are such that $i \neq j$, then $a(cl(\{i, j\})) > b$. From Proposition 12 we know that for all $i \in C$ there exists $j(i) \in B[x^*, i]$ such that $x^*_i = x^*_{j(i)}$. Let $\bar{C} = \{j(i) : i \in C\}$. Note that $|\bar{C}| \leq |C|$. Since either $j(i) = i$ or $(j(i), i) \in A$ for all $i \in C$ we know that $cl(i) \subseteq cl(j(i))$. This implies that for $i, k \in \bar{C}$ such that $i \neq k$, we have $a(cl\{i, k\}) > b$. Thus $x(\bar{C}) \leq 1$ defines a clique inequality. To see that $x(\bar{C}) \leq 1$ is at least as violated, observe that by Proposition 9 it is not possible to have two vertices i and i' in C such that $j(i) = j(i')$. Hence, $x^*(C) = \sum_{i\in C} x^*_i = \sum_{i\in C} x^*_{j(i)} = x^*(\bar{C})$, and so, $\frac{x^*(C)-1}{|C|} \leq \frac{x^*(\bar{C})-1}{|\bar{C}|}$. □

From Propositions 16 and 17 the following theorem directly follows.

Theorem 18. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. There exist maximally violated clique inequalities and maximally violated MIC inequalities that are weight-balanced.*

Since both clique and MIC inequalities have no negative coefficients, Theorem 18 implies that, given a fractional solution x^*, the support of maximally violated constraints will be contained in $B[x^*]$. As we will see later in the computational section this is a very important result, because the set $B[x^*]$ tends to be very small.

4. Shrinking and PCKP

In applying separation algorithms to combinatorial optimization problems defined over graphs it is very helpful to preprocess these graphs in order to reduce the optimization problem to one that is equivalent, but defined over a smaller graph. This serves to reduce the number of operations that must be performed. Given a graph G and a set of vertices $S \subseteq V$ let G/S denote the graph obtained by contracting the set S into a single node s and deleting any self-loop edges that appear. This operation is called *shrinking* S in G, and has been used with great success in solving large-scale integer programming problems such as the Traveling Salesman Problem [19,10]. In this section we describe how shrinking can be defined in the context of PCKP. As we will see, the notion of weight-balanced inequalities plays a key role in shrinking.

Definition 19. Let (G, a, b) define an instance of PCKP satisfying our working assumptions and let $\mathcal{S}$ be a family of sets defining a partition of V. Define $\bar{G} = (\bar{V}, \bar{A})$ as the graph obtained from contracting the sets $S \in \mathcal{S}$, and define $\bar{a} \in \mathbb{Z}^{\bar{V}}$ by $\bar{a}_S = a(S)$. We say that $\bar{G}$ and $\bar{a}$ are obtained from G and a by *shrinking* the partition $\mathcal{S}$. Note that $(\bar{G}, \bar{a}, b)$ defines an instance of PCKP.

Theorem 20. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, let $\mathcal{S}$ be a family of sets defining a partition of V and let $(\bar{G}, \bar{a}, b)$ the instance obtained by shrinking the partition $\mathcal{S}$. Consider an inequality $\alpha x \leq \beta$ valid for $P(G, a, b) \cap \{0, 1\}^V$. Define $\bar{\alpha}_S = \alpha(S)$ for each $S \in \mathcal{S}$. Then, $\bar{\alpha} z \leq \beta$ is valid for $P(\bar{G}, \bar{a}, b) \cap \{0, 1\}^{\bar{V}}$.*

Proof. Let $z \in P(\bar{G}, \bar{a}, b) \cap \{0, 1\}^{\bar{V}}$ and define $y \in [0, 1]^V$ as $y_i = z_S$ for all $i \in S$, $S \in \mathcal{S}$. We first prove that $y \in P(G, a, b) \cap \{0, 1\}^V$. In fact,

$$ay = \sum_{i\in V} a_i y_i = \sum_{S\in\mathcal{S}} \sum_{i\in S} a_i y_i = \sum_{S\in\mathcal{S}} \left(\sum_{i\in S} a_i \right) z_S = \sum_{S\in\mathcal{S}} \bar{a}_S z_S = \bar{a} z \leq b. \tag{4}$$

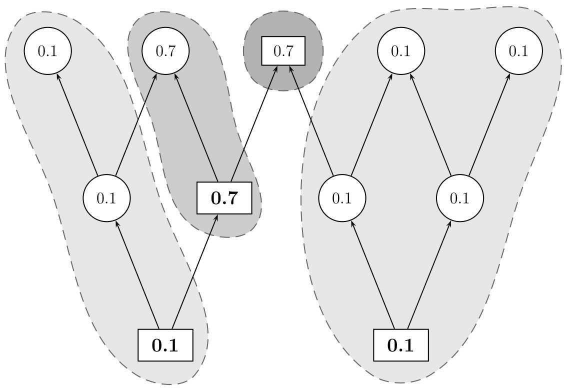

Fig. 2. Example of a canonical partition of V implied by G and x^*. The values of x^* are written inside the vertices. The rectangular vertices correspond to representing elements for each set in the partition. The vertices in $B[x^*]$ have bold letters.

Additionally, for each $(i, j) \in A$ either i and j are in the same set $S \in \mathcal{S}$, or in different sets. In the first case, $y_i = y_j = z_S$, so in particular it satisfies $y_i \leq y_j$. In the second case, suppose that $i \in S$ and $j \in T$. Since $(i, j) \in A$, we also have $(S, T) \in \bar{A}$, so $y_i = z_S \leq z_T = y_j$.

Now, since $y \in P(G, a, b)$ it follows that $\alpha y \leq \beta$. Thus, replacing a and $\bar{a}$ by α and $\bar{\alpha}$ respectively in equation (4), we obtain:

$$\bar{\alpha} z = \sum_{S\in\mathcal{S}} \bar{\alpha}_S z_S = \sum_{S\in\mathcal{S}} \left(\sum_{i\in S} \alpha_i\right) z_S = \sum_{S\in\mathcal{S}} \sum_{i\in S} \alpha_i y_i = \sum_{i\in V} \alpha_i y_i = \alpha y \leq \beta. \quad \square$$

Corollary 1. *Let $\mathcal{S}$ be a family of sets defining a partition of V and x^* such that if $i, j \in S$ then $x^*_i = x^*_j$. Let $z^* \in \mathbb{R}^{\bar{V}}$ be defined such that $z^*_S = x^*_i$, for all $i \in S$, $S \in \mathcal{S}$. Then, $z^* \in P(\bar{G}, \bar{a}, b)$. Moreover, $\alpha x^* = \bar{\alpha} z^*$.*

Proof. This result is obtained by replacing $a, \bar{a}, z$ and y by $\alpha, \bar{\alpha}, z^*$ and x^* respectively in equation (4). $\square$

The previous corollary implies that if there exists a violated inequality for PCKP, then there exists a violated inequality for the shrunken graph. This result, however, does not imply that every violated constraint in the shrunken problem can be mapped back to a violated constraint in the original variable space. This issue is addressed in the following theorem. Before, however, we need to show how it is possible to define a partition of V such that it is safe to perform shrinking for PCKP.

