Constraint programming approach
to a bilevel scheduling problem

András Kovács · Tamás Kis

Abstract Bilevel optimization problems involve two decision makers who make their choices sequentially, either one according to its own objective function. Many problems arising in economy and management science can be modeled as bilevel optimization problems. Several special cases of bilevel problem have been studied in the literature, e.g., linear bilevel problems. However, up to now, very little is known about solution techniques of discrete bilevel problems. In this paper we show that constraint programming can be used to model and solve such problems. We demonstrate our first results on a simple bilevel scheduling problem.

Keywords Scheduling · Bilevel programming · Constraint modeling · QCSP

1 Introduction

Bilevel programming deals with decision and optimization problems whose outcome is determined by the interplay of two self-interested decision makers who decide sequentially. First, the decision maker called the *leader* makes its choice. Then, in view of the leader's decision, the *follower* chooses its response. Either decision maker aims at minimizing (maximizing) its own objective function. In the general case, the objective values mutually depend on the choices of the other party. Technically, the follower's role can be seen as solving a parametric optimization problem, whose parameters are determined by the leader. The particularly interesting situation is that of the leader, who is assumed to have a complete knowledge of the follower's constraints, objective, and input data. He endeavors to find his best choice subject to

A. Kovács (✉) · T. Kis
Computer and Automation Research Institute,
Hungarian Academy of Sciences, Budapest, Hungary
e-mail: akovacs@sztaki.hu

T. Kis
e-mail: tamas.kis@sztaki.hu

🙋 Springer

the response that he can expect from the self-interested follower. In the optimistic (pessimistic) case the leader assumes that the follower chooses from the set of its optimal responses the one that is the most (least) favorable for the leader.

Formally, the set of all variables in the problem is partitioned into two sets: the leader's variables X, and the follower's variables Y. The leader can assign values to X, while the follower decides about Y, and it is assumed that all variables have finite domains. The leader aims at minimizing f subject to the constraint set C and the follower's optimality condition, which states that the follower will minimize g subject to D. Also, the leader must avoid the values of X for which the follower's response does not satisfy C. Throughout the paper we assume that both the leader and the follower try to *minimize* their objectives, though, the same techniques can be used for maximization or mixed problems as well. Hence, the optimistic bilevel problem can be formulated as:

$$\min_{X,Y} \ f(X, Y) \tag{1}$$

subject to

$$C(X, Y) \tag{2}$$

$$Y \in \arg\min_{Y'}(g(X, Y') \mid D(X, Y')) \tag{3}$$

In formula (3), the operator *arg min* refers to the set of all optimal solutions of the problem at hand. Moreover, the pessimistic case of the problem is described as:

$$\min_X \max_Y \ f(X, Y) \tag{4}$$

subject to

$$C(X, Y) \tag{5}$$

$$Y \in \arg\min_{Y'}(g(X, Y') \mid D(X, Y')). \tag{6}$$

Bilevel programming techniques can be applied to model various decision problems of actors in customer-producer relations, in competition, or at various levels of an organizational hierarchy. Despite this, well-founded theoretical results are known for special cases of bilevel problems only. These include various exact and heuristic approaches to linear bilevel problems (where all constraints and both objective functions are linear expressions over continuous variables), and mostly heuristic methods for other cases, such as bilinear problems [11]. The papers [10, 12] address problems where the follower's variables can take discrete values. For (fully) discrete bilevel problems, which are in the focus of this study, only sporadic application results are available, see [21, 25]. Also, to the best of our knowledge, this paper is the first to investigate the solution of bilevel optimization problems using constraint programming (CP) techniques.

1.1 A motivating example

The classical approach in management science assumes that the different departments of the same company, although have individual decision roles and responsibilities, subsume their interest to the same global objective. This objective is related to maximizing the long-term profit of the company. The reality is often different: the

performance of each department is evaluated using, and rewarded based on, a different performance measure. These performance measures are only distantly related to the global objective of the company, and are often conflicting. Hence, a relevant alternative model of the joint operation of several departments is using multilevel programming techniques [2]. A simple case study is presented below.

Consider the bilevel scheduling problem where the management of the company (the leader) is responsible for *order acceptance* and the workshop foreman (the follower) decides on the *execution sequence* of the tasks corresponding to accepted orders. The leader has no direct influence on the sequencing decisions. Formally, there is a set of tasks T, some of which will have to be scheduled on a single unary resource. Task j is characterized by its processing time p_j, release time r_j, and deadline d_j. The difference between the profit if j is executed on time and the loss of reputation if it is rejected is captured by the cost (or task weight) w_j^1 to be paid if the task is rejected. A solution is acceptable for the leader only if all the accepted tasks are completed on time. The leader must select the tasks that will be actually executed: the binary variable x_j is 1 if task j is accepted and 0 if rejected. The objective of the leader is to minimize $\sum_j w_j^1 (1 - x_j)$ subject to the temporal constraints.

