

On the exact separation of mixed integer knapsack cuts

Ricardo Fukasawa · Marcos Goycoolea

Received: 29 October 2007 / Accepted: 18 April 2009 / Published online: 6 June 2009
© Springer and Mathematical Programming Society 2009

Abstract During the last decades, much research has been conducted on deriving classes of valid inequalities for mixed integer knapsack sets, which we call *knapsack cuts*. Bixby et al. (The sharpest cut: the impact of Manfred Padberg and his work. MPS/SIAM Series on Optimization, pp. 309–326, 2004) empirically observe that, within the context of branch-and-cut algorithms to solve mixed integer programming problems, the most important inequalities are knapsack cuts derived by the mixed integer rounding (MIR) procedure. In this work we analyze this empirical observation by developing an algorithm to separate over the convex hull of a mixed integer knapsack set. The main feature of this algorithm is a specialized subroutine for optimizing over a mixed integer knapsack set which exploits dominance relationships. The exact separation of knapsack cuts allows us to establish natural benchmarks by which to evaluate specific classes of them. Using these benchmarks on MIPLIB 3.0 and MIPLIB 2003 instances we analyze the performance of MIR inequalities. Our computations, which are performed in exact arithmetic, are surprising: In the vast majority of the instances in which knapsack cuts yield bound improvements, MIR cuts alone achieve over 87% of the observed gain.

Keywords Cutting plane algorithms · Mixed integer knapsack problem · Mixed integer programming

Mathematics Subject Classification (2000) 90C10 · 90C11 · 90C57

R. Fukasawa (✉)
Mathematical Sciences Department, IBM Research, Yorktown Heights, NY, USA
e-mail: rfukasaw@us.ibm.com

M. Goycoolea
School of Business, Universidad Adolfo Ibáñez, Santiago, Chile
e-mail: marcos.goycoolea@uai.cl

1 Introduction

Consider a positive integer n , $b \in \mathbb{Q}$, $a \in \mathbb{Q}^n$, $l \in \{\mathbb{Q} \cup \{-\infty\}\}^n$, $u \in \{\mathbb{Q} \cup \{+\infty\}\}^n$ and $I \subseteq N := \{1, \dots, n\}$. In this article we consider the mixed integer knapsack set

$$K = \{x \in \mathbb{R}^n : ax \leq b, l \leq x \leq u, x_i \in \mathbb{Z}, \forall i \in I\}.$$

More specifically, we are interested in solving the separation problem for $\text{conv}(K)$. That is, given a point $x^* \in \mathbb{Q}^n$, we are interested in either (a) identifying an inequality (π, π_0) which is valid for $\text{conv}(K)$ (equivalently valid for K) and violated by x^* , or (b) proving no such inequality exists.

By obtaining valid inequalities for mixed integer knapsack sets we can also obtain valid inequalities for general Mixed Integer Programming (MIP) problems. In fact, consider the mixed integer set

$$P = \{x \in \mathbb{R}^n : Dx \leq d, l \leq x \leq u, x_i \in \mathbb{Z}, \forall i \in I\}.$$

where D is an $m \times n$ matrix, and $d \in \mathbb{Q}^m$. If (a, b) can be obtained as a non-negative linear combination of rows from (D, d) , then $P \subseteq K$ and we say that K is an implied knapsack set of P . If this is the case, any inequality which is valid for K will also be valid for P and we henceforth call it *knapsack inequality* or *knapsack cut* of P .

Deriving strong knapsack inequalities is of great practical importance to the solution of MIPs. In fact, most cutting planes known for general MIPs are knapsack cuts. For example, Gomory Mixed Integer (GMI) cuts [30, 43] are knapsack cuts derived from the tableaus of linear programming relaxations, and lifted cover inequalities [19, 34] are knapsack cuts derived from the original rows of P . Other classes of knapsack cuts include mixed integer rounding (MIR) cuts and their variations [17, 41, 44], split cuts [16], lift-and-project cuts [8], and group cuts [21, 31]—to name but a few.

In practice, the most successful class of knapsack cuts for MIPs are the GMI/MIR cuts [12]. Dash and Günlük [22] have recently done an empirical study to better understand this practical success. They show that, in a significant number of instances, after adding GMI cuts for all the optimal tableau rows, there are no more violated group cuts for those tableau rows. Their results are quite surprising, and already suggest how strong are GMI cuts within the context of group cuts.

There are, however, natural questions left to answer: In the case where there are some violated group cuts, do these cuts give any significant improvement in the continuous relaxation bound? In the case where there are no violated group cuts, are there any other classes of knapsack cuts that are violated and can yield significantly better bounds? In fact, S. Dash (personal communication) proposed a closely related conjecture that for many problems in MIPLIB 3.0 the combined effect of GMI cuts from different tableau rows is to push the resulting solution x into the intersection of the convex hulls of the mixed integer knapsack sets defined by the tableau rows.

To help resolve these questions, we empirically assess the performance of MIR and GMI inequalities relative to that of all possible knapsack cuts using the objective function of the continuous relaxation as a measure of quality. Formally, consider P as

Table 1 Computational studies of classes of cuts

Paper	Instance set	Implied Knapsack Set (\mathcal{K})	Class of cuts
Boyd [15]	Pure 0-1	Formulation	Knapsack
Yan and Boyd [48]	Mixed 0-1	Formulation	Knapsack
Fischetti and Lodi [26]	Pure integer	All	CG
Bonami et al. [14]	General MIP	All	Pro-CG
Fischetti and Lodi [25]	Pure 0-1	All	Knapsack
Balas and Saxena [9]	General MIP	All	MIR
Dash et al. [22]	General MIP	All	MIR
Kaparis and Letchford [39]	Pure 0-1	Formulation	Knapsack
This paper	General MIP	Formulation and Tableaus	Knapsack

defined above, $c \in \mathbb{Q}^n$, and \mathcal{C} a set of valid inequalities for P . Define,

$$z^*(\mathcal{C}) = \min\{cx : Dx \leq d, l \leq x \leq u, \pi x \leq \pi_o \ \forall (\pi, \pi_o) \in \mathcal{C}\}.$$

Observe that the value $z^*(\mathcal{C})$ defines a benchmark by which to evaluate classes of cuts that are subsets of \mathcal{C} . This idea will be applied in our context in the following way: Given a family of implied knapsack sets \mathcal{K} , let $\mathcal{C}^{\mathcal{K}}$ and $\mathcal{M}^{\mathcal{K}}$ represent, respectively, the set of all knapsack cuts and MIR cuts which can be derived from sets $K \in \mathcal{K}$. Since $\mathcal{M}^{\mathcal{K}} \subseteq \mathcal{C}^{\mathcal{K}}$ it is easy to see that $z^*(\mathcal{C}^{\mathcal{K}}) \geq z^*(\mathcal{M}^{\mathcal{K}})$ and that the proximity of these two values gives an indication of the strength of MIR inequalities derived from that particular family \mathcal{K} .

In this paper we compute the values $z^*(\mathcal{C}^{\mathcal{F}})$ and $z^*(\mathcal{C}^{\mathcal{T}})$, where \mathcal{F} is the set of original formulation rows and \mathcal{T} is the set of first tableau rows. Note that MIR cuts obtained from rows of a simplex tableau are simply GMI cuts. In this study we are able to obtain results for a large subset of MIPLIB 3.0 and MIPLIB 2003 instances.

There have been several related papers computing the value of the continuous relaxation using all cuts in a certain class of knapsack cuts. Boyd [15], Yan and Boyd [48] and Kaparis and Letchford [39] compute $z^*(\mathcal{C}^{\mathcal{F}})$, where \mathcal{F} is the set of original formulation rows. They perform these tests on a subset of pure and mixed 0–1 instances in MIPLIB 3.0 [13]. Balas and Saxena [9] and Dash et al. [22] compute $z^*(\mathcal{M}^{\mathcal{A}})$ for all MIPLIB 3.0 problems, where \mathcal{A} is the set of all implied knapsack sets. This generalizes the results of Fischetti and Lodi [26] and Bonami et al. [14] where they consider Chvátal-Gomory cuts and projected Chvátal-Gomory cuts. Fischetti and Lodi [25] compute $z^*(\mathcal{C}^{\mathcal{A}})$ for a very reduced test set of pure 0–1 problems. We summarize these results and compare them to ours in Table 1.

Finally, note that, since our method establishes benchmarks for knapsack cuts, we use exact arithmetic to ensure that the results obtained are accurate. Indeed, Boyd [15] already cites numerical difficulties in P2756, a pure 0–1 instance, and as a consequence, he gets different bounds than us and Kaparis and Letchford [39].

