



Discrete Optimization

Decomposition approaches for recoverable robust optimization problems



J. M. van den Akker^{a,*}, P. C. Bouman^{b,1}, J. A. Hoogeveen^a, D. D. Tönissen^{c,1}

^a Department of Information and Computing Sciences Utrecht University, Princetonplein 5, 3584 CC Utrecht, The Netherlands

^b Rotterdam School of Management Erasmus University, Burgemeester Oudlaan 50, 3062 PA Rotterdam, The Netherlands

^c School of Industrial Engineering Eindhoven University of Technology, Den Dolech 2, 5612 AZ Eindhoven, The Netherlands

ARTICLE INFO

Article history:

Received 16 October 2014

Accepted 1 December 2015

Available online 4 January 2016

Keywords:

Recoverable robustness

Column generation

Branch-and-price

Knapsack

Shortest path

ABSTRACT

Real-life planning problems are often complicated by the occurrence of disturbances, which imply that the original plan cannot be followed anymore and some recovery action must be taken to cope with the disturbance. In such a situation it is worthwhile to arm yourself against possible disturbances by including recourse actions in your planning strategy. Well-known approaches to create plans that take possible, common disturbances into account are robust optimization and stochastic programming. More recently, another approach has been developed that combines the best of these two: recoverable robustness. In this paper, we solve recoverable robust optimization problems by the technique of branch-and-price. We consider two types of decomposition approaches: separate recovery and combined recovery. We will show that with respect to the value of the LP-relaxation combined recovery dominates separate recovery. We investigate our approach for two example problems: the size robust knapsack problem, in which the knapsack size may get reduced, and the demand robust shortest path problem, in which the sink is uncertain and the cost of edges may increase. For each problem, we present elaborate computational experiments. We think that our approach is very promising and can be generalized to many other problems.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Most optimization algorithms rely on the assumption that all input data are deterministic and known in advance. However, in many practical optimization problems, such as planning in public transportation or health care, data may be subject to changes. To deal with this uncertainty, different approaches have been developed. In case of *robust optimization* (see Ben-Tal, Ghaoui, and Nemirovski (2009); Bertsimas and Sim (2004)) we choose the solution with minimum cost that remains feasible for a given set of disturbances in the parameters. In case of *stochastic programming* (Birge & Louveaux, 1997), we take *first stage decisions* on basis of the current information and, after the true value of the uncertain data has been revealed, we take the *second stage* or *recourse decisions*. The objective here is to minimize the cost of the first stage decisions plus the expected cost of the recourse decisions. The recourse decision variables may be restricted to a polyhedron

through the so-called technology matrix (Birge & Louveaux, 1997). Summarized, robust optimization wants the initial solution to be completely immune for a predefined set of disturbances, while stochastic programming includes a lot of options to postpone decisions to a later stage or change decisions in a later stage.

Recently, the notion of *recoverable robustness* (Liebchen, Lübbecke, Möhring, & Stiller, 2009) has been developed, which combines robust optimization and second-stage recovery options. Recoverable robust optimization computes solutions, which for a given set of scenarios can be recovered to a feasible solution according to a set of pre-described, fast, and simple recovery algorithms. The main difference between recoverable robustness and stochastic programming is the way in which recourse actions are limited. The property of recoverable robustness that recourse actions must be achieved by applying a simple algorithm instead of being bounded by a polyhedron makes this approach very suitable for combinatorial problems. As an example, consider the planning of buses and drivers in a large city. We may expect that during rush hours buses may be delayed, and hence may be too late to perform the next trip in their schedule. In case of robust optimization, we can counter this only by making the time between two consecutive trips larger than the maximum delay that we want to

* Corresponding author. Tel.: +31 302533989.

E-mail addresses: J.M.vandenAkker@uu.nl (J.M. van den Akker), PBouman@rsm.nl (P.C. Bouman), J.A.Hoogeveen@uu.nl (J.A. Hoogeveen), D.D.Tonissen@tue.nl (D.D. Tönissen).

¹ The research was performed while this author was at Utrecht University.

take into account. This may lead to a very conservative schedule. In case of recoverable robustness, we are allowed to change, if necessary, the bus schedule, but this is limited by the choice of the recovery algorithm. For example, we may schedule a given number of stand-by drivers and buses, which can take over the trip of a delayed driver/bus combination. Especially in the area of railway optimization recoverable robust optimization methods have gained a lot of attention (see e.g. Caprara, Galli, Kroon, Maróti, and Toth (2010); Cicerone et al. (2009)).

In this paper we present a *Branch-and-Price* approach for solving recoverable robust optimization problems. We present two new types of solution approaches: *Separate Recovery Decomposition* (SRD) and *Combined Recovery Decomposition* (CRD). These approaches can be used to model many problems; we will test them on the size robust knapsack problem and on the demand robust shortest path problem.

This paper extends our conference paper (Bouman, van den Akker, Hoogeveen, Demetrescu, & Halldorsson, 2011) by presenting a general definition of our decomposition approaches, a proof that the LP-relaxation of the CRD model dominates the LP-relaxation of the SRD model, and a further study of the solution algorithm for the demand robust shortest path problem. This study includes different column generation strategies and elaborate computational experiments.

To the best of our knowledge, Bouman et al. (2011) and this paper are the first ones applying column generation to recoverable robust optimization. Another decomposition approach, namely Benders decomposition, is used by Cacchiani, Caprara, Galli, Kroon, and Maróti (2008) to assess the Price of Recoverability for recoverable robust rolling stock planning in railways.

The remainder of the paper is organized as follows. In Section 2, we define the concept of recoverable robustness. In Section 3, we present our two different decomposition approaches, and we show the general result that the LP-relaxation of the CRD Model is stronger than the LP-relaxation of the SRD model. In Section 4, we consider the size robust knapsack problem. We investigate the two decomposition approaches in a branch-and-price framework and we present computational experiments in which we compare different solution algorithms. Besides algorithms based on Separate and Combined Recovery Decomposition, we test hill-climbing, dynamic programming, and branch-and-bound. The experiments indicate that Separate Recovery Decomposition performs best. Section 5 is devoted to the demand robust shortest path problem. Since Separate Recovery Decomposition does not seem to be appropriate for this problem, we focus on Combined Recovery Decomposition and consider the settings of the branch-and-price algorithm in more detail. In our experiments we show that the column generation strategy has a significant influence on the computation time. Finally, Section 6 concludes the paper.

2. Recoverable robustness

In this section we formally define the concept of recoverable robustness. We are given an optimization problem

$$P = \min\{f(x) | x \in F\},$$

where $x \in \mathbb{R}^n$ are the decision variables, f is the objective function, and F is the set of feasible solutions.

Disturbances are modeled by a set of discrete scenarios S . We use F_s to denote the set of feasible solutions for scenario $s \in S$, and we denote the decision variables for scenarios s by y^s . The set of algorithms that can be used for recovery are denoted by \mathcal{A} , where $A(x, s) \in \mathcal{A}$ determines a feasible solution y^s from a given initial solution x in case of scenario s . In case of planning buses and drivers a scenario corresponds to a set of bus trips that are delayed, and the algorithms in \mathcal{A} decide about the use of standby drivers.

The *recovery robust optimization problem* is now defined as:

$$\mathcal{RRP}_{\mathcal{A}} = \min\{f(x) + g(\{c^s(y^s) | s \in S\}) | x \in F, A \in \mathcal{A}, \forall_{s \in S} y^s = A(x, s)\}.$$

Here, $c^s(y^s)$ denotes the cost associated with the recovery variables y^s , and g denotes the function to combine these cost into the objective function. There are many possible choices for g . A few examples are as follows:

1. $g(\{c^s(y^s)\}) = 0$. This models the situation where our only concern is the feasibility of the recovered solutions.
2. $g(\{c^s(y^s)\}) = \max_{s \in S} c^s(y^s)$, that is, it models the maximal cost of the recovered solutions y^s . This corresponds to minimizing the worst-case cost. If $c^s(y^s)$ measures the deviation of the solution y^s from x , we minimize the maximum deviation from the initial solution. Note that this deviation may also be limited by the recovery algorithms.
3. $g(\{c^s(y^s)\}) = \sum_{s \in S} p_s c^s(y^s)$, where p_s denotes the probability that scenarios s occurs. This corresponds to minimizing the expected value of the solution after recovery.

In the remainder of the paper we will consider applications that use either a function of sum or max type for $g(\{c^s(y^s)\})$.

Although earlier papers on recoverable robustness (e.g. Liebchen et al. (2009)) consider the latter type of definition of g as two-stage stochastic programming, we think that the requirement of a pre-described easy recovery algorithm makes this definition fit into the framework of recoverable robustness.

