

# Pruning Moves

Matteo Fischetti

Dipartimento di Ingegneria dell'Informazione, University of Padova,  
35131 Padova, Italy, matteo.fischetti@unipd.it

Domenico Salvagnin

Dipartimento di Matematica Pura ed Applicata, University of Padova,  
35121 Padova, Italy, salvagni@math.unipd.it

The concept of dominance among nodes of a branch-and-bound tree, although known for a long time, is typically not exploited by general-purpose mixed-integer linear programming (MILP) codes. The starting point of our work was the general-purpose dominance procedure proposed in the 1980s by Fischetti and Toth, where the dominance test at a given node of the branch-and-bound tree consists of the (possibly heuristic) solution of a restricted MILP only involving the fixed variables. Both theoretical and practical issues concerning this procedure are analyzed, and important improvements are proposed. In particular, we use the dominance test not only to fathom the current node of the tree, but also to derive variable configurations called “nogoods” and, more generally, “improving moves.” These latter configurations, which we rename “pruning moves” so as to stress their use in a node-fathoming context, are used during the enumeration to fathom large sets of dominated solutions in a computationally effective way. Computational results on a testbed of MILP instances whose structure is amenable to dominance are reported, showing that the proposed method can lead to a considerable speedup when embedded in a commercial MILP solver.

*Key words:* mixed-integer programming; constraint programming; dominance procedure; nogoods

*History:* Accepted by John N. Hooker, Jr., Area Editor for Constraint Programming and

Optimization; received May 2008; revised November 2008, January 2009; accepted March 2009.

Published online in *Articles in Advance*.

## 1. Introduction

The mixed-integer linear programming (MILP) paradigm is among the most powerful and most used methods for modeling and solving both real-life and theoretical optimization problems; see, e.g., Papadimitriou and Steiglitz (1982), Schrijver (1998), and Garfinkel and Nemhauser (1972). Almost 50 years of active research in the field has produced huge improvements in the solving capability of MILP codes, as reported in, e.g., Bixby et al. (2000).

A concept playing a potentially relevant role in trying to keep the size of the branch-and-bound tree as small as possible is that of *dominance*. Following, e.g., Papadimitriou and Steiglitz (1982), a branching node  $\alpha$  is said to be dominated by another node  $\beta$  if every feasible solution in the subtree of  $\alpha$  corresponds to a solution in the subtree of  $\beta$  having a better cost (tie breaks being handled by a suitable lexicographic rule). This concept seems to have been studied first by Kohler and Steiglitz (1974) and has been developed in the subsequent years, most notably by Ibaraki (1977). However, although known for a long time, dominance criteria are not exploited in general-purpose MILP codes because of a number of important limitations of the classical definition. In fact, as stated, the dominance relationship is too vague to be

applicable in a generic MILP context—in principle, every node not leading to an optimal solution could be declared as being dominated.

In this paper we build on the general-purpose dominance procedure proposed in the late 1980s by Fischetti and Toth (1988), where the dominance test at a given node consists of the (possibly heuristic) solution of an auxiliary MILP involving only the fixed variables. As far as we know, no attempt to actually use the above dominance procedure within a general-purpose MILP scheme is reported in the literature. This is because the approach tends to be excessively time consuming—the reduction in the number of nodes does not compensate for the large overhead introduced. In an attempt to find a viable way to implement the original scheme, we found that the concept of *nogood*, borrowed from constraint programming (CP), can play a crucial role for the practical effectiveness of the overall dominance procedure. Roughly speaking, a *nogood* is a partial assignment of the variables such that every completion is either infeasible (for constraint satisfaction problems) or nonoptimal (for constraint optimization problems). Nogoods are widely used by the CP community (Dechter 1990; Katsirelos and Bacchus 2003, 2005; Lecoutre et al. 2007) and proved to be crucial

for the effectiveness of last-generation SAT solvers (Moskewicz et al. 2001, Zhang et al. 2001, Eén and Sörensson 2004). However, the concept of nogood is seldom used in mathematical programming. One of the first uses of nogoods in ILP was to solve verification problems (Hooker and Yan 1995) and fixed-charge network flow problems (Kim and Hooker 2002). Further applications can be found in Codato and Fischetti (2006), where nogoods are used to generate cuts for a MILP problem. Recently, attempts to apply the concept of nogood to MILPs, deriving nogoods from nodes discarded by the branch-and-bound (B&B) algorithm, have been presented in Achterberg (2007) and Sandholm and Shields (2006), with mixed results. In the present paper, we depart from the “classical” use of nogoods—namely, the representation of the result of some conflict analysis algorithm—and we derive them from dominance considerations. In particular, in the context of dominance, a nogood configuration is available, as a by-product, whenever the auxiliary problem is solved successfully. More generally, we show how the auxiliary problem can be used to derive “improving moves” that capture in a much more effective way the presence of general integer (as opposed to binary) variables.

Improving moves are the basic elements of *test sets*. For an integer programming problem, a test set is defined as a set  $T$  of vectors such that every feasible solution  $x$  to the integer program is nonoptimal if and only if there exists an element  $t \in T$  (the “improving move”) such that  $x + t$  is a feasible solution with strictly better objective function value; see, e.g., Scarf (1986), Graver (1975), and Thomas and Weismantel (1996). Test sets are often computed for a family of integer linear programs, obtained by varying the right-hand-side vector. This makes test sets useful for sensitivity analysis or for solving stochastic programs with a large number of scenarios. Improving moves are meant to be used to convert any feasible solution to an optimal one by a sequence of moves maintaining solution feasibility and improving in a strictly monotone way the objective value. Computing and using test sets for NP-hard problems is, however, far too expensive in practice. In our approach, instead, improving moves are exploited heuristically within a node-fathoming procedure—hence, the name “pruning moves.” More specifically, we generate pruning moves on the fly and store them in a *move pool* that is checked in a rather inexpensive way at each node of the branch-and-bound tree. Computational results show that this approach can be very successful for problems whose structure is amenable to dominance.

The rest of this paper is organized as follows. In §2 we describe the Fischetti-Toth technique in more detail and address some issues related to the tie-break rule to be used to get a mathematically correct overall method. In §3 we introduce the concepts

of nogood and pruning move, and show how to derive them in an effective way. In §4 we deal with some extensions of the basic scheme, and important implementation issues are addressed in §5. Computational results obtained on hard knapsack and network design instances are given in §6. We show that for those problems, the proposed method can lead to a significant speedup when embedded within a general-purpose MILP solver (e.g., ILOG CPLEX 11). Some conclusions are finally drawn in §7.

## 2. The Fischetti-Toth Dominance Procedure

In the standard B&B or branch-and-cut (B&C) framework, a node is fathomed in two situations:

1. The LP relaxation of the node is infeasible, or
2. The optimal value of LP relaxation is not better than the value of the incumbent optimal solution.

There is, however, a third way of fathoming a node: by using the concept of dominance. According to Papadimitriou and Steiglitz (1982), a dominance relation is defined as follows: If we can show that a descendant of a node  $\beta$  is at least as good as the best descendant of a node  $\alpha$ , then we say that node  $\beta$  *dominates* node  $\alpha$ , meaning that the latter can be fathomed (in case of ties, an appropriate rule has to be taken into account to avoid fathoming cycles). Unfortunately, this definition may become useless in the context of general MILPs, where we do not actually know how to perform the dominance test without storing huge amounts of information for all the previously generated nodes—which is often impractical.

Fischetti and Toth (1988) proposed a different dominance procedure that overcomes many of the drawbacks of the classical definition, and resembles somehow the *logic cuts* introduced by Hooker et al. (1994) and the *isomorphic pruning* introduced recently by Margot (2002, 2003). Here is how the procedure works.