The organization of this paper is as follows. In the next section, we discuss how to solve the problem of separating over a single mixed integer knapsack set. In Sect. 3 we

describe our methodology for a practical implementation of the separation algorithm. Computational results are presented in Sect. 4, while final remarks are given in Sect. 5.

2 Identifying a violated knapsack cut

Consider $x^* \in \mathbb{Q}^n$ and a feasible mixed integer knapsack set K . In this section we are concerned with developing an effective algorithm for answering the following questions: Is $x^* \in \text{conv}(K)$? If not, can we find an inequality $\pi x \leq \pi_o$ which is valid for K , and such that $\pi x^* > \pi_o$?

Throughout this section we assume that x^* satisfies all constraints defining K except for the integrality requirements. We also assume that for each $i \in N$ either $l_i \neq -\infty$ or $u_i \neq \infty$. Further, we assume that the bound constraints are tight. That is, for every finite bound there exists a point in K which meets that bound at equality. Finally, we assume that $l_i < u_i$ for all $i \in 1, \dots, n$.

Let $\{x^1, x^2, \dots, x^q\}$ and $\{r^1, r^2, \dots, r^t\}$ represent the extreme points and rays of $\text{conv}(K)$. Applegate et al. [3] suggest solving the following linear program to solve the separation problem.

$$\begin{aligned}
 LP_1 : \min \quad & \sum_{i=1}^n (u_i + v_i) \\
 \text{s.t.} \quad & \pi x^k - \pi_o \leq 0 \quad \forall k = 1 \dots q \quad (C1) \\
 & \pi r^k \leq 0 \quad \forall k = 1 \dots t \quad (C2) \\
 & \pi x^* - \pi_o = 1 \quad (C3) \\
 & \pi + u - v = 0 \quad (C4) \\
 & u \geq 0, v \geq 0.
 \end{aligned}$$

Notice that if LP_1 is infeasible, then $x^* \in \text{conv}(K)$, and thus there exists no knapsack cut violated by x^* . Otherwise, the optimal solution (u, v, π, π_o) gives a valid knapsack cut $\pi x \leq \pi_o$ separating x^* from $\text{conv}(K)$ and with maximum value of $\frac{\pi x^* - \pi_o}{\|\pi\|_1}$.

Because LP_1 has an exponential number of rows, it cannot be solved directly. Rather, it must be solved with a dynamic row generation algorithm which separates constraints (C1) and (C2) by solving, for fixed π , the sub-problem $\max\{\pi x : x \in K\}$. We develop a specialized solver for this problem in Sect. 3. We also note that the straightforward dynamic cut generation algorithm that repeatedly optimizes over a relaxation of LP_1 and adds a violated cut (C1) or (C2) to the relaxation is prohibitively slow. For that reason, we present in Sects. 2.1–2.2 steps which were used to help speed up the procedure. We summarize the procedure in Algorithm 1.

2.1 Eliminating variables from LP_1 .

Say that a knapsack cut for K is *trivial* if it is implied by the linear programming relaxation of K . A proof of the following result concerning non-trivial knapsack cuts can be found in Atamtürk [5].

Algorithm 1: KNAPSACK_SEPARATOR(K, x^*)

Input: A mixed integer knapsack set K , and a vector x^* .

Output: An answer to the question “Is $x^* \in \text{conv}(K)$?”. In case that the answer is FALSE, the algorithm also returns a cut separating x^* from $\text{conv}(K)$.

- 1 Simplify the problem by fixing variables which are at their bounds, as indicated in Sect. 2.2, and reduce to a smaller dimensional knapsack separation problem;
 - 2 Formulate problem LP_1 in the reduced variable space without adding any of the constraints in $(C1) - (C2)$;
 - 3 Apply Propositions 1 and 2 to simplify LP_1 by eliminating variables from consideration and adding bounds;
 - 4 Solve LP_1 using the dynamic row generation algorithm;
 - 5 If solving LP_1 reveals that $x^* \in \text{conv}(K)$, then STOP and report that $x^* \in \text{conv}(K)$;
 - 6 Let $(\hat{\pi}, \hat{\pi}_o)$ represent the optimal solution to LP_1 ;
 - 7 Lift cut $\hat{\pi}x \leq \hat{\pi}_o$ to obtain a cut $\pi x \leq \pi_o$ which is valid for K ;
 - 8 STOP and return the cut (π, π_o) ;
-

Proposition 1 *Every non-trivial facet-defining knapsack cut $\pi x \leq \pi_o$ of $\text{conv}(K)$ satisfies the following properties:*

- (i) *If $a_i > 0$, then $\pi_i \geq 0$.*
- (ii) *If $a_i < 0$, then $\pi_i \leq 0$.*
- (iii) *$\pi_i = 0$ for all $i \notin I$ such that $a_i > 0$ and $u_i = +\infty$.*
- (iv) *$\pi_i = 0$ for all $i \notin I$ such that $a_i < 0$ and $l_i = -\infty$.*
- (v) *There exists a constant $\alpha > 0$ such that $\pi_i = \alpha a_i$ for all $i \notin I$ such that $a_i > 0$ and $l_i = -\infty$, and for all $i \notin I$ such that $a_i < 0$ and $u_i = +\infty$.*

The following result concerning violated and non-trivial knapsack cuts is a simple generalization of a remark made in Boyd [15].

Proposition 2 *Consider $x^* \in K^{LP} \setminus \text{conv}(K)$, where K^{LP} stands for the LP relaxation of K . Let $H^+ = \{i \in 1, \dots, n : a_i > 0, x_i^* = l_i\}$ and $H^- = \{i \in 1, \dots, n : a_i < 0, x_i^* = u_i\}$. Then there exists a knapsack cut $\pi x \leq \pi_o$ separating x^* from $\text{conv}(K)$ such that $\pi_i = 0, \forall i \in H^+ \cup H^-$.*

Proof Since $x^* \in K^{LP} \setminus \text{conv}(K)$, then there exists a valid non-trivial inequality $\hat{\pi}x \leq \hat{\pi}_o$ for $\text{conv}(K)$ violated by x^* . We may assume that $\hat{\pi}x \leq \hat{\pi}_o$ is nontrivial, since if there is no such cut, we conclude the proof. Define $\pi_o = \hat{\pi}_o - \sum_{i \in H^+} \hat{\pi}_i l_i - \sum_{i \in H^-} \hat{\pi}_i u_i$ and

$$\pi_i = \begin{cases} 0 & \text{if } i \in H^+ \cup H^- \\ \hat{\pi}_i & \text{otherwise.} \end{cases}$$

We now prove $\pi x \leq \pi_o$ is valid for K . For this, consider $\bar{x} \in \text{conv}(K)$. Consider $i \in H^+$. From Proposition 1 we know that $a_i > 0$ implies $\hat{\pi}_i \geq 0$, and thus, $\hat{\pi}_i \bar{x}_i \geq \hat{\pi}_i l_i$. Likewise, consider $i \in H^-$. From Proposition 1 we know that $a_i < 0$ implies $\hat{\pi}_i \leq 0$, and thus, $\hat{\pi}_i \bar{x}_i \geq \hat{\pi}_i u_i$. Hence, $\hat{\pi}_o \geq \hat{\pi} \bar{x} \geq \pi \bar{x} + \sum_{i \in H^+} \hat{\pi}_i l_i + \sum_{i \in H^-} \hat{\pi}_i u_i$, and so, $\pi \bar{x} \leq \pi_o$. On the other hand, $\hat{\pi}_o < \hat{\pi} x^* = \pi x^* + \sum_{i \in H^+} \hat{\pi}_i l_i + \sum_{i \in H^-} \hat{\pi}_i u_i$. Thus, $\pi x^* > \pi_o$, and we conclude the result. □

From Propositions 1 and 2 we can see that variables π_i with $i \in 1, \dots, n$ can be assumed to have value zero whenever any of the following conditions are met:

- $i \notin I, a_i > 0$ and $u_i = +\infty$,
- $i \notin I, a_i < 0$ and $l_i = -\infty$,
- $a_i > 0$ and $x_i^* = l_i$,
- $a_i < 0$ and $x_i^* = u_i$.

In these cases the corresponding variables can simply be eliminated from LP_1 . Further, from Proposition 1 it can be seen that if we add a non-negative continuous variable α to LP_1 , we can replace all variables π_i satisfying condition (v) of Proposition 1 by the corresponding terms αa_i . This can effectively reduce the variable space by eliminating all remaining continuous unbounded variables. Finally, observe that Proposition 1 allows us to add non-negativity bounds on all variables π_i such that $a_i > 0$, and non-positivity bounds on all variables π_i such that $a_i < 0$.