Definition 21. Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. For each $i \in V$ define,

$$\Pi_i = \{j \in V : B[x^*, i] = B[x^*, j]\}.$$

As immediately follows from Proposition 12, the family of sets $\mathcal{S} = \{\Pi_i\}_{i\in V}$ defines a partition of V. Moreover, every $S \in \mathcal{S}$ is such that $i, j \in S$ implies $x^*_i = x^*_j$. We say that the family of sets $\mathcal{S}$ defines the *canonical* partition of V implied by G and x^*.

Note that this partition can be constructed in $\mathcal{O}((|V| + |E|)^2)$ using a topological order (see [9]).

Theorem 22. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions, and consider $x^* \in P(G, a, b)$. Let $\mathcal{S}$ represent the canonical partition of V defined by G and x^*. Let $\bar{G} = (\bar{V}, \bar{A})$, $\bar{a}$ and z^* be obtained from G, a and x^* by shrinking partition $\mathcal{S}$ as described in Theorem 20. For each set $S \in \mathcal{S}$ let us choose $i_S \in S$ (a representing element, see Fig. 2 for an example) such that either $i_S \in B[x^*]$ or if $(j, i_S) \in A$ then $j \notin S$. Finally, let $\bar{\alpha} z \leq \beta$ be a valid inequality for $P(\bar{G}, \bar{a}, b) \cap \{0, 1\}^{\bar{V}}$ that is weight-balanced with respect to z^*. Define,*

$$\alpha_i = \begin{cases} \bar{\alpha}_S & \text{if } i = i_S \text{ for some } S \in \mathcal{S}, \\ 0 & \text{otherwise.} \end{cases}$$

Then, inequality $\alpha x \leq \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$ and weight-balanced with respect to x^*. Moreover, $\alpha x^* = \bar{\alpha} z^*$.

Proof. It is easy to see that $\alpha x^* = \bar{\alpha} z^*$. First, we prove that $\alpha x \leq \beta$ is weight-balanced with respect to x^*. We only need to prove that if $i \notin B[x^*]$ and $i = i_S$ for some $S \in \mathcal{S}$ then $\alpha_i \leq 0$. In fact, there must exist $j \in T \neq S$ such that $(j, i) \in A$ and $x^*_i = x^*_j$. But, this means that $(T, S) \in \bar{A}$ and $z^*_S = z^*_T$, so $S \notin B[z^*]$. Since $\bar{\alpha} z \leq b$ is weight-balanced with respect to z^*, we must have $\alpha_i = \bar{\alpha}_S \leq 0$.

Second, we show that $\alpha x \le \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$. For this consider $x \in P(G, a, b) \cap \{0, 1\}^V$ and let $F = \{i \in B[x^*] : x_i = 1\}$. Define:

$$\hat{x}_i = \begin{cases} 1 & \text{if } i \in cl(F), \\ 0 & \text{otherwise.} \end{cases}$$

Also define $\hat{z} \in \{0, 1\}^{\hat{V}}$ so that $\hat{z}_S = \hat{x}_{i_S}$ for each $S \in \mathcal{S}$. An outline of the proof is as follows: We first prove (a) $\hat{x} \in P(G, a, b)$. From this follows (b) $\hat{z} \in P(\bar{G}, \bar{a}, b)$. We then show (c) $\alpha\hat{x} \ge \alpha x$ and (d) $\alpha\hat{x} = \bar{\alpha}\hat{z}$. From (c), (d) and (b) it follows that $\alpha x \le \alpha\hat{x} = \bar{\alpha}\hat{z} \le \beta$, with which we conclude.

(a) By definition $\hat{x}$ is the incidence vector of the smallest closure containing F. Since it is the incidence vector of a closure, $\hat{x}$ satisfies all precedence constraints. Further, since x is also the incidence vector of a closure containing F we know that $\hat{x} \le x$. Together with $a \ge 0$, this implies that $a\hat{x} \le b$.

(b) We first show that $(S, T) \in \bar{A}$ and $S \ne T$ imply $\hat{z}_S \le \hat{z}_T$. We only need to consider the case $\hat{z}_S = 1$. In this case $\hat{x}_{i_S} = 1$, and so, by how $\hat{x}$ was constructed, there exists $k \in B[x^*, i_S]$ such that $\hat{x}_k = 1$. On the other hand, $(S, T) \in \bar{A}$ implies there exists $(i, j) \in A$ such that $i \in S$ and $j \in T$. Since $S \ne T$ and $(i, j) \in A$ we know $B[x^*, i] \subsetneq B[x^*, j]$. This, however, implies that $k \in B[x^*, i_S] = B[x^*, i] \subsetneq B[x^*, j] = B[x^*, i_T]$ and so $\hat{x}_{i_T} = 1$. With this we conclude $\hat{z}_T = 1$. Second, we must show that $\bar{a}\hat{z} \le b$. For this, it suffices to show that $\hat{x}_{i_S} \le \hat{x}_i$ for all $i \in S$, because this implies that

$$\bar{a}\hat{z} = \sum_{S \in \mathcal{S}} \bar{a}_S \hat{z}_S = \sum_{S \in \mathcal{S}} \left(\sum_{i \in S} a_i \right) \hat{x}_{i_S} \le \sum_{S \in \mathcal{S}} \sum_{i \in S} a_i \hat{x}_i = a\hat{x} \le b.$$

Consider $S \in \mathcal{S}$. To prove that $\hat{x}_{i_S} \le \hat{x}_i$ for all $i \in S$ we only need to consider the case $\hat{x}_{i_S} = 1$. From the definition of $\hat{x}$ we know $\hat{x}_{i_S} = 1$ implies the existence of $k \in B[x^*, i_S]$ such that $\hat{x}_k = 1$. Since $B[x^*, i_S] = B[x^*, i]$ for all $i \in S$, it follows that $k \in B[x^*, i]$ for all $i \in S$. From this we conclude $\hat{x}_i = 1$ for all $i \in S$.

(c) Note that $D = \{i \in V : x_i > \hat{x}_i\} \subseteq V \setminus B[x^*]$. Since $\alpha x \le \beta$ is weight-balanced with respect to x^*, it follows that $\alpha_i \le 0$ for all $i \in D$. Thus, $\alpha\hat{x} \ge \alpha x$.

(d) Finally,

$$\sum_{i \in V} \alpha_i \hat{x}_i = \sum_{S \in \mathcal{S}} \alpha_{i_S} \hat{x}_{i_S} = \sum_{S \in \mathcal{S}} \bar{\alpha}_S \hat{z}_S. \quad \square$$

5. A simple mechanism for strengthening valid inequalities

Given a valid inequality $\alpha x \le \beta$ for $P(G, a, b) \cap \{0, 1\}^V$ with $\alpha \in \mathbb{Z}^V$, $\beta \in \mathbb{Z}_+$, in this section we are concerned with quickly finding a stronger inequality $\alpha' x \le \beta'$ that is also valid for $P(G, a, b) \cap \{0, 1\}^V$. The mechanism we introduce requires solving a single instance of a maximum closure problem, as described in the following theorem.

Theorem 23. *Let (G, a, b) define an instance of PCKP satisfying our working assumptions. Let $\alpha x \le \beta$ be a valid inequality for $P(G, a, b) \cap \{0, 1\}^V$, and let $S = \{i \in V : \alpha_i \ne 0\}$. Let $\hat{x}$ be an optimal solution to problem,*

$$\begin{array}{lll} \max & \alpha x & \\ s.t. & x_i \le x_j & \forall (i, j) \in A \\ & x_i \in \{0, 1\} & \forall i \in V, \end{array} \tag{5}$$

such that $I(\hat{x}) = \{i \in V : \hat{x}_i = 1\}$ is minimal inclusion-wise among all optimal solutions of (5). Then, inequality

$$\sum_{i \in I(\hat{x})} \alpha_i x_i \le \beta \tag{6}$$

is valid for $P(G, a, b) \cap \{0, 1\}^V$. Moreover, if $I(\hat{x}) \subsetneq S$, then constraint (6) is stronger than constraint $\alpha x \le \beta$.