Let  $P$  be the MILP problem

$$P: \min \{c^T x : x \in F(P)\},$$

whose feasible solution set is defined as

$$F(P) := \{x \in \mathbb{R}^n : Ax \leq b, l \leq x \leq u, \\ x_j \text{ integer for all } j \in J\}, \quad (1)$$

where  $J \subseteq N := \{1, \dots, n\}$  is the index set of the integer variables. For any  $J' \subseteq J$  and for any  $x' \in \mathbb{R}^{J'}$ , let

$$c(J', x') := \sum_{j \in J'} c_j x'_j$$

denote the contribution of the variables in  $J'$  to the overall cost. Now, suppose we are solving  $P$  by an

enumerative (B&B or B&C) algorithm whose branching rule fixes some of the integer-constrained variables to certain values. For every node  $k$  of the branch-and-bound tree, let  $J^k \subseteq J$  denote the set of indices of the variables  $x_j$  fixed to a certain value  $x_j^k$  (say). Every solution  $x \in F(P)$  such that  $x_j = x_j^k$  for all  $j \in J^k$  (i.e., belonging to the subtree rooted at node  $k$ ) is called a *completion* of the partial solution associated at node  $k$ .

DEFINITION 1. Let  $\alpha$  and  $\beta$  be two nodes of the branch-and-bound tree. Node  $\beta$  dominates node  $\alpha$  if

1.  $J^\beta = J^\alpha$ ;
2.  $c(J^\beta, x^\beta) \leq c(J^\alpha, x^\alpha)$ , i.e., the cost of the partial solution  $x^\beta$  at node  $\beta$  is not worse than that at node  $\alpha$ , namely,  $x^\alpha$ ;
3. Every completion of the partial solution  $x^\alpha$  associated with node  $\alpha$  is also a completion of the partial solution  $x^\beta$  associated with node  $\beta$ .

According to the classical dominance theory, the existence of a node  $\beta$  unfathomed that dominates node  $\alpha$  is a sufficient condition to fathom node  $\alpha$ . A key question at this point is given the current node  $\alpha$ , how can we check the existence of a dominating node  $\beta$ ? Fischetti and Toth (1988) answered this question by modeling the search of dominating nodes as a structured optimization problem, to be solved exactly or heuristically. For generic MILP models, this leads to the following *auxiliary problem*  $XP^\alpha$ :

$$\min \sum_{j \in J^\alpha} c_j x_j$$

$$\sum_{j \in J^\alpha} A_j x_j \leq b^\alpha := \sum_{j \in J^\alpha} A_j x_j^\alpha, \quad (2)$$

$$l_j \leq x_j \leq u_j, \quad j \in J^\alpha, \quad (3)$$

$$x_j \text{ integer}, \quad j \in J^\alpha. \quad (4)$$

If a solution  $x^\beta$  (say) of the auxiliary problem having a cost strictly smaller than  $c(J^\alpha, x^\alpha)$  is found, then it defines a dominating node  $\beta$  and the current node  $\alpha$  can be fathomed.

It is worth noting that the auxiliary problem is of the same nature as the original MILP problem but with a smaller size, and thus it is often easily solved (possibly in a heuristic way) by a general-purpose MILP solver. In a sense, we are using here the approach of “MIPping the dominance test” (i.e., of modeling it as a MILP; see Fischetti et al. 2009) in a vein similar to the recent approaches of Fischetti and Lodi (2003) (the so-called *local branching* heuristic, where a suitable MILP model is used to improve the incumbent solution) and of Fischetti and Lodi (2007) (where an ad hoc MILP model is used to generate violated Chvátal-Gomory cuts). Also note that, as discussed in §4, the auxiliary problem gives a sufficient but not necessary condition for the existence of a

dominating node, in the sense that some of its constraints could be relaxed without affecting the validity of the approach. In addition, inequalities (2) could be converted to equalities to reduce the search space and get a simpler, although possibly less effective, auxiliary problem.

The Fischetti-Toth dominance procedure, called LD (for *local dominance*) in the sequel, has several useful properties:

- There is no need to store any information about the set of previously generated nodes;
- There is no need to make any time-consuming comparison of the current node with other nodes;
- A node can be fathomed even if the corresponding dominating one has not been generated yet;
- The correctness of the enumerative algorithm does not depend on the branching rule; this is a valuable property because it imposes no constraints on the B&B parameters—although an inappropriate branching strategy could prevent several dominated nodes to be fathomed; and
- The LD test needs not be applied at every node; this is crucial from a practical point of view, because the dominance test introduces some overhead and it would make the algorithm too slow if applied at every node.

An important issue to be addressed when implementing the LD test is to avoid fathoming cycles arising when the auxiliary problem actually has a solution  $x^\beta$  different from  $x^\alpha$  but of the same cost, in which case one is allowed to fathom node  $\alpha$  only if a tie-break rule is used to guarantee that node  $\beta$  itself is not fathomed for the same reason. To prevent these “tautological” fathoming cycles, the following criterion (among others) has been proposed in Fischetti and Toth (1988): in case of cost equivalence, define as unfathomed the node  $\beta$  corresponding to the solution found by the *deterministic*<sup>1</sup> (exact or heuristic) algorithm used to solve the auxiliary problem. However, even very simple “deterministic” algorithms may lead to a wrong result, as shown in the following example.

Let  $P$  be the problem

$$\left\{ \begin{array}{l} \min \quad -x_1 - x_2 - x_3 - x_4 - 99x_5 \\ \text{s.t.} \quad x_1 + x_2 \leq 1, \\ \quad \quad x_3 + x_4 \leq 1, \\ \quad \quad x_4 + x_5 \leq 1, \\ \quad \quad x \in \{0, 1\}^5, \end{array} \right.$$

whose optimal solutions are  $[1, 0, 1, 0, 1]$  and  $[0, 1, 1, 0, 1]$ , and let us consider the branch-and-bound tree depicted in Figure 1. The deterministic

<sup>1</sup> In this context, an algorithm is said to be *deterministic* if it always provides the same output solution for the same input.




Figure 1 A Nasty Situation for LD Test

algorithm used to perform the LD test is as follows: If the number of variables in the auxiliary problem is smaller than 3, use a greedy heuristic trying to fix variables to 1 in decreasing index order; otherwise, use the same greedy heuristic but in increasing index order.

When node  $\alpha$  (that corresponds to the partial solution  $x_1 = 1, x_2 = 0$  with cost  $-1$ ) is processed, the following auxiliary model is constructed:

$$\begin{cases} \min & -x_1 - x_2 \\ \text{s.t.} & x_1 + x_2 \leq 1, \\ & x_1, x_2 \in \{0, 1\}, \end{cases}$$

and the deterministic heuristic returns the partial solution  $x_2 = 1, x_1 = 0$  of cost  $-1$  associated with node  $\beta$ , so node  $\alpha$  is declared to be dominated by  $\beta$  and node  $\alpha$  is fathomed assuming (correctly) that node  $\beta$  will survive the fathoming test. However, when the descendant node  $\alpha'$  (that corresponds to the partial solution  $x_1 = 0, x_2 = 1, x_3 = 1, x_4 = 0$  with cost  $-2$ ) is processed, the following auxiliary model is constructed:

$$\begin{cases} \min & -x_1 - x_2 - x_3 - x_4 \\ \text{s.t.} & x_1 + x_2 \leq 1, \\ & x_3 + x_4 \leq 1, \\ & x_4 \leq 1, \\ & x_1, x_2, x_3, x_4 \in \{0, 1\}, \end{cases}$$

and our deterministic heuristic returns the partial solution  $x_1 = 1, x_2 = 0, x_3 = 1, x_4 = 0$  of cost  $-2$  associated with node  $\beta'$ , so node  $\alpha'$  is declared to be dominated by  $\beta'$  and node  $\alpha'$  is fathomed as well.