The sequencing decisions are made by the follower, who aims at minimizing the total weighted completion time of the tasks selected by the leader, i.e., $\{j \mid x_j = 1\}$. The start and completion times of tasks j are denoted by S_j and C_j, respectively, and the relation $S_j + p_j = C_j$ holds. The task weights w_j^2 that express the importance of tasks for the follower are independent from the leader's task weights w_j^1. We assume that the follower observes the release times, but the organizational relations within the company are such that the leader cannot force the follower to obey the deadlines. Hence, it might happen that a set of tasks could be scheduled on time, but the follower prefers to execute them in a sequence that violates some deadlines. Such task sets do not lead to feasible solutions of the bilevel problem.

Using the classical three-field scheduling notation [17], the follower's problem corresponds to a parametric version of $1|r_j| \sum_j w_j^2 C_j$. The first field of the notation specifies the machine environment; in our case number 1 stands for a single machine problem. The second field defines the constraints on activities; they are subject to individual release dates (r_j) in this problem. Finally, the third field states that the optimization criterion is the total weighted completion time of the tasks ($\sum_j w_j^2 C_j$). In our bilevel problem, this widely studied problem is parameterized with variables x_j, which decide the set of tasks to be considered by the follower.

This sample problem is a special type of bilevel problems where the leader's objective depends only on the leader's variables. However, the feasibility of a solution depends on the follower's response as well. For further examples from the scheduling domain, see Section 2.

1.2 Structure of this paper

The remainder of this paper is organized as follows. First, we review the related literature. After making the necessary definitions and presenting some basic theoretical results (Section 3), we introduce a generic CP approach to discrete bilevel optimization problems (Section 4). In Section 5 we illustrate the use of those techniques on the sample scheduling problem. Finally, we present experimental results in Section 6, and then conclude the paper.

2 Related literature

A number of different approaches in optimization deal with situations where the decision maker has only limited control of the problem at hand. Stochastic programming [30] considers random events occurring with known probability, and aims at optimizing the expected performance. Quantified problem solving looks at finding strategies for all possible actions of an adversary. In contrast, bilevel programming assumes a self-interested adversary with completely known objectives, and wishes to find a solution with the assumption that the adversary acts rationally.

2.1 Applications of bilevel programming

Probably the earliest example and a motivation of bilevel optimization problems came from economic game theory. In a two-player *Stackelberg game* two competing firms, the market leader and a follower company, for example a new entrant, produce equivalent goods. The firms decide their production quantities sequentially, which together determine the market price, with the aim of maximizing their own profit [11].

The application of bilevel programming to the coordination of multi-divisional organizations has been proposed in [2]. The approach is illustrated on a case study of three divisions of a paper company. The divisions are responsible for different stages of processing the paper, hence, the end product of one division serves as raw material for another division. Each division can decide to buy or sell on the outside market or from/to another division. The objective of the corporate unit is to set the internal transfer prices in such a way that the optimal decisions on the divisional level coincide with the corporate optimum. This problem can be encoded into a linear bilevel problem, and solved by known algorithms from the literature.

There exist a few application areas of discrete bilevel problems, and especially bilevel scheduling. In [25], the production planning problem of a pharmaceutical company is considered, while [21] studies a bilevel problem that may arise in flow shop scheduling. These papers take a relatively straightforward solution approach: they enumerate (a part of) the leader's possible choices, and for each choice, compute the follower's response. Brown et al. [6] investigate a bilevel project scheduling problem where the objective of the decision maker is to cause maximal delay of its adversary's project, which is given as a PERT network. The interdictor can buy delays, while the project owner can buy speed-ups on some arcs of the network from their limited budget. In [22], we present basic complexity and algorithmic results for bilevel scheduling problems.

Various other bilevel optimization problems arise naturally in economy and management science. Perhaps the most widely discussed example is the *toll setting* problem in a network, e.g., in a system of regional highways [23]. The owner of the network (the leader) seeks for the optimal pricing of each link in the network so as to maximize its profit. The follower corresponds to the ensemble of the users of the network. A fixed amount of users belong to each origin-destination pair, and each user selects the path that minimizes his costs, composed of the travel time and the tolls to pay. Many variations of this basic problem have been investigated, including problems where tolls or traffic signs are set by the local authorities who wish to control the movement of hazardous materials or consider other environmental

effects [26]. Another typical application is the optimization of chemical processes. Here, the follower's optimality condition describes that the steady-state result of a chemical reaction is an equilibrium where the reacting substances reach their energy minimum [8].