Proof. We first show that (6) is valid for $P(G, a, b) \cap \{0, 1\}^V$. In fact, consider $x \in P(G, a, b) \cap \{0, 1\}^V$ and let $I(x) = \{i \in V : x_i = 1\}$. Since $I(\hat{x})$ and $I(x)$ are both closures, we know that $I(x) \cap I(\hat{x})$ is also a closure. Define $x' \in \{0, 1\}^V$ so that $x'_i = 1$ iff $i \in I(x) \cap I(\hat{x})$. Since $x' \le x$ and $a \ge 0$ we know that $ax' \le b$ and so $x' \in P(G, a, b) \cap \{0, 1\}^V$. From all this it follows that,

$$\sum_{i \in I(\hat{x})} \alpha_i x_i = \sum_{i \in I(\hat{x}) \cap I(x)} \alpha_i x_i = \sum_{i \in V} \alpha_i x'_i \le \beta.$$

We now show that if $I(\hat{x}) \subsetneq S$, then constraint (6) is stronger than constraint $\alpha x \le \beta$. For this we first show that there exists $z \in P(G, a, b) \cap \{0, 1\}^V$ such that $\alpha z < \sum_{i \in I(\hat{x})} \alpha_i z_i$. Since $I(\hat{x}) \subsetneq S$, we know that there exists $i \in S \setminus I(\hat{x})$ such that $\alpha_i \ne 0$. Since $\hat{x}$ is an optimal solution of (5), we know that $\alpha(cl(i) \setminus I(\hat{x})) \le 0$. Moreover, we can assume that this $i \in S \setminus I(\hat{x})$ is such that for all $(i, j) \in A$ we have $j \in I(\hat{x})$ or $\alpha_j = 0$. With this assumption we have that $\alpha(cl(i) \setminus I(\hat{x})) = \alpha_i < 0$.

Define $z \in \{0,1\}^V$ so that $z_i = 1$ iff $i \in cl(i)$. Since $cl(i)$ is by definition a closure, and since $az \leq b$, we know that $z \in P(G,a,b) \cap \{0,1\}^V$. On the other hand,

$$\alpha z = \sum_{i \in cl(i) \cap I(\hat{x})} \alpha_i + \sum_{i \in cl(i) \setminus I(\hat{x})} \alpha_i < \sum_{i \in cl(i) \cap I(\hat{x})} \alpha_i = \sum_{i \in I(\hat{x})} \alpha_i z_i.$$

Finally, we show that, for any $x \in P(G,a,b) \cap \{0,1\}^V$ such that $\alpha x = \beta$, we have $\sum_{i \in I(\hat{x})} \alpha_i x_i = \beta$. Let $I(x) = \{i \in V : x_i = 1\}$. Note that $\alpha(I(x) \setminus I(\hat{x})) \leq 0$. Otherwise, we would have that $\alpha(I(x) \cup I(\hat{x})) > \alpha(I(\hat{x}))$, which, because $I(\hat{x}) \cup I(x)$ is a closure, contradicts the optimality assumption of $\hat{x}$. From this we have that,

$$\alpha x = \sum_{i \in S \cap I(x)} \alpha_i x_i = \sum_{i \in I(\hat{x})} \alpha_i x_i + \sum_{i \in I(x) \setminus I(\hat{x})} \alpha_i x_i \leq \sum_{i \in I(\hat{x})} \alpha_i x_i \leq \beta.$$

Thus $\alpha x = \beta$ implies $\sum_{i \in I(\hat{x})} \alpha_i x_i = \beta$, and we conclude our result. □

First note that finding a minimal support optimal solution to (5) is solvable using a flow algorithm. This theorem also implies that; if $\alpha x \leq \beta$ is a strong inequality for $P(G,a,b) \cap \{0,1\}^V$; then $\alpha(rcl(G,S')) > 0$ for any $S' \subset S$. Moreover, given any valid inequality for $P(G,a,b) \cap \{0,1\}^V$, a strong version of it can be found by maximizing a single closure problem.

6. Lifting and PCKP

Definition 24. Let $P(G,a,b)$ define an instance of PCKP satisfying our working assumptions. Consider disjoint sets $O, I \subseteq V$ and define,

$$P(G,a,b,O,I) = \{x \in P(G,a,b) : x_i = 1 \ \ \forall i \in I,\ \ x_i = 0 \ \ \forall i \in O\}.$$

Assume $I, O \subseteq V$ are such that $P(G,a,b,O,I) \cap \{0,1\}^V$ is non-empty, and consider an inequality $\alpha x \leq \beta$ valid for $P(G,a,b,O,I) \cap \{0,1\}^V$. We say that *lifting* inequality $\alpha x \leq \beta$ consists in computing coefficients $\gamma, \eta \geq 0$ such that $\gamma \neq 0$ or $\eta \neq 0$, and such that inequality

$$\alpha x + \sum_{i \in I} \gamma_i (1 - x_i) + \sum_{j \in O} \eta_j x_j \leq \beta, \tag{7}$$

is valid for $P(G,a,b) \cap \{0,1\}^V$.

For an introduction to lifting, see [23,24,14]. Lifting in the context of PCKP can be used to serve two important purposes:

- **Strengthening valid inequalities.** Consider an inequality $\alpha x \leq \beta$ that is valid for $P(G,a,b) \cap \{0,1\}^V$, but that is not facet-defining. If $P(G,a,b,I,O) \cap \{0,1\}^V$ is non-empty for sets $I, O \subseteq V$, lifting can be used to *strengthen* inequality $\alpha x \leq \beta$. Such a strengthening procedure would correspond to tilting [14] inequality $\alpha x \leq b$ in directions $(1 - e_i)$ and e_j for $i \in I$ and $j \in O$. This use of lifting has been studied by Boyd [7], Park and Park [20], van de Leensel et al. [22] and Boland et al. [6] in the context of induced cover inequalities, K-covers, $(1, K)$-configurations and induced clique inequalities.
- **Speeding up the computation of cutting planes.** Consider a non-integral solution $\bar{x} \in P(G,a,b)$. Let $I = \{i \in V : \bar{x}_i = 1\}$ and $O = \{i \in V : \bar{x}_i = 0\}$. Note that a cutting plane separating $\bar{x}$ from $P(G,a,b,O,I) \cap \{0,1\}^V$ exists if and only if a cutting plane separating $\bar{x}$ from $P(G,a,b)$ exists. From a computational point of view, it could be easier to compute a cutting plane separating $\bar{x}$ from $P(G,a,b,O,I) \cap \{0,1\}^V$, since this results in solving a separation problem of lower dimension. If the separation algorithm in $P(G,a,b,O,I)$ is successful we obtain an inequality $\alpha x \leq b$ that is valid for $P(G,a,b,O,I) \cap \{0,1\}^V$, and violated by $\bar{x}$. Lifting this inequality, we obtain an inequality of form (7) that is valid for $P(G,a,b) \cap \{0,1\}^V$ and violated by $\bar{x}$.