Therefore, in this case the enumerative algorithm cannot find any of the two optimal solutions; i.e., the LD tests produced a wrong answer.

In view of the considerations above, in our implementation we used a different tie-break rule, also described in Fischetti and Toth (1988), that consists in ranking cost-equivalent solutions in lexicographical order ( $<$ ). To be more specific, in case of cost ties we fathom node  $\alpha$  if and only if  $x^\beta < x^\alpha$ , meaning that the partial solution  $x^\beta$  associated with the dominating node  $\beta$  is lexicographically smaller<sup>2</sup> than  $x^\alpha$ . Using this tie-break rule, it is possible to prove the correctness of the overall enumerative method.

**PROPOSITION 1.** *Assuming that the projection of the feasible set  $F(P)$  on the integer variable space is a bounded set, a B&B algorithm exploiting LD with the lexicographical tie-break rule returns the same optimal value as the classical B&B algorithm.*

**PROOF.** Let  $x^*$  be the lexicographically minimal optimal solution, whose existence is guaranteed by the boundedness assumption and by the fact that  $<$  is a well-order. We need to show that no node  $\alpha$  having  $x^*$  among its descendants (i.e., such that  $x_j^* = x_j^\alpha$  for all  $j \in J^\alpha$ ) can be fathomed by the LD test. Assume by contradiction that a node  $\beta$  dominating  $\alpha$  exists, and define

$$z_j := \begin{cases} x_j^\beta & j \in J^*, \\ x_j^* & j \notin J^*, \end{cases}$$

where  $J^* := J^\beta (=J^\alpha)$ . In other words,  $z$  is a new solution obtained from  $x^*$  by replacing its dominated part with the dominating one. Two cases can arise:

1.  $c(J^*, x^\beta) < c(J^*, x^\alpha)$ : We have

$$c^T z = \sum_{j \in J^*} c_j x_j^\beta + \sum_{j \notin J^*} c_j x_j^* < \sum_{j \in J^*} c_j x_j^\alpha + \sum_{j \notin J^*} c_j x_j^* = c^T x^*$$

and

$$\sum_{j=1}^n A_j z_j = \sum_{j \in J^*} A_j x_j^\beta + \sum_{j \notin J^*} A_j x_j^* \leq \sum_{j \in J^*} A_j x_j^\alpha + \sum_{j \notin J^*} A_j x_j^* \leq b;$$

so  $z$  is a feasible solution with a cost strictly smaller than  $x^*$ , which is impossible.

2.  $c(J^*, x^\beta) = c(J^*, x^\alpha)$ : using the same argument as in the previous case, one can easily show that  $z$  is an alternative optimal solution with  $z < x^*$ , also impossible.  $\square$

It is important to notice that the above proof of correctness uses just two properties of the lexicographic order, namely,

<sup>2</sup> We use the standard definition of lexicographic order on vectors of fixed size over a totally ordered set.

P1 (*well-order*): Required for the existence of a minimum optimal solution;

P2 (*inheritance*): If  $x^\alpha$  and  $x^\beta$  are two partial assignments such that  $x^\alpha < x^\beta$ , then the lexicographic order is not changed if we apply the same completion to both of them.

This observation will be used in §4 to derive a more efficient tie-break rule.

### 3. Nogoods and Pruning Moves

The computational overhead related to the LD test can be reduced considerably if we exploit the notion of nogoods taken from CP. A *nogood* is a partial assignment of the problem variables such that every completion is either infeasible (for constraint satisfaction problems) or nonoptimal (for constraint optimization problems). The key observation here is that whenever we discover (through the solution of the auxiliary problem) that the current node  $\alpha$  is dominated, we have in fact found a *nogood configuration*  $[J^\alpha, x^\alpha]$  that we want to exclude from being reanalyzed at a later time.

Actually, when the LD test succeeds we have not just a dominated partial assignment ( $x^\alpha$ ) but also a dominating one ( $x^\beta$ ). Combining the two, we get a *pruning move* ( $\Delta = [J^\alpha, x^\beta - x^\alpha]$ ), i.e., a list of variable changes that we can apply to a (partial) assignment to find a dominating one, provided that the new values assigned to the variables stay within the prescribed bounds. For binary MILPs, the concept of pruning move is equivalent to that of nogood, in the sense that a pruning move can be applied to a partial assignment if and only if the same assignment can be ruled out by a corresponding (single) nogood. For general-integer MILPs, however, pruning moves can be much more effective than nogoods. This is easily seen by the following example. Suppose we have two integer variables,  $x_1, x_2 \in [0, U]$ , and that the LD test produces the pair

$$x^\alpha = (1, 0), \quad x^\beta = (0, 1).$$

It is easy to see that, in this case, all the following (dominating, dominated)-pairs are valid:

$$\{(a, b), (a - 1, b + 1)\}: a \in [1, U], b \in [0, U - 1].$$

The standard LD procedure would need to derive each such pair by solving a different auxiliary problem, whereas they can be derived all together by solving a single MILP leading to the pruning move  $(-1, 1)$ . As such, pruning moves can be seen as compact representations of sets of nogoods, a topic also studied in Katsirelos and Bacchus (2005).

In our implementation, we maintain explicitly a pool of previously found pruning moves and solve the following problem (akin to separation for cutting-plane methods) at each branching node  $\alpha$ : find, if

any, a move  $\Delta = [J', \delta]$  stored in the pool such that  $J' \subseteq J^\alpha$  and  $l_j \leq x_j^\alpha + \delta_j \leq u_j$  for all  $j \in J'$ . If the test is successful, we can of course fathom node  $\alpha$  without needing to construct and solve the corresponding auxiliary problem  $XP^\alpha$ . In our implementation, pruning moves are stored in sparse form, whereas the pool is simply a list of moves sorted by length. It is worth noting that we are interested in minimal (with respect to set inclusion) pruning moves so as to improve effectiveness of the method. To this end, before storing a pruning move in the pool we remove its components  $j$  such that  $x_j^\alpha = x_j^\beta$  (if any).

It is also worth noting that a move  $\Delta^1 = [J^1, \delta^1]$  implies (absorbs) a move  $\Delta^2 = [J^2, \delta^2]$  in case

$$J^1 \subseteq J^2, \text{ and } |\delta_j^1| \leq |\delta_j^2| \quad \forall j \in J^1.$$

This property can be exploited to keep the pool smaller without affecting its fathoming power.

At first glance, the use of a move pool resembles classical state-based dominance tests, but this is really not the case because the amount of information stored is much smaller—actually, it could even be limited to be of polynomial size by exploiting techniques such as relevance- or length-bounded nogood recording (see Bayardo and Miranker 1996).

### 4. Improving the Auxiliary Problem

The effectiveness of the dominance test presented in the previous section heavily depends on the auxiliary problem that is constructed at a given node  $\alpha$ . In particular, it is advisable that its solution set is as large as possible so as to increase the chances of finding a dominating partial solution. Moreover, we aim to find a partial solution different from (and, we hope, lexicographically better than) the one associated with the current node—finding the same solution  $x^\alpha$  is of no use within the LD context. For these reasons, we next propose a number of improvements over the original auxiliary problem formulation.

#### 4.1. Objective Function

The choice of the lexicographic order as a mean to resolve ties, although natural and simple, is not well suited in practice.

In the most naïve implementation, there is a good chance of not finding a lexicographically better solution even if this exists, because we do not convey to the solver in any way the information that we are interested in lexicographically minimal solutions. This is unfortunate, because we risk wasting a great computational effort.