3. Decomposition approaches

We discuss two decomposition approaches for recovery robust optimization problems. In both cases we reformulate the problem such that we have to select one solution for the initial problem and one for each scenario. The difference consists of the way we deal with the scenarios.

3.1. Separate Recovery Decomposition

In *Separate Recovery Decomposition*, we select an initial solution and *separately* we select a solution for each scenario. This means that for each feasible initial solution $k \in F$ we have a decision variable x_k signaling if this solution is selected; similarly for each feasible solution for each scenario $q \in F_s$ we have a decision variable y_q^s . In the formulation we enforce that we select exactly one initial solution and one solution for each scenario. The recovery constraints enforces that for each scenario the initial solution can be transformed into a feasible solution by the given recovery algorithm. We assume that the recovery constraint and the objective function can be expressed *linearly*. We now obtain an Integer Linear Programming formulation which is formulated as follows (for maximization objective):

$$\max \sum_{k \in F} c_k x_k + \sum_{s \in S} \sum_{q \in F_s} c_q^s y_q^s$$

subject to

$$\sum_{k \in F} x_k = 1 \quad (1)$$

$$\sum_{q \in F_s} y_q^s = 1 \text{ for all } s \in S \quad (2)$$

$$A_1 x + A_2^s y^s \leq b_s \text{ for all } s \in S \quad (3)$$

$$x_k \in \{0, 1\} \text{ for all } k \in F \quad (4)$$

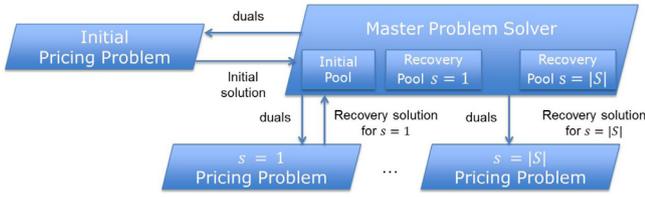


Fig. 1. Column generation process Separate Recovery Decomposition.

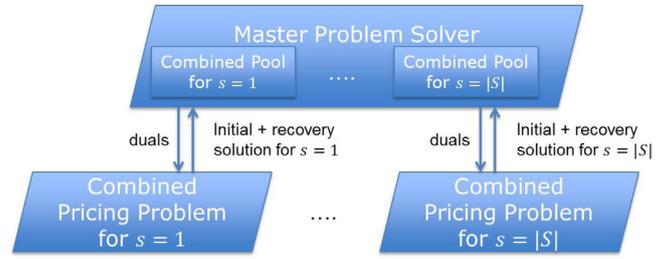


Fig. 2. Column generation process Combined Recovery Decomposition.

$$y_q^s \in \{0, 1\} \text{ for all } q \in F_s, s \in S \quad (5)$$

Here y^s denotes the vector of all entries y_q^s . Constraints (1) and (2) state that exactly one solution is selected in the initial phase and for each scenario, respectively. Constraints (3) model the fact that for each scenario $s \in S$ the solution y^s should be obtained from x by the given recovery algorithm. We want to solve this ILP formulation using Branch-and-Price (Barnhart, Johnson, Nemhauser, Savelsbergh, & Vance, 1998).

The column generation process is depicted in Fig. 1.

Since the variables x_k model initial solutions, we have that for these variables the pricing problem boils down to an optimization problem over the set of initially feasible solutions F , i.e. a variant of the initial problem. Similarly, for variables y_q^s for scenario s the pricing problem becomes an optimization problem over F_s , i.e. a variant of the original problem in case of scenario s .

3.2. Combined Recovery Decomposition

In *Combined Recovery Decomposition*, we select for each scenario a combination of an initial solution together with the optimal recovery of a solution for that single scenario. This means that for each scenario $s \in S$ we have for each combination of an initial solution $k \in F$ and the corresponding solution $q \in F_s$ obtained by the recovery algorithm a binary variable z_{kq}^s signaling if this combination is selected. Obviously, the combinations selected for the different scenarios must all correspond to the same initial solution. To model this we first assume that any initial solution k can be described by the parameters a_{ik} ($i = 1, \dots, l$), where $a_{ik} = 1$ if the initial solution k possesses characteristic i , and $a_{ik} = 0$ otherwise. Next, we introduce binary variables v_i ($i = 1, \dots, l$) to indicate whether the chosen initial solution possesses characteristic i . This leads to the following Integer Linear Programming formulation, where again we assume that we can express the functions f and g in a linear way.

$$\max \sum_{i=1}^l f_i v_i + \sum_{s \in S} \sum_{(k,q) \in F \times F_s} c_q^s z_{kq}^s$$

subject to

$$\sum_{(k,q) \in F \times F_s} z_{kq}^s = 1 \text{ for all } s \in S \quad (6)$$

$$v_i = \sum_{k \in F} a_{ik} \left(\sum_{q \in F_s} z_{kq}^s \right) \text{ for all } i \in \{1, \dots, l\}, s \in S \quad (7)$$

$$z_{kq}^s \in \{0, 1\} \text{ for all } k \in F, q \in F_s, q = A(k, s), s \in S. \quad (8)$$

Recall that $A(k, s)$ denotes the result of the recovery algorithm applied to solution k for scenario s . The binary character of the v variables is guaranteed by the combination of Constraints (7) and (8). Constraint (7) ensures that the same initial solution is selected for each scenario; this initial solution can be constructed using the binary variables v_i .

We also solve this ILP formulation by Branch-and-Price. The column generation process is depicted in Fig. 2. The Pricing Problem now boils down to finding an optimal combination of an initial solution and a recovery solution for a given scenario.

Our approach can be applied to different recoverable robust optimization problems. In the next two sections, we will demonstrate the approach for two different problems. First, we will compare the two models with respect to the strength of the LP-relaxation.

3.3. Comparing the formulations

Since we solve both formulations through Branch-and-Price, we solve the LP-relaxation to find an upper bound. We will show that the LP-relaxation of combined recovery decomposition gives a stronger bound than the LP-relaxation of the separate recovery decomposition. In the proof, we need the following lemma.

Lemma 1. For any feasible solution of the LP-relaxation of the CRD formulation we have that

$$\sum_{q \in F_s} z_{kq}^s \equiv z_k^s = z_k$$

for all $k \in F$.

Proof. Let A denote the $(l \times |F|)$ matrix with entries a_{ik} ($i = 1, \dots, l; k = 1, \dots, |F|$). Then Constraints (7) can be expressed as $v = Az^s$, for all $s \in S$, where $z^s = (z_1^s, \dots, z_{|F|}^s)^T$ and $v = (v_1, \dots, v_l)^T$. Using the linearity of A , we find that $A(z^{s_1} - z^{s_2}) = 0$ for any $s_1, s_2 \in S$. Without loss of generality, we may assume that A has rank l , since we can remove the redundant constraints otherwise. Hence, we may conclude that $z^1 = z^2 = \dots = z^{|S|}$. \square

Theorem 1. Let Z_{LP}^{SRD} and Z_{LP}^{CRD} be the solution value of the LP-relaxation of the Separate and Combined Recovery Decomposition model, respectively. Then

$$Z_{LP}^{SRD} \geq Z_{LP}^{CRD},$$

i.e. the LP-relaxation of combined recovery decomposition is stronger.

Proof. Let (\tilde{v}, \tilde{z}) be an optimal solution to the LP-relaxation of the Combined Recovery Decomposition model. We will give a solution (\bar{x}, \bar{y}) to the Separate Recovery Decomposition model, and show that it is feasible and has the same objective value as (\tilde{v}, \tilde{z}) . We define:

$$\bar{x}_k = \sum_{q \in F_1} \tilde{z}_{kq}^1 \text{ for all } k \in F$$

$$\bar{y}_q^s = \sum_{k \in F} \tilde{z}_{kq}^s \text{ for all } s \in S, q \in F_s$$

Remark that in our definition of \bar{x}_k we could have used the values \tilde{z}_{kq}^s for any scenario, because of Lemma 1. We first check the constraints.

We have that $\sum_{k \in F} \bar{x}_k = \sum_{k \in F} \sum_{q \in F_1} \tilde{z}_{kq}^1 = 1$ because of Constraint (6). Hence, Constraint (1) holds. In the same way,

$\sum_{q \in F_s} \bar{y}_q^s = \sum_{q \in F_s} \sum_{k \in F} \bar{z}_{kq}^s = 1$ because of Constraint (6). Hence, Constraint (2) holds for each $s \in S$.