In order for lifting to be computationally effective it should be quick. In this section we describe simple techniques for quick lifting in the context of PCKP. In Section 6.1 we describe simple relaxations of the lifting problem that are easy to solve. In Section 6.2 we characterize those variables that can be lifted and those that cannot. Finally, in Section 6.3 we present an argument for lifting in a greedy manner so as to obtain highly-violated inequalities.

6.1. Optimal lifting and relaxed lifting of a single variable

Let (G,a,b) define an instance of PCKP satisfying our working assumptions, and consider $i \in V$. Assume inequality $\alpha x \leq \beta$ is valid for $P(G,a,b) \cap \{x : x_i = 1\} \cap \{0,1\}^V$. *Down-lifting* variable i in constraint $\alpha x \leq \beta$ consists in computing a coefficient γ_i such that

$$\alpha x + \gamma_i (1 - x_i) \leq \beta \quad \forall x \in P(G,a,b) \cap \{0,1\}^V. \tag{8}$$

Optimally down-lifting variable i consists in computing the largest possible lifting coefficient γ_i. This can be done by solving

$$\begin{aligned} z_i = \max \quad & \alpha x \\ s.t. \quad & x_i = 0 \\ & x \in P(G, a, b) \cap \{0, 1\}^V \end{aligned} \tag{9}$$

and defining $\gamma_i = \beta - z_i$. Observe that if $\alpha x \leq \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$, then $\gamma_i \geq 0$. Moreover, if $\gamma_i > 0$, then constraint (8) is strictly stronger than $\alpha x \leq \beta$.

The problem with optimal down-lifting is that solving each instance of (9) could potentially be very difficult. An alternative is to lift in a non-optimal way by solving a relaxation of (9). That is, by replacing constraint $x \in P(G, a, b) \cap \{0, 1\}^V$ by a weaker constraint $x \in R(G, a, b)$ where $P(G, a, b) \cap \{0, 1\}^V \subseteq R(G, a, b)$, and solving the *relaxed down-lifting problem* with respect to $R(G, a, b)$

$$\begin{aligned} w_i = \max \quad & \alpha x \\ s.t. \quad & x_i = 0 \\ & x \in R(G, a, b). \end{aligned} \tag{10}$$

If $w_i > 0$, then $\gamma_i = \beta_i - w_i$ is still a valid lifting coefficient in the sense that the resulting inequality (8) is valid for $P(G, a, b) \cap \{0, 1\}^V$. Although γ_i might be a weaker coefficient, if $R(G, a, b)$ is chosen appropriately, it should be computationally easier to obtain. The first relaxation $R(G, a, b)$ we consider is

$$R_{LP}(G, a, b) = \left\{x \in [0, 1]^V : ax \leq b,\ x_i \leq x_j\ \forall (i, j) \in A\right\}. \tag{11}$$

In relaxation (11) we relax the integrality condition $x \in \{0, 1\}^V$. Solving (9) subject to this relaxation can be done in $O(mn \log n)$ (see [8]).

Remark. In some cases, the use of relaxation (11) may result in a coefficient γ that is not integral. If the original constraint $\alpha x \leq \beta$ is such that α and β are both integer (as is the case for induced cover and induced clique inequalities) this results in an opportunity to further strengthen the resulting inequality. In fact, if γ_i is fractional and all other coefficients are integral, γ_i can be rounded up to obtain a stronger, integral coefficient. That this new coefficient is valid can be proved using a simple argument: Add a constraint $\epsilon x_i \leq \epsilon$ to the lifted inequality, choosing ϵ so that the fractional left-hand side coefficient becomes integer. Then, round the right-hand side down applying Gomory's rounding procedure to obtain the desired result.

A second relaxation proposed by Park and Park [20] is

$$R_{\alpha,\beta}(G, a, b) = \left\{x \in \{0, 1\}^V : \alpha x \leq \beta,\ x_i \leq x_j \quad \forall (i, j) \in A,\ i, j \in cl(\alpha)\right\}, \tag{12}$$

where $\alpha x \leq \beta$ is the inequality for $P(G, a, b) \cap \{0, 1\}^V$ to be lifted and $cl(\alpha) = cl(\{i \in V : \alpha_i \neq 0\})$. Using this relaxation amounts to solving,

$$\begin{aligned} w_i = \max \quad & \alpha x \\ s.t. \quad & \alpha x \leq \beta \\ & x_i = 0 \\ & x_k \leq x_j \quad \forall (k, j) \in A,\ k, j \in cl(\alpha) \\ & x \in \{0, 1\}^V. \end{aligned} \tag{13}$$

It is not difficult to see that the solution to (13) satisfies $w_i \leq \hat{w}_i$, where

$$\hat{w}_i = \min\{\beta, \max\{\alpha x : x_i = 0,\ x_k \leq x_j\ \forall (k, j) \in A,\ k, j \in cl(\alpha),\ x \in [0, 1]^V\}\},$$

thus we can use $\gamma_i = \beta_i - \hat{w}_i$ as a lifting coefficient. Note that computing this coefficient only requires solving a single max-closure problem.

Further note that if the α coefficients are all 0 or 1 (e.g., clique inequalities), then $\hat{w}_i = w_i$. That is, we can solve (13) to optimality by solving a max-closure problem. Moreover, in the specific case of Lifted Cover Inequalities, Park and Park [20] show that $z_i = w_i = \hat{w}_i$.

Let $P(G, a, b)$ define an instance of PCKP satisfying our working assumptions, and consider $i \in V$. Assume inequality $\alpha x \leq \beta$ is valid for $P(G, a, b) \cap \{x : x_i = 0\} \cap \{0, 1\}^V$. *Up-lifting* variable i in constraint $\alpha x \leq \beta$ consists in computing a coefficient η_i such that

$$\alpha x + \eta_i x_i \leq \beta \quad \forall x \in P(G, a, b) \cap \{0, 1\}^V. \tag{14}$$

Optimally up-lifting variable i consists in computing the largest possible lifting coefficient η_i. This can be done by solving

$$\begin{aligned} w_i = \max \quad & \alpha x \\ s.t. \quad & x_i = 1 \\ & x \in P(G, a, b) \cap \{0, 1\}^V \end{aligned} \tag{15}$$

and defining $\eta_i = \beta - w_i$. Observe that if $\alpha x \leq \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$, then $\eta_i \geq 0$. Moreover, if $\eta_i > 0$, then constraint (14) is strictly stronger than $\alpha x \leq \beta$.

As in optimal down-lifting, solving each instance of (15) can be very difficult, and an alternative is to solve a relaxed lifting problem using $R_{LP}(G, a, b)$ or $R_{\alpha,\beta}(G, a, b)$.

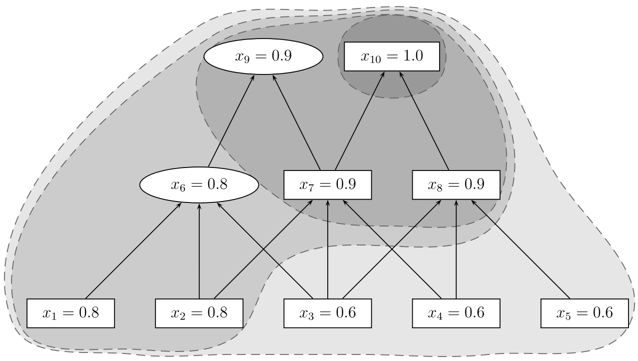

Fig. 3. Example of a fractional solution for the PCKP instance with the graph shown in the figure and knapsack constraint $\sum\{x_i : i = 1, \ldots, 10\} \le 8$.

