

On the Exact Separation of Mixed Integer Knapsack Cuts

Ricardo Fukasawa¹ and Marcos Goycoolea²

¹ H. Milton Stewart School of Industrial and Systems Engineering
Georgia Institute of Technology
`rfukasaw@isye.gatech.edu`

² School of Business
Universidad Adolfo Ibáñez
`marcos.goycoolea@uai.cl`

Abstract. During the last decades, much research has been conducted deriving classes of valid inequalities for single-row mixed integer programming polyhedrons. However, no such class has had as much practical success as the MIR inequality when used in cutting plane algorithms for general mixed integer programming problems. In this work we analyze this empirical observation by developing an algorithm which takes as input a point and a single-row mixed integer polyhedron, and either proves the point is in the convex hull of said polyhedron, or finds a separating hyperplane. The main feature of this algorithm is a specialized subroutine for solving the Mixed Integer Knapsack Problem which exploits cost and lexicographic dominance. Separating over the entire closure of single-row systems allows us to establish natural benchmarks by which to evaluate specific classes of knapsack cuts. Using these benchmarks on Miplib 3.0 instances we analyze the performance of MIR inequalities. Computations are performed in exact arithmetic.

Keywords: cutting plane algorithms, integer programming

1 Introduction

Consider positive integers n, m and let $d \in \mathbb{Q}^m$, $D \in \mathbb{Q}^{m \times n}$, $l \in \{\mathbb{Q} \cup \{-\infty\}\}^n$ and $u \in \{\mathbb{Q} \cup \{+\infty\}\}^n$. Let $I \subseteq N := \{1, \dots, n\}$ and consider the mixed integer set:

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

We say that a mixed integer knapsack set of the form,

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

with $b \in \mathbb{Q}$, $a \in \mathbb{Q}^n$ is *implied* by P if (a, b) is a non-negative linear combination of rows obtained from (D, d) . Observe that if K is implied by P , then $P \subseteq K$. Hence, any inequality which is valid for K is also valid for P . We henceforth call such inequalities *knapsack cuts* derived from K .

Deriving strong knapsack cuts is of great practical importance to Mixed Integer Programming (MIP). In fact, most cutting planes known for general mixed integer programming are knapsack cuts. For example, Gomory Mixed Integer cuts [19, 28] are knapsack cuts derived from the tableaus of linear programming relaxations, and Lifted Cover Inequalities [12, 23] are knapsack cuts derived from the original rows of P . Other classes of knapsack cuts include mixed-integer-rounding (MIR) cuts and their variations [11, 26, 28], split cuts [10], lift-and-project cuts [4], and group cuts [15, 20] – to name but a few.

In this paper we discuss an empirical methodology for evaluating sub-classes of knapsack cuts. Formally, consider P as 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} . For example, consider a family of implied knapsack sets \mathcal{K} and let $\mathcal{C}^{\mathcal{K}}$ represent the set of all knapsack cuts which can be derived from some set $K \in \mathcal{K}$. Likewise, let $\mathcal{M}^{\mathcal{K}}$ represent the set of all MIR inequalities which can be derived from some set $K \in \mathcal{K}$. Given that $\mathcal{M}^{\mathcal{K}} \subseteq \mathcal{C}^{\mathcal{K}}$ 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 our computational experiments we will consider two specific families of implied knapsack sets: The set \mathcal{F} of all formulation rows of P ; and, given a basic solution of the simplex algorithm, the set \mathcal{T} of all tableau rows.

Boyd [8] and Yan and Boyd [30] compute $z^*(\mathcal{C}^{\mathcal{F}})$ for a subset of pure and mixed 0-1 instances in MIPLIB 3.0 [7]. Fischetti and Lodi [18] extend this result by computing $z^*(\mathcal{C}^{\mathcal{A}})$, where \mathcal{A} is the set of all implied knapsack polyhedra, for a similar test set of pure 0-1 problems.