Next, we have to show that \bar{x}_k and \bar{y}_q^s satisfy the recovery constraints, which are defined as

$$A_1 x + A_2^s y^s \leq b_s \text{ for all } s \in S.$$

Let s be any scenario; we will show that the constraint holds for this scenario. We define e_k^x as the unit vector of length $|F|$ with a 1 on position k ; similarly, we define e_q^s as the unit vector of length $|F_s|$ with a 1 on position q . Since within a variable \bar{z}_{kq}^s , we have that solution q is obtained by applying the recovery algorithm to initial solution k we must have that

$$A_1 e_k^x + A_2^s e_q^s \leq b_s \text{ for all } k \in F; q \in F_s.$$

For each $k \in F$ and $q \in F_s$ we multiply these inequalities with \bar{z}_{kq}^s and add them up to obtain a convex combination of the recovery constraints. This yields

$$\sum_{k \in F} \sum_{q \in F_s} \bar{z}_{kq}^s [A_1 e_k^x + A_2^s e_q^s] \leq \sum_{k \in F} \sum_{q \in F_s} \bar{z}_{kq}^s b_s = b_s.$$

Since the recovery constraints are linear, this can be rewritten as

$$A_1 \sum_{k \in F} \left[\sum_{q \in F_s} \bar{z}_{kq}^s \right] e_k^x + A_2^s \sum_{q \in F_s} \left[\sum_{k \in F} \bar{z}_{kq}^s \right] e_q^s \leq b_s.$$

Using Lemma 1, we have that

$$\bar{x}_k = \sum_{q \in F_1} \bar{z}_{kq}^1 = \sum_{q \in F_s} \bar{z}_{kq}^s;$$

and hence we can substitute \bar{x}_k and \bar{y}_q^s to obtain

$$A_1 \sum_{k \in F} \bar{x}_k e_k^x + A_2^s \sum_{q \in F_s} \bar{y}_q^s e_q^s \leq b_s.$$

Consequently $A_1 \bar{x} + A_2^s \bar{y} \leq b_s$, implying that (\bar{x}, \bar{y}) satisfies Constraint (3). Clearly \bar{x}_k and \bar{y}_q^s are in $[0, 1]$, which implies that the solution is feasible. Now we consider the objective value.

Remark that $c_k = \sum_{i=1}^l f_i a_{ik}$. It follows that

$$\begin{aligned} \sum_{k \in F} c_k \bar{x}_k + \sum_{s \in S} \sum_{q \in F_s} c_q^s \bar{y}_q^s &= \sum_{k \in F} c_k \sum_{q \in F_1} \bar{z}_{kq}^1 + \sum_{s \in S} \sum_{q \in F_s} c_q^s \sum_{k \in F} \bar{z}_{kq}^s \\ &= \sum_{k \in F} \sum_{i=1}^l f_i a_{ik} \left(\sum_{q \in F_1} \bar{z}_{kq}^1 \right) + \sum_{s \in S} \sum_{k \in F} \sum_{q \in F_s} c_q^s \bar{z}_{kq}^s \\ &= \sum_{i=1}^l f_i \sum_{k \in F} a_{ik} \left(\sum_{q \in F_1} \bar{z}_{kq}^1 \right) + \sum_{s \in S} \sum_{k \in F} \sum_{q \in F_s} c_q^s \bar{z}_{kq}^s. \end{aligned}$$

By Constraint (7), the above expression equals

$$\sum_{i=1}^l f_i \bar{v}_i + \sum_{s \in S} \sum_{k \in F} \sum_{q \in F_s} c_q^s \bar{z}_{kq}^s,$$

which is equal to the cost of (\bar{v}, \bar{z}) . \square

4. Size robust knapsack problem

We consider the following knapsack problem. We are given n items, where item j ($j = 1, \dots, n$) has revenue c_j and weight a_j . Each item can be selected at most once. The knapsack size is b . We define the *size robust knapsack problem* as the knapsack problem where the knapsack size b is subject to uncertainty. We denote by $b_s < b$ the size of the knapsack in scenario $s \in S$. We assume that the knapsack will keep its original size with probability p_0 and that scenario s will occur with probability p_s . Our objective is to maximize the expected revenue after recovery. We study the situation in which recovery has to be performed by removing items.

As soon as it becomes clear which scenario s appears, recovery is performed by removing items in such a way that the remaining items give a knapsack with maximal revenue and size at most b_s . This boils down to solving a knapsack problem where the item set is the set of items selected in the initial solution and the knapsack size is b_s . Hence, our set of recovery algorithms is given by the dynamic programming algorithm for solving these knapsacks. For ease of exposition, we consider the binary knapsack problem. Our approach can directly be generalized to the integral knapsack problem.

Recently, Büsing, Koster, and Kutschka (2011a) have studied a different version of our knapsack problem, where they focus on uncertainty in the weights and in the items. They show \mathcal{NP} -hardness of several variants of the problem and develop a polyhedral approach to solve these problems. In a follow-up paper, Büsing et al. (2011b) present an integer linear programming formulation of quadratic size and evaluates the gain of recovery. With respect to approximation algorithms, Disser et al. (2014) consider policies for packing a knapsack with unknown capacity and Goerigk et al. (2014) consider the knapsack problem in which there is a limited budget to decrease item weights.

In this section, we discuss our two decomposition approaches for the size robust knapsack problem and present elaborate computational experiments in which we compare our method with other algorithms.

4.1. Separate Recovery Decomposition

We define $K(b)$ as the set of feasible knapsack fillings with size at most b . For $k \in K(b)$, we denote its revenue by $C_k = \sum_{i \in k} c_i$. In the same way, we denote the revenue of $q \in K(b_s)$ by $C_q^s = \sum_{i \in q} c_i$.

We define decision variables:

$$x_k = \begin{cases} 1 & \text{if knapsack } k \in K(b) \text{ is selected,} \\ 0 & \text{otherwise.} \end{cases}$$

and

$$y_q^s = \begin{cases} 1 & \text{if knapsack } q \in K(b_s) \text{ is selected for scenario } s, \\ 0 & \text{otherwise.} \end{cases}$$

The problem can now be formulated as follows:

$$\max p_0 \sum_{k \in K(b)} C_k x_k + \sum_{s \in S} p_s \sum_{q \in K(b_s)} C_q^s y_q^s$$

subject to

$$\sum_{k \in K(b)} x_k = 1 \tag{9}$$

$$\sum_{q \in K(b_s)} y_q^s = 1 \text{ for all } s \in S \tag{10}$$

$$\sum_{k \in K(b)} a_{ik} x_k - \sum_{q \in K(b_s)} a_{iq}^s y_q^s \geq 0 \text{ for all } i \in \{1, 2, \dots, n\}, s \in S \tag{11}$$

$$x_k \in \{0, 1\} \text{ for all } k \in K(b) \tag{12}$$

$$y_q^s \in \{0, 1\} \text{ for all } q \in K(b_s), s \in S. \tag{13}$$

where the parameters a_{ik} and a_{iq}^s are defined as follows:

$$a_{ik} = \begin{cases} 1 & \text{if item } i \text{ is in knapsack } k \in K(b), \\ 0 & \text{otherwise.} \end{cases}$$

and

$$a_{iq}^s = \begin{cases} 1 & \text{if item } i \text{ is in knapsack } q \in K(b_s), \\ 0 & \text{otherwise.} \end{cases}$$

In the above model Constraint (9) states that exactly one knapsack is selected for the original situation and Constraints (10) that exactly one knapsack is selected for each scenario. Constraints (11) ensures that recovery is done by removing items.

Recall that we solve this ILP formulation using Branch-and-Price; thereto, we relax the integrality Constraints (12) and (13) into $x_k \geq 0$ and $y_k^s \geq 0$, and solve this LP-relaxation by column generation.

The pricing problem

From the theory of linear programming it is well-known that for a maximization problem increasing the value of a variable will improve the current solution if and only if its reduced cost is positive. The pricing problem then boils down to maximizing the reduced cost.

Let λ , μ_s , and π_{is} be the dual variables of Constraints (9)–(11), respectively. Now the reduced cost $c^{\text{red}}(x_k)$ of x_k is given by

$$c^{\text{red}}(x_k) = p_0 \sum_{i \in k} c_i - \lambda - \sum_{i=1}^n \sum_{s \in S} a_{ik} \pi_{is}$$

$$= \sum_{i=1}^n a_{ik} \left(p_0 c_i - \sum_{s \in S} \pi_{is} \right) - \lambda.$$

Observe that λ and μ are free variables, and $\pi_{is} \leq 0$. The pricing problem is to find a feasible knapsack for the original scenario, where the revenue of item i , equals $(p_0 c_i - \sum_{s \in S} \pi_{is})$. This is just the original knapsack problem with modified objective coefficients. Similarly the reduced cost $c^{\text{red}}(y_q^s)$ are given by $c^{\text{red}}(y_q^s) = \sum_{i=1}^n a_{iq}^s (p_s c_i + \pi_{is}) - \mu_s$. Therefore, the pricing is exactly the knapsack problem with knapsack size b_s and modified objective coefficients. Note that in the pricing problem an item may have a negative revenue. Clearly such items can be discarded.