6.2. Selecting which variables to lift

The following lemma says that if $\alpha x \le \beta$ is a face-defining inequality, it only makes sense to down-lift a very specific set of variables.

Lemma 25. *Let $P(G, a, b)$ define an instance of PCKP satisfying our working assumptions, and consider $i \in V$. Assume inequality $\alpha x \le \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$ and that there exists $\bar{x} \in P(G, a, b) \cap \{0, 1\}^V$ such that $\alpha\bar{x} = \beta$. Finally, let $S_+ = \{s \in V : \alpha_s > 0\}$. Let γ_i be the optimal down-lifting coefficient for variable i in constraint $\alpha x \le \beta$. If $i \notin cl(S_+)$ then $\gamma_i = 0$.*

Proof. Define $\hat{x} \in \{0, 1\}^V$ as follows,

$$\hat{x}_j = \begin{cases} 0 & \text{if } j \in rcl(i) \\ \bar{x}_j & \text{otherwise.} \end{cases}$$

First, note that the support of $\hat{x}$ defines a closure, since it is obtained by removing a reverse closure from a closure. Second, note that $a\hat{x} \le b$ since $a \ge 0$ and $\hat{x} \le \bar{x}$. These two facts imply that $\hat{x} \in P(G, a, b) \cap \{0, 1\}^V$, and $\alpha\hat{x} \le \beta$. Next, note that $\alpha\hat{x} = \beta$. In fact, let $S = \{s \in V : \alpha_s \neq 0\}$. It is easy to see that if $i \notin cl(S_+)$ then $rcl(i) \cap S_+ = \emptyset$ and $rcl(i) \cap S \subseteq S \setminus S_+$. Hence,

$$\alpha\hat{x} = \sum_{j \in S} \alpha_j \hat{x}_j = \sum_{j \in S \setminus rcl(i)} \alpha_j \hat{x}_j = \sum_{j \in S \setminus rcl(i)} \alpha_j \bar{x}_j \ge \sum_{j \in S} \alpha_j \bar{x}_j = \beta.$$

Thus, since $\hat{x} \in P(G, a, b) \cap \{0, 1\}^V$ and since $\alpha x \le \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$, we conclude $\alpha\hat{x} = \beta$. Moreover, since $\hat{x}_i = 0$ and $\alpha\hat{x} = \beta$, we conclude that $\hat{x}$ is an optimal solution of Problem (9) with objective β, where we conclude that $\gamma_i = 0$. □

The following lemma gives a similar result for up-lifting.

Lemma 26. *Let $P(G, a, b)$ define an instance of PCKP satisfying our working assumptions, and consider $i \in V$. Let $S = \{s \in V : \alpha_s \neq 0\}$. Assume inequality $\alpha x \le \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^V$ and that for all $j \in S$ there exists $x^j \in P(G, a, b) \cap \{x_j = 1\} \cap \{0, 1\}^V$ such that $\alpha x^j = \beta$. Let η_i be the optimal up-lifting coefficient for variable i in constraint $\alpha x \le \beta$. If $i \in cl(S)$ then $\eta_i = 0$.*

Proof. Since $i \in cl(S)$, there exists $j \in S$ such that $i \in cl(\{j\})$. Consider $x^j \in P(G, a, b) \cap \{x_j = 1\} \cap \{0, 1\}^V$ such that $\alpha x^j = \beta$ (as in the hypothesis). Since $i \in cl(\{j\})$ we also have that x^j is in $P(G, a, b) \cap \{x_i = 1\} \cap \{0, 1\}^V$. Thus, x^j is feasible for problem (15), and since it has objective function value β we conclude it is also optimal. Hence $\eta_i = 0$. □

Note that minimal induced cover inequalities and induced clique inequalities satisfy the conditions required by the two previous Lemmas.

6.3. Choosing a lifting order

We are now ready to address the following two important questions: Is the lifting order important when strengthening valid inequalities through lifting? If so, which is the *best* ordering of the variables in which to do lifting? To answer the first question consider the instance of PCKP and the corresponding fractional solution depicted in Fig. 3. Suppose we would like to strengthen the following inequality that is valid for $P(G, a, b) \cap \{0, 1\}^V$:

$$x_1 + x_2 + x_3 + x_4 + x_5 \le 3. \tag{16}$$

Note that this inequality is violated by the fractional solution by 0.4. Lemma 26 shows that we should only consider down-lifting variables x_6, x_7, x_8, x_9 and x_{10}. Moreover, optimally down-lifting variable x_k is the same as solving the problem

$$\gamma_k = 3 - \begin{array}{ll} \max & x_1 + x_2 + x_3 + x_4 + x_5 \\ s.t. & x \in P(G, a, b) \cap \{x_k = 0\}. \end{array}$$

If we optimally down-lift (in order) variables x_6, x_7, x_8, x_9 and x_{10}, we obtain the (facet-defining) inequality

$$x_1 + x_2 + x_3 + x_4 + x_5 - x_6 - x_7 - x_8 \leq 0,$$

which is violated by 0.8. If, instead, we optimally down-lift in the same variables in reversed order (e.g., $x_{10}, x_9, x_8, x_7, x_6$), we obtain the valid inequality

$$x_1 + x_2 + x_3 + x_4 + x_5 - x_9 - 2x_{10} \leq 0 \tag{17}$$

which is violated by 0.5. What this shows is that, in general, the resulting inequality, and the amount by which it is violated, is dependent of the *lifting order*. So the question of choosing an appropriate lifting order is very relevant. Again, consider the same example as before. We know that if we optimally down-lift x_{10} and then x_9 we obtain inequality (17), which is violated by 0.5. If, instead, we optimally down-lift variables x_9 and x_{10} we obtain inequality,

$$x_1 + x_2 + x_3 + x_4 + x_5 - 2x_9 - x_{10} \leq 0. \tag{18}$$

This equality is violated by 0.6. That is, lifting x_9 before x_{10} results in an inequality that is more violated. This is because we are down-lifting and because $(1 - x^*_{10}) < (1 - x^*_9)$. It turns out that when lifting pairs of variables, this greedy way of lifting always results in an inequality that is more violated. This is explained in the following lemma.