In this paper we compute the values $z^*(\mathcal{C}^{\mathcal{F}})$ and $z^*(\mathcal{C}^{\mathcal{T}})$ for a larger subset of MIPLIB 3.0 instances, including general mixed integer problems. We compare these values to estimates of $z^*(\mathcal{M}^{\mathcal{F}})$ and $z^*(\mathcal{M}^{\mathcal{T}})$ (i.e., the bounds obtained by using MIR inequalities) and attempt to address the well acknowledged observation that it is difficult to identify classes of knapsack inequalities which systematically outperform the MIR inequality in broad test sets. Recently, Dash and Günlük [15] also try to analyze this issue in terms of cuts from the cyclic group problem.

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. This methodology described requires the use of a subroutine for solving the mixed integer knapsack problem. An algorithm for solving this problem is discussed in Sect. 3. Computational results are presented in Sect. 4, while final remarks and a discussion ensues in Sect. 5.

2 Identifying violated knapsack cuts

Consider $x^* \in \mathbb{R}^n$ and a mixed integer knapsack set K . In this section we address 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 $\pi x^* > \pi_o$?

We assume that K has no free variables, since it is easy to substitute a free variables by two non-negative variables. Let $\{x^1, x^2, \dots, x^q\}$ and $\{r^1, r^2, \dots, r^t\}$ represent the extreme points and extreme rays of $\text{conv}(K)$. The following proposition, which follows from the work of Applegate et. al [1], allows us to address this question.

Proposition 1. *Consider the following linear programming (LP) problem with variables $u, v, \pi \in \mathbb{R}^n$, and $\pi_o \in \mathbb{R}$:*

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

If this problem is infeasible, then $x^ \in \text{conv}(K)$, and thus there exists no knapsack cut violated by x^* . Otherwise, this problem admits an optimal solution (u, v, π, π_o) such that inequality $\pi x \leq \pi_o$ is a valid knapsack cut maximizing:*

$$\frac{\pi x^* - \pi_o}{\|\pi\|_1}$$

That is, the hyperplane defined by (π, π_o) maximizes the L_1 distance to x^ .*

Because LP_1 has an exponential number of constraints, we use a dynamic cut generation algorithm to solve the problem. We begin with constraints (C3)–(C4) and a subset of constraints (C1) – (C2). The cut generation algorithm requires solving the problem $\max\{\pi x : x \in K\}$ at each iteration. If this problem is unbounded at any given iteration, then there exists an extreme ray r^j of $\text{conv}(K)$ such that $\pi r^j > 0$. That is, we have identified a violated constraint. If this problem is not unbounded, then there exists an optimal solution corresponding to an extreme point x^k of $\text{conv}(K)$. If $\pi x^k > \pi_o$ then we have found a violated constraint. Otherwise, it means that all constraints of the problem are satisfied. Solving the oracle problem is discussed in Sect. 3.

Notice that in general, it is not possible to assure that the solution of $\max\{\pi x : x \in K\}$ given by the oracle will correspond to an extreme point or ray of $\text{conv}(K)$. However, constraints (C1) – (C2) can be re-defined in terms of all points/rays of K without affecting the correctness of Proposition 1. Even though this would result in an infinite number of constraints, under very mild assumptions [17], the dynamic cut generation algorithm will still converge in a finite

number of iterations.

In order to speed up the solution of LP_1 we make use of certain characterizations of violated knapsack cuts.

Let $K = \{x \in \mathbb{R}^n : ax \leq b, l \leq x \leq u, x_i \in \mathbb{Z}, \forall i \in I\}$. We may assume without loss of generality [21] that the bound constraints are tight. 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 [3].

Proposition 2. *Every non-trivial facet-defining knapsack cut $\pi x \leq \pi_o$ of $\text{conv}(K)$ satisfies the following properties:*

- (i) *If $a_i > 0$, $\pi_i \geq 0$*
- (ii) *If $a_i < 0$, $\pi_i \leq 0$*
- (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 technique employed in Boyd [8].

Proposition 3. *Consider $x^* \notin \text{conv}(K)$. Let $H^+ = \{i \in N : a_i > 0, x_i^* = l_i\}$ and $H^- = \{i \in N : a_i < 0, x_i^* = u_i\}$. If there does not exist a trivial inequality separating x^* from $\text{conv}(K)$, then there exists a knapsack cut $\pi x \leq \pi_o$ such that $\pi_i = 0, \forall i \in H^+ \cup H^-$.*

We make use of Propositions 2 – 3 in the following way: We restrict the signs of coefficients according to Proposition 2 items (i) and (ii). Coefficients π_i with $i = 1, \dots, n$ which can be assumed to be zero are eliminated from LP_1 . Further, a single variable is used for all coefficients π_i with $i = 1, \dots, n$ for which we know that $\pi_i = \alpha a_i$. Note that this last reduction is equivalent to aggregating the unbounded continuous variables into a single variable.