2.2 Related problems in CP

In constraint programming, a problem class strongly related to bilevel programming is the class of *quantified constraint satisfaction problems* (QCSPs), and their optimization versions, *quantified constraint optimization problems* (QCOPs) [3, 16]. While a classical constraint program corresponds to evaluating a formula that contains existentially quantified variables only (e.g., $\exists x \exists y \; C(x, y)$), in QCSP it is allowed to have universally quantified variables as well (e.g., $\exists x \forall y \; C(x, y)$). In papers [3] and [5], the basic QCSP and QCOP language has been extended with restricted quantification, resulting in the QCSP+ and QCOP+ languages. A sample QCSP+ formula is $\exists x \forall y [L(x, y)] \; C(x, y)$, which contains the restricted quantifier $\forall y [L(x, y)]$. This reads *"for all y such that $L(x, y)$ it holds that..."*. It is easy to show that a QCSP+ formula can be translated into a QCSP formula with negation and disjunction.

A number of QCSP solvers have been proposed in the literature, including the open-source QCOP+ solver called QeCode by Benedetti et al. [31], built on the top of Gecode [14]; QCSP-Solve, a solver partly motivated by ideas from QBF-solving by Gent et al. [15]; and the bottom-up solver BlockSolve by Verger and Bessière [36].

We are aware of two applications of QCSP to scheduling problems. Benedetti et al. [4] present a QCSP+ model of a scheduling game in which an adversary can change some task parameters—e.g., the resource requirement of some tasks subject to a limit on the overall increase of requirements. The objective is to find a robust schedule that remains feasible whatever actions the adversary takes. Nightingale [28] presents a QCSP model for job-shop scheduling with the risk of machine breakdowns. We will investigate the relation of bilevel programming to QCSP in detail in the next section.

Another related problem is the class of *adversarial constraint satisfaction problems* (ACSP) [7]. ACSP can be used to model games played by n agents with potentially conflicting interests that consist of a fixed number of rounds. The number of rounds equals the number of variables in the ACSP, which is typically much larger that the number of agents. The main difference between ACSP and bilevel programming is that in ACSP in each round the forthcoming agent is free to choose an arbitrary variable to instantiate, i.e., variables are not assigned to agents a priori. Also, in ACSP, all agents must satisfy the same set of constraints, although in theory it is possible to incorporate a measure of constraint violations into the optimization criteria of each agent.

2.3 Bilevel problems in game theory

The presence of two self-interested decision makers make bilevel programming interesting from the game theoretical point of view as well. The optimal solution of an optimistic (pessimistic) bilevel program corresponds to a weak (strong) Stackelberg equilibrium [34]. As opposed to the classical Nash equilibrium for continuous games, Stackelberg equilibrium refers to games with turns. Also, Stackelberg equilibrium

is different from subgame perfect equilibrium (SPE) [37], since SPE requires the strategy to cover all possible moves of the opponent, not only its optimal ones. The original concept of Stackelberg has been extended to an oligopolistic market with one leader and N followers by Sherali et al. [33]. In that model the followers reach an equilibrium solution in the market of a single homogeneous product and the leader, supplying the same product without any collusion with the other firms, sets the production levels in an optimal (profit maximizing) fashion by explicitly considering the reaction of the other firms to its output variations. This model leads to a mathematical program with equilibrium constraints, and the latter area has a rich literature.

2.4 Solution techniques applied

In our work we rely on known techniques of constraint programming and operations research, and adapt these to bilevel problems. For a detailed presentation of the applied constraint propagation algorithms and search techniques the reader is referred to [1, 32].

We note that the applied decomposition to a master problem and a subproblem resembles the (logic-based) Benders decomposition approach to single-level problems [18]. However, a substantial difference is that in the single level Benders case one is free to choose the separation of the master and the subproblem as it is the most efficient computationally, whereas in the bilevel case the separation comes from the problem definition. This implies that for bilevel problems it can be rather challenging to feedback strong cuts (or constraints) from the subproblem to the master problem.

3 Basic properties of discrete bilevel problems

In this section we analyze basic properties of discrete bilevel problems. First, we give a closer look at the potential definitions of the follower's optimality condition. Then, we demonstrate that bilevel problems differ substantially from single level problems with a single or multiple objectives, whereas they are more related to quantified constraint satisfaction problems. Finally, we investigate how the bilevel problem can be relaxed to a single level problem, and address the computational complexity of bilevel problems.