2.2 Fixing variables and lifting back again

Consider a mixed integer set P , such as the one defined in Sect. 1:

$$P = \{x \in \mathbb{R}^n : Dx \leq d, l \leq x \leq u, x_i \in \mathbb{Z}, \forall i \in I\}.$$

Let U, L be disjoint subsets of $\{1, \dots, n\}$ such that $u_i \neq \infty$ for all $i \in U$ and such that $l_i \neq -\infty$ for all $i \in L$. Define

$$P(L, U) = P \cap \{x \in \mathbb{R}^n : x_i = l_i \forall i \in L\} \cap \{x \in \mathbb{R}^n : x_i = u_i \forall i \in U\}.$$

Given an inequality $\pi x \leq \pi_o$ which is valid for $P(L, U)$, *lifting* consists in obtaining from this an inequality $\hat{\pi} x \leq \hat{\pi}_o$ which is valid for P .

Given a mixed integer knapsack set K and a point x^* we are concerned in determining if $x^* \in \text{conv}(K)$, or if we can find a separating hyperplane. For this we use lifting in the following way. Define $U = \{i \in 1, \dots, n : x_i^* = u_i\}$ and $L = \{i \in 1, \dots, n : x_i^* = l_i\}$. It is easy to see that $x^* \in \text{conv}(K)$ if and only if $x^* \in \text{conv}(K(L, U))$. However, this latter problem is easier to solve since its dimension is much smaller. In practice, this means we can apply the dynamic cut generation methodology on the smaller knapsack constraint:

$$\sum_{i=1, i \notin L \cup U}^n a_i x_i \leq b - \sum_{i \in L} a_i l_i - \sum_{i \in U} a_i u_i.$$

If solving this problem reveals that $x^* \in \text{conv}(K(L, U))$ then we conclude that $x^* \in \text{conv}(K)$. However, if it reveals that $x^* \notin \text{conv}(K(L, U))$ we will obtain an inequality $\pi x \leq \pi_o$ separating x^* from $\text{conv}(K(L, U))$ but which is not necessarily valid for K . In this case we simply apply lifting to obtain a new inequality which is valid for K .

Lifting methodologies typically come in two flavors: Sequential lifting [10,47,49] and sequence-independent lifting [6,35]. In our study we used the sequential lifting procedure described in Fukasawa [29], which allowed us to lift both continuous and general integer variables. The sequence in which we lifted the variables was completely arbitrary (from lowest to highest index).

3 Solving the mixed integer knapsack problem

Consider K as defined in Sect. 1, let $c \in \mathbb{Q}^n$, and assume l_i is finite for each $i \in 1, \dots, n$. In this section we are concerned with solving the problem

$$\max\{cx : x \in K\}, \quad (1)$$

which we call the Mixed Integer Knapsack Problem (MIKP).

Variants of MIKP have long been studied in the research literature. In these it is typically assumed that all coefficients defining the problem are integer, that all variables must take integer values (i.e. no continuous variables are allowed), and that $l_i = 0$ for all $i = 1, \dots, n$. In addition: In the Knapsack Problem (KP) $u_i = 1$ for all $i = 1, \dots, n$, in the Bounded Knapsack Problem (BKP) $u_i < \infty$ for all $i = 1, \dots, n$, and in the Unbounded Knapsack Problem (UKP) $u_i = \infty$ for all $i = 1, \dots, n$. Most modern algorithms for solving KP, BKP, and UKP are based either on branch and bound (following the work of Horowitz and Sahni [36]) and on dynamic programming (following the work of Bellman [11]). However, the most efficient codes seldom make explicit use of Linear Programming and in addition, they never consider the use of both integer and continuous variables. For excellent surveys describing the rich literature on this topic, see Kellerer et al. [40] and Martello and Toth [42].

While many of these algorithms could be adapted for solving 1, they require scaling so as to obtain all-integer coefficients. Our MIKP instances come from solving the separation problem over tableau rows, which are numerically very bad, requiring large scaling coefficients. Because of this we chose to develop an LP-based branch and bound approach. In what follows we describe our algorithm for solving MIKP.

3.1 Detecting unbounded solutions

For each $i \in 1, \dots, n$ define the *efficiency* of variable x_i as $e_i = c_i/a_i$ if $a_i \neq 0$, as $e_i = +\infty$ if $a_i = 0$ and $c_i > 0$, and as $e_i = -\infty$ if $a_i = 0$ and $c_i < 0$.

Proposition 3 *If MIKP is feasible, then it is unbounded if and only if one of the following conditions hold,*

(a) *There exists $i \in 1, \dots, n$ such that*

$$(a_i \leq 0, c_i > 0, u_i = +\infty)$$

(b) *There exist $i, j \in 1, \dots, n$ such that $e_i > e_j$,*

$$(a_i > 0, c_i > 0, u_i = +\infty) \text{ and } (a_j < 0, c_j \leq 0, u_j = +\infty).$$

Proof It is clear that if either condition holds then the problem must be unbounded. We prove the converse. Let $U = \{i \in 1, \dots, n : u_i = +\infty\}$, and consider the recession cone of $\text{conv}(K)$, which is given by $C = \{x \in \mathbb{R}^n : ax \leq 0, x_i \geq 0 \forall i \in U, x_i = 0 \forall i \notin U\}$.

The extreme rays of $\text{conv}(K)$ are the extreme rays of C and must satisfy $n - 1$ of the linearly independent constraints defining C at equality. Thus, all extreme ray vectors have at most two-nonzero values. Let x correspond to an extreme ray satisfying $cx > 0$. We may assume that (a) does not hold. Therefore, for any nonzero component x_i we have that $c_i > 0 \Rightarrow a_i > 0$ and $a_i \leq 0 \Rightarrow c_i \leq 0$. Given that $cx > 0$ we may assume that $c_i > 0$. Thus $a_i > 0$ and since $ax \leq 0$, it follows that there must exist j such that $x_j > 0$ and $a_j < 0$. Finally, $cx > 0$ implies that $c_i x_i > -c_j x_j$, and $ax = 0$ implies that $a_i x_i = -a_j x_j$. From this we have that $e_i > e_j$, and this concludes our proof. \square

Observe that Proposition 3 implies that it can be determined if MIKP is unbounded in linear time. In fact, this can be achieved by looping through the unbounded variables and keeping track of the most efficient one satisfying $(a_i > 0, c_i > 0)$, the least efficient one satisfying $(a_j < 0, c_j \leq 0)$, and checking if any satisfy $(a_i \leq 0, c_i > 0)$.

3.2 Preprocessing

We utilize a simple eight-step pre-processing procedure in order to speed our algorithm. Let $N^+ = \{i \in 1, \dots, n : a_i > 0\}$, and $N^- = \{i \in 1, \dots, n : a_i < 0\}$. We assume that for each $i \in I$ we have $l_i \in \mathbb{Z}$ and $u_i \in \mathbb{Z} \cup \{\infty\}$. If there exists $i \in N^-$ such that $u_i = +\infty$, define $L_{min} = -\infty$. Otherwise, define,

$$L_{min} = \sum_{i \in N^-} a_i u_i + \sum_{i \in N^+} a_i l_i$$

If there exists $i \in N^+$ such that $u_i = +\infty$, define $L_{max} = +\infty$. Otherwise,

$$L_{max} = \sum_{i \in N^-} a_i l_i + \sum_{i \in N^+} a_i u_i$$

Finally, let $sign(x) = 1$ if $x > 0$, $sign(x) = -1$ if $x < 0$ and $sign(x) = 0$ if $x = 0$.

1. **Detect infeasibility.** The problem is infeasible if and only if one of the following conditions hold: (a) $L_{min} \neq -\infty$ and $L_{min} > b$, or (b) there exists $i \in 1, \dots, n$ such that $u_i \neq \infty$ and $l_i > u_i$.
2. **Detect unboundedness.** Use the algorithm discussed in Sect. 3.1 to detect if the problem is unbounded. Observe that if the problem is bounded, all variables x_i satisfying $c_i \geq 0$ and $a_i \leq 0$ will be such that u_i is finite.