Two other techniques are used to speed up the separation process. The first one uses the fact that MIR inequalities are knapsack cuts. With that in mind, we first apply an MIR separation heuristic to try to find violated knapsack cuts and only use the above separation procedure if the MIR heuristic fails.

The other technique relies on the following simple observation. Let $U = \{i \in N : x_i^* = u_i\}$ and $L = \{i \in N : x_i^* = l_i\}$. If we define,

$$K^* = K \cap \{x : x_i = u_i \forall i \in U\} \cap \{x : x_i = l_i \forall i \in L\},$$

we know that $x^* \in \text{conv}(K)$ iff $x^* \in \text{conv}(K^*)$. Thus, answering the question: “Is $x^* \in \text{conv}(K)$?” can be done in a space of usually much smaller dimension by testing instead if $x^* \in \text{conv}(K^*)$.

If our test shows that $x^* \in \text{conv}(K^*)$, we are done with the separation since we know that in this case $x^* \in \text{conv}(K)$. However, if $x^* \notin \text{conv}(K^*)$ we still need to get a cut separating x^* from $\text{conv}(K)$ and thus we have to run our separation algorithm in the original space. Notice, however, that if $x^* \notin \text{conv}(K^*)$, our separation algorithm will return a cut separating x^* from $\text{conv}(K^*)$, so one could potentially lift this cut to obtain a cut separating x^* from $\text{conv}(K)$. We have not implemented this feature yet, but we expect that it will significantly speed up our algorithm.

To summarize, we outline the complete algorithm below:

Algorithm 1: Outline of knapsack separation process

Input: x^* and K
Output: $x^* \in \text{conv}(K)$ or a cut separating x^* from $\text{conv}(K)$

```

begin
  Run the MIR separation heuristic
  if cut found then
    | return the MIR cut separating  $x^*$  from  $\text{conv}(K)$ 
  else
    | Apply Propositions 1 and 2 to simplify  $LP_1$ 
    | Solve  $LP_1$  in a reduced space to separate  $x^*$  from  $\text{conv}(K^*)$ 
    | if  $x^* \in \text{conv}(K^*)$  then
    | | return  $x^* \in \text{conv}(K)$ 
    | else
    | | Solve  $LP_1$  in the original variable space to separate  $x^*$  from  $\text{conv}(K)$ 
  end

```

3 Solving the Mixed Integer Knapsack problem

In this section we are concerned with the problem of solving the Mixed Integer Knapsack Problem (MIKP),

$$\max\{cx : x \in K\} \tag{1}$$

We will assume that the problem is feasible, and are interested in either (a) proving that the problem is unbounded by finding an extreme ray r^* of $\text{conv}(K)$, or (b) computing the optimal value of the problem by finding the optimal solution $x^* \in K$.

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 [24]) and on dynamic programming (following the work of Bellman [6]). 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, the reader is advised to consult Kellerer et al [25] and Martello and Toth [27].

While it is reasonable to expect that many of these algorithms could be adapted for solving our general case with a mix of continuous, integer, bounded and unbounded variables, the fact that they are designed to work with integer coefficients raises certain concerns with regards to the application discussed in this paper. In fact, part of our motivation is to study the efficacy of cuts derived from tableau rows. However, these rows are rarely made up of integer coefficients, and what's more, they are typically very ill conditioned. Thus, scaling them so as to obtain integers may result in extremely large numbers. Considering this important shortcoming, and the need to further study these algorithms in order to account for the mixed use of bounded, unbounded, continuous and integer variables, our approach has been to pursue an LP-based branch and bound approach, which seems naturally suited to mixed integer programming problems. This issue, however, is one which merits further research. In what follows we describe our algorithm for solving MIKP.