3.1 On the optimistic and pessimistic cases

The optimistic and pessimistic formulations of the bilevel problem, shown in formulae (1–3) and (4–6), respectively, capture two different standpoints of the leader. This difference is relevant in problems where the follower can have several optimal solutions, and these differ essentially from the viewpoint of the leader: only some of them satisfy C or they incur different costs f. The *optimistic* formulation assumes that the leader is allowed to choose one from the follower's optimal solutions, or, equivalently, the follower is friendly enough to choose an optimal response that satisfies C and minimizes f, if there exists one.

In contrast, in the *pessimistic* case, the leader wishes to safeguard against the risks of an unfavorable follower response by assuming that the follower selects its optimal response that is the least favorable for the leader. There are two possible interpretations of the pessimistic formulation (4–6) used in the literature. Both approaches consider a response from the follower's optimal set that maximizes f. However, the first interpretation allows the follower to choose a response that violates C [35], whereas the second interpretation assumes that the follower must select a response satisfying C, if there exists one [11]. In the first case, which we call the *hard pessimistic* formulation, the leader must select values for X in such a way that *all* optimal solutions of the follower satisfies C. This is not necessary in the second, so-called *soft pessimistic* formulation. It must be noted that in most of the existing applications of bilevel programming the constraint set C is empty, and therefore the two pessimistic cases are equivalent.

In the core of this paper, we focus on the optimistic case. At the same time, we note that the similar techniques can be used for the pessimistic case. The necessary, minor changes in the algorithmic details will be discussed in Section 4.5.
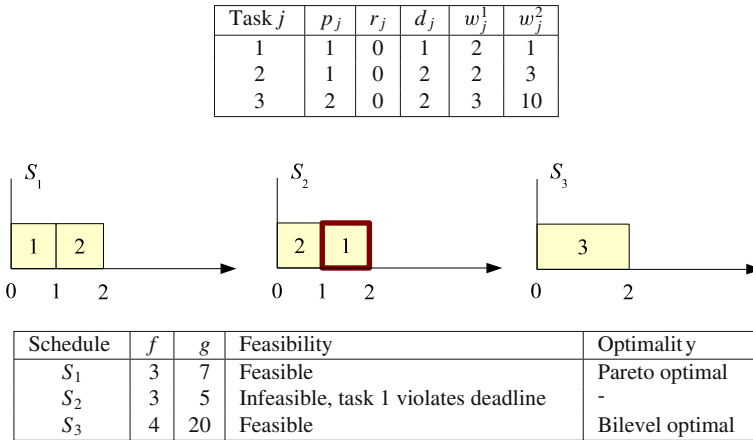
3.2 Bilevel versus single level problems

Below we demonstrate the difference of the single level and the bilevel problems on an instance of our sample problem presented in Fig. 1. In the single level case, the leader could accept all the four tasks and process them, e.g., in the order (1, 2, 3, 4). In the bilevel case, the leader only chooses the tasks to process, but the follower sequences them. If the leader selects all tasks, then the follower's response is the solution of the corresponding $1|r_j|\sum_j w_j^2 C_j$ problem, i.e., the sequence (4, 3, 2, 1). This solution is infeasible, because task 1 violates its deadline. In fact, the optimal bilevel solution is selecting the tasks $\{1, 2, 3\}$, and processing them in the order (1, 3, 2), which respects all deadlines. An interesting, seemingly paradoxical situation is that the strictly smaller set of tasks $\{1, 2\}$ cannot be scheduled, because the follower's response, (2, 1), violates the deadline of task 1. This also warns us that inference methods that work for the single level case might not generalize to the bilevel problem.

Task j	p_j	r_j	d_j	w_j^1	w_j^2
1	1	0	1	2	1
2	2	0	100	2	4
3	1	1	100	2	20
4	1	0	100	1	5

Fig. 1 A bilevel problem instance and the follower's response for various choices of the leader. The tasks marked with a thick frame in the schedules violate their deadlines

Task j	p_j	r_j	d_j	w_j^1	w_j^2
1	1	0	1	2	1
2	1	0	2	2	3
3	2	0	2	3	10

S_1

| 1 | 2 |
0   1   2

S_2

| 2 | 1 |
0   1   2

S_3

| 3 |
0   2

Schedule	f	g	Feasibility	Optimality
S_1	3	7	Feasible	Pareto optimal
S_2	3	5	Infeasible, task 1 violates deadline	-
S_3	4	20	Feasible	Bilevel optimal

Fig. 2 Difference of the bilevel and the bicriteria Pareto optimal solutions. The figure presents a problem instance and its three different solutions

3.3 Bilevel versus bicriteria approaches

Although both bilevel programming and single level bicriteria approaches seek for solutions that are attractive w.r.t. two different objective functions, the two approaches differ essentially. They model two different situations: bicriteria optimization looks for the best compromise in a centralized way, while bilevel optimization follows a simple, hierarchical protocol with two autonomous partners, each interested in optimizing its own objective value. Indeed, the optimal solution of the bilevel problem might not be Pareto optimal for the corresponding single level bicriteria problem, and vice versa. Below we illustrate this phenomenon on our sample problem.