3. **Fix variables at their bounds.** Fix to u_i all variables x_i such that $c_i \geq 0$ and $a_i \leq 0$. Fix to l_i to all variables x_i such that $c_i \leq 0$ and $a_i \geq 0$. Observe that after fixing, all variables will be such that $(a_i > 0$ and $c_i > 0)$ or $(a_i < 0$ and $c_i < 0)$.
4. **Tighten upper bounds.** If $L_{min} \neq -\infty$, then the for each $i \in N^+$ we can re-define $u_i := \min\{u_i, (b - L_{min} + a_i l_i)/a_i\}$. If $i \in I$ we can further strengthen this by letting $u_i := \lfloor u_i \rfloor$.
5. **Tighten lower bounds.** If $L_{min} \neq \infty$, then the for each $i \in N^-$ we can re-define $l_i := \max\{l_i, (b - L_{min} + a_i u_i)/a_i\}$. If $i \in I$ we can further strengthen this by letting $l_i := \lceil l_i \rceil$.
6. **Sort variables** Sort variables in order of increasing efficiency. Break ties by letting integer variables precede continuous variables. Break second ties in order of increasing a_i .
7. **Aggregate integer variables.** If two integer variables x_i, x_j are such that $a_i = a_j$ and $c_i = c_j$ aggregate them into a new variable x_k of the same type such that $a_k = a_i = a_j$, $c_k = c_i = c_j$, $l_k = l_i + l_j$ and $u_k = u_i + u_j$.
8. **Aggregate continuous variables.** If two continuous variables x_i, x_j are such that $\frac{c_i}{a_i} = \frac{c_j}{a_j}$ and $sign(a_i) = sign(a_j)$ aggregate them into a new variable x_k of the same type such that $a_k = a_i$, $c_k = c_i$, $l_k = l_i + \frac{a_j}{a_i} l_j$ and $u_k = u_i + \frac{a_j}{a_i} u_j$.

Note that steps one through five can be performed together in a single loop, and that step seven can be performed in a single pass after sorting. For more information and for possible extensions see Savelsbergh [45]. From this point on, we will assume that we are always dealing with instances that have been preprocessed according to the steps above.

3.3 Branch and bound

We use a depth-first-search branch and bound algorithm which always branches on the unique fractional variable. We use a simple linear programming algorithm, a variation of Dantzig's algorithm [20], which runs in linear time by taking advantage of the fact that variables are sorted by decreasing efficiency. We do not use any cutting planes in the algorithm, nor any heuristics to generate feasible solutions. The algorithm uses variable reduced-cost information to improve variable bounds at each node of the tree.

3.4 Domination

Consider x^1 and x^2 , two feasible solutions of MIKP. We say that x^1 *cost-dominates* x^2 if $cx^1 > cx^2$ and $ax^1 \leq ax^2$. On the other hand, we say that x^1 *lexicographically-dominates* x^2 if $cx^1 = cx^2$, $ax^1 \leq ax^2$, and if in addition there exists $i \in \{1, \dots, n\}$ such that $x_i^1 < x_i^2$ and $x_k^1 = x_k^2$ for all $k < i$. We say that a solution is *dominated* if it is cost-dominated or lexicographically-dominated. Observe that there exists a unique non-dominated optimal solution (or none at all).

Traditional branch and bound algorithms work by pruning nodes when (a) they are proven infeasible, or (b) when the LP relaxation of the nodes have value worse than

the best known feasible solution. In our implementation, we additionally prune nodes when (c) it can be shown that every optimal solution in those nodes is dominated.

Using dominance to improve the branch and bound search can have an important impact on the effectiveness of the search [37]. In fact, lexicographic and cost dominance allow us to disregard feasible solutions that are not the unique lexicographically smallest optimum solution, hence significantly reducing the search space.

In general, the problem of detecting if a solution is dominated can be extremely difficult. Fischetti et al. [27,28] detect domination by means of solving an MIP. In what follows we describe a simple methodology for identifying specific cases of domination arising in instances of the mixed integer knapsack problem. Note however, that the notion of dominance as defined above is much more general, and could be used for general mixed integer programming problems as well [27,28,32]. As a final comment, note that Andonov et al [2] have also used domination in the context of dynamic programming algorithms, having had great success in tackling the unbounded knapsack problem.

Throughout this section we will adopt the convention that $\infty + r = \infty, \forall r \in \mathbb{R}$.

3.4.1 Domination between pairs of integer variables

Consider indices $i, j \in I$, and non-zero integers k_i, k_j . If $a_i k_i + a_j k_j \geq 0$ and $c_i k_i + c_j k_j < 0$ we say that (i, j, k_i, k_j) defines an integer cost-domination tuple. If $i < j, k_i > 0, a_i k_i + a_j k_j \geq 0$ and $c_i k_i + c_j k_j = 0$ we say that (i, j, k_i, k_j) defines an integer lexicographic-domination tuple. The propositions below show how domination tuples allow for the easy identification of dominated solutions.

Proposition 4 Consider an integer cost/lexicographic-domination tuple (i, j, k_i, k_j) and let x be a feasible MIKP solution. Define $\delta \in \mathbb{Z}^n$ such that $\delta_i = k_i, \delta_j = k_j$ and $\delta_q = 0$ for all $q \in \{1, \dots, n\} \setminus \{i, j\}$. If

$$l_i + k_i \leq x_i \leq u_i + k_i \text{ and } l_j + k_j \leq x_j \leq u_j + k_j \quad (2)$$

Then x is cost/lexicographically-dominated by $x - \delta$.

To see that this proposition is true, it is simply a matter of observing that condition 2 implies that $(x - \delta)$ is a feasible solution in terms of the bounds, and that $a_i k_i + a_j k_j \geq 0$ and $c_i k_i + c_j k_j \leq 0$ imply the domination.

The following Theorem follows directly from Proposition 4, and illustrates how domination tuples can be used to strengthen a branch and bound algorithm.

Theorem 1 Consider two integer type variables x_i and x_j and a domination tuple (i, j, k_i, k_j) . If in some node of the branch and bound tree we have that $l_i + k_i \leq x_i \leq u_i + k_i$ is satisfied by all solutions in that node, then:

- If $k_j > 0$ we can impose the constraint $x_j \leq l_j + k_j - 1$ in that node,
- If $k_j < 0$ we can impose the constraint $x_j \geq u_j + k_j + 1$ in that node,

and in either case we will not cut off the unique non-dominated optimal solution to the original MIP.

Observe that whenever (c_i, c_j) and (a_i, a_j) are linearly independent there exist an infinite number of cost-domination tuples. Likewise, there exist an infinite number of lexicographic-domination tuples in the linear dependent case.

The two following propositions state that, in each case, there always exists a *minimal* domination tuple. That is, a domination tuple (i, j, k_i, k_j) such that all other domination tuples (i, j, k'_i, k'_j) defined for the same variables, satisfy $|k_i| \leq |k'_i|$ and $|k_j| \leq |k'_j|$. The proof follows from the observation that since (c_i, c_j) and (a_i, a_j) are in the same orthant (they have the same sign in each coordinate), the intersection of the half spaces $c_i x_i + c_j x_j < 0$ and $a_i x_i + a_j x_j \geq 0$ must also be contained in an orthant. Thus, it easily follows that there must be a minimal domination tuple. We now prove this more formally.

Proposition 5 *If (a_i, a_j) and (c_i, c_j) are linearly independent, let,*

$$\mathcal{D} = \{(k_i, k_j) \in \mathbb{Z} \times \mathbb{Z} : k_i, k_j \neq 0, a_i k_i + a_j k_j \geq 0 \text{ and } c_i k_i + c_j k_j < 0\}$$

The following two properties hold:

- (i) \mathcal{D} is contained in the interior of an orthant in $\mathbb{N} \times \mathbb{N}$. More precisely, for every $(k_i, k_j) \in \mathcal{D}$ we have that: $\left(c_i - \frac{c_j}{a_j} a_i\right) < 0$ if and only if $k_i > 0$ and $\left(c_j - \frac{c_i}{a_i} a_j\right) < 0$ if and only if $k_j > 0$.
- (ii) There exists (k_i^o, k_j^o) in \mathcal{D} such that (k_i, k_j) in \mathcal{D} implies $|k_i^o| \leq |k_i|$ and $|k_j^o| \leq |k_j|$.

Proof Consider $(k_i, k_j) \in \mathcal{D}$. Observe that since $sign(a_i) = sign(c_i)$ we can multiply constraint $a_i k_i + a_j k_j \geq 0$ by $-\frac{c_j}{a_j}$ and add it to $c_i k_i + c_j k_j < 0$ and we get that $\left(c_i - \frac{c_j}{a_j} a_i\right) k_i < 0$. Likewise $\left(c_j - \frac{c_i}{a_i} a_j\right) k_j < 0$. Thus, (i) follows.