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 $c_i < 0$. In addition, we say that x_i is a *potentiator* if,

$$(a_i \leq 0, c_i > 0, u_i = +\infty) \text{ or } (a_i \geq 0, c_i < 0, l_i = -\infty).$$

We say that x_i is an *incrementor* if,

$$(a_i > 0, c_i > 0, u_i = +\infty) \text{ or } (a_i < 0, c_i < 0, l_i = -\infty).$$

We say that x_i is a *decrementor* if,

$$(a_i > 0, c_i \geq 0, l_i = -\infty) \text{ or } (a_i < 0, c_i \leq 0, u_i = +\infty).$$

By identifying a potentiator, or instead, by identifying the most efficient incrementor and the least efficient decrementor, it is possible to easily establish if a problem is unbounded, as shown by the following Proposition:

Proposition 4. *MIKP is unbounded if and only if one of the following conditions hold,*

- MIKP admits a potentiator x_j .
- MIKP admits an incremator x_i and a decremator x_j such that $e_i > e_j$.

Note that Proposition 4 implies that it can be determined if MIKP is unbounded in linear time. Note also that once the potentiator, or instead, the incremator and decremator have been identified, it is easy to construct an extreme ray of $\text{conv}(K)$.

Preprocessing

We consider the following four-step preprocessing algorithm (see [21],[29]) which assumes the problem is not unbounded.

1. 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 $a_i \geq 0$.
2. Make all bounds as tight as possible.
3. Aggregate variables. If two variables x_i, x_j of the same type (integer or continuous) 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 $u_k = u_i + u_j$.
4. Sort variables in order of decreasing efficiency. Break ties checking for variable types (integer or continuous).

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 [13], 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.

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 $cx^1 = cx^2$ and $ax^1 \leq ax^2$, and if in addition, there exists $i \in 1, \dots, n$ such that $x_i^1 < x_i^2$ and $x_k^1 = x_k^2, \forall k \in 1, \dots, (i-1)$. 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 it can be shown that the optimal solution in those nodes has value worse than a bound previously obtained. 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. 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. In what follows we describe a simple methodology for identifying specific cases of domination.

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-dominance tuple. If $k_i \geq 0$, $a_i k_i + a_j k_j \geq 0$ and $c_i k_i + c_j k_j = 0$ we say that (i, j, k_i, k_j) defines an integer lexicographic-dominance tuple. Observe that whenever (c_i, a_i) and (c_j, a_j) are linearly independent there exist an infinite amount of cost-dominance pairs. Likewise, there exist an infinite amount of lexicographic-dominance tuples in the linear dependence case. However, in each case, there always exists a *minimal* dominance tuple. That is, a dominance tuple (i, j, k_i, k_j) such that all other dominance 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 propositions below show how dominance tuples allow for the easy identification of dominated solutions.

Proposition 5. *Consider an integer cost-dominance tuple (i, j, k_i, k_j) and let x be a feasible MIKP solution. If any of the following three conditions hold:*

- $k_i > 0$, $k_j > 0$, $x_i \geq l_i + k_i$ and $x_j \geq l_j + k_j$,
- $k_i < 0$, $k_j > 0$, $x_i \leq u_i + k_i$ and $x_j \geq l_j + k_j$,
- $k_i < 0$, $k_j < 0$, $x_i \leq u_i + k_i$ and $x_j \leq u_j + k_j$.

Then x is cost-dominated.

Proposition 6. *Consider an integer lexicographic-dominance tuple (i, j, k_i, k_j) and let x be a feasible MIKP solution. If either of the following conditions hold:*

- $k_j > 0$, $x_i \geq l_i + k_i$, and $x_j \geq l_j + k_j$,
- $k_j < 0$, $x_i \geq l_i + k_i$, and $x_j \leq u_j + k_j$,

then x is lexicographically-dominated.

To see that these propositions are true, it is simply a matter of observing that if the conditions hold for a feasible x , then defining x' so that $x'_i = x_i - k_i$, $x'_j = x_j - k_j$ and $x'_k = x_k$ for $k \neq i, j$, we have x' is feasible and dominates x .

The following propositions illustrate how dominance tuples can be used to strengthen branch and bound algorithm. This is achieved by preventing nodes with dominated solutions from being created through additional enforced bound changes.