Consider the problem instance presented in Fig. 2. The candidate task sets to be scheduled are {1}, {2}, {3}, and {1, 2}, and it is easy to see that no other task set can be scheduled to meet the deadlines. A Pareto optimal schedule is (1, 2), denoted by S_1, which leads to objective values $f = 3$ and $g = 7$. Observe that schedule S_1 is not a feasible solution of the bilevel problem: if the leader decided to accept tasks {1, 2}, then the follower would sequence these according to (2, 1), resulting in schedule S_2. However, S_2 violates the leader's deadline constraint on task 1, and hence, it is not a feasible solution of the bilevel problem. In fact, the optimal bilevel solution is the schedule (3), called S_3, which has $f = 4$ and $g = 20$. The leader prefers S_3 to {1} and {2} as well. Note that S_1 Pareto dominates S_3, which means that the bilevel optimal solution is Pareto dominated.

3.4 Bilevel programming versus QCSP

As it has been described above, the main difference between QCSP and bilevel programming is that in QCSP, one wishes to find a strategy that covers all possible actions of the adversary, whereas in bilevel programming we assume that the follower will act rationally according to its known objectives. Now we show that bilevel programs can be translated into a QCOP+ with a single pair of quantifiers $\exists \diamond \forall \diamond$

and vice versa. We assume that the function symbols f and g and the relation $\leq$ is available in the constraint language.

First, note that the optimistic bilevel problem corresponds to the QCOP+

$$\min_{X,Y}\{f(X, Y) \mid C(X, Y) \wedge D(X, Y)$$

$$\wedge\, \forall Y'\big[D(X, Y')\big]\, g(X, Y) \leq g(X, Y')\}.$$

Here, the first line of the formula describes that $\langle X, Y \rangle$ is a feasible solution, while the second line states that $\langle X, Y \rangle$ is an optimal response of the follower, because all the alternative responses Y' would result in a greater or equal value of g. Furthermore, the hard pessimistic bilevel problem can be rewritten as

$$\min_{X,Y}\{f(X, Y) \mid C(X, Y) \wedge D(X, Y)$$

$$\wedge\, \forall Y'\big[D(X, Y')\big]\, g(X, Y) \leq g(X, Y')$$

$$\wedge\, \forall Y''\big[D(X, Y'') \wedge g(X, Y) = g(X, Y'')\big]$$

$$C(X, Y'') \wedge f(X, Y) \geq f(X, Y'')\}.$$

Similarly to the optimistic case, the first and second lines describe that the solution $\langle X, Y \rangle$ is feasible and optimal for the follower. The third and fourth lines encode the hard pessimistic assumption, i.e., that all the optimal responses of the follower Y'' must satisfy C and result in a value of f not worse than $f(X, Y)$. Note that the QCOP+ equivalent of a soft pessimistic bilevel program can be derived from the above formula by omitting $C(X, Y'')$ from the last line.

Although translation to QCOP+ is a theoretically sound approach to solving bilevel problems, it can be rather inefficient. The main deficiency of the approach is that the computed strategy must cover all possible decisions of the follower explicitly. To verify these claims, we have implemented our sample problem in QeCode. Experimental results are presented in Section 6.

3.5 The single level relaxation

Various components of the solution algorithms for bilevel problems rely on well understood techniques for single level problems. Therefore, it seems natural to look for relations between bilevel and single level problems. The simplest way of reduction is to let the leader decide on every variable, and completely disregard the existence of the follower. The resulting problem will be called the *single level relaxation* of the bilevel problem, and its solution value is obviously a lower bound on the bilevel solution cost:

Definition 1 The single level relaxation of a bilevel program is, using the set of all variables $X' = (X, Y)$, the problem $\min\{f(X') \mid C(X') \wedge D(X')\}$.