Moreover, the lexicographic order cannot be expressed as a linear objective without resorting to huge coefficients: the only way to enforce the discovery of lexicographically better solutions is through

ad hoc branching and node selection strategies, which are quite intrusive and greatly degrade the efficiency of the solution process.

The solution we propose is to use an alternative randomly generated objective function, according to the following scheme:

- An *alternative* random objective function for the auxiliary problems is generated at the beginning of the computation and is kept unchanged for the whole search so as to satisfy the two properties P1 and P2 of §2 needed for the correctness of the algorithm.

- To guarantee that the optimal solution of the auxiliary problem will be not worse than the original partial assignment, we add the following *optimality* constraint:

$$\sum_{j \in J^\alpha} c_j x_j \leq \sum_{j \in J^\alpha} c_j x_j^\alpha.$$

- After having solved the auxiliary problem with the alternative objective function, we compare its solution  $x^\beta$  (if any) with the original partial assignment  $x^\alpha$ , first by using the original objective function, then by using the alternative function, and finally, in the unlikely case that both functions yield the same value, by lexicographic order.

#### 4.2. Local Branching Constraints

As the depth of the nodes in the B&B increases, the auxiliary problem grows in size and becomes harder to solve. Moreover, we are interested in detecting moves involving only a few variables, because these are more likely to help prune the tree and are more efficient to search. For these reasons one can heuristically limit the search space of the auxiliary problem to alternative assignments not too far from the current one. To this end, we use a *local branching* (Fischetti and Lodi 2003) constraint defined as follows.

For a given node  $\alpha$ , let  $B^\alpha \subseteq J^\alpha$  be the (possibly empty) set of fixed binary variables, and define

$$U = \{j \in B^\alpha \mid x_j^\alpha = 1\}, \quad L = \{j \in B^\alpha \mid x_j^\alpha = 0\}.$$

Then we can guarantee a solution  $x$  of the auxiliary problem to be different from  $x^\alpha$  in at most  $k$  binary variables through the following *local branching* constraint:

$$\sum_{j \in U} (1 - x_j) + \sum_{j \in L} x_j \leq k.$$

A similar reasoning could be extended to deal with general integer variables as well, although in this case the constraint is not as simple as before and requires the addition of certain auxiliary variables (Fischetti and Lodi 2003). According to our computational experience, a good compromise is to consider a local branching constraint involving only the (binary

or general integer) variables fixed to their lower or upper bound, namely,

$$\sum_{j \in U} (u_j - x_j) + \sum_{j \in L} (x_j - l_j) \leq k,$$

where

$$U = \{j \in J^\alpha \mid x_j^\alpha = u_j\}, \quad L = \{j \in J^\alpha \mid x_j^\alpha = l_j\}.$$

#### 4.3. Right-Hand-Side Improvement

One could observe that we have been somehow overly-conservative in the definition of the auxiliary problem  $XP^\alpha$ . In particular, as noticed already in Fischetti and Toth (1988), in some cases the condition

$$\sum_{j \in J^\alpha} A_j x_j \leq b^\alpha$$

could be relaxed without affecting the correctness of the method.

To illustrate this possibility, consider a simple knapsack constraint  $4x_1 + 5x_2 + 3x_3 + 2x_4 \leq 10$  and suppose we are given the partial assignment  $[1, 0, 1, *]$ . The corresponding constraint in the auxiliary problem then reads  $4x_1 + 5x_2 + 3x_3 \leq 7$ . However, because the maximum load achievable with the free variables is 2, one can safely consider the relaxed requirement  $4x_1 + 5x_2 + 3x_3 \leq 10 - 2 = 8$ . Notice that the feasible partial solution  $[0, 1, 1, *]$  is forbidden by the original constraint but allowed by the relaxed one; i.e., the relaxation does improve the chances of finding a dominating node. Another example arises for set-covering problems, where the  $i$ th constraint reads  $\sum_{j \in Q_i} x_j \geq 1$  for some  $Q_i \subseteq \{1, \dots, n\}$ . Suppose we have a partial assignment  $x_j^\alpha$  ( $j \in J^\alpha$ ) such that  $k := \sum_{j \in J^\alpha \cap Q_i} x_j^\alpha > 1$ . In this case, the corresponding constraint in the auxiliary problem would be  $\sum_{j \in J^\alpha \cap Q_i} x_j \geq k$ , although its relaxed version  $\sum_{j \in J^\alpha \cap Q_i} x_j \geq 1$  is obviously valid as well.

The examples above show, however, that the improvement of the auxiliary problem may require some knowledge of the particular structure of its constraints. Even more important, the right-hand-side strengthening procedure above can interfere and become incompatible with the postprocessing procedure that we apply to improve the moves. For this reason, in our implementation we decided to avoid any right-hand-side improvement.

#### 4.4. Local Search on Incumbents

A drawback of the proposed scheme is that node fathoming is very unlikely at the very beginning of the search. Indeed, at the top of the tree only few variables are fixed, and the LD test often fails (this is also true if short moves exist, because their detection depends on the branching strategy used) while the move pool is empty.

To mitigate this problem, each time a new incumbent is found we invoke the following local search phase aimed at feeding the move pool.

• Given the incumbent  $x^*$ , we search the neighborhood  $N_k(x^*)$  defined through the following constraints:

$$\sum_{j \in J} A_j x_j = \sum_{j \in J} A_j x_j^*, \quad (5)$$

$$\sum_{j \in J: x_j^* = u_j} (u_j - x_j) + \sum_{j \in J: x_j^* = l_j} (x_j - l_j) \leq k, \quad (6)$$

$$\sum_{j \in J} c_j x_j \leq \sum_{j \in J} c_j x_j^*, \quad (7)$$

$$l_j \leq x_j \leq u_j, \quad j \in J,$$

$$x_j \text{ integer}, \quad j \in J.$$

In other words, we look for alternative values of the integer variables  $x_j$  ( $j \in J$ ) using each constraint—variable bounds excluded—in the same way as  $x^*$  (constraint (5)), having a Hamming distance from  $x^*$  not larger than  $k$  (constraint (6)), and with an objective value not worse than  $x^*$  (constraint (7)).

• We populate a solution list by finding multiple solutions to the MILP

$$\min \left\{ \sum_{j \in J} c_j x_j : x \in N_k(x^*) \right\}$$

by exploiting the multiple-solution mode available in our MILP solver (ILOG 2007).

• Given the solution list  $L = (x^1, \dots, x^p)$ , we compare the solutions pairwise and generate a pruning move accordingly, to be stored in the move pool.

• If we find a better solution during the search, we use it to update the incumbent.

It is worth noting that the use of equalities (5) allows us to generate a pruning move for *every* pair of distinct solutions in the list, because for every pair we have a dominating and a dominated solution whose difference produces the pruning move.

## 5. Implementation

The enhanced dominance procedure presented in the previous sections was implemented in C++ on a Linux platform and applied within a commercial MILP solver. Here are some implementation details that deserve further description.

An important drawback of LD tests is that their use can postpone the finding of a better incumbent solution, thus increasing the number of nodes needed to solve the problem. This behavior is quite undesirable, particularly in the first phase of the search, when we have no incumbent and no nodes can be fathomed through bounding criteria. Our solution to this problem is to skip the dominance test until the first feasible solution is found.