To find an integral solution, we are going to apply Branch-and-Price. We branch on items that are fractional in the current solution, i.e. items i for which $\sum_{k \in K(b)} a_{ik} x_k$ is fractional. This is easily combined with column generation, since enforcing that an item is taken in or omitted from the knapsack can easily be included in the pricing problem. If all values $\sum_{k \in K(b)} a_{ik} x_k$ are integral, then a single initial knapsack is selected with value 1. Now consider a basic solution to the LP for scenario s . It is easy to see that this solution contains an optimal subset of the initial knapsack with total weight at most b_s which is selected with value 1. Consequently, the solution is integral.

4.2. Combined Recovery Decomposition

In contrast to the *Separate Recovery Decomposition*, we consider fillings of the initial knapsack in combination with the optimal recovery for *one* scenario. Consequently, we introduce decision variables:

$$z_{kq}^s = \begin{cases} 1 & \text{if the combination of initial solution } k \in K(b) \\ & \text{and recovery solution } q \in K(b_s) \text{ is selected} \\ & \text{for scenario } s, \\ 0 & \text{otherwise.} \end{cases}$$

Clearly, z_{kq}^s is only defined if q is a subset of k . The ILP model further contains the original variable x_i signaling if item i is contained in the initial knapsack. We can formulate the problem as follows:

$$\max p_0 \sum_{i=1}^n c_i x_i + \sum_{s \in S} p_s \sum_{(k,q) \in K(b) \times K(b_s)} C_q^s z_{kq}^s$$

subject to

$$\sum_{(k,q) \in K(b) \times K(b_s)} z_{kq}^s = 1 \quad \text{for all } s \in S \tag{14}$$

$$x_i - \sum_{(k,q) \in K(b) \times K(b_s)} a_{ik} z_{kq}^s = 0 \quad \text{for all } i \in \{1, 2, \dots, n\}, s \in S \tag{15}$$

$$x_i \in \{0, 1\} \quad \text{for all } i \in \{1, 2, \dots, n\} \tag{16}$$

$$z_{kq}^s \in \{0, 1\} \quad \text{for all } k \in K(b), q \in K(b_s), s \in S, \tag{17}$$

Constraints (14) enforce that exactly one combination is selected for each scenario; Constraints (15) ensure that the same initial knapsack filling is selected for all scenarios.

Again, we are going to solve the LP-relaxation by column generation. We include the variables x_i in the restricted master LP and, hence pricing is only performed for the variables z_{kq}^s . We denote the dual variables of Constraints (14) and (15) by ρ_s and σ_{is} , respectively. The reduced cost of z_{kq}^s is now equal to:

$$c^{\text{red}}(z_{kq}^s) = \sum_{i=1}^n a_{iq}^s p_s c_i + \sum_{i=1}^n a_{ik} \sigma_{is} - \rho_s.$$

We solve the pricing problem for each scenario separately. We have to find an initial and recovery solution. This can be solved by dynamic programming. The main observation is that there are three types of items: items included in both the initial and recovery knapsack, items selected for the initial knapsack, but removed by the recovery, and non-selected items. We define state variables $D(i, w_0, w_s)$ as the best value for a combination of an initial and recovery knapsack for scenario s , such that the initial knapsack is a subset of $\{1, 2, \dots, i\}$, the recovery knapsack is a subset of the initial knapsack, and the initial and recovery knapsack have weight w_0 and w_s , respectively. The recurrence relation is as follows:

$$D(i, 0, 0) = 0 \quad \forall i$$

$$D(0, w_0, w_s) = -\infty \quad \text{for } w_0, w_s > 0$$

$$D(i, w_0, w_s) = \max \begin{cases} D(i-1, w_0, w_s) \\ D(i-1, w_0 - a_i, w_s) + \sigma_{is} \\ D(i-1, w_0 - a_i, w_s - a_i) + \sigma_{is} + p_s c_i \end{cases}$$

4.3. Computational results

We performed extensive computational experiments with the knapsack problem. The algorithms were implemented in the Java Programming language and the Linear Programs were solved using ILOG CPLEX 11.0. All experiments were run on a PC with an Intel®Core™Duo 6400 @ 2.13 GHz processor with 1 GB of RAM. Our experiments were performed in three phases. Since we want to focus on difficult instances, in the first phase we tested 12 different instance types to find out which types are the most difficult. We ran two algorithms: the separate recovery branch-and-price and a hill climbing algorithm. Our instance types are based on the instance types by Pisinger (2005), where we have to add the knapsack weight b_s and the probability p_s for each of the scenarios. In the second phase, we tested five different algorithms on relatively small instances with up to 25 items and up to 8 scenarios. In the third phase we tested the best two algorithms from the second phase on larger instances with 50 or 100 items and up to 20 scenarios. In this section, we will present the most important observations from the second and third phase. For further details we refer to Bouman (2011).

Table 1
Second phase results.

Algorithm	Failed	Avg t(ms)	Max t(ms)	Avg $\frac{c}{c^*}$	Min $\frac{c}{c^*}$	Avg nodes	Max nodes
Separate Recovery	128	107	2563	–	–	3.27	122
Combined Recovery	1407	417	2969	–	–	1.12	17
Branch and Bound	190	111	2906	–	–	1281	33321
DP	2840	347	2984	–	–	–	–
Hill climbing	n.a.	17.3	422	0.99	0.85	–	–

In the second phase we tested the following 5 instance classes with $R = 30$:

- The almost strongly correlated instances: pick a_i uniformly from $[1, R]$ and c_i uniformly from $[a_i + R/10 - R/500, a_i + R/10 + R/500]$.
- The inverse strongly correlated instances: pick c_i uniformly from $[1, R]$ and a_i from $[c_i + R/10]$.
- The circle(d) instances: pick a_i uniformly from $[1, R]$ and $c_i = d\sqrt{4R^2 - (a_i - 2R)^2}$, where d is set to $\frac{2}{3}$.
- The span(v, m) uncorrelated instances. Create v items with a'_j and c'_j from $[1, R]$ and divide a'_j and c'_j by $m + 1$. The knapsack instance is then generated by repeatedly choosing one of the v combinations (a'_j, c'_j) and a random value μ_i from $[1, m]$: the created item gets $a_i = \mu_i a'_j$ and $c_i = \mu_i c'_j$.
- The subset sum instances, where $a_i = c_i$ drawn from $[1, R]$.

We considered instances with 5, 10, 15 and 25 items and with 2, 4, 6 and 8 scenarios (except for the instances with 5 items, for which we only considered 2 and 4 scenarios). For each combination we generated 100 item sets (20 from each instance class). For each item set we generated 3 sets of scenarios, with large, middle, and small values of b_s relative to b , where large, middle, and small mean that the values b_s are drawn uniformly from $[\frac{1}{2}b, b]$, $[\frac{3}{8}b, \frac{6}{8}b]$, and $[\frac{1}{4}b, \frac{1}{2}b]$, respectively. This means that we considered 4200 $((2 + 3 \times 4) \times 100 \times 3)$ instances in total.

We report results on the following algorithms:

- Separate Recovery Decomposition with Branch-and-Price, where we branch on the fractional item with largest $\frac{c_j}{a_j}$ ratio and first evaluate the node which includes the item.
- Combined Recovery Decomposition with Branch-and-Price, where we branch in the same way as in Separate Recovery Decomposition.
- Branch-and-Bound where we branch on the fractional item with smallest $\frac{c_j}{a_j}$ ratio and first evaluate the node which includes the item. In every node, the optimal recovery for the current set of items is computed using dynamic programming for the regular knapsack problem. When a node adds a new item, the DP-table of the parent node is extended with only this item in order to keep the computation efficient. The upper bounds are computed by solving the LP-relaxation, which is solved by including the items in order of $\frac{c_j}{a_j}$ ratio in the initial solution and in each scenario.
- Dynamic programming. Our DP resembles the DP for the pricing problem in case of Combined Recovery Decomposition. We use state variables $D(i, w_0, w_1, \dots, w_{|S|})$, where w_j ($j = 1, \dots, |S|$) corresponds to scenario j . The recurrence relation gets exponential size now, since, if item k is included in the initial solution, we must compute for each possible subset of S the value of including or omitting this item. Hence, the running time becomes $O(n2^{|S|})$.
- Hill Climbing. We apply neighborhood search on the initial knapsack only and compute for each initial knapsack the optimal recovery by Dynamic Programming. Hill climbing performs 100 restarts.

Table 2
Third phase results for Separate Recovery decomposition.