3.6 Computational complexity

Bilevel problems are complex optimization problems, they often belong to a higher complexity class than their corresponding single level relaxations. For example, linear bilevel problems (where both the single level relaxation and the follower's

subproblem is a linear program) are known to be NP-complete [11]. Here, we focus on the complexity of decision versions of discrete bilevel problems, especially in the case where the (decision version of the) single level relaxation is NP-hard. It is easy to observe that a bilevel problem is—except for degenerate cases—at least as complex as its single level relaxation, hence, NP-hard. On the other hand, discrete bilevel problems are in PSPACE, because all instantiations of the variables can be enumerated and evaluated in polynomial space using a recursive algorithm, similarly to the algorithm defined for solving quantified boolean formulae in [13].

Now, a discrete bilevel program may or may not belong to NP. It is easy to define discrete bilevel problems with an NP-compete follower's subproblem that are outside NP. Consider an unconventional, but valid bilevel scheduling problem where all variables belong to the follower. Both the leader and the follower aim at sequencing the set of tasks on a single machine subject to release times r_j and strict deadlines d_j. However, the leader would be interested in minimizing $\sum C_j$, whereas the follower minimizes the weighted earliness penalty $\sum w_j(d_j - C_j)$. This problem is NP-hard, because with follower weights $w_j \equiv 0$, all solutions are equivalent for the follower, and hence, the optimistic bilevel problem corresponds to the classical $1|r_j, d_j| \sum C_j$ problem, which is NP-hard [24]. On the other hand, it is also co-NP-hard, because with follower weights $w_j \equiv 1$ the two criteria are the negatives of each other. Hence, verifying the feasibility of a bilevel solution is equivalent to proving that no better solution exists for the follower's NP-hard subproblem, $1|r_j, d_j| \sum w_j(d_j - C_j)$. Since the bilevel problem is both NP-hard and co-NP-hard, it is outside NP (unless P=NP).

The above complexity results indicate that no direct encoding of discrete bilevel problems into CP or MIP can be expected. For the case of our main sample problem, we were not able to prove that it is outside NP, but we conjecture that it is, since no trivial certificate seems to exist for a positive answer.

4 Modeling and solving bilevel problems by CP

In this section we first present a generic approach to solving discrete bilevel optimization problems by CP. Then, we introduce several algorithmic techniques to improve the efficiency of the solver. Each of the subsections has a counterpart in the Section 5, where the use of the given technique is illustrated on the sample scheduling problem. Unless stated otherwise, we consider the optimistic bilevel problem.

We use the notation $\mathrm{Dom}(X)$ for the current domain of a CP variable Z. Furthermore, the minimum and maximum values in $\mathrm{Dom}(X)$ will be denoted by $\check{Z}$ and $\hat{Z}$, respectively.

4.1 The basic constraint model

Given the discrete optimistic bilevel problem as described in formulae (1–3), let us define an equivalent constraint program that encodes the problem from the leader's point of view. We also call it the *master problem*. The decision variables are both X and Y. They are subject to constraints C and C_{Opt}, where C is the set of the leader's constraints, and C_{Opt} describes the follower's optimality condition:

$$\min_{X,Y}\{ f(X, Y) \mid C \wedge C_{\mathrm{Opt}}\},$$

where

$$C_{\text{Opt}}: \quad Y \in \arg\min_{Y'}\{g(X, Y') \mid D(X, Y')\}.$$

The optimization problem contained in C_{Opt} is called the *follower's subproblem*. We assume that C is a set of classical constraints over finite-domain variables, which have appropriate propagation algorithms defined in the literature. In contrast, in the generic case, constraint C_{Opt} contains the parametric version of an NP-hard discrete optimization problem. Hence, there is little hope that algorithms readily available in the literature can be applied to propagate it, or generalized arc-consistency can be achieved efficiently. Therefore, we propose to settle for a generate-and-test approach for propagating C_{Opt}, i.e., to propagate only when all of the leader's variables X become bound. The pseudo-code of the propagation algorithm is presented in Fig. 3. The algorithm first determines the follower's minimum cost, g^* (line 3), or returns the symbol `'no_solution'` if no feasible solution exists. Then, it computes the follower's response according to the optimistic assumptions, Y^+ (line 6). Both of these steps require solving the follower's subproblem with known parameters X. Exact solution approach, e.g., CP search must be used, since the bilevel problem formulation requires finding exact optimum. When solving the follower's subproblem in lines 3 and 6, the domains of Y in the master problem must be ignored, since those domains are corrupted by the propagators of C, i.e., constraints that the follower disregards.

Regarding *search techniques* for the master problem, we propose to perform any kind of search, exact or non-exact, in the space of the instantiations of the leader's variables. It is not necessary to consider the follower's variables, since values will be assigned to them by constraint C_{Opt}.