Consider (x_i, x_j) and (y_i, y_j) in \mathcal{D} . Let $z_i = x_i$ if $|x_i| \leq |y_i|$. Otherwise, let $z_i = y_i$. Let $z_j = x_j$ if $|x_j| \leq |y_j|$. Otherwise, let $z_j = y_j$. Since \mathcal{D} is contained in a pointed cone in one of the four orthants, it is easy to see that $(z_i, z_j) \in \mathcal{D}$. Thus there must exist (k_i^o, k_j^o) in \mathcal{D} satisfying the required condition. □

Proposition 6 *If (a_i, a_j) and (c_i, c_j) are linearly dependent, let,*

$$\mathcal{D} = \{(k_i, k_j) \in \mathbb{N} \times \mathbb{Z} : k_i, k_j \neq 0, a_i k_i + a_j k_j = 0 \text{ and } c_i k_i + c_j k_j = 0\}$$

There exists (k_i^o, k_j^o) in \mathcal{D} such that (k_i, k_j) in \mathcal{D} implies $|k_i^o| \leq |k_i|$ and $|k_j^o| \leq |k_j|$.

Proof Given the linear dependence of (a_i, a_j) and (c_i, c_j) we have that,

$$\mathcal{D} = \left\{ (k_i, k_j) \in \mathbb{N} \times \mathbb{Z} : k_i, k_j \neq 0, \frac{k_i}{k_j} = -\frac{a_j}{a_i} \right\}$$

That this set is non-empty follows from the fact that a_i and a_j are rational numbers. Let (k_i^o, k_j^o) be the point in \mathcal{D} with the smallest k_i^o . Clearly, $|k_j^o| = \left|\frac{a_i}{a_j}\right| |k_i^o| \leq \left|\frac{a_i}{a_j}\right| |k_i| = |k_j|$ for all $(k_i, k_j) \in \mathcal{D}$. □

From Proposition 6 we see that if (a_i, a_j) and (c_i, c_j) are linearly dependent, then it is trivial to compute a minimal domination pair. From Proposition 5 it follows that if (a_i, a_j) and (c_i, c_j) are linearly independent, a minimal domination tuple can be obtained by solving the following two-dimensional integer programming problem:

$$\begin{aligned}
 & \text{minimize } |k_i| + |k_j| \\
 & \text{subject to,} \\
 & \quad c_i k_i + c_j k_j \geq 0 \\
 & \quad a_i k_i + a_j k_j \leq 0 \\
 & \quad |k_i| \geq 1, |k_j| \geq 1, \\
 & \quad k_i, k_j \in \mathbb{Z}
 \end{aligned} \tag{3}$$

Further, given that we know the feasible region of Problem 3 is contained in an orthant of $\mathbb{Z} \times \mathbb{Z}$, we can remove the absolute values and make the problem linear.

While in theory Problem 3 can be solved in polynomial time [24, 38, 46], in practice we found it easier to solve the problem by simple enumeration. For this, loop through all possible values of k_i , starting from $k_i = 1$. Once k_i is fixed it is trivial to solve the problem for k_j (or determine infeasibility). From the propositions above it suffices to stop at the first value of k_i for which the problem is feasible.

In order to use the above propositions in the branch and bound algorithm we compute what we call a *domination table*. This table is defined as a list of all possible (minimal) domination tuples. In order to obtain a domination table we must loop through all possible pairs of variables before solving the problem. Observe that because we will only use these tables to strengthen bounds after branching, as in Theorem 1, we only need to store domination tuples (i, j, k_i, k_j) such that $|k_i| \leq (u_i - l_i)$ and $|k_j| \leq (u_j - l_j)$.

As a final note, we observe that once domination is enforced through branching, the branch-and-bound tree will be of finite size, regardless of there existing unbounded variables. In fact, in any branch-and-bound tree of infinite size there must exist a branching-path of infinite length. Further, because there are only a finite number of integer variables, there must exist two unbounded ones on which the tree branches an infinite number of times. Since for every pair of integer variables there exists a minimal domination tuple, it is not difficult to see that if we enforce domination, at some point branching on one of these variables would impose a bound on the other variable, thus making the impossible path. Goycoolea [32] analogously shows that using domination tables one could explicitly impose bounds a priori, which also guarantees finite termination of branch-and-bound.

3.4.2 Domination between integer and continuous variables

Let $C_{j,k} = \{i \in C : j \leq i \leq k\}$. In this section we show how branching on integer variable x_i can be used to strengthen the bounds on continuous variables x_j with $j \in C_{1,i}$ or those with $j \in C_{i+1,n}$. Our attention will focus on instances of MIKP satisfying $a_i > 0$ for all $i \in 1, \dots, n$. This assumption is made only to simplify the

notation and analysis. All results can easily be extended to the more general case in which coefficients can be positive or negative.

Proposition 7 *Consider a pre-processed instance of MIKP, $\max\{cx : x \in K\}$. For any $x \in K$ and $i \in I$ we have that x is cost-dominated if there exists $k \in C_{i,n}$ such that $e_k < e_i$, $l_k < x_k$, and,*

$$x_i \leq u_i - 1 \quad \text{and} \quad \sum_{j \in C_{i,n}} a_j(x_j - l_j) \geq a_i \tag{4}$$

We have that x is dominated if the following condition holds:

$$x_i \geq l_i + 1 \quad \text{and} \quad \sum_{j \in C_{1,i}} a_j(u_j - x_j) \geq a_i \tag{5}$$

Proof Assume that $x_i \leq u_i - 1$ and $\sum_{j \in C_{i,n}} a_j(x_j - l_j) \geq a_i$. Define $\delta \in \mathbb{R}^n$ as follows: let $\delta_i = -1$, $\delta_j = 0 \forall j \in I \setminus \{i\}$, and $\delta_j = 0 \forall j \in C \setminus C_{i,n}$. For $j \in C_{i,n}$ define δ_j so that (i) $0 \leq \delta_j \leq x_j - l_j$, (ii) $\sum_{j \in C_{i,n}} a_j \delta_j = a_i$, and (iii) $\delta_k > 0$. Define $x' = x - \delta$. It is easy to see that x' satisfies the integrality constraints. Further, (i) implies that x' satisfies all of the bound constraints, and (ii) implies that $ax' = ax$. Thus, $x' \in K$. Finally, since the variables are sorted by decreasing efficiency, we have that $cx' \geq cx$, and from (iii) we obtain that this inequality is strict. Hence, x' cost-dominates x . The proof that 5 implies x is dominated is analogous. \square

Consider $i \in I$. Define,

$$f(i) = \min\{k \in C_{i,n} : \sum_{j \in C_{i,k}} a_j(u_j - l_j) > a_i\}$$

Likewise, define:

$$g(i) = \max\{k \in C_{1,i} : \sum_{j \in C_{k,i}} a_j(u_j - l_j) > a_i\}$$

Proposition 8 *Consider a pre-processed instance of MIKP, $\max\{cx : x \in K\}$, and $i \in I$. Let $f = f(i)$ and $g = g(i)$. If in some node of the branch and bound tree we have that $x_i \leq u_i - 1$ is satisfied by all solutions in that node, and in addition, $e_f < e_i$, then we may also impose that:*

$$x_j = l_j \quad \forall j \in C_{f+1,n} \tag{6}$$

and,

$$x_f \leq l_f + \frac{a_i - \sum_{j \in C_{i,f-1}} a_j(u_j - l_j)}{a_f} \tag{7}$$

Likewise, if in some node of the branch and bound tree we have that $x_i \geq l_i + 1$ is satisfied by all solutions in that node, then we may also impose that:

$$x_j = u_j \quad \forall j \in C_{1,g-1}, \tag{8}$$

and,

$$x_g \geq u_g - \frac{a_i - \sum_{j \in C_{g,i}} a_j(u_j - l_j)}{a_g}. \tag{9}$$

In either case we will only be cutting off dominated solutions to the original MIP.

Proof We will prove that if $e_f < e_i$, then any solution $x \in K$ that satisfies $x_i \leq u_i - 1$ and does not satisfy 6 or 7 is cost-dominated.