Lemma 27. *Consider an inequality $\alpha x \leq \beta$ that is valid for $P(G, a, b) \cap \{0, 1\}^n$. Let R be an integral polyhedral relaxation of $P(G, a, b) \cap \{0, 1\}^n$. That is, R is a polytope such that $P(G, a, b) \cap \{0, 1\}^n \subseteq R$ and such that the extreme points of R are integral. Let $x^* \in P(G, a, b)$, and suppose we sequentially compute the optimal R-lifting coefficients, i.e.*

$$\eta_i^R := \beta - \max\{\alpha x : x_i = 1,\ x \in R\}, \tag{19}$$

and

$$\gamma_i^R := \beta - \max\{\alpha x : x_i = 0,\ x \in R\}. \tag{20}$$

Then

1. *If $(1 - x^*_i) \geq (1 - x^*_j)$, down-lifting x_i first and then down-lifting x_j results in an inequality with violation that is equal to or greater than the violation of the inequality that would be obtained by down-lifting x_j first and then down-lifting x_i.*
2. *If $x^*_i \geq x^*_j$, up-lifting x_i first and then up-lifting x_j results in an inequality with violation that is equal to or greater than the violation of the inequality that would be obtained by up-lifting x_j first and then up-lifting x_i.*
3. *If $(1 - x^*_i) \geq x^*_j$, down-lifting x_i first and then up-lifting x_j results in an inequality with violation that is equal to or greater than the violation of the inequality that would be obtained by down-lifting x_j first and then up-lifting x_i.*

Proof. We assume that $\alpha x \leq \beta$ is valid for $P(G, a, b) \cap \{0, 1\}^n$ and that $i, j \in \{1, \ldots, n\}$ are two indices of variables in $P(G, a, b)$. We define

$$Z_{pq} = \min\left\{\beta, \max\{\alpha x : x \in R, x_i = p, x_j = q\}\right\}.$$

Note that Z_{pq} is at most the maximum value of αx attainable in $P(G, a, b) \cap \{0, 1\}^n \cap \{x_i = p, x_j = q\}$. Hence, the resulting coefficients would be valid for our original problem. With this, we prove each part of Lemma 27 by showing that the sum of the lifting coefficients remains constant when we exchange the order of lifting between two consecutive variables.

In order to prove 1, we call γ_i, γ_j the optimal R-lifting coefficients for x_i, x_j obtained when first down-lifting x_i, and then x_j, and δ_i, δ_j the optimal R-lifting coefficients for x_i, x_j obtained when first down-lifting x_j and then x_i. We will prove that $\delta_i + \delta_j = \gamma_i + \gamma_j$.

For this, note that

$$\begin{aligned} \gamma_i &= (\beta - \max\{Z_{00}, Z_{01}\}), \\ \gamma_j &= (\beta - \gamma_i - \max\{Z_{00}, Z_{10} - \gamma_i\}), \\ \delta_j &= (\beta - \max\{Z_{00}, Z_{10}\}), \\ \delta_i &= \left(\beta - \delta_j - \max\{Z_{00}, Z_{01} - \delta_j\}\right). \end{aligned}$$

Hence,

$$\begin{aligned} \gamma_i + \gamma_j &= \beta - \max\{Z_{00}, Z_{10} - \gamma_i\} \\ &= \beta - \max\{Z_{00}, Z_{10} - \beta + \max\{Z_{00}, Z_{01}\}\} \\ &= \beta - \max\{Z_{00}, Z_{10} - \beta + Z_{00}, Z_{10} - \beta + Z_{01}\} \\ &= \beta - \max\{Z_{00}, Z_{00} - (\beta - Z_{10}), Z_{10} + Z_{01} - \beta\}. \end{aligned}$$

But, since $\beta \geq Z_{10}$, we conclude that

$$\gamma_i + \gamma_j = \beta - \{\max\{Z_{00}, Z_{10} + Z_{01} - \beta\}\}.$$

Since δ_i and δ_j are obtained from γ_i and γ_j by swapping Z_{10} and Z_{01}, we conclude that $\delta_i + \delta_j = \gamma_i + \gamma_j$.

Statements 2 and 3 of the lemma can be shown analogously. □

An important limitation of Lemma 27 is that it is only true when it comes to lifting pairs of variables. That is, if we want to lift three or more variables, it is not clear if lifting in a greedy order will result in an inequality that is most violated with respect to all possible permutations of lifting orders. Park and Park [20] prove that in the case of lifting induced cover inequalities it is indeed true that greedy lifting orders (for down-lifting variables) result in maximally violated inequalities. However, it remains an open question to determine whether it is true in general or not.

7. Computational experiments and comparisons

In this section we analyze the performance of our proposed methodologies. That is, we aim to show, via computational experimentation, that (1) our shrinking methodologies are effective in reducing problem size, (2) separating weight-balanced inequalities is computationally advantageous (i.e., the number of fractional break-points in practice is small), and (3) that lifting can have a strongly positive effect on improving LP bounds of problems having PCPKP substructures. For this purpose we develop a prototype implementation of our proposed algorithmic methodologies to show these points. A detailed analysis of parameter calibration, and further integration in a branch-and-bound scheme, are beyond the scope of this article and a topic for future research. In Section 7.1 we describe the algorithms that were implemented, the platform on which the experiments took place, and the experiments themselves. In Section 7.2 we describe the instances that we analyzed. In Section 7.3 we describe our computational results.

7.1. Detailed implementation

We developed a simple cutting plane algorithm that separates MIC inequalities as follows. Given an instance of PCKP described by (G, a, b), and a fractional solution x^*, perform the following steps:

1. Eliminate all vertices $i \in V$ such that $x_i^* = 0$.
2. Compute all break-points $B[x^*]$ as described in Definition 11.
3. Compute the canonical partition associated with (G, a, b) and x^* as described in Definition 21.
4. Use the canonical partition to obtain a contracted instance $(\bar{G}, \bar{a}, b)$ and the corresponding solution $\bar{x}^*$, as described in Definition 19.
5. Compute the break-points $B[\bar{x}^*]$ of instance $P(\bar{G}, \bar{a}, b)$ and solution $\bar{x}^*$, as described in Definition 11.
6. For each fractional break-point $f \in B[\bar{x}^*]$ construct a minimal induced cover inequality as follows:
   (a) Define $C := \{f\} \cup (B[\bar{x}^*] \cap \{i \in \bar{V} : \bar{x}_i^* = 1\})$.
   (b) Add the remaining break-points, in decreasing order of value, until $\bar{a}(cl(C)) > b$.
   (c) Enforce minimality of C by eliminating non-minimal elements in reverse order of inclusion to the set.
7. Apply the down-lifting algorithm detailed in Section 6 to obtain a stronger valid inequality in $P(\bar{G}, \bar{a}, b) \cap \{0, 1\}^{\bar{V}}$. For this purpose use the relaxation detailed in Eq. (12).
8. Apply the up-lifting algorithm detailed in Section 6 to obtain a stronger valid inequality in $P(\bar{G}, \bar{a}, b) \cap \{0, 1\}^{\bar{V}}$. For this purpose we use the relaxation detailed in Eq. (11).
9. Map the resulting inequality of instance $(\bar{G}, \bar{a}, b)$ to the corresponding inequality of instance (G, a, b), as described in Theorem 22.

The steps described above were repeated many times, in rounds. In each round we generated one inequality for each fractional break-point, as described above. We only kept cuts that were violated by more than 2^{-7}. Of these cuts, we selected the three most violated, added them to the LP, and resolved. We iterated until we were not able to find any inequalities meeting the minimum violation requirement.

In order to measure the effect of using different features, we ran all instances five times, each time with one of the following configurations:

1. (CPX) We let CPLEX generate cuts with its default settings.
2. (CPX + MIC) After letting CPLEX generate cuts, we generated MIC inequalities, but we did not lift them.
3. (CPX + D-MIC) After letting CPLEX generate cuts, we generated down-lifted MIC inequalities. We did not do any up-lifting.
4. (CPX + U-MIC) After letting CPLEX generate cuts, we generated up-lifted MIC inequalities. We did not do any down-lifting.
5. (CPX + DU-MIC) After letting CPLEX generate cuts, we generated down-and-then-up-lifted MIC inequalities.