4.2 Lifting the follower's constraints and dominance rules into the master problem

The above basic CP formulation can be strengthened by adding redundant constraints that propagate even when a part of the variables X is not bound. First, observe that the constraint set D can be added to the basic model, because it reduces the search space to values of X that have at least one feasible follower response. Furthermore, assume that there are weak dominance rules known for the follower's sub-problem, i.e., properties that *all* optimal solutions of the sub-problem must satisfy. Then, the conjunct of these rules can be encoded into a constraint C_{Dom}, and

```
PROCEDURE Propagate_C_Opt()
1    IF X is not bound
2       RETURN
3    LET g* := minY'{g(X, Y') | D(X, Y')}
4    IF g* = 'no_solution'
5       Fail
6    LET Y+ := arg minY'{f(X, Y') | g(X, Y') = g* ∧ C(X, Y') ∧ D(X, Y')}
7    IF Y+ = 'no_solution'
8       Fail
9    Instantiate Y ← Y+
```

Fig. 3 Algorithm for propagating constraint C_{Opt}

added to the CP model as a redundant constraint. Hence, the following CP model is a sound representation of the bilevel problem, and it leads to stronger propagation than the basic model:

$$\min_{X,Y}\{ f(X, Y) \mid C \wedge D \wedge C_{\text{Opt}} \wedge C_{\text{Dom}}\}.$$

4.3 Bounds on the follower's cost

Below we present a novel technique that prunes the search tree based on the difference of the constraint sets that the leader and the follower must satisfy. We will characterize the values that g can take in solutions that are feasible for the leader, as well as the values that g can take in optimal responses of the follower. Clearly, if the two ranges do not overlap, then there is no feasible bilevel solution in the current branch of the search tree.

Let UB denote the value of the best known solution, and let us characterize the current branch of the search tree by the domain of X, denoted by $\text{Dom}(X)$. Any feasible improving solution of the master problem must obey C, D, and $f < UB$. Hence, a valid lower bound $g^{L\min}$ on the values g in the solutions that are acceptable for the leader in the current branch is:

$$g^{L\min} = \min_{X' \in \text{Dom}(X)} \min_{Y}\{g(X', Y) \mid C(X', Y) \wedge D(X', Y) \wedge f(X', Y) < UB\}.$$

On the other hand, for any fixed leader's choice in this branch, the follower will return a response that minimizes g subject to D. By taking the maximum of these minimum values, we get an upper bound $g^{F\max}$ on g in the solutions that are acceptable for the follower in the current branch:

$$g^{F\max} = \max_{X' \in \text{Dom}(X)} \min_{Y}\{g(X', Y) \mid D(X', Y)\}.$$

Note that the constraints in the definition of $g^{F\max}$ are a subset of the constraints for $g^{L\min}$, and therefore $g^{F\max} < g^{L\min}$ can occur. This means that no solution in the current branch of the search tree is both feasible for the leader and optimal for the follower. At the same time $g^{F\max} \geq g^{L\min}$ is also possible, since $g^{F\max}$ is a maximin, whereas $g^{L\min}$ is a minimum.

Lemma 1 *If $g^{F\max} < g^{L\min}$, then the current search branch contains no feasible improving bilevel solution, and therefore it can be fathomed.*

In general, it is difficult to compute the exact values of $g^{F\max}$ and $g^{L\min}$. Instead, an upper estimate of $g^{F\max}$, denote by $\hat{g}^{F\max}$ can be used, and similarly, a lower estimate of $g^{L\min}$, denoted by $\check{g}^{L\min}$ can be applied.

Finally, note that the application of the above bounds makes sense in the leaves of the search tree as well, since they may prove the infeasibility of the leaf faster than solving the follower's subproblem in an exact way.

4.4 Lower bounds on the leader's cost

In theory, it is straightforward to apply the classical lower bounding technique of operations research to bilevel problems: let $\check{f}$ be a lower bound and $\hat{f}$ an upper bound,

typically the value of the best known solution. Now, if $\check{f} \geq \hat{f}$ then the current branch of the search tree does not contain a feasible improving solution. In practice, the effective use of this technique is challenging, because good lower bounds for bilevel problems are rarely available from the literature. A possible approach is using the single level relaxation, whose solution imposes a lower bound on the bilevel problem. If the single level relaxation is still intractable, then it can be relaxed further. We note that the value of the single level relaxation is often far from the optimal solution of the bilevel problem.