Items	Scenarios	Failed	Avg ms	Max ms	Avg nodes	Max nodes
50	2	2	686	56312	1.56	68
50	3	12	2724	53454	1.7	25
50	4	46	3799	58688	2.6	35
50	10	125	3295	53483	2.29	35
50	20	144	1473	38766	1.4	17
100	2	114	1695	47531	1.05	5
100	3	173	703	24781	1.16	11
100	4	176	964	46172	2.03	59
100	10	213	469	34547	1.39	25
100	20	210	103	2703	1.13	13

For the branching algorithms we tested different branching strategies. All our algorithms branch on whether an item should be included in the initial knapsack. The different branching strategies select which item should be branched on first and whether the item should first be included or excluded. The first item is a combination of best/worst according to weight/ratio/profit. In total this yields 12 branching strategies. In the Branch-and-Price algorithms the difference in performance turned out to be minor and we report on the strategy that performed best in Separate Recovery Decomposition. However, in the Branch-and-Bound algorithm some difference could be observed and we report on the strategy that shows the best performance for this algorithm.

The results of the second phase are given in Table 1. For each instance and each algorithm, we allowed at most 3000 milliseconds of computation time. For each algorithm, we report on the number of instances (out of 4200) that could not be solved within 3000 milliseconds, and the average and maximum computation time over the successful instances. For Hill Climbing we compute for each instance the performance ratio, which is equal to the value of the obtained solution c divided by the value of the optimum solution c^* ; in the table we report the average and worst performance ratio. For the branching algorithms we further report the average and maximum number of evaluated nodes. For Hill Climbing 'Failed' is not applicable.

The results indicate that for this problem Separate Recovery Decomposition outperforms Combined Recovery Decomposition. DP is inferior to Branch-and-Bound and Hill Climbing. The fact that Branch-and-Bound requires more nodes than the Separate and Combined Recovery Branch-and-Price indicated that the decomposition models have a stronger LP-relaxation than the standard ILP based on variables for each item and scenario.

In the third phase we experimented with larger instances for the two best algorithms. We present results for instances with 50 and 100 items and 2, 3, 4, 10, or 20 scenarios. Again, for each combination of number of items, number of scenarios, we generated 100 item sets (20 from each instance class) with 3 scenario sets each. This results in 300 instances per combination of number of items and number of scenarios, where the maximum computation time per instance per algorithm is 4 minutes. The results are depicted in Tables 2 and 3.

The results suggest that the computation time of Separate Recovery Decomposition scales very well with the number of

Table 3
Third phase results for hill climbing.

Items	Scenarios	Avg ms	Max ms	Avg $\frac{c}{c}$	Min $\frac{c}{c}$
50	2	104	969	0.98	0.68
50	3	173	1204	0.98	0.84
50	4	180	1203	0.98	0.83
50	10	268	1407	1	0.94
50	20	309	1515	1	0.84
100	2	887	19656	0.98	0.66
100	3	1257	25578	1	0.86
100	4	1783	32625	1	0.8
100	10	3546	34703	1	0.8
100	20	4546	37312	1	0.94

scenarios. As may be expected, Hill Climbing shows a significant increase in the computation time when the number of scenarios is increased. Moreover, the small number of nodes indicates that Separate Recovery Decomposition is well-suited for instances with a larger number of scenarios. On average the quality of the solutions from Hill Climbing is very high. However, a worst performance ratio of about 0.66 shows that there is no guarantee of quality. Hill Climbing succeeded to complete the algorithm with 100 restarts for all instances.

5. The demand robust shortest path problem

The demand robust shortest path problem is an extension of the shortest path problem that has been introduced by Dhamdhere, Goyal, Ravi, and Singh (2005). We are given a graph (V, E) with cost c_e on the edges $e \in E$, and a source node $v_{source} \in V$. The question is to find the cheapest path from source to the sink, but the exact location of the sink is subject to uncertainty. Moreover, the cost of the edges may change over time. More formally, there are multiple scenarios $s \in S$ that each define a sink v_{sink}^s and a factor $f^s > 1$ by which the cost of the edges are scaled. To work with the same problem as Dhamdhere et al. (2005), we choose as objective to minimize the cost of the worst case scenario. It is not difficult to see that this problem is \mathcal{NP} -hard, as it generalises the Steiner Tree Problem. When we pick each f^s high enough, the optimal solution is to buy a minimum cost tree that connects the source and all sinks during the first phase.

In contrast to Büsing (2012), we can buy any set of edges in the initial planning phase. In the recovery phase, we have to extend the initial set such that it contains a path from the source to the sink v_{sink}^s , while paying increased cost for the additional edges. Remark that, when the sink gets revealed, the recovery problem can be solved as a shortest path problem, where the edges already bought get zero cost. Hence, the recovery algorithm is a shortest path algorithm.

Observe that the recovery problem has the constraint that the union of the edges selected during recovery and the initially selected edges contains a path from source v_{source} to sink v_{sink}^s . It is quite involved to express this constraint using linear inequalities, and hence to apply Separate Recovery Decomposition. However, the constraint fits very well into Combined Recovery Decomposition.

Our Combined Recovery Decomposition model contains the variable x_e signaling if edge $e \in E$ is selected initially. Moreover, for each scenario, it contains variables indicating which edges are selected initially and which edges are selected during the recovery:

$$z_{kq}^s = \begin{cases} 1 & \text{if the combination of initial edge set } k \subseteq E \\ & \text{and recovery edge set } q \subseteq E \text{ is selected for scenario } s, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that z_{kq}^s is only defined if k and q are feasible, i.e., their intersection is empty and their union contains a path from v_{source} to v_{sink}^s . Finally, it contains Z_{max} defined as the maximal recovery cost.

We can formulate the problem as follows:

$$\begin{aligned} \min & \sum_{e \in E} c_e x_e + Z_{max} \\ \text{subject to} & \\ & \sum_{(k,q) \subseteq E \times E} z_{kq}^s = 1 \quad \text{for all } s \in S \end{aligned} \tag{18}$$

$$x_e - \sum_{(k,q) \subseteq E \times E} a_{ek} z_{kq}^s = 0 \quad \text{for all } e \in E, s \in S \tag{19}$$

$$Z_{max} - \sum_{e \in E} f^s c_e \sum_{(k,q) \subseteq E \times E} a_{eq}^s z_{kq}^s \geq 0 \quad \text{for all } s \in S \tag{20}$$

$$x_e \in \{0, 1\} \quad \text{for all } e \in E \tag{21}$$

$$z_{kq}^s \in \{0, 1\} \quad \text{for all } k \subseteq E, q \subseteq E, s \in S, \tag{22}$$

where the binary parameters a_{ek} signal whether edge e is in edge set k and the binary parameters a_{eq}^s signal whether edge e is in edge set q for scenario s .

Constraints (18) ensure that exactly one combination of initial and recovery edges is selected for each scenario; Constraints (19) enforce that the same set of initial edges is selected for each scenario. Finally, Constraints (20) make sure that Z_{max} represents the cost of the worst case scenario.

5.1. Solving the LP by column generation

We first relax the integrality Constraints (21) and (22) into $x_e \geq 0$ and $z_{kq}^s \geq 0$, and solve this LP-relaxation. To deal with the huge number of variables we are going to solve the problem by *column generation*.

The pricing problem Since we have a minimization problem, the pricing problem boils down to minimizing the reduced cost. Let λ_s , ρ_{es} , and π_s be the dual variables of the Constraints (18)–(20) respectively. The reduced cost of z_{kq}^s is now equal to:

$$c^{\text{red}}(z_{kq}^s) = -\lambda_s + \sum_{e \in E} \rho_{es} a_{ek} + \sum_{e \in E} \pi_s f^s c_e a_{eq}^s.$$

We solve the pricing problem for each scenario *separately*. For a given scenario s , the pricing problem boils down to minimizing $c^{\text{red}}(z_{kq}^s)$ over all feasible a_{ek} and a_{eq}^s . This means that we have to select a subset of edges that contains a path from v_{source} to v_{sink}^s . This subset consists of edges which have been bought initially and edges which are attained during recovery. The first type corresponds to $a_{ek} = 1$ and has cost ρ_{es} and the second type to $a_{eq}^s = 1$ and has cost $\pi_s f^s c_e$. The pricing problem is close to a shortest path problem, but we have two binary decision variables for each edge. However, we can apply the following *preprocessing steps*:

1. First, we select all edges with negative cost. From LP theory it follows that all dual variables π_s are nonnegative, and hence, all recovery edges have nonnegative cost. So we only select initial phase edges with negative cost ρ_{es} . From now on, the cost of these edges is considered to be 0.
2. The other edges can either be selected in the initial phase or in the recovery phase. To minimize the reduced cost, we have to choose the cheapest option. This means that we can set the cost of an edge equal to $\min(\rho_{es}, \pi_s f^s c_e)$.

The pricing problem now boils down to a *shortest path* problem with nonnegative cost on the edges and hence can be solved by Dijkstra (1959) algorithm. We implemented the algorithm by a min heap with running time $O(|E|\log(|V|))$.