Proposition 7. *Consider two integer type variables x_i and x_j and a dominance tuple (i, j, k_i, k_j) such that $k_i > 0$. If in some node of the branch and bound tree we impose $x_i \geq l_i + \alpha_i$, where $\alpha_i \geq k_i$, 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.

The case $k_i < 0$ is analogous.

In order to use the above propositions in the branch and bound algorithm we compute what we call a *domination table* before initiating the solve. This table is defined as a list of all possible (minimal) domination tuples. Observe that we only need store domination tuples (i, j, k_i, k_j) such that $|k_i| \leq (u_i - l_i)$ and $|k_j| \leq (u_j - l_j)$. In order to compute domination tuples we perform a simple enumeration algorithm which uses bounds to identify where to start and stop the enumerations.

4 Computational experiments

In this section, our computational experiments are described. All implementations were compiled using the “C” and “C++” programming languages, using the Linux operating system (v2.4.27) and Intel Xeon dual-processor computers (2GB of RAM, at 2.66GHz). Since generating cuts which are invalid is a real point of concern, we found it appropriate to use the exact arithmetic, both for solving LP_1 , and for the MIKP oracle. Thus, we used Applegate et al. [2] exact LP solver for LP_1 , and the GNU Multiple Precision (GMP) Arithmetic library [22] to implement the MIKP algorithm.

4.1 The optimization oracle

We first compare the performance of our MIKP algorithm (“kbb”) with the performance of CPLEX 9.0 (“cpx”), the only alternative for MIKP we know of to date. Note that CPLEX was ran with all its default settings, except for the tolerance, which was set to 10^{-6} . Note also that our MIKP algorithm was ran using double floating arithmetic, with a tolerance of 10^{-6} .

In our first implementation of the separation algorithm we had incorporated a version of kbb which did not use domination branching. We quickly realized that this algorithm was not efficient enough. When running this version of the code, we saved all problems which took our algorithm more than 2.0 seconds to solve. These are the 1,556 problems that we now use to compare cpx with the full version of kbb. It is important to note that by the nature of the way these instances were generated, there might be some of instances that are very similar to each other.

In Fig. 1 we present a histogram summarizing the running times of kbb and cpx. Each point in the curves represents the number of instances which were solved within a given maximum time. For instance, note that the number of instances solved to optimality by kbb within a second is roughly 1150, whereas the number of instances solved to optimality by cpx is roughly 700. Note that the time is represented in logarithmic scale. Further, observe that the hardest instance for kbb takes several hundred seconds – roughly ten times less than the hardest instance for cpx.




Fig. 1. Histogram comparing KBB algorithm with CPLEX.

It is clear from this histogram that the kbb algorithm outperforms cpx in the instance set. Note that this does not necessarily mean that kbb solves every instance faster than cpx, but rather, that cumulatively, kbb performs better. In fact, on average, kbb takes 81% less time than cpx, and explores 37.5% less branch-and-bound nodes. Moreover, in 49 instances, CPLEX fails to find the optimum solution since it runs out of memory after creating too large a branch and bound tree.

4.2 Knapsack cuts

We next use an implementation of the algorithms presented in Sect. 2 and Sect. 3 to compare the practical performance of MIR cuts against the performance of separating all possible knapsack cuts. As detailed in Sect. 1, given a family \mathcal{K} of knapsack sets implied by P , such a comparison can be made by comparing the values $z^*(\mathcal{C}^{\mathcal{K}})$ and $z^*(\mathcal{M}^{\mathcal{K}})$. In this article we only consider the set $\mathcal{K} = \mathcal{F}$, i.e., the family of knapsack sets induced by the original formulation rows, and the set $\mathcal{K} = \mathcal{T}$, i.e., the family of knapsack sets induced by the simplex tableau rows of the optimal LP solution for the original LP relaxation.

Computing $z^*(\mathcal{M}^{\mathcal{K}})$ is NP-hard [9], so instead we approximate this value using an MIR separation heuristic. Given a point x^* , for every $K \in \mathcal{K}$ we try to find MIR inequalities that are violated by x^* . We add these inequalities to the LP relaxation of P and repeat the process until no more MIR inequalities are found. The MIR inequalities for each K are derived by a separation heuristic which combines scaling and variable complementation techniques (for details see [21], [26], and [14]). Denote by $z_M^{\mathcal{K}}$ the objective function value at the end of the procedure. Since this is just a heuristic, after completing a run, there may be violated MIR inequalities which have not been identified. Therefore $z_M^{\mathcal{K}}$ should be considered an estimate of $z^*(\mathcal{M}^{\mathcal{K}})$.