The definition and solution of the auxiliary problem at every node of the branch-and-bound tree can

become too expensive in practice. We face here a situation similar to that arising in B&C methods, where new cuts are typically not generated at every node—although the generated cuts are exploited at each node. A specific LD consideration is that we had better skip the auxiliary problem on nodes close to the top or the bottom of the branch-and-bound tree. Indeed, in the first case only a few variables have been fixed; hence, there is little chance of finding dominating partial assignments. In the latter case, instead, it is likely that the node would be fathomed anyway by standard bounding tests. Moreover, at the bottom of the tree the number of fixed variables is quite large and the auxiliary problem may be quite hard to solve. In our implementation, we provide two thresholds on tree depth, namely,  $depth_{\min}$  and  $depth_{\max}$ , and solve the auxiliary problem for a node  $\alpha$  only if  $depth_{\min} \leq depth(\alpha) \leq depth_{\max}$ . Moreover, we decided to solve the auxiliary problem at a node only if its depth is a multiple of a given parameter, say,  $depth_{\text{interval}}$ .

In addition, because it is undesirable to spend a large computing time on the auxiliary problem for a node that would have been pruned anyway by the standard B&B rules, we decided to apply our technique just before branching—applying the LD test before the LP relaxation is solved turned out to be less effective.

To avoid spending too much computing time on pathologically hard auxiliary MILPs, we also set a node limit (say)  $N_1$  for the auxiliary problem solution, and a node limit (say)  $N_2$  for the local search on incumbents.

Finally, because the discovery of new pruning moves decreases as we proceed with the search, we set an upper bound  $M$  on the number of times the LD test is called: after this limit is reached, the pruning effect is left to the move pool only.

It is important to stress that although the auxiliary problem is solved only at certain nodes, we check the current partial assignment against the move pool at every node, because this check is relatively cheap.

## 6. Computational Results

In our computational experiments we used the commercial solver ILOG CPLEX 11.0 (ILOG 2007) with default options. All runs were performed on an Intel Q6600 2.4 GHz PC with 4 GB of RAM, under Linux.

The definition of the test bed for testing the potential of our approach is, of course, a delicate issue. In fact, one cannot realistically expect any dominance relationship to be effective on all types of MILPs. This is confirmed by a preliminary test we performed on the MIPLIB 2003 (Achterberg et al. 2006) test bed, where pruning moves are seldom generated. We here face a situation similar to that arising when testing techniques designed for highly symmetric problems, such

as the *isomorphic pruning* proposed recently by Margot (2002, 2003)—although remarkably effective on some classes of problems, the approach is clearly of no use for problems that do not exhibit any symmetry.

Therefore, we looked for classes of practically relevant problems whose structure can trigger the dominance relationship, and measured the speedup that can be achieved by using our specific LD procedure. In particular, we next give results on two combinatorial problems: knapsack problems (Martello and Toth 1990) and network loading problems (van Hoesel et al. 2002, Atamtürk and Rajan 2002). Whereas the first class of problems is quite natural for dominance tests (and could in principle be solved much more effectively by using specialized codes), the second one is representative of very important applications where the dominance property is hidden well inside the solution structure.

### 6.1. Knapsack Problem

We generated hard single knapsack instances according to the so-called *spanner instance* method in combination with the *almost-strongly correlated* profit generation technique; see Pisinger (2005) for details.

The parameters of our LD procedure were set to  
 $depth_{\min}$ : 5,  
 $depth_{\max}$ : 0.8 times the number of integer variables,  
 $depth_{\text{interval}}$ : 6,  
 $k$ : 0.2 times the number of integer variables,  
 $N_1$ : 10,  
 $N_2$ : 5,000,  
 $M$ : 1,000.

The results on hard single knapsack instances with 60 to 90 items are given in Table 1, where labels “Dominance” and “Standard” refer to the

performance of the B&C scheme with and without the LD tests, respectively, and label “Ratio” refers to the ratios Standard/Dominance. For a fair comparison, the same seed was used to initialize the random generator in all runs. The performance figures used in the comparison are the number of nodes of the resulting branch-and-bound tree and the computing time (in CPU seconds). For these problems, the LD tests provided an overall speedup of 23 times and with substantially fewer nodes (the ratio being approximately 1:33). Note that, in some cases, the ratios reported in Table 1 are just lower bounds on the real ones, because the standard algorithm was stopped before completion because of the time limit.

Additional statistics are reported in Table 2, where we provide, for each instance, the final size of the move pool, the percentage of the whole solution time spent either on pool management (Pool time) or LD tests (LD time) along with their success rates (Pool success and LD success, respectively). The statistics in Table 2 indicate that the number of pruning moves stored in the pool is always manageable. In addition, the pool checks and the LD tests are rather successful because they both allow for node fathoming approximately 1/3 of the times they are applied—the total effect being of fathoming approximately 2/3 of the nodes where the test is applied, which results in a dramatic reduction of the total number of branch-and-bound nodes because we often fathom early nodes and hence large subtrees. Although the relative overhead introduced by the LD procedure is significant (because of the fact that node relaxations are particularly cheap for knapsack problems), the overall benefit is striking, with an average speedup of about one to two orders of magnitude.

**Table 1** Computational Results for Hard Knapsack Instances

Problem	Standard			Dominance			Ratio	
	Nodes	Time (s)	Gap	Nodes	Time (s)	Gap	Nodes	Time
<i>kp60_1</i>	311,490	14.70	0.00	1,793	3.45	0.00	173.73	4.26
<i>kp60_2</i>	831,319	43.72	0.00	3,718	3.05	0.00	223.59	14.35
<i>kp60_3</i>	865,469	45.32	0.00	3,995	2.15	0.00	216.64	21.11
<i>kp60_4</i>	1,012,287	47.54	0.00	19,720	6.42	0.00	51.33	7.41
<i>kp70_1</i>	>12,659,538	>1,200.00	0.41	1,634,517	138.68	0.00	7.75	8.65
<i>kp70_2</i>	783,092	41.21	0.00	4,466	6.43	0.00	175.35	6.41
<i>kp70_3</i>	830,794	41.97	0.00	6,396	4.93	0.00	129.89	8.51
<i>kp70_4</i>	>13,226,464	>1,200.00	0.48	27,591	4.49	0.00	479.38	267.18
<i>kp80_1</i>	403,396	18.71	0.00	2,599	9.41	0.00	155.21	1.99
<i>kp80_2</i>	559,447	28.11	0.00	3,118	1.87	0.00	179.42	15.04
<i>kp80_3</i>	576,885	23.46	0.00	2,962	3.40	0.00	194.76	6.90
<i>kp80_4</i>	277,981	13.35	0.00	5,690	3.29	0.00	48.85	4.06
<i>kp90_1</i>	>16,013,282	>1,200.00	0.07	803,333	65.57	0.00	19.93	18.30
<i>kp90_2</i>	18,330,528	863.77	0.00	5,024	3.56	0.00	3,648.59	242.74
<i>kp90_3</i>	>15,273,264	>1,200.00	0.27	37,017	5.61	0.00	412.60	213.78
<i>kp90_4</i>	2,136,389	116.21	0.00	8,056	3.82	0.00	265.19	30.46
Average	>5,255,727	>381.13	—	160,165	16.63	—	398.89	54.45
Geom. mean	>1,761,164	>98.78	—	11,625	6.00	—	151.50	16.47



**Table 2 Internal Statistics for Hard Knapsack Instances**

Problem	Pool size	Pool time (%)	Pool success (%)	LD time (%)	LD success (%)
kp60_1	404	20.08	38.95	13.82	35.00
kp60_2	196	1.61	34.97	29.82	47.06
kp60_3	360	14.86	37.66	24.45	16.67
kp60_4	148	3.48	38.18	58.24	16.80
kp70_1	330	26.00	29.62	1.10	50.50
kp70_2	464	18.62	39.97	20.67	40.59
kp70_3	299	15.11	40.58	25.28	30.43
kp70_4	324	16.61	29.31	37.60	59.90
kp80_1	384	20.67	42.89	4.99	37.05
kp80_2	286	8.45	37.87	35.88	25.45
kp80_3	833	24.78	43.96	9.54	35.50
kp80_4	294	15.95	43.02	28.13	25.13
kp90_1	268	22.32	37.55	2.08	28.70
kp90_2	705	17.76	38.56	34.19	58.15
kp90_3	286	16.55	34.20	20.88	41.20
kp90_4	181	3.85	33.97	44.78	42.20
Average	360.13	15.42	37.58	24.47	36.90
Geom. mean	326.24	12.49	37.33	16.80	34.59

**6.2. Network Loading Problem**

Network loading problems arise in telecommunications applications where demand for capacity for multiple commodities has to be realized by allocating capacity to the arcs of a given network. Along with a capacity plan, a routing of all commodities has to be determined, and each commodity must be routed from source to destination on a single path through the network. The objective is to minimize the cost of the installed capacity in the network, ensuring that all commodities can be routed from source to destination simultaneously.