First, let $L_f = l_f + \frac{a_i - \sum_{j \in C_{i,f-1}} a_j(u_j - l_j)}{a_f}$ and $L_j = l_j, \forall j \in C_{f+1,n}$. Assume, for the sake of contradiction, that $x_t > L_t$ for some $t \in C_{f,n}$. If $x_j < u_j$ for any $j \in C_{1,t-1}$, then we can define \bar{x} so that $\bar{x}_k = x_k, \forall k \neq j, t, \bar{x}_t = x_t - \epsilon, \bar{x}_j = x_j + \epsilon a_t/a_j$ and have that for some $\epsilon > 0$ small enough, $\bar{x} \in K$. Since the variables are sorted by decreasing efficiency, and observing that no two continuous variables can have the same efficiency, \bar{x} cost-dominates x .

Therefore we may assume that $x_j = u_j$ for all $j \in C_{1,t-1}$. But then since $t \geq f$, we have that $\sum_{j \in C_{i,n}} a_j(x_j - l_j) = \sum_{j \in C_{i,t-1}} a_j(u_j - l_j) + \sum_{j \in C_{i,n}} a_j(x_j - l_j) > a_i$ and that $e_t \leq e_f < e_i$. Therefore, by Proposition 7 we will have that x is cost-dominated.

The second part of the Proposition is analogous. □

4 Computational experiments

In this section we describe the results of our computational experiments. All implementations were developed in the ‘‘C’’ and ‘‘C++’’ programming languages, using the Linux operating system (v2.4.27) on Intel Xeon dual-processor computers (2GB of RAM, at 2.66GHz). Since generating cuts which are invalid is a real point of concern, we use exact arithmetic rather than the floating point arithmetic typically used in computer codes. Specifically, we used the exact LP solver of Applegate et al. [4] for solving LP_1 , and the GNU Multiple Precision (GMP) Arithmetic library [33] for implementing all other computations.

In order to test our results we use the MIPLIB 3.0 [13] and MIPLIB 2003 [1] data sets. Results are organized as follows. In Sect. 4.1 we utilize the methodology developed in Sect. 2 to benchmark the MIR inequalities. In Sect. 4.2 we analyze the performance of the mixed integer knapsack algorithm developed in Sect. 3.

4.1 Benchmarking the MIR inequalities

We now describe our experiments computing the values $z^*(\mathcal{C}^K)$ and $z^*(\mathcal{M}^K)$ and our use of these values to benchmark the performance of MIR inequalities, as detailed

Table 2 Benchmarks for formulation rows

Instance	ORIG-GAP (%)	MIR-PERF (%)	KNAP-PERF(%)	MIR-REL (%)
arki001	0.02	12.99	13.12	99.01
fiber	61.55	91.07	93.82	97.06
gen	0.16	99.78	99.78	100
gesa2	1.18	69.96	71.03	98.48
gesa3	0.56	47.80	49.33	96.90
gt2	36.41	92.56	94.52	97.93
1152lav	1.39	0.01	1.36	0.41
lseu	25.48	67.91	76.09	89.25
mitre	0.36	89.14	100.00	89.14
mod008	5.23	71.23	89.21	79.84
mod010	0.24	18.34	18.34	100
nsrand-ipx	4.53	47.94	48.75	98.32
p0033	18.40	76.33	87.42	87.31
p0201	9.72	33.78	33.78	100
p0282	31.56	94.08	98.59	95.42
p0548	96.37	52.84	84.34	62.66
p2756	13.93	44.46	86.35	51.49
qnet1	10.95	50.48	89.06	56.68
qnet1_o	24.54	84.32	95.12	88.65
rgn	40.63	57.49	57.49	100
roll3000	13.91	58.27	61.62	94.56
sp97ar	1.38	2.19	2.49	88.08
timtab1	96.25	21.88	22.07	99.13
timtab2	93.49	13.82	13.96	99

in Sect. 1. In these experiments we only consider the sets $\mathcal{K} = \mathcal{F}$, i.e., the family of knapsack sets induced by the original formulation rows, and $\mathcal{K} = \mathcal{T}$, i.e., the family of knapsack sets induced by the tableau rows of the optimal LP relaxation solution.

As a first step, we implemented the MIR separation heuristic described in Goycoolea [32] to compute an approximation of $z^*(\mathcal{M}^{\mathcal{K}})$ which we call $z_M^{\mathcal{K}}$. We compute $z_M^{\mathcal{K}}$ by repeatedly adding MIR inequalities valid for each $K \in \mathcal{K}$ to the LP relaxation and re-optimizing it until no more cuts are found.

As a second step, we compute $z^*(\mathcal{C}^{\mathcal{K}})$ by repeatedly adding knapsack cuts found for each $K \in \mathcal{K}$ to the LP relaxation and re-optimizing it until no more cuts are found. We note that before using the exact separation routine for some $K \in \mathcal{K}$ one can check if there is an MIR inequality (or any other known knapsack cut) valid for K that is violated by the current LP relaxation. This helps speed up computations significantly.

Tables 2 and 3 present the results for $\mathcal{K} = \mathcal{F}$ and $\mathcal{K} = \mathcal{T}$ respectively. For each problem instance let z_{UB}^* represent the value of the optimal (or best known) solution and z_{LP}^* the LP relaxation value. For each set \mathcal{K} and each instance we compute the following performance measures:

Table 3 Benchmarks for Tableau rows

Instance	ORIG-GAP (%)	MIR-PERF (%)	KNAP-PERF (%)	MIR-REL (%)
air03	0.38	100.00	100.00	100
a1c1s1	91.33	18.94	18.94	100
aflow30a	15.10	13.41	14.77	90.78
aflow40b	13.90	7.91	8.02	98.71
bell3a	1.80	60.15	60.15	100
bell5	3.99	14.53	14.68	98.94
blend2	8.99	20.63	20.63	100
dcmulti	2.24	50.46	50.49	99.94
egout	73.67	55.33	55.33	100
fiber	61.55	75.89	77.27	98.21
fixnet6	69.85	11.08	11.08	100
flugpl	2.86	11.74	11.74	100
gen	0.16	61.67	61.97	99.52
gesa2	1.18	28.13	28.13	99.98
gesa2_o	1.18	29.55	29.65	99.67
gesa3	0.56	45.76	45.83	99.85
gesa3_o	0.56	49.96	49.99	99.94
khb05250	10.31	75.14	75.14	100
liu	69.60	27.02	27.02	100
lseu	25.48	61.21	61.21	100
manna81	1.01	100	100	100
misc03	43.15	7.24	7.24	100
misc06	0.07	26.98	26.98	100
misc07	49.64	0.72	0.72	100
mod008	5.23	22.57	22.59	99.92
mod011	13.86	22.17	22.17	100
modglob	1.49	18.05	18.05	100
nsrand-ipx	4.53	38.94	39.26	99.20
nw04	3.27	22.37	22.37	100
p0033	18.40	74.71	74.71	100
p0201	9.72	34.36	34.36	100
pp08a	62.61	50.97	50.97	100
qiu	601.15	3.47	3.47	100
rentacar	5.11	27.42	27.42	100
rgn	40.63	9.78	9.78	100
set1ch	41.31	39.18	39.18	100
swath	28.44	33.22	33.58	98.91
timtab1	96.25	24.74	24.80	99.76
timtab2	93.49	15.57	15.76	98.79
tr12-30	89.12	60.27	60.27	100
vpm1	22.92	47.27	49.09	96.30
vpm2	28.08	19.17	19.39	98.85

ORIG-GAP: The value of the LP relaxation gap:

$$\frac{z_{UB}^* - z_{LP}^*}{|z_{UB}^*|}.$$

MIR-PERF: How much of the LP gap was closed by the MIR heuristic:

$$\frac{z_M^{\mathcal{K}} - z_{LP}^*}{z_{UB}^* - z_{LP}^*}.$$

KNAP-PERF: How much of the LP gap was closed by the knapsack cuts:

$$\frac{z^*(\mathcal{C}^{\mathcal{K}}) - z_{LP}^*}{z_{UB}^* - z_{LP}^*}.$$

MIR-REL: Relative performance of MIR separation heuristic. That is, how much of the LP gap closed by the knapsack cuts was closed by the MIR cuts:

$$\frac{z_M^{\mathcal{K}} - z_{LP}^*}{z^*(\mathcal{C}^{\mathcal{K}}) - z_{LP}^*}.$$

4.1.1 Knapsack cuts derived from formulation rows

We now discuss the computational results obtained for $\mathcal{K} = \mathcal{F}$. Of the 92 instances in MIPLIB 3.0 and MIPLIB2003, we were able to compute the values $z_M^{\mathcal{K}}$ for 79 of them. The 13 problems which we were unable to solve were atlanta-ip, cap6000, dano3mip, harp2, momentum2, momentum3, markshare1, markshare2, mkc, msc98-ip, net12, rd-rplusc-21 and stp3d. Of the 79 problems solved, 4 of them were such that ORIG-GAP was equal to 0.0 (disctom, dsbmip, enigma, and noswot), and 51 of them were such that KNAP-PERF and MIR-PERF were both equal to 0.0. In Table 2 we present the results for the remaining 24 problems.