Note that though the MIR separation problem is NP-hard, one could use the approaches of Balas and Saxena [5] or Dash, Günlük and Lodi [16] to better approximate $z^*(\mathcal{M}^{\mathcal{K}})$.

To compute $z^*(\mathcal{C}^{\mathcal{K}})$, we proceed as follows. Given a fractional solution, we loop through all of the mixed integer knapsack sets $K \in \mathcal{K}$. For each of these we invoke the procedure outlined in Sect. 2 and identify a violated cut if such exists. After completing this loop we add the cuts to the problem and repeat. The procedure ends when for every K we can prove that there is no violated knapsack cut.

Computational tests are performed on all MIPLIB 3.0 instances using the mixed integer knapsack sets $\mathcal{K} = \mathcal{F}$ and $\mathcal{K} = \mathcal{T}$. 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:

LP-PERF: Performance of the original LP formulation. That is, the value of the LP relaxation gap:

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

KNAP-PERF: Performance of the knapsack cuts. That is, 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-PERF: Performance of MIR separation heuristic. That is, how much of the LP gap closed by the knapsack cuts was closed by the MIR cuts:

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

Knapsack cuts derived from formulation rows

In this section we analyze the performance of knapsack and MIR inequalities on formulation rows of MIPLIB 3.0 instances. Results are summarized in Table 1. Of the 59 instances in the library, we eliminated eight instances which were unfinished at the time of writing the article (arki001, cap6000, dano3mip, harp2, mitre, mod008, pk1 and rout), three for which LP-PERF was equal to 0.0 (dsbmip, enigma, and noswot), and thirty two for which KNAP-PERF and MIR-PERF were both equal to 0.0.

First, note that knapsack cuts alone can considerably close the remaining LP gap in some problems (column KNAP-PERF). In fact, in 9 problems out of the 16 problems in which knapsack cuts improved the gap, over 84% of the gap was closed, and in 14 out of 16 problems, over 50 % of the gap was closed. On average, the GAP closed by the knapsack cuts among these 16 instances is around 71%. It is interesting, however, that in thirty two instances knapsack

Table 1. Benchmarks for Formulation Closure

Instance	LP-PERF	KNAP-PERF	MIR-PERF
fiber	61.55%	93.82%	97.06 %
gen	0.16%	99.78%	100.00 %
gesa2	1.18%	71.03%	98.48 %
gesa3	0.56%	49.33%	96.90 %
gt2	36.41%	94.52%	97.93 %
l152lav	1.39%	1.36%	0.41 %
lseu	25.48%	76.09%	88.25 %
mod010	0.24%	18.34%	100.00 %
p0033	18.40%	87.42%	87.31 %
p0201	9.72%	33.78%	100.00 %
p0282	31.56%	98.59%	95.42 %
p0548	96.37%	84.34%	62.76 %
p2756	13.93%	86.35%	51.49 %
qnet1	10.95%	89.06%	56.68 %
qnet1_o	24.54%	95.12%	88.65 %
rgn	40.63%	57.49%	100.00 %

cuts should do nothing to improve the gap. If in addition we consider in our average the thirty two instances for which KNAP-PERF is 0.0%, this drops to 23.66%.

Second, consider the column MIR in which we can get an idea of how well the mixed integer rounding cut closure compares to the knapsack cut closure. Observe that of the 16 problems, in 12 of them, by using the MIR cuts alone, we close over 87% of the GAP closed by the knapsack cuts. This indicates that MIR inequalities are a very important subset of knapsack inequalities; at least 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 and qnet1. In these, we managed to improve the value of MIR-PERF from 87.31% to 100% and from 56.68% to 77.27% respectively.