Since we solve the pricing problem for each scenario separately, the following questions arise: ‘For which scenarios do we actually solve the pricing problem?’ and ‘Which columns do we actually add to the restricted LP?’. We investigate the following strategies:

- *Interleaved*: goes through the pricing problems of the different scenarios one by one. As soon as a variable with negative reduced cost is identified, the corresponding column is added and the master problem gets resolved. After that, it goes to the next scenario. When the pricing problem has a solution with non-negative reduced cost for every scenario the column generation process is stopped.
- *Best*: solves the pricing problem for all scenarios, but only a column z_{kq}^s with overall minimal reduced cost is added to the master problem. The master problem is solved again and this repeats itself until the minimal reduced cost is nonnegative.
- *All*: solves the pricing problem for all scenarios and adds a column for all scenarios for which a variable z_{kq}^s with negative reduced cost was found; after adding all those columns it resolves the master problem.

Within the first few experiments it became very clear that the LP problem is very degenerate. Certainly for larger graphs with a lot of scenarios this tends to slow down the computation enormously. Observe that every solution needs at least $|S|$ columns. To get a complete solution, because of Constraint (19), we need a collection of columns such that for each edge e the total amount by which it is selected in the initial solution $\sum_{(k,q) \subseteq E \times E} a_{ek} z_{kq}^s$ is the same for every scenario s . This has the consequence that, although it is included in the basis, a promising new column often does not influence the primal solution. To deal with this problem, we use the following method: When a column is added, we always guarantee that it can be selected for the solution by generating for every scenario the best column with the same initial edges. Those columns are generated by fixing the set of initial edges and finding the best recovery edges by running Dijkstra’s algorithm for all other $|S| - 1$ scenarios.

As a starting solution we take the column in which all edges are taken in the initial solution. Other strategies were tested but the differences were small and instance dependent.

Moreover, we have investigated column deletion, i.e. deletion of columns with too positive reduced cost. However, this does not seem to work well in combination with including additional columns.

5.2. Solving the ILP

If the solution of the LP-relaxation is integral, our problem is solved to optimality. Otherwise, we proceed by *Branch-and-Price* (Barnhart et al., 1998), i.e. Branch-and-Bound, where we generate additional columns in the nodes of the search tree.

In a Branch-and-Price algorithm the branching strategy has to be designed in such a way that we are still able to solve the pricing problem in each node of the tree. In our algorithm we branch on the variables x_e . In a node with $x_e = 1$ we only generate columns where edge e is bought in the initial phase. This implies that in the first preprocessing step of the pricing we buy edge e at cost ρ_{es} and then set its cost to 0. In a node with $x_e = 0$ we are not allowed to buy edge e in the initial phase. Therefore, we have to define the cost of the edge as $\pi_s^f c_e$ instead of $\min(\rho_{es}, \pi_s^f c_e)$.

Concerning the choice of the edge for branching, besides considering the edges in order of their index, we implemented branching on the most doubtful edge. This means that the we branch on

the edge for which $|x_e - \frac{1}{2}|$ is minimal. This strongly speeds up the computation.

Moreover, we investigated different node selection strategies. We considered best bound branching, i.e. branching on the node with the minimal lower bound, breadth first search, depth first search and also best depth first, which from the deepest nodes in the tree selects the one with the best lower bound. In our experiments best depth first did not improve depth first search very much. Although depth first search sometimes slightly improved best bound search, it showed a less stable behavior. The same is true for breadth first search. Therefore we chose to use best bound branching in our algorithm.

To compute an upper bound three rounding heuristics were tested. The first heuristic was to select for the initial phase only edges with $x_e = 1$ in the LP-solution. In the second heuristic, all edges with $x_e \geq \frac{1}{2}$ were selected in the initial solution. As a third alternative we applied a randomized strategy: each edge was selected in the initial solution with a probability equal to the value of x_e in the optimal solution of the LP-relaxation. In all three cases, for each scenario the best recovery solution was determined by Dijkstra’s algorithm. There did not seem to be much difference in performance between the heuristics and we applied the second one since we thought it to be the most intuitive one.

5.3. Computational results

We have implemented our column generation and branch-and-price algorithms in Java and used ILOG CPLEX 12.4 as linear programming solver. We ran experiments on an Intel®Core™Duo 2.66 GHz processor with 4 GB of RAM.

Again, our experiments were performed in three phases. We first investigated all column generation strategies, to determine the best one. Secondly, we performed a sensitivity analysis. Finally, we ran our algorithm on some larger instances.

We first present results for linear programming to illustrate the effect of the different column addition strategies from Section 5.1. The strategies InterA, BestA, and AllA denote extensions of the strategies Inter, Best and All, in which, when we add a column, we also add for each scenario the best column with the same initial edges. In Table 4 we give results for 4 different randomly generated relatively small instances, where $G_{n,e}$ has a graph with n nodes and e edges. The recovery factor f for these instances is fixed at 2.0 and every non source node is a possible sink and thus a scenario. For each instance, we give the number of iterations (it), the number of added columns (col), and the computation time in milliseconds (t).

Our results reveal that the strategies with additional columns strongly speed up the computation. In most of our cases the number of columns is also reduced, but as may be expected, the reduction is not that strong. In some cases ($G_{17,31}$ Inter and All) the number of columns increased.

We also solved the ILP for these instances, where we applied all combinations of strategies in the root and in the tree. The strategies without additional columns resulted in large running times.

The improvement from the additional columns is explained by the fact that these columns enable the solver to actually use every generated ‘interesting’ column in a solution. This is especially important for the combined recovery decomposition model, where each column is a combination of an initial solution and a recovery solution for one scenario. In this model, a feasible solution requires a set of columns that constitute the same initial solution for each scenario. So the additional columns prevent a known problem in column generation: a very good column is generated but can not be used because other columns to complete the solution are lacking.

Table 4
Results for linear programming.

Method	$G_{4,5}$			$G_{14,13}$			$G_{17,23}$			$G_{17,31}$		
	it	col	t	it	col	t	it	col	t	it	col	t
Inter	32	40	10	613	639	1014	3701	3740	25.941	7132	7179	134213
Best	22	30	10	545	571	1038	2835	2874	19.278	6065	6112	136585
All	17	39	8	114	699	554	577	3632	8069	1101	7453	54145
InterA	5	21	6	33	443	55	64	1048	393	537	8624	13132
BestA	6	24	4	24	326	39	73	1192	583	257	4144	7503
AllA	8	21	3	40	443	55	81	1048	407	652	8624	16649

Table 5
Results for branch-and price with random instances.

ed	InterA			BestA			BestANoSort			AllA		
	it	nd	t	it	nd	t	it	nd	t	it	nd	t
11	37	1	20	30	1	19	36	2	24	49	1	20
14	153	5	135	104	6	129	143	10	158	195	5	139
17	385	11	791	209	10	510	358	19	689	472	11	824
20	715	17	2842	414	19	1510	1042	77	2417	888	17	2951
23	2493	41	24669	1353	49	10208	2782	119	14832	3016	41	26606
26	10121	157	174213	4520	154	64500	18514	1238	119645	11921	157	193029
29	20630	325	488442	9641	256	184075	21837	638	297424	23939	272	571894

Therefore, we do not feel it necessary to perform more experiments without additional columns. Moreover, it did not pay off to use a different strategy in the root than in the remainder of the tree. Consequently, from now on we only consider strategies with additional columns that use the same strategy for the complete tree.

Recall that we branch on the most doubtful edge with $|x_e - \frac{1}{2}|$ minimal and select the node with the best lower bound. We first report results for a set \mathcal{G}_{500} of 500 random instances. They are based on graphs with 10 to 29 edges, where for each number of edges we vary the number of nodes. For every number of edges a total of 25 graphs are generated. All graphs are connected, the cost of the edges is drawn uniformly randomly from the interval $[0; 100]$, every non-source is a possible sink, and f has a random value in the interval $[1; 10]$.

The average solution times for those instances are 53.6, 21.2, 66.7 and 34.8 seconds for InterA, BestA, and AllA and BestANoSort, where in the latter strategy we branch on the edges in order of their numbering instead of on the most doubtful edge. The BestA method performs significantly better than the other methods according the Wilcoxon signed-rank test (done with R version 3.1.1, with $p = 3.823e^{-13}$ as the highest p). Table 5 shows results for a subset of the set of random instances \mathcal{G}_{500} . For each number of edges and each strategy, we report on the average total number of iterations of column generation (it), the average number of nodes in the branch-and-bound tree (nodes), and the average computation time in milliseconds (t).

In Fig. 3 we plot on a logarithmic scale the computation time for each number of edges.