Given a directed graph  $G = (V, A)$ , a set of commodities  $K$  (each commodity being described by a source node  $s^k$ , a destination node  $t^k$ , and a demand

size  $d^k$ ), a base capacity unit  $C$ , and capacity installation costs  $c_{ij}$ , our network loading problem can be formulated as

$$\min \sum_{(i,j) \in A} c_{ij} y_{ij}$$

$$\sum_{j \in V} x_{ij}^k - \sum_{j \in V} x_{ji}^k = \begin{cases} 1 & i = s^k, \\ -1 & i = t^k, \\ 0 & \text{otherwise,} \end{cases}$$

$$k \in K, i \in V, \quad (8)$$

$$\sum_k d^k x_{ij}^k \leq C y_{ij}, \quad (i, j) \in A, \quad (9)$$

$$x_{ij}^k \in \{0, 1\}, y_{ij} \in \mathbb{Z}_0^+, \quad k \in K, (i, j) \in A. \quad (10)$$

Random instances of the above network loading problem were generated as follows:

- To obtain a gridlike structure, we generated grid networks of dimensions  $3 \times 5$  and  $4 \times 4$  and randomly deleted arcs with probability 0.1. Each arc has base unit capacity of value 4 and cost 10.

- We generated a commodity of flow 10 for each pair of nodes with probability 0.2.

Some parameters of the LD procedure were changed with respect to the knapsack test bed, because of the greatly increased size of the instances:


- $k$ : 0.01 times the number of integer variables,
- $N_1$ : 100,
- $N_2$ : 25,000.

The results of our experiments are given in Table 3. In this class of problems, the dominance procedure is still quite effective, with an overall speedup of four times and a node ratio of more than 5.

**Table 3 Computational Results for Network Loading Problems**

Problem	Standard			Dominance			Ratio	
	Nodes	Time (s)	Gap	Nodes	Time (s)	Gap	Nodes	Time
g_15_17_43	1,954,292	797.01	0.00	242,693	163.80	0.00	8.05	4.87
g_15_17_45	>8,711,335	>3,600.00	0.34	1,544,646	845.00	0.00	5.64	4.26
g_15_17_51	>8,022,870	>3,600.00	0.29	963,576	545.14	0.00	8.33	6.60
g_15_18_35	3,764,325	1,559.38	0.00	286,539	172.48	0.00	13.14	9.04
g_15_18_37	3,959,652	1,525.63	0.00	567,899	279.64	0.00	6.97	5.46
g_15_18_39	752,035	251.52	0.00	303,667	146.55	0.00	2.48	1.72
g_15_18_40	>10,156,564	>3,600.00	0.48	1,071,922	493.57	0.00	9.48	7.29
g_15_19_43	1,609,434	886.51	0.00	415,472	294.61	0.00	3.87	3.01
g_16_18_48	581,268	226.13	0.00	122,824	86.65	0.00	4.73	2.61
g_16_18_53	6,425,061	3,183.84	0.00	6,489	56.14	0.00	990.15	56.71
g_16_19_51	>7,222,780	>3,600.00	0.43	3,774,093	2,158.96	0.00	1.91	1.67
g_16_20_47	5,593,517	3,436.69	0.00	587,773	449.83	0.00	9.52	7.64
g_16_20_51	2,229,792	1,394.58	0.00	272,355	257.26	0.00	8.19	5.42
g_16_21_40	>6,187,221	>3,600.00	0.37	2,334,537	1,524.71	0.00	2.65	2.36
g_16_21_44	1,079,588	717.43	0.00	151,869	182.00	0.00	7.11	3.94
g_16_21_52	>4,565,496	>3,600.00	0.27	1,279,007	1,186.96	0.00	3.57	3.03
Average	>4,550,951.88	>2,223.67	—	870,335.06	552.71	—	67.86	7.85
Geom. mean	>3,354,264	>1,621.25	—	436,484	339.43	—	7.68	4.78

Copyright: INFORMS holds copyright to this *Articles in Advance* version, which is made available to institutional subscribers. The file may not be posted on any other website, including the author's site. Please send any questions regarding this policy to permissions@informs.org.



**Figure 2** Typical Moves Captured by the LD Tests on Network Design Instances

Typical pruning moves discovered by the LD test for this class of problems are shown in Figure 2. In particular, Figure 2 illustrates two moves taken from instance  $g_{15\_17\_43}$ . For each of them we present two equivalent configurations. Flows of different commodities are drawn with different line styles, and the values reported by the edges (if relevant) indicate the amount  $y_{ij}$  of unit capacities installed on the arcs. In move (a) we have a simple swap of the routing of two different flows with the same value: the corresponding move involves only binary variables  $x_{ij}^k$ , whereas variables  $y_{ij}$  (not shown in the figure) are unaffected. In move (b) we have a more complex swap, involving also arc capacities (and thus general integer variables as well). Here, we reroute two commodities so as to use a different arc and move an appropriate amount of unit capacities accordingly. This last example shows the effectiveness of pruning moves: the move can be applied whenever it is possible to transfer five units of capacity from the bottom arc to the upper arc, regardless of their actual values.

Additional statistics are reported in Table 4. The format of the table is the same as for knapsack problems. As expected, both the computational overhead of the tests and their rate of success are significantly smaller than in the knapsack case, but still very satisfactory.

### 6.3. Pool Effectiveness

Finally, we tested the effectiveness of the main improvements we proposed to the original Fischetti-Toth scheme. Results are reported in Tables 5 and 6 for knapsack and network loading instances, respectively. We compared our final code (Dominance) against two versions of the same code obtained by disabling the use of the move pool (versions LD1 and LD2)

**Table 4** Internal Statistics for Network Loading Instances

Problem	Pool size	Pool time ratio (%)	Pool success (%)	LD time ratio (%)	LD success (%)
$g_{15\_17\_43}$	48	0.73	13.39	26.20	1.50
$g_{15\_17\_45}$	61	1.08	17.80	8.43	0.80
$g_{15\_17\_51}$	167	2.19	17.14	5.79	1.30
$g_{15\_18\_35}$	55	0.94	16.64	15.84	1.10
$g_{15\_18\_37}$	53	1.15	4.04	3.27	0.50
$g_{15\_18\_39}$	108	1.52	12.22	16.93	1.20
$g_{15\_18\_40}$	89	1.66	18.90	5.17	2.60
$g_{15\_19\_43}$	135	1.68	18.51	8.97	3.30
$g_{16\_18\_48}$	78	0.95	9.81	28.68	0.90
$g_{16\_18\_53}$	84	0.17	3.19	59.68	0.60
$g_{16\_19\_51}$	178	2.45	22.92	1.49	1.70
$g_{16\_20\_47}$	56	0.92	8.63	8.76	0.80
$g_{16\_20\_51}$	69	0.81	15.32	18.28	2.60
$g_{16\_21\_40}$	67	0.92	11.46	2.77	1.90
$g_{16\_21\_44}$	108	0.78	7.25	25.27	1.00
$g_{16\_21\_52}$	71	0.77	13.24	5.58	2.30
Average	89.19	1.17	13.15	15.07	1.51
Geom. mean	82.13	1.01	11.70	9.89	1.31