First, note that knapsack cuts alone can considerably close the remaining LP gap in some problems (column KNAP-PERF). In fact, in 11 problems out of the 24 problems in which knapsack cuts improved the gap, over 84% of the gap was closed, and in 15 out of 24 problems, over 50% of the gap was closed. On average, the GAP closed by the knapsack cuts among these 24 instances is around 62%. Note, however, that in 51 instances knapsack cuts do nothing to improve the gap. If we consider the average GAP closed including these 51 instances, the average drops to 19.84%.

Second, consider the column MIR-REL in which we can get an idea how close is the bound obtained by using only MIR cuts compared to the one obtained by using all knapsack cuts. Observe that in 19 out of the 24 problems we close over 87% of the GAP closed by the knapsack cuts by using the MIR cuts alone. This indicates that MIR inequalities are a very important subset of knapsack inequalities for the instances considered. A natural question is the following: How much could we improve the value of MIR-PERF if we used an exact MIR separation algorithm as opposed to a heuristic?

In an attempt to answer this question we fine-tuned the settings of the MIR heuristic for the problems p0033, p0548 and qnet1. In these, we managed to improve the value of MIR-PERF from 87.31 to 100%, from 62.66 to 63.66% and from 56.68 to 77.27% respectively. This indicates that the true value of MIR-REL may be much closer to 100% than suggested by the table.

4.1.2 Knapsack cuts derived from tableau rows

In order to obtain \mathcal{T} for a given problem instance, we first solve the linear programming relaxation of the corresponding problem and let \mathcal{T} be the implied knapsack sets induced by the optimal tableau rows in the augmented variable space consisting of both structural and slack variables. When attempting to generate a cut from a tableau row our algorithms all worked in this augmented space. Whenever a cut was generated from a tableau row, we would first substitute out slack variables before adding it back to the LP. Slack variables were always assumed to be non-negative and continuous.

We now discuss the computational results obtained for $\mathcal{K} = \mathcal{T}$. Of the 92 instances in MIPLIB 3.0 and MIPLIB2003, we were able to compute the values $z_M^{\mathcal{T}}$ for 48 of them. Of these 48 problems, 4 of them were such that ORIG-GAP was equal to 0.0 (disctom, dsbmip, enigma, and noswot), and 2 of them were such that KNAP-PERF and MIR-PERF were both equal to 0.0 (stein27 and stein45). In Table 3 we present the results for the remaining 42 problems.

First, it is important to remark that separating knapsack cuts from tableau rows was much more difficult than from original formulation rows. This is because tableau rows are typically more dense, with more ill-conditioned coefficients, and with more continuous variables. This can be seen in that only 48 of 92 instances were solved, as opposed to the 79 which were solved for formulation rows.

Second, it is interesting to note that the value KNAP-PERF is very erratic, uniformly ranging in values from 100 to 0.0%. In contrast to the case of formulation rows, only 2 instances are such that KNAP-PERF is 0.0%.

The last, and perhaps most startling observation, is that the MIR-REL is very often close to 100%. If this result were true in general, it would be very surprising. However, because there are still 44 instances for which we were not able to obtain results, one must be careful. It could be the case that those 44 instances are precisely the ones which have MIR-REL much lower than 100%.

4.2 Performance of the mixed integer knapsack solver

We now compare the performance of our MIKP algorithm (*kbb*) with the performance of CPLEX 9.0 [18] (*cpix*), the only alternative for MIKP we know to date. CPLEX was used with all of its default settings, except for the tolerance, which was set to 10^{-6} . Our MIKP algorithm was used with double floating arithmetic, with a tolerance of 10^{-6} . Finally, note that we also tested a hybrid algorithm (*cpx*) which combined the *kbb* pre-processor (as in Sect. 3.2) with the CPLEX solver.

In our first implementation of the separation algorithm we had incorporated a version of *kbb* which did not use domination branching nor reduced cost bound




Fig. 1 Performance profile comparing *kbb*, *cpx* and *cpp*

improvements. However, this algorithm was not efficient enough to achieve our aim of separating knapsack cuts derived from MIPLIB instances. When running this version of the code, we saved around 130,000 instances which took our algorithm more than 2.5 s to solve. Most of those instances were very similar to each other, since many would originate from separating over the convex hull of the same implied knapsack set. When a single implied knapsack set had more than one knapsack instance associate to it, we eliminated all of these but the one which was generated last. This was due to the observation that these were usually the hardest ones to solve. We thus reduced our instances to only 288. Most of these were posed in exact arithmetic. After converting them to double floating point, some of them became very easy to solve. Hence, from these 288, we removed all instances in which *kbb*, *cpx* and *cpp* all took less than 1 s to solve in double arithmetic. This left us with 67 instances, which are the ones used in the following experiments. These instances are available upon request.

In Fig. 1 we present the performance profiles of running times of *kbb*, *cpx* and *cpp*. Each point (x, y) of the curves tells us that in y percent of the instances, this particular solver was at most x times slower than the fastest solver for that instance (for a more detailed explanation of performance profiles, see [23]). For instance, by looking at the point where $x = 1$ we see that *kbb* is the fastest in about 77% of the instances and by looking at the point where $x = 10$, it is at most 10 times slower than either *cpx* or *cpp* in over 85% of the instances. Also, note that giving the preprocessed instance to CPLEX does not seem to make a big difference. This can be explained since CPLEX has its own preprocessing, which probably makes most of our preprocessing unnecessary.




Fig. 2 Performance profile comparing *kbb*, *kbd*, *kbr*, and *kbo*

It is clear from this Figure that the *kbb* algorithm outperforms *cpx* and *cpp* in this instance set. Note that this does not necessarily mean that *kbb* solves every instance faster than *cpx*, but rather, that cumulatively, *kbb* performs better. Moreover, *cpx/cpp* fails to find the optimum solution in 11/8 instances, since it runs out of memory after creating too large a branch and bound tree.

In Fig. 2 we compare the different versions of our MIKP solver:

- **kbo**: Without domination branching or reduced-cost bound improvement.
- **kbr**: With reduced-cost bound improvement but without domination branching.
- **kbd**: With domination branching but without reduced-cost bound improvement
- **kbb**: With both domination branching and reduced-cost bound improvement

It is interesting to see in Fig. 2 how domination branching and reduced-cost fixing interact to improve the performance of the algorithm. Domination branching, when used as a single feature, clearly helps decrease solution times. This is not the case of reduced-cost fixing, however, which actually makes the algorithm perform worse. What is very surprising, however, is that when both features are used together, the performance is altogether markedly improved.

It makes sense that these two features should complement each other, since, whenever bounds are changed, more domination conditions can be applied which leads to additional pruning of the tree. On the other hand, the extra computational effort of improving bounds by reduced-cost does not translate in any pruning of the branch-and-bound tree. Indeed, in all instances, the size of the branch-and-bound tree for *kbo* and *kbr* were exactly the same.

5 Final remarks

One of the goals of this study has been to assess the overall effectiveness of MIR inequalities relative to knapsack cuts. We observe that in most test problems from MIPLIB 3.0 and MIPLIB 2003, the bound obtained by using just MIR inequalities is very similar in value (if not equal) to the bound obtained using all possible knapsack cuts. This observation helps explain the difficulty in outperforming MIRs by using other cuts from tableau and formulation rows, and suggests that for further bound improvements on this instance set we might have to consider new row aggregation schemes, or cuts derived from multiple row systems. We would like to point out, however, that it may still be possible to obtain a significant improvement on the practical performance of MIR inequalities using other knapsack cuts, for instance if the performance is measured in terms of total time it takes to reach a similar gap level.

We put great care into ensuring that the generated cuts are valid and that the procedure runs correctly, but this makes the methodology very slow. For example, some of the unsolved instances ran for over a week without a final answer being reported. Part of the difficulty arises from the fact that exact arithmetic is being employed. We have observed that on the average performing exact arithmetic computations takes several orders of magnitudes longer than floating-point computations.

As a final remark, we note that recently, knapsack cuts have been used in practice to help solve some combinatorial optimization problems [7]. Following up on this, it would be interesting to see if we can also use our knapsack cut generation procedure in practice to help solve some classes of mixed integer programming problems.