These results suggest that especially for larger graphs BestA outperforms the other column addition strategies. Even when BestA is combined with the inferior branching strategy of branching on edges in lexicographical order, this is faster than the other column addition strategies.

We also did some sensitivity experiments on the influence of the edge cost and the recovery factor f . To test the influence of the costs of the edges we used $G_{14,13}$ and $G_{17,31}$ with fixed recovery factor $f = 2$. We created 500 random cost versions by generating the edge costs uniformly randomly from the interval $[0; 50]$. For

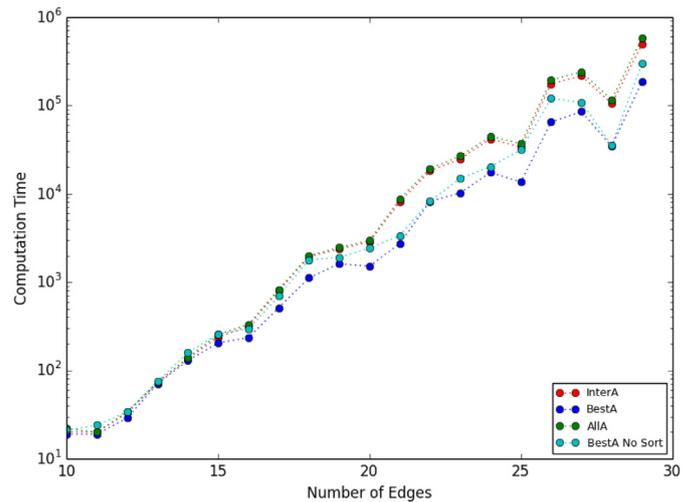


Fig. 3. Logarithmic time results of the methods.

the recovery factor f we generated similar instances only now we fixed the cost and varied the recovery factor in the interval $[1; 10]$. Because the $G_{17,31}$ graph with random recovery factor f was solved relatively slowly, we only solved 25 instances. Because these experiments are about sensitivity we report the quartile points, which divide the data into four equal groups, instead of averages. We looked at the iterations, nodes and time separately, and also report the minimum and maximum. These results can be found in Table 6. All instances are solved with the BestA method.

These results suggest that cost as well as recovery factor have a large influence on the iterations, nodes and solution time of the instance, this difference can be a factor of more than 100. This might be explained by the fact that some combinations of cost and recovery factor result in alternative solutions with approximately the same value, which have an impact on the size of the search tree. We consider a small example with 3 nodes: one source s and two possible sink t_1 and t_2 each occurring with probability $\frac{1}{2}$. There are

Table 6
Results for varying the cost (c) and the recovery factor (f) of the edges.

	$G_{14, 13} c$			$G_{17, 31} c$			$G_{14, 13} f$			$G_{17, 31} f$		
	it	nd	t	it	nd	t	it	nd	t	it	nd	t
Minimum	30	3	57	146	3	281	27	1	161	483	1	42651
First quartile	134	13	235	1044	38	6323	51	1	325	3183	39	100754
Median	185	29	325	1922	85	13289	74	1	416	6964	99	173167
Third quartile	235	49	445	3177	166	18687	116	9	606	9404	153	238108
Maximum	675	211	1475	9753	1567	51236	279	73	1112	51923	1221	520106

Table 7
Effect of number of scenarios on average solution time (ms).

Edges	Number of scenarios			
	5	10	15	20
11	9	-	-	-
14	19	123	-	-
23	102	1390	7103	-
26	120	3495	19941	59429
29	573	16296	107511	249716

two edges (s, t_1) and (s, t_2) with the same initial cost Q . If $f = 2$, then it does not make a difference if you buy all edges, one edge, or no edges in the initial phase.

The experiments from now on, were performed with a better computer with an Intel®Core™i5 3.40 GHz processor and 8 GB of RAM. This computer is approximately twice as fast. Until now we considered instances where every non-source node could be the sink, which are instances with a relatively large number of scenarios. Since the size of the ILP model is linear in the number of scenarios, we may expect that instances with fewer scenarios can be solved faster. From the set G_{500} used before we generated 42747 new instances by varying the number of scenarios.

For each instance in the set G_{500} , we generate $\min\{\binom{n-1}{i}, 10\}$ instances with i scenarios, i.e., i possible sink nodes, where n is the number of nodes of the instance. Hence, for each instance, in total $\sum_{i=1}^{n-1} \min\{\binom{n-1}{i}, 10\}$ instances are generated. The maximum of 10 instances which can be generated per instance and number of scenarios is used to limit the total number of instances. If $\binom{n-1}{i} > 10$, we randomly select 10 subsets of size i . For example, when an instance of the G_{500} set has 4 nodes, 7 new instances are generated. Three instances, with 1 scenario are generated, three with 2 scenarios and one instances with 3 scenarios. Note that the instance which has 3 scenarios is an exact copy from the G_{500} instance. For each instance the nodes, edges and the cost of the edges are copied from the G_{500} instance. For each instance the corresponding recovery values f^s from the G_{500} instance are used for the scenarios.

These instances were solved with the BestA method. Solving all 42747 instances took 56.4 hours. In Table 7 we show the average computation time in milliseconds for different numbers of edges and different numbers of scenarios.

Moreover, in Figs. 4 and 5, we plot the computation time on a logarithmic scale, per number of edges as a function of the number of scenarios and per number of scenarios as a function of the number of edges, respectively. A larger version of the figures can be found in the appendix.

Our results suggest that the number of edges has a strong impact on the computation time, which indicates exponential behavior. The impact of the number of scenarios, seems very strong in the beginning but then somewhat flattens out in the logarithmic scale. We conclude that both have a strong influence.

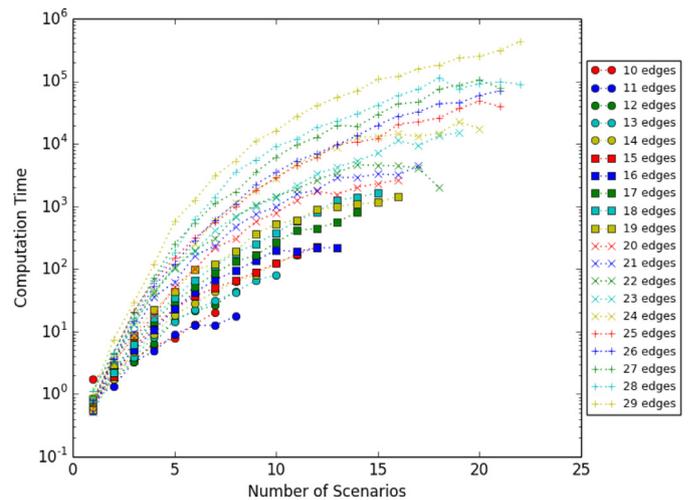


Fig. 4. Time results per number of edges.

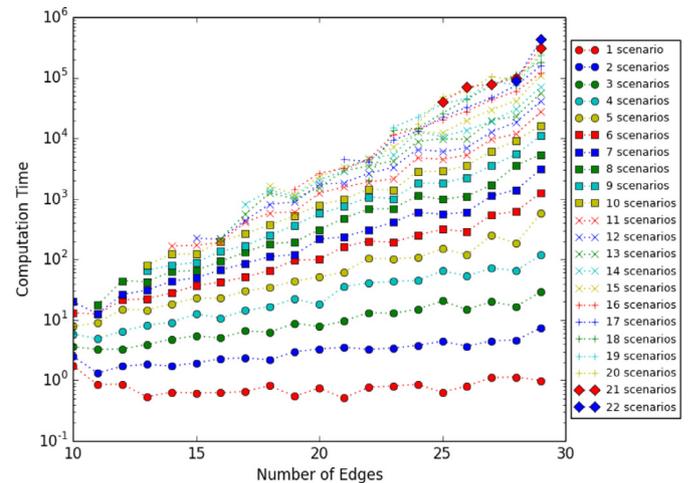


Fig. 5. Time results per number of scenarios.

Finally, we performed experiments for four larger instances, which are generated to get an indication of the boundaries of what can be solved. Because these instances take a long time to solve, only four instances were generated. These graphs are generated in exactly the same way as the G_{500} graphs. Every non source node is a possible sink, the graphs are connected, the cost of the edges are uniformly random from the interval $[0; 100]$, and f has a random value in the interval $[1; 10]$. In Table 8 we show the number of nodes and edges, the time needed to solve the LP, the time needed to solve the ILP, the total solution time, and the number of nodes in the branch-and-bound tree, together with the node in which the best solution was found. For G_4 we ran into

Table 8
Results for large instances.

Graph	V	E	t_{LP} (min)	t_{ILP} (min)	t	n	n_{sol}
G_1	13	50	3.6	3.6	7.2 min	3	2
G_2	25	50	19.3	2083	35 h	1597	1332
G_3	15	60	27.2	167	3.2 h	207	142
G_4	30	60	167	–	–	–	–

memory problems while solving the ILP. These results indicate that instances which have 25 nodes and 50 edges or more can take a long time to solve and cannot always be solved with 8 GB of RAM available.