Knapsack cuts derived from tableau rows

In this section we analyze the performance of knapsack and MIR inequalities on tableau rows of MIPLIB 3.0 instances. For this we compute z_{LP}^* and store the tableau rows in the set of knapsack polyhedra $\mathcal{K} = \mathcal{T}$, which we use for all subsequent computations. Results are summarized in Table 2. Of the 59 instances in the library, we eliminated thirty two instances which were unfinished at the time of writing the article, three for which LP-PERF was equal to 0.0 (dsbmip, enigma, and noswot), and two for which KNAP-PERF and MIR-PERF were both equal to 0.0 (stein27 and stein45).

Table 2. Benchmarks for Tableau Closure

Instance	LP-PERF	KNAP-PERF	MIR-PERF
air03	0.38 %	100.00 %	100.00%
bell3a	1.80 %	60.15 %	100.00%
bell5	3.99 %	14.68 %	98.94%
dcmulti	2.24 %	50.49 %	99.94%
egout	73.67 %	55.33 %	100.00%
fixnet6	69.85 %	11.08 %	100.00%
flugpl	2.86 %	11.74 %	100.00%
gesa2	1.18 %	28.13 %	99.98%
gesa2_o	1.18 %	29.65 %	99.67%
khb05250	10.31 %	75.14 %	100.00%
misc03	43.15 %	7.24 %	100.00%
misc06	0.07 %	26.98 %	100.00%
misc07	49.64 %	0.72 %	100.00%
modglob	1.49 %	18.05 %	100.00%
p0033	18.40 %	74.71 %	100.00%
p0201	9.72 %	34.36 %	100.00%
pp08a	62.61 %	50.97 %	100.00%
qiu	601.15 %	3.47 %	100.00%
rgn	40.63 %	9.78 %	100.00%
set1ch	41.31 %	39.18 %	100.00%
vpm1	22.92 %	49.09 %	96.30%
vpm2	28.08 %	19.39 %	98.85%

First, it is important to remark that separating knapsack cuts from tableau rows is considerable more difficult than separating knapsack cuts from original formulation rows. This is due to several reasons: Tableau rows are typically much more dense, coefficients tend to be numerically very bad, and rows tend to have a lot of continuous variables. This added difficulty is reflected in the fact that out of 59 instances, in two days of runs we just managed to solve 24 instances to completion, as opposed to the 48 which we solved when considering 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 two instances are such that KNAP-PERF is 0.0%.

The last, and perhaps most startling observation, is that the MIR-PERF is always at 100%, if not very close. If this result were true in general, it would be very surprising. However, because there are still thirty two instances which have not been solved one must be very careful. Because of the way in which we computed these numbers, it could be the case that those instances with MIR-PERF close to 100% are easier for our methodology to solve. It is very reasonable to expect that instances with MIR-PERF well below 100% are more difficult to solve as they require more iterations of the knapsack separation algorithm as opposed to iterations of the MIR separation heuristic.

5 Final remarks

It is important to note that these results are very preliminary. 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 K NAP-PERF values computed took as much as 5 days to obtain. Some of the unsolved instances have been ran for over a week without a final answer being reported. We are currently developing further techniques by which these computations can be accelerated. Part of the difficulty arises from the fact that exact arithmetic is being employed. In average, we have observed that performing exact arithmetic computations take 100 times longer than floating point arithmetic computations.

One of the main goals of this study has been to assess the overall effectiveness of MIR inequalities relative to knapsack cuts. The motivation being the empirical observation that though much research has been conducted studying inequalities derived from single row systems, no such class of inequalities has been able to systematically improve upon the performance of MIRs. In this regard, the results we present are surprising. We observe that in most test problems, the bound obtained by optimizing over the MIR closure is very similar in value (if not equal) to the bound obtained optimizing over the knapsack closure. Though it is important to note that this observation is limited in the number of test problems considered, it does help explain the lack of success in generating other cuts from tableau and formulation rows, and, suggests that for further bound improvements we might have to consider new row aggregation schemes, or cuts derived from multiple row systems.