Acknowledgments The authors would like to thank William J. Cook for his support and encouragement in doing this work. They are also grateful to Sanjeeb Dash and Oktay Günlük for hosting them at the IBM T.J. Watson Research Center as student interns during the summers of 2004 and 2006. The discussions and questions raised during these internships were key factors in motivating this research.

References

1. Achterberg, T., Koch, T., Martin, A.: MIPLIB 2003. *Oper. Res. Lett.* **34**(4), 1–12 (2006). doi: [10.1016/j.orl.2005.07.009](https://doi.org/10.1016/j.orl.2005.07.009). <http://www.zib.de/Publications/abstracts/ZR-05-28/>. See <http://miplib.zib.de>
2. Andonov, R., Poirriez, V., Rajopadhye, S.: Unbounded knapsack problem: dynamic programming revisited. *Eur. J. Oper. Res.* **123**, 394–407 (2000)
3. Applegate, D., Bixby, R.E., Chvátal, V., Cook, W.: TSP cuts which do not conform to the template paradigm. In: *Computational Combinatorial Optimization, Optimal or Provably Near-Optimal Solutions [based on a Spring School]*, pp. 261–304. Springer-Verlag GmbH, London (2001)
4. Applegate, D., Cook, W., Dash, S., Espinoza, D.: Exact solutions to linear programming problems. *Oper. Res. Lett.* **35**, 693–699 (2007)
5. Atamtürk, A.: On the facets of the mixed-integer knapsack polyhedron. *Math. Program.* **98**, 145–175 (2003)
6. Atamtürk, A.: Sequence independent lifting for mixed-integer programming. *Oper. Res.* **52**, 487–490 (2004)
7. Avella, P., Boccia, M., Vasilyev, I.: A computational study of exact knapsack separation for the generalized assignment problem. Technical report available at Optimization Online (2006)
8. Balas, E., Perregaard, M.: A precise correspondence between lift-and-project cuts, simple disjunctive cuts, and mixed integer Gomory cuts for 0–1 programming. *Math. Program.* **94**, 221–245 (2003)
9. Balas, E., Saxena, A.: Optimizing over the split closure. *Math. Program.* **113**(2), 219–240 (2008)
10. Balas, E., Zemel, E.: Facets of knapsack polytope from minimal covers. *SIAM J. Appl. Math.* **34**(1), 119–148 (1978)

11. Bellman, R.E.: *Dynamic Programming*. Princeton University Press, Princeton (1957)
12. Bixby, R., Gu, Z., Rothberg, E., Wunderling, R.: Mixed integer programming: a progress report. In: *The sharpest cut: the impact of Manfred Padberg and his work*. MPS/SIAM Series on Optimization, pp. 309–326
13. Bixby, R.E., Ceria, S., McZeal, C.M., Savelsbergh, M.W.P.: An updated mixed integer programming library: MIPLIB 3.0. *Optima* **58**, 12–15 (1998)
14. Bonami, P., Cornuéjols, G., Dash, S., Fischetti, M., Lodi, A.: Projected Chvátal Gomory cuts for mixed integer linear programs. *Math. Program.* **113**(2), 241–257 (2008)
15. Boyd, A.E.: Fenchel cutting planes for integer programs. *Oper. Res.* **42**, 53–64 (1992)
16. Cook, W., Kannan, R., Schrijver, A.: Chvátal closures for mixed integer programming problems. *Math. Program.* **47**, 155–174 (1990)
17. Cornuéjols, G., Li, Y., Vanderbussche, D.: K-cuts: a variation of Gomory mixed integer cuts from the LP tableau. *INFORMS J. Comput.* **15**(4), 385–396 (2003)
18. CPLEX: <http://www.ilog.com/products/cplex>
19. Crowder, H., Johnson, E., Padberg, M.: Solving large-scale zero-one linear-programming problems. *Oper. Res.* **31**(5), 803–834 (1983)
20. Dantzig, G.B.: Discrete variable extremum problems. *Oper. Res.* **5**(2), 266–277 (1957)
21. Dash, S., Günlük, O.: On the strength of gomory mixed-integer cuts as group cuts. IBM research report RC23967 (2006)
22. Dash, S., Günlük, O., Lodi, A.: On the MIR closure of polyhedra. In: Fischetti, M., Williamson, D. (eds.) *IPCO*, Lecture notes in computer science, vol. 4513, pp. 337–351 (2007)
23. Dolan, E., Moré, J.: Benchmarking optimization software with performance profiles. *Math. Program.* **91**(2), 201–213 (2001)
24. Eisenbrand, F., Laue, S.: A linear algorithm for integer programming in the plane. *Math. Program.* **102**(2), 249–259 (2005)
25. Fischetti, M., Lodi, A.: On the knapsack closure of 0–1 integer linear problems. In: *Presentation at 10th International Workshop on Combinatorial Optimization*, Aussois (2006). Available at <http://www-id.imag.fr/IWCO2006/slides/Fischetti.pdf>
26. Fischetti, M., Lodi, A.: Optimizing over the first Chvátal closure. *Math. Programm. B* **110**(1), 3–20 (2007)
27. Fischetti, M., Salvagnin, D.: A local dominance procedure for mixed-integer linear programming. <http://www.dei.unipd.it/~fisch/papers/MIPdominance.pdf>
28. Fischetti, M., Toth, P.: A new dominance procedure for combinatorial optimization problems. *Oper. Res. Lett.* **7**, 181–187 (1988)
29. Fukasawa, R.: *Single-row mixed-integer programs: Theory and computations*. Ph.D. thesis, Georgia Institute of Technology (2008)
30. Gomory, R.E.: Early integer programming (reprinted). *Oper. Res.* **50**(1), 78–81 (2002)
31. Gomory, R.E., Johnson, E.: Some continuous functions related to corner polyhedra I. *Math. Program.* **3**, 23–85 (1972)
32. Goycoolea, M.: *Cutting planes for large mixed integer programming models*. Ph.D. thesis, Georgia Institute of Technology (2006)
33. Granlund, T.: The GNU multiple precision arithmetic library. Available on-line at <http://www.swox.com/gmp/>
34. Gu, Z., Nemhauser, G.L., Savelsbergh, M.W.P.: Lifted cover inequalities for 0–1 integer programs: computation. *INFORMS J. Comput.* **10**, 427–437 (1998)
35. Gu, Z., Nemhauser, G.L., Savelsbergh, M.W.P.: Sequence independent lifting in mixed integer programming. *J. Combin Optimiz.* **4**(1), 109–129 (2000)
36. Horowitz, E., Sahni, S.: Computing partitions with applications to the knapsack problem. *J. ACM* **21**, 277–292 (1974)
37. Ibaraki, T.: The power of dominance relations in branch-and-bound algorithms. *J. ACM* **24**, 264–279 (1977)
38. Kannan, R.: A polynomial algorithm for the two-variable integer programming problem. *J. ACM* **27**, 118–122 (1980)
39. Kaparis, K., Letchford, A.: Separation algorithms for 0–1 knapsack polytopes. Technical report available at *Optimization Online* (2007)
40. Kellerer, H., Pferschy, U., Pisinger, D.: *Knapsack Problems*. Springer, Berlin (2004)

41. Marchand, H., Wolsey, L.: Aggregation and mixed integer rounding to solve MIPs. *Oper. Res.* **49**, 363–371 (2001)
42. Martello, S., Toth, P.: *Knapsack Problems: Algorithms and Computer Implementations*. Wiley, New York (1990)
43. Nemhauser, G.L., Wolsey, L.A.: A recursive procedure for generating all cuts for 0-1 mixed integer programs. *Math. Program.* **46**, 379–390 (1990)
44. Nemhauser, G.L., Wolsey, L.A.: *Integer and Combinatorial Optimization*. Discrete Mathematics and Optimization. Wiley, New York (1999)
45. Savelsbergh, M.: Preprocessing and probing for mixed integer programming problems. *ORSA J. Comput.* **6**, 445–454 (1994)
46. Scarf, H.E.: Production sets with indivisibilities—part II: the case of two activities. *Econometrica* **49**, 395–423 (1981)
47. Wolsey, L.: Facets and strong valid inequalities for integer programs. *Oper. Res.* **24**(2), 367–372 (1976)
48. Yan, X.Q., Boyd, E.A.: Cutting planes for mixed-integer knapsack polyhedra. *Math. Program.* **81**(2), 257–262 (1998)
49. Zemel, E.: Lifting the facets of zero-one polytopes. *Math. Program.* **15**(1), 268–277 (1978)