and a version without the local search on incumbents (version LD3). We provide the absolute performance figures for Dominance, and we give relative performance for the other versions—the numbers in Tables 5 and 6 give the slowdown factors of the various versions with respect to our final code (the larger the worse). According to Tables 5 and 6, disabling the move pool while retaining the limit  $M$  on the number of times the LD test is actually called (version LD1) is disastrous: not only do we lose the fathoming effect on a large part of the tree, but we waste a large computing time in solving auxiliary problems discovering a same pruning move (or even a dominated one) over and over. Better results can be obtained if we remove limit  $M$  (version LD2): in this way we retain much of the fathoming effect but at a much greater computational effort (LD2 is about five times slower than the default version on knapsack problems, and reaches the one-hour time limit in 11 out of 16 instances on network problems). The contribution of the local search on incumbents (version LD3) is more difficult to evaluate: although the number of nodes is always reduced by its use, the overall computing time is reduced only for network problems. A closer look at the individual table entries shows, however, that local search is ineffective only for easy problems, where its overhead is not balanced by the increased fathoming power, but local search turns out to be very useful in harder instances (e.g.,  $kp70\_1$  or  $kp90\_1$ ).

## 7. Conclusions

In this paper we have presented a dominance procedure for general MILPs. The technique is an elaboration of an earlier proposal of Fischetti and Toth (1988),

**Table 5 Comparison of Different LD Versions on Knapsack Problems**

Problem	Dominance		LD1 ratios		LD2 ratios		LD3 ratios	
	Nodes	Time (s)	Nodes	Time	Nodes	Time	Nodes	Time
<i>kp60_1</i>	1,793	3.45	104.23	3.09	3.66	1.36	1.20	0.23
<i>kp60_2</i>	3,718	3.05	179.47	11.96	4.61	4.77	1.56	0.76
<i>kp60_3</i>	3,995	2.15	143.03	14.41	5.01	4.67	1.97	0.82
<i>kp60_4</i>	19,720	6.42	19.34	3.07	4.33	10.45	1.67	0.85
<i>kp70_1</i>	1,634,517	138.68	>8.16	>8.65	>1.26	>8.65	7.21	6.25
<i>kp70_2</i>	4,466	6.43	133.94	5.05	5.58	2.70	1.26	0.30
<i>kp70_3</i>	6,396	4.93	117.94	7.94	5.31	4.43	1.35	0.44
<i>kp70_4</i>	27,591	4.49	>500.77	>267.18	2.85	17.85	2.33	1.36
<i>kp80_1</i>	2,599	9.41	105.96	1.46	4.97	0.72	1.38	0.11
<i>kp80_2</i>	3,118	1.87	103.60	9.35	4.19	4.70	1.30	0.69
<i>kp80_3</i>	2,962	3.40	119.87	4.49	5.14	1.79	1.44	0.26
<i>kp80_4</i>	5,690	3.29	39.23	3.56	4.78	4.36	1.54	0.70
<i>kp90_1</i>	803,333	65.57	>20.06	>18.30	10.66	16.30	14.49	10.52
<i>kp90_2</i>	5,024	3.56	3,354.31	221.01	4.45	4.02	1.58	0.67
<i>kp90_3</i>	37,017	5.61	>432.23	>213.78	3.67	20.43	1.58	0.94
<i>kp90_4</i>	8,056	3.82	240.35	27.55	3.32	4.58	1.00	0.64
Average	—	—	>351.41	>51.30	>4.61	>6.99	2.68	1.60
Geom. mean	—	—	>116.83	>13.14	>4.26	>4.85	1.88	0.74

*Notes.* LD1 is the version without the move pool and with the same limit  $M$  on the number of times the LD test is called, whereas LD2 is still without the move pool, but with no such limit. LD3 is the version without local search on the incumbents. As in Table 1, label Dominance refers to the default LD version; > indicates a reached time limit.

with important improvements aimed at making the approach computationally more attractive in the general MILP context. In particular, the use of nogoods and of pruning moves that we propose in this paper turned out to be crucial for an effective use of dominance tests within a general-purpose MILP code.

In our view, a main contribution of our work is the innovative use of improving moves. In a classical (yet

computationally impractical) test set approach, these moves are used within a primal solution scheme, leading eventually to an optimal solution. In our approach, instead we heuristically generate improving moves on small subsets of variables by solving, on the fly, small MILPs. These moves are not used to improve the incumbent, as in the classical test set environment, but rather to fathom nodes in the

**Table 6 Comparison of Different LD Versions on Network Problems**

Problem	Dominance		LD1 ratios		LD2 ratios		LD3 ratios	
	Nodes	Time (s)	Nodes	Time	Nodes	Time	Nodes	Time
<i>g_15_17_43</i>	242,693	164	4.36	2.73	>1.97	>21.98	0.82	0.90
<i>g_15_17_45</i>	1,544,646	845	>5.38	>4.26	>0.22	>4.26	1.26	1.22
<i>g_15_17_51</i>	963,576	545	>8.38	>6.60	>0.55	>6.60	1.99	1.82
<i>g_15_18_35</i>	286,539	172	6.80	4.89	1.48	15.78	1.08	1.02
<i>g_15_18_37</i>	567,899	280	2.13	1.77	1.25	12.43	0.47	0.55
<i>g_15_18_39</i>	303,667	147	2.19	1.63	1.11	7.99	0.97	0.89
<i>g_15_18_40</i>	1,071,922	494	>9.46	>7.29	>0.65	>7.29	1.44	1.35
<i>g_15_19_43</i>	415,472	295	3.39	2.68	1.04	10.83	0.98	0.93
<i>g_16_18_48</i>	122,824	87	9.51	5.49	>4.17	>41.55	4.01	2.63
<i>g_16_18_53</i>	6,489	56	>1,126.00	>64.13	28.33	33.69	3.84	1.07
<i>g_16_19_51</i>	3,774,093	2,159	>1.93	>1.67	>0.11	>1.67	1.65	1.57
<i>g_16_20_47</i>	587,773	450	>10.07	>8.00	>1.05	>8.00	2.96	2.61
<i>g_16_20_51</i>	272,355	257	3.46	2.43	>1.83	>13.99	1.63	1.35
<i>g_16_21_40</i>	2,334,537	1,525	>2.66	>2.36	>0.17	>2.36	2.42	2.36
<i>g_16_21_44</i>	151,869	182	7.18	4.25	>3.04	>19.78	1.76	1.31
<i>g_16_21_52</i>	1,279,007	1,187	>3.57	>3.03	>0.23	>3.03	1.01	0.97
Average	—	—	>75.40	>7.70	>2.95	>13.20	1.77	1.41
Geom. mean	—	—	>6.49	>4.14	>1.00	>9.24	1.51	1.29

*Notes.* LD1 is the version without the move pool and the same  $M$  as the default LD version, whereas LD2 is still without the move pool, but with no such limit. LD3 is the version without local search on the incumbents. As in Table 1, label Dominance refers to the default LD version; > indicates a reached time limit (for LD1 and LD2, the time limit is reached so often that the means are seriously underestimated).

branch-and-bound tree—hence, the name pruning moves. If implemented in a proper way, this approach introduces an acceptable overhead even if embedded in a highly efficient commercial MILP solver such as ILOG CPLEX 11 and may produce a drastic reduction in the number of nodes. Indeed, computational results show that the method can lead to a speedup of up to two orders of magnitude on hard MILPs whose structure is amenable to dominance, in the sense that it allows for local variable adjustments that produce equivalent solutions. An example of this kind of problem has been discussed—namely, a network loading problem arising in telecommunication.