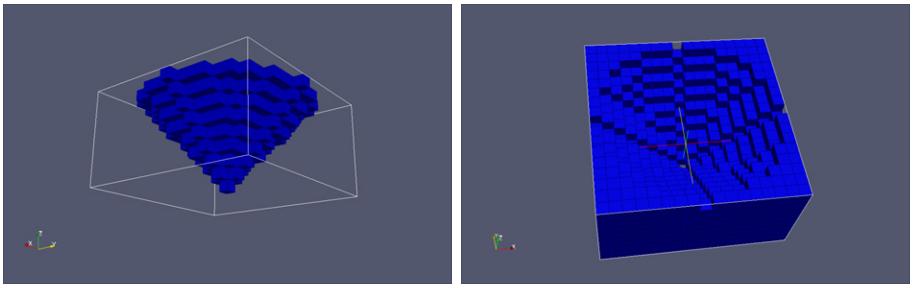

Fig. 4. Precedence example, with angle of 45°.

All of the cuts were added at the root node of the branch and bound tree. We did not analyze branch and cut algorithms. Instead, for each configuration and for each instance, we analyzed the number of fractional break-points found per round, the amount of time that our code dedicated to adding cuts and resolving LPs after letting CPLEX first finish adding its cuts, and the integrality gap closed. All of the code was developed using the C programming language. The underlying Linear Programming solver used was CPLEX 12.2 [15]. All runs were made using a single thread on the machine with an address space limit and data limit of 8 Gb. The running time limit was 40 h. The integer tolerance was set to 10^{-6}. The target relative gap was set to 10^{-5}. The cuts-factor (maximum number of cuts to add relative to the original number of constraints) was set to 10. Finally, the root-node LP was solved using the barrier solver with crossover. All other CPLEX settings were left at their default values.

7.2. Instances

To test the algorithms presented above, we considered randomly generated instances inspired by open pit mining applications. For a given base height h, we considered a three dimensional regular grid of size $\lfloor \frac{5}{2}h \rfloor \times \lfloor \frac{5}{2}h \rfloor \times h$ comprised of units we call *blocks*. For a given block b with coordinates (x, y, z), its precedences were defined as those blocks $b' = (x', y', z')$ in the L_2 cone with slope 45° above b (see Fig. 4 for an illustration). For a description of how such precedence structures relate to open pit mining, see [8]. In our implementation, we only incorporated precedence constraints in the transitively reduced set. That is, we only added constraints $x_i \le x_j$ for arcs $(i, j) \in \bar{A}$, where, $\bar{A} \subset A$ is the maximal set satisfying the following condition: if (i, j) and (j, k) are in $\bar{A}$, then (i, k) is not in $\bar{A}$.

For each instance we added k knapsack constraints (with positive integer coefficients). These coefficients were randomly generated with a discrete uniform distribution of range $\{0, \dots, 4094\}$. The right-hand side for all constraints was defined as $2144h^3$. This is approximately $\frac{\pi}{3}\mathbb{E}(a_i)$, and can be interpreted as the expected weight of the largest L_2 cone completely included in the set. The objective function of each element $b = (x, y, z)$ was randomly generated with a discrete uniform distribution of range $\{-36(h-z), \dots, 48(h-z)\}$, thus making, in average, elements deeper in the grid more valuable than those on top. We considered values of h in $\{3, \dots, 7\}$, and values of k in $\{1, \dots, 9\}$. The resulting instances had between 147 nodes and 490 precedences ($h = 3$) to 2023 nodes and 31 790 precedences ($h = 7$).

7.3. Results

In order to compare the different configurations of our cut separation algorithm we used two metrics: the time to solve the LP and generate the cuts, and the percentage of the gap closed at the root node. Formally, the latter is defined as follows:

$$\text{gap closed (percentage)} = \frac{z_{\text{cuts}} - z_{LP}}{z_{IP} - z_{LP}},$$

where z_{IP} is the optimal integer programming solution value, z_{LP} is the value of the root *LP* relaxation without any cuts, and z_{cuts} is the value obtained at the root node after adding cuts.

In total we generated 450 random instances. Specifically, for each $h \in \{3, \dots, 7\}$ and for each $k \in \{1, \dots, 9\}$ we generated 10 different instances (using different seeds). As previously mentioned, each instance was ran five times, once with each configuration. For each run with the CPX configuration, we took note of two things: The total time it took to run this configuration, and the percentage gap closed by CPLEX with its cuts. For each of the other four configurations, we took note of six things: The percentage gap closed with the CPLEX cuts together with the cuts generated by our algorithm, the total time taken to run the instance with the given configuration (including the time it took CPLEX to solve the root node), the total number of cuts that were added by our code, the relative size of the shrunken graph (relative to the original graph), the total number of break-points (relative the original number of variables in the problem), and the total number of break-points that had fractional values (relative the original number of variables in the problem). The average of these values over the 90 instances corresponding to each h-configuration pair, are reported in Table 1.

From the table we can make four first conclusions: (1) By comparing the CPX and CPX + MIC configurations, we can see that adding MIC inequalities without lifting contributes little value to what CPLEX can achieve with its own cutting plane

Table 1
Computational results corresponding to the application of 5 different configurations of our cutting-plane algorithm to 450 instances.

		$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$
CPX	Closed gap (%)	49.62	22.73	0.00	0.00	0.00
	Time (s)	0.10	0.57	0.75	3.01	9.60
CPX + MIC	Closed gap (%)	52.25	28.13	0.13	0.01	0.00
	Time (s)	0.13	1.17	2.14	7.30	10.10
	Number of cuts	24.57	31.04	0.01	0.00	0.00
	Shrinkage (%)	18.78	23.91	46.75	52.76	60.29
	Break-points (%)	54.69	34.79	19.04	13.07	9.82
	Fract. break-points (%)	3.28	1.34	1.66	1.51	3.34
CPX + D-MIC	Closed gap (%)	56.95	35.97	18.56	17.13	26.67
	Time (s)	0.14	1.44	6.86	25.33	62.94
	Number of cuts	38.13	105.28	113.38	107.75	94.61
	Shrinkage (%)	18.49	24.22	47.47	52.37	59.07
	Break-points (%)	56.07	37.09	25.93	18.64	14.50
	Fract. break-points (%)	3.01	1.18	0.85	0.59	0.54
CPX + U-MIC	Closed gap (%)	58.77	40.88	19.00	25.02	38.32
	Time (s)	0.52	22.45	892.55	25 931.93	172 591.71
	Number of cuts	37.67	62.28	42.83	153.23	131.51
	Shrinkage (%)	17.86	22.19	42.86	45.69	49.75
	Break-points (%)	57.04	38.44	26.47	19.87	16.24
	Fract. break-points (%)	2.98	1.29	0.75	0.37	0.37
CPX + DU-MIC	Closed gap (%)	62.08	44.96	29.94	31.34	45.51
	Time (s)	0.60	25.78	897.69	18 768.45	104 194.77
	Number of cuts	42.72	100.27	122.18	138.12	93.43
	Shrinkage (%)	18.11	21.80	43.26	45.30	49.36
	Break-points (%)	57.17	39.45	27.48	20.80	16.78
	Fract. break-points (%)	2.62	1.16	0.78	0.49	0.54