6. Concluding remarks

We have presented a new approach, based on column generation, for solving recoverable robust optimization problems. The key to this approach are two decomposition methods: Separate Recovery Decomposition (SRD) and Combined Recovery Decomposition (CRD). In the SRD approach, we work with separate solutions for the initial problem and recovery solutions for the different scenarios; in the CRD approach, we work with combined solutions for the initial problem and the recovery problem for a single scenario. We have shown that the LP-relaxation of the CRD model provides a stronger bound than the LP-relaxation of the SRD model.

We have tested our approach on two problems. The first one is the size robust knapsack problem, where the knapsack size can decrease, and where we can recover by removing items. Our computational experiments revealed that for this problem Separate Recovery Decomposition outperformed Combined Recovery Decomposition and that the SRD method scaled very well with the number of scenarios. Moreover, SRD outperforms the standard branch-and-bound and dynamic programming approach.

We further considered the demand robust shortest path problem. Here we need to buy the edges of a shortest path from a known source to an unknown sink that will be revealed later. We buy an initial set of edges, and when the sink gets revealed we can recover by buying additional edges (at a higher cost). For this problem the CRD approach seems to be very appropriate; the SRD model is hard to implement. The computational experiments revealed that it is crucial to add columns that correspond to the same initial solution for all scenarios.

There are several directions for future research. An obvious question is to find out what type of problems can be solved by this kind of approach; we are currently investigating our framework on scheduling problems. Next, it is important to formulate conditions to determine beforehand which decomposition model will work best for a certain problem. Recently, we have extended the size robust knapsack problem to multiple knapsacks (Tönissen, van den Akker, & Hoogeveen (2015)). For the single knapsack problem the SRD model outperforms CRD, but this changes if the number of knapsacks increases, and for four and more knapsacks CRD outperforms SRD. Another important question concerns the reduction of the running time; for example, in our experiments we have seen that adding the right columns or improving the primal heuristic can lead to a large improvement. Interesting issues for further research are restrictions on the recovery solution such as a limited budget for the cost of the recovery solution.

Appendix A. Time results per amount of edges/scenarios on a logarithmic scale

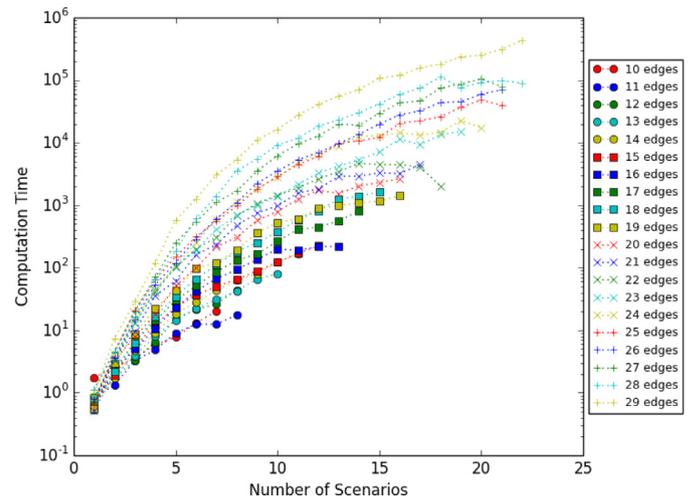


Fig. A.6. Time results per amount of edges on a logarithmic scale.

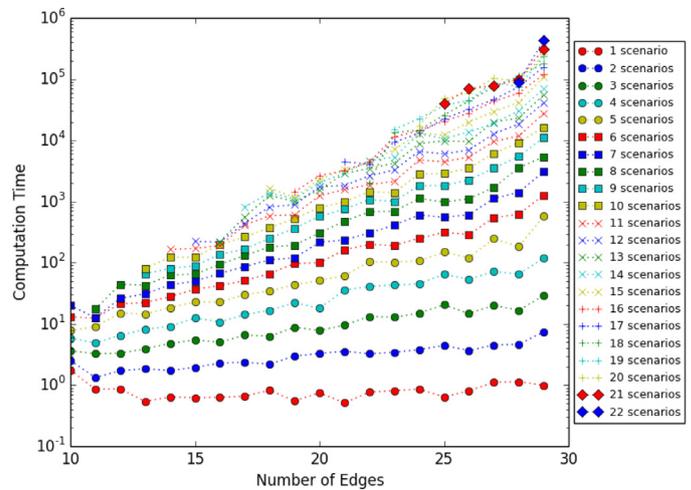


Fig. A.7. Time results per amount of scenarios on a logarithmic scale.

References

Barnhart, C., Johnson, E. L., Nemhauser, G. L., Savelsbergh, M. W. P., & Vance, P. H. (1998). Branch-and-price: column generation for solving huge integer programs. *Operations Research*, 46, 316–329.

Ben-Tal, A., Ghaoui, L. E., & Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.

Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations Research*, 52, 35–53.

Birge, J. R., & Louveaux, F. (1997). *Introduction to stochastic programming*. Springer Series in Operations Research and Financial Engineering. Springer.

Bouman, P. C. (2011). *A column generation framework for recoverable robustness*. Department of Information and Computing Sciences, Utrecht University Master's thesis.

Bouman, P. C., van den Akker, J. M., & Hoogeveen, J. A. (2011). Recoverable robustness by column generation. In C. Demetrescu, & M. M. Halldorsson (Eds.), *Algorithms ESA 2011, volume 6942 of Lecture Notes in Computer Science* (pp. 215–226). Berlin, Heidelberg: Springer.

Büsing, C. (2012). Recoverable robust shortest path problems. *Networks*, 59, 181–189.

Büsing, C., Koster, A. M. C. A., & Kutschka, M. (2011). Recoverable robust knapsacks: the discrete scenario case. *Optimization Letters*, 5, 1–14.

Büsing, C., Koster, A. M. C. A., & Kutschka, M. (2011). Recoverable robust knapsacks: γ -scenarios. In J. Pahl, T. Reiners, & S. Voß (Eds.), *Network Optimization, volume 6701 of Lecture Notes in Computer Science* (pp. 583–588). Berlin, Heidelberg: Springer.

- Cacchiani, V., Caprara, A., Galli, L., Kroon, L. G., & Maróti, G. (2008). Recoverable robustness for railway rolling stock planning. In *Proceedings of ATMOS 2008 - 8th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems, Karlsruhe, Germany*.
- Caprara, A., Galli, L., Kroon, L. G., Maróti, G., & Toth, P. (2010). Robust train routing and online re-scheduling. In *ATMOS 2010 - 10th Workshop on Algorithmic Approaches for Transportation Modeling, Optimization, and Systems, Liverpool, United Kingdom* (pp. 24–33).
- Cicerone, S., D'Angelo, G., Stefano, G. D., Frigioni, D., Navarra, A., Schachtebeck, M., & Schöbel, A. (2009). Recoverable robustness in shunting and timetabling. *Robust and Online Large-Scale Optimization: Models and Techniques for Transportation Systems*, 28–60.
- Dhamdhere, K., Goyal, V., Ravi, R., & Singh, M. (2005). How to pay, come what may: Approximation algorithms for demand-robust covering problems. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science, FOCS '05* (pp. 367–378). Washington, DC, USA: IEEE Computer Society.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik*, 1, 269–271.
- Disser, Y., Klimm, M., Megow, N., & Stiller, S. (2014). Packing a knapsack of unknown capacity. In E. W. Mayr, & N. Portier (Eds.), *Proceedings of the 31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014)*, volume 25 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Dagstuhl, Germany (pp. 276–287).
- Goerigk, M., Sabharwal, Y., Schöbel, A., & Sen, S. (2014). Approximation algorithms for the weight-reducible knapsack problem. In T. V. Gopal, M. Agrawal, A. Li, & S. B. Cooper (Eds.), *Proceedings of Theory and Applications of Models of Computation*, volume 8402 of *Lecture Notes in Computer Science* (pp. 203–215). Springer International Publishing.
- Liebchen, C., Lübbecke, M. E., Möhring, R. H., & Stiller, S. (2009). The concept of recoverable robustness, linear programming recovery, and railway applications. *Robust and Online Large-Scale Optimization: Models and Techniques for Transportation Systems*, 1–27.
- Pisinger, D. (2005). Where are the hard knapsack problems? *Computers & Operations Research*, 32, 2271–2284.
- Tönissen, D. D., van den Akker, J. M., & Hoogeveen, J. A. (2015). Column generation strategies and decomposition approaches to the size robust multiple knapsack problem. *Technical Report UU-CS-2015-010*. Department of Information and Computing Sciences, Utrecht University.