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:

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