routines. Also, by looking at these configurations we can see that the cuts generated by CPLEX and our methodology are only effective on small, shallow, problems ($h \leq 4$). (2) By comparing the CPX + MIC, CPX + DMIC, CPX + UMIC, and CPX + DUMIC configurations, we can see that lifting is critical to getting MIC inequalities to close the integrality gap. Moreover, the gap can be closed anywhere from 20% to 60%, considering instances of all sizes. (3) By comparing the CPX + DMIC, CPX + UMIC, and CPX + DUMIC configurations, we can see that up-lifting results in stronger inequalities than down-lifting, but that up-lifting is considerably slower. Down-lifting is remarkably fast, but up-lifting is prohibitively slow. It should be noted that in our implementation we use CPLEX to solve the lifting problem over the relaxation defined in (11). It would probably be significantly faster to use the techniques described in [8,3] to solve these lifting problems. This should be the first thing to change in any new implementation of our up-lifting mechanism; first because the LP problems are very degenerate; and second because polynomial and practical algorithms are available, where even hot-start could be used. (4) By comparing the CPX and CPX + MIC configurations, we can see that the time needed to generate the MIC inequalities is minimal. By looking at CPX + DMIC and CPX + UMIC, we can see that lifting is what takes most of the time—especially up-lifting. Why is generating the MIC so fast? Shrinking certainly helps. As can be seen from the table, by shrinking we are obtaining a graph that is one-fifth the size of the original graph on small instances (20% for small h) and one-half the size on larger instances (50% for larger h values). However, what is really striking, is the small number of fractional break-points. This is really what is making MIC separation so fast. With such a small number of fractional break-points it is even conceivable that one could use integer programming to separate the cuts.

8. Final remarks

In this paper we have considered the problem of separating maximally violated inequalities for the precedence constrained knapsack problem. Our contributions include a new partial characterization of maximally violated inequalities (the concept of weight-balanced inequalities), a new safe shrinking technique, and new insights on strengthening and lifting. In our computational results sections we applied the methodologies described in this paper to an implementation of a cutting-plane algorithm for separating Minimally Induced Cover (MIC) inequalities. The computational experiments clearly illustrated the practicality of our results. In fact, the experiments showed that the number of fractional break-points in practice tends to be very small. This is very important because it greatly eases the separation of weight-balanced inequalities. For clique and MIC inequalities this is specially useful, because the support of these inequalities is made up entirely of positive coefficients, and thus, is limited to variables corresponding to fractional break-points. Moreover, our experiments show that shrinking also serves to greatly reduce the problem size, and that our proposed lifting mechanisms significantly strengthen the MIC inequalities obtained. More so, without lifting, our computations show that the inequalities are useless in practice. We expect that these results can significantly contribute to solving larger integer programming formulations of PCKP generalizations, especially those that appear in open pit mining problems, which are characterized by being extremely large.

Acknowledgments

This work was partially funded by FONDECYT grants 1110024 and 1110674, ANILLO grant ACT-88, Basal project CMM (Universidad de Chile) and ICM grant P10-024-F. We would especially like to thank the anonymous referees who handled this paper for their meticulous attention to detail, and for their generous and very useful suggestions.

References

[1] J. Amaya, D. Espinoza, M. Goycoolea, E. Moreno, T. Prevost, E. Rubio, A scalable approach to optimal block scheduling, in: Proceedings of APCOM 2009, The Canadian Institute of Mining, Metallurgy and Petroleum, 2009, pp. 567–575.
[2] E. Balas, R. Jeroslow, Canonical cuts on the unit hypercube, SIAM J. Appl. Math. 23 (1) (1972) 61–69.
[3] D. Bienstock, M. Zuckerberg, Solving LP relaxations of large-scale precedence constrained problems, in: Proceedings of IPCO 2010, Vol. 6080, in: LNCS, 2010, pp. 1–14.
[4] A. Bley, N. Boland, C. Fricke, G. Froyland, A strengthened formulation and cutting planes for the open pit mine production scheduling problem, Comput. Oper. Res. 37 (9) (2010) 1641–1647.
[5] N. Boland, A. Bley, C. Fricke, G. Froyland, R. Sotirov, Clique-based facets for the precedence constrained knapsack problem, Math. Program. 133 (2012) 481–511.
[6] N. Boland, I. Dumitrescu, G. Froyland, A.M. Gleixner, LP-based disaggregation approaches to solving the open pit mining production scheduling problem with block processing selectivity, Comput. Oper. Res. 36 (2009) 1064–1089.
[7] E.A. Boyd, Polyhedral results for the precedence-constrained knapsack problem, Discrete Appl. Math. 41 (1993) 185–201.
[8] R. Chicoisne, D. Espinoza, M. Goycoolea, E. Moreno, E. Rubio, A new algorithm for the open-pit mine scheduling problem, Oper. Res. 60 (2012) 517–528.
[9] W.J. Cook, W.H. Cunningham, W.R. Pulleyblank, A. Schrijver, Combinatorial Optimization, Springer, 2009.
[10] W. Cook, D. Espinoza, M. Goycoolea, Computing with domino-parity inequalities for the TSP, INFORMS J. Comput. 19 (3) (2007) 356–365.
[11] J. Correa, A. Schulz, Single-machine scheduling with precedence constraints, Math. Oper. Res. 30 (4) (2005) 1005–1021.
[12] C. Cullenbine, R. Wood, A. Newman, A sliding time window heuristic for open pit mine block sequencing, Opt. Lett. 5 (3) (2011) 365–377.
[13] K. Dagdelen, T. Johnson, Optimum open pit mine production scheduling by Lagrangian parameterization, in: 19th APCOM Symposium of the society of mining engineers, AIME.
[14] D. Espinoza, R. Fukasawa, M. Goycoolea, Lifting, tilting and fractional programming revisited, Oper. Res. Lett. 38 (2010) 559–563.
[15] IBM ILOG CPLEX, IBM®ILOG®CPLEX®Optimization Studio, Version 12 Release 2 Information Center, 2010. URL http://publib.boulder.ibm.com/infocenter/cosinfoc/v12r2/index.jsp.
[16] T. Johnson, Optimum open-pit mine production scheduling, in: A. Weiss (Ed.), A Decade of Digital Computing in the Mining Industry, AIME, New York, 1969, (Chapter 4).
[17] D. Johnson, K. Niemi, On knapsacks, partitions, and a new dynamic programming technique for trees, Math. Oper. Res. 8 (1) (1983) 1–14.
[18] M.W. Padberg, (1,k)-Configurations and facets for packing problems, Math. Program. 18 (1) (1980) 94–99.
[19] M.W. Padberg, G. Rinaldi, A branch-and-cut algorithm for the resolution of large-scale symmetric traveling salesman problems, SIAM Rev. 33 (1991) 60–100.
[20] K. Park, S. Park, Lifting cover inequalities for the precedence-constrained knapsack problem, Discrete Appl. Math. 72 (1997) 219–241.
[21] M. Queyranne, A. Schulz, Approximation bounds for a general class of precedence constrained parallel machine scheduling problems, SIAM J. Comput. 35 (5) (2006) 1241–1253.
[22] R.L.M.J. van de Leensel, C.P.M. van Hoesel, J.J. van de Klundert, Lifting valid inequalities for the precedence constrained knapsack problem, Math. Program. 86 (1999) 161–185.
[23] L. Wolsey, Faces for a linear inequality in 0-1 variables, Math. Program. 8 (2) (1975) 165–178.
[24] L. Wolsey, Facets and strong valid inequalities for integer programs, Oper. Res. 24 (2) (1976) 367–372.