A drawback of our approach is of course the overhead introduced when addressing problems that turn out not to produce any effective pruning move. A possible remedy is to use a conservative adaptive scheme that tries to generate those moves only in the first part of the search and deactivates the solution of auxiliary problem if its success rate is not satisfactory. The detailed design and implementation of this idea is, however, out of the scope of the present paper and is left to future investigation.

### Acknowledgments

This work was supported by the University of Padova “Progetti di Ricerca di Ateneo,” under Contract CPDA051592 (project: Integrating Integer Programming and Constraint Programming) and by the Future and Emerging Technologies unit of the EC (IST priority), under Contract FP6-021235-2 (Project ARRIVAL).

### References

Achterberg, T. 2007. Conflict analysis in mixed integer programming. *Discrete Optim.* 4(1) 4–20.

Achterberg, T., T. Koch, A. Martin. 2006. MIPLIB 2003. *Oper. Res. Lett.* 34(4) 361–372. [Problems available at <http://miplib.zib.de>.]

Atamtürk, A., D. Rajan. 2002. On splittable and unsplittable flow capacitated network design arc-set polyhedra. *Math. Programming* 92(2) 315–333.

Bayardo, R. J., D. P. Miranker. 1996. A complexity analysis of space-bounded learning algorithms for the constraint satisfaction problem. *Proc. 13th National Conf. Artificial Intelligence (AAAI-96)*, Vol. 1. Association for the Advancement of Artificial Intelligence, Menlo Park, CA, 298–304.

Bixby, R. E., M. Fenelon, Z. Gu, E. Rothberg, R. Wunderling. 2000. MIP: Theory and practice—Closing the gap. M. J. D. Powell, S. Scholtes, eds. *Proc. 19th IFIP TC7 Conf. System Modeling Optim. IFIP Conf. Proc.*, Vol. 174. Kluwer, B. V., Dordrecht, The Netherlands, 19–50.

Codato, G., M. Fischetti. 2006. Combinatorial Benders’ cuts for mixed-integer linear programming. *Oper. Res.* 54(4) 756–766.

Dechter, R. 1990. Enhancement schemes for constraint processing: Backjumping, learning, and cutset decomposition. *Artificial Intelligence* 41(3) 273–312.

Eén, N., N. Sörensson. 2004. An extensible SAT-solver. E. Giunchiglia, A. Tacchella, eds. *SAT 2003. Lecture Notes in Computer Science*, Vol. 2919. Springer-Verlag, Berlin, 333–336.

Fischetti, M., A. Lodi. 2003. Local branching. *Math. Programming* 98(1–3) 23–47.

Fischetti, M., A. Lodi. 2007. Optimizing over the first Chvátal closure. *Math. Programming* 110(1) 3–20.

Fischetti, M., P. Toth. 1988. A new dominance procedure for combinatorial optimization problems. *Oper. Res. Lett.* 7(4) 181–187.

Fischetti, M., A. Lodi, D. Salvagnin. 2009. Just MIP it! V. Maniezzo, T. Stautzle, S. Voss, eds. *MATHEURISTICS: Hybridizing Metaheuristics and Mathematical Programming. Operations Research/Computer Science Interfaces Series*. Springer, New York. Forthcoming.

Garfinkel, R. S., G. L. Nemhauser. 1972. *Integer Programming (Series in Decision and Control)*. John Wiley & Sons, New York.

Graver, J. E. 1975. On the foundations of linear and integer linear programming I. *Math. Programming* 9(1) 207–226.

Hooker, J. N., H. Yan. 1995. Verifying logic circuits by Benders’ decomposition. V. Saraswat, P. Van Hentenryck, eds. *Principles and Practice of Constraint Programming. The Newport Papers*, MIT Press, Cambridge, MA, 267–288.

Hooker, J. N., H. Yan, I. E. Grossmann, R. Raman. 1994. Logic cuts for processing networks with fixed charges. *Comput. Oper. Res.* 21(3) 265–279.

Ibaraki, T. 1977. The power of dominance relations in branch-and-bound algorithms. *J. ACM* 24(2) 264–279.

ILOG. 2007. CPLEX: ILOG CPLEX 11.0 user’s manual and reference manual. ILOG, Sunnyvale, CA. <http://www.ilog.com>.

Katsirelos, G., F. Bacchus. 2003. Unrestricted nogood recording in CSP search. F. Rossi, ed. *Principles and Practice of Constraint Programming (CP 2003). Lecture Notes in Computer Science*, Vol. 2833. Springer-Verlag, Berlin, 873–877.

Katsirelos, G., F. Bacchus. 2005. Generalized nogoods in CSPs. M. M. Veloso, S. Kambhampati, eds. *AAAI-05. AAAI Press/The MIT Press*, Cambridge, MA, 390–396.

Kim, H.-J., J. N. Hooker. 2002. Solving fixed-charge network flow problems with a hybrid optimization and constraint programming approach. *Ann. Oper. Res.* 115(1–4) 95–124.

Kohler, W. H., K. Steiglitz. 1974. Characterization and theoretical comparison of branch-and-bound algorithms for permutation problems. *J. ACM* 21(1) 140–156.

Lecoutre, C., L. Sais, S. Tabary, V. Vidal. 2007. Recording and minimizing nogoods from restarts. *J. Satisfiability, Boolean Modeling Comput.* 1 147–167.

Margot, F. 2002. Pruning by isomorphism in branch-and-cut. *Math. Programming* 94(1) 71–90.

Margot, F. 2003. Exploiting orbits in symmetric ILP. *Math. Programming* 98(1–3) 3–21.

Martello, S., P. Toth. 1990. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, New York.

Moskewicz, M. W., C. F. Madigan, Y. Zhao, L. Zhang, S. Malik. 2001. Chaff: Engineering an efficient SAT solver. *Proc. 38th Design Automation Conf. (DAC 2001), Las Vegas, NV*. ACM, New York, 530–535.

Papadimitriou, C. H., K. Steiglitz. 1982. *Combinatorial Optimization: Algorithms and Complexity*. Prentice Hall, Englewood Cliffs, NJ.

Pisinger, D. 2005. Where are the hard knapsack problems? *Comput. Oper. Res.* 32(9) 2271–2284.

Sandholm, T., R. Shields. 2006. Nogood learning for mixed-integer programming. Technical Report CMU-CS-06-155, Carnegie Mellon University, Pittsburgh.

Scarf, H. E. 1986. Neighborhood systems for production sets with indivisibilities. *Econometrica* 54(3) 507–532.

Schrijver, A. 1998. *Theory of Linear and Integer Programming*. John Wiley & Sons, New York.

Thomas, R. R., R. Weismantel. 1996. Test sets and inequalities for integer programs. *Proc. 5th Internat. IPCO Conf. Integer Programming Combin. Optim. Lecture Notes in Computer Science*, Vol. 1084. Springer-Verlag, Berlin, 16–30.

van Hoesel, S. P. M., A. M. C. A. Koster, R. L. M. J. van de Leensel, M. W. P. Savelsbergh. 2002. Polyhedral results for the edge capacity polytope. *Math. Programming* 92(2) 335–358.

Zhang, L., C. F. Madigan, M. H. Moskewicz, S. Malik. 2001. Efficient conflict driven learning in boolean satisfiability solver. *Proc. ICCAD-01, San Jose, CA*. IEEE Computer Society, Washington, DC, 279–285.