References

1. D. Applegate, R. E. Bixby, V. Chvátal, and W. Cook. 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]*, pages 261–304, London, UK, 2001. Springer-Verlag GmbH.
2. D. Applegate, W. Cook, S. Dash, and D. Espinoza. Exact solutions to linear programming problems. *Submitted to Operations Research Letters*, 2006.
3. A. Atamtürk. On the facets of the mixed-integer knapsack polyhedron. *Mathematical Programming*, 98:145–175, 2003.
4. E. Balas and M. Perregaard. A precise correspondence between lift-and-project cuts, simple disjunctive cuts, and mixed integer Gomory cuts for 0-1 programming. *Mathematical Programming*, 94:221–245, 2003.
5. E. Balas and A. Saxena. Optimizing over the split closure. *Mathematical Programming*, To appear.
6. R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
7. R. E. Bixby, S. Ceria, C. M. McZeal, and M. W. P Savelsbergh. An updated mixed integer programming library: MIPLIB 3.0. *Optima*, (58):12–15, June 1998.
8. A. E. Boyd. Fenchel cutting planes for integer programs. *Operations Research*, 42:53–64, 1992.

9. A. Caprara and A. Letchford. On the separation of split cuts and related inequalities. *Mathematical Programming*, 94(2-3):279–294, 2003.
10. W. Cook, R. Kannan, and A. Schrijver. Chvátal closures for mixed integer programming problems. *Mathematical Programming*, 47:155–174, 1990.
11. G. Cornuéjols, Y. Li, and D. Vanderbussche. K-cuts: A variation of gomory mixed integer cuts from the LP tableau. *Inform Journal On Computing*, 15:385–396, 2003.
12. H. Crowder, E.L. Johnson, and M. Padberg. Solving large-scale zero-one linear-programming problems. *Operations Research*, 31:803–834, 1983.
13. G. B. Dantzig. Discrete variable extremum problems. *Operations Research*, 5(2):266–277, 1957.
14. S. Dash, M. Goycoolea, and O. Günlük. Two-step mir inequalities for mixed-integer programs. *Optimization Online*, Jul 2006.
15. S. Dash and O. Günlük. On the strength of gomory mixed-integer cuts as group cuts. *IBM research report RC23967*, 2006.
16. S. Dash, O. Günlük, and A. Lodi. MIR closures of polyhedral sets. Available online at http://www.optimization-online.org/DB_HTML/2007/03/1604.html.
17. D. G. Espinoza. *On Linear Programming, Integer Programming and Cutting Planes*. PhD thesis, School of Industrial and Systems Engineering, Georgia Institute of Technology, March 2006.
18. M. Fischetti and A. Lodi. On the knapsack closure of 0-1 integer linear problems. Presentation at 10th International Workshop on Combinatorial Optimization, Aussois (2006). Available at <http://www-id.imag.fr/IWCO2006/slides/Fischetti.pdf>.
19. R. E. Gomory. Early integer programming (reprinted). *Operations Research*, 50:78–81, Jan 2002.
20. R. E. Gomory and E.L. Johnson. Some continuous functions related to corner polyhedra I. *Mathematical Programming*, 3:23–85, 1972.
21. M. Goycoolea. *Cutting Planes for Large Mixed Integer Programming Models*. PhD thesis, Georgia Institute of Technology, 2006.
22. T. Granlund. The GNU multiple precision arithmetic library. Available on-line at <http://www.swox.com/gmp/>.
23. Z. Gu, G. L. Nemhauser, and M. W. P. Savelsbergh. Lifted cover inequalities for 0-1 integer programs: Computation. *INFORMS Journal on Computing*, 10:427–437, 1998.
24. E. Horowitz and S. Sahni. Computing partitions with applications to the knapsack problem. *Journal of the ACM*, 21:277–292, 1974.
25. H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, Berlin, Germany, 2004.
26. H. Marchand and L.A. Wolsey. Aggregation and mixed integer rounding to solve MIPs. *Operations Research*, 49:363–371, 2001.
27. S. Martello and P. Toth. *Knapsack Problems: Algorithms and Computer Implementations*. J. Wiley, New York, 1990.
28. G. L. Nemhauser and L. A. Wolsey. A recursive procedure for generating all cuts for 0-1 mixed integer programs. *Mathematical Programming*, 46:379–390, 1990.
29. M.W.P. Savelsbergh. Preprocessing and probing for mixed integer programming problems. *ORSA Journal on Computing*, 6:445–454, 1994.
30. X. Q. Yan and E. A. Boyd. Cutting planes for mixed-integer knapsack polyhedra. *Mathematical Programming*, 81:257–262, 1998.