4.5 Extension to the pessimistic case

The models and algorithms for the optimistic case can be extended easily to the pessimistic case. The extension requires modifying the propagator of C_{Opt}. In the soft pessimistic case, the function *min* in line 6 of Fig. 3 must be replaced with *max* as follows:

6   `LET` $Y^+ := \arg\max_{Y'}\{f(X, Y') \mid g(X, Y') = g^* \wedge C(X, Y') \wedge D(X, Y')\}$

In the hard pessimistic case, in addition to the above change, the following lines must be inserted after line 5 (outside the IF-THEN branch of lines 4 and 5) to ensure that the follower does not have an optimal solution that violates C:

5a `IF there exists an` Y' `such that` $\{g(X, Y')=g^* \wedge \bar{C}(X, Y') \wedge D(X, Y')\}$
5b   `Fail`

In line 5a, the expression $\bar{C}$ stands for the negation of C, i.e., $\bar{C} = \vee_{c \in C} \neg c$. Hence, $\bar{C}$ corresponds to a reified constraint. We note that the use of reified constraints in $\bar{C}$ makes the CP approach substantially less efficient for hard pessimistic problems, and may limit the applicable constraints depending on the solver used. The enhancements presented in Sections 4.2–4.4 can be applied without any change.

5 Modeling and solving the scheduling problem

5.1 The basic constraint model

The basic constraint model of our scheduling problem contains n binary variables x_j to denote if task j is scheduled, and n optional activities with start and end time variables. The activities are subject to a unary resource and time window constraints, and the follower's optimality constraint. The objective function is expressed as $f = \sum_j w_j^1(1 - x_j)$ using a weighted sum constraint. Our search strategy selects in each node the task j whose x_j variable is unbound and has the greatest w_j^1. Then, it creates two children of the node according to $x_j = 1$ (left branch) or $x_j = 0$ (right branch).

    The follower's optimality constraint is a custom developed constraint, which embeds a constraint-based solver for the $1|r_j| \sum w_j^2 C_j$ problem. The naive constraint model with a unary resource constraint, release time constraint, and the cost

expressed using a weighted sum constraint is used. A classical chronological schedule-or-postpone search strategy (called *setTimes* in Ilog) is used, and the subproblem solver also includes dominance rules from [19].

5.2 Lifting the follower's dominance rules into the master problem

A number of efficient dominance rules are known for the $1|r_j| \sum w_j^2 C_j$ problem, see, e.g., [19]. However, the condition side of most of these rules is too complex to fire when only the x_j variables are bound, and very little is known about the task start times or the order. We lifted the following simple dominance rule to the master problem. Without loss of generality we can assume that tasks are indexed in the *weighted shortest processing time* order (WSPT, non-increasing w_j^2/p_j), with ties broken by *earliest due date* (EDD, non-decreasing d_j).

Lemma 2 *For any fixed leader's choice, i.e., assignment of the variables x_j, there exists an optimal follower's response according to the optimistic bilevel assumption such the tasks that start after $r_{max} = \max_{\{j|x_j=1\}} r_j$ are ordered by the above defined task index.*

Proof The proof essentially matches the proofs of the optimality of the WSPT order for the $1|| \sum_j w_j C_j$ problem and the EDD order for the feasibility problem subject to deadlines, but with uniform release times. Let j and k be two subsequent tasks, both starting after r_{max} in any feasible bilevel solution, i.e., in a schedule that is both optimal for the follower and feasible for the leader. Now, $w_j^2/p_j \geq w_k^2/p_k$ holds, because otherwise the schedule would be suboptimal for the follower: swapping j and k would decrease $\sum_j w_j^2 C_j$. Furthermore, if $w_j^2/p_j = w_k^2/p_k$ and $d_j > d_k$, then swapping j and k preserves the optimality of the follower's response and does not cause infeasibility for the leader. Repeating the swaps until no further such task pairs exists results in an optimal feasible response according to the optimistic bilevel assumption. □

In theory, it would be possible to encode the above dominance rule directly into $\frac{1}{2}n(n-1)$ reified constraints, i.e., one constraint for each pair of tasks. However, to achieve more efficient propagation, we implemented a new algorithm for propagating the above dominance rule as a single global constraint. The algorithm, displayed in Fig. 4, tightens the bounds of the domains of start and end time variables. Recall that the minimum and maximum start (end) times are denoted by $\check{S}_j$ and $\hat{S}_j$ ($\check{C}_j$ and $\hat{C}_j$), respectively. Tightening the bounds of a start time variable S_j means updating its initial domain of $[\check{S}_j, \hat{S}_j]$ to $[\check{S}'_j, \hat{S}'_j]$, where $\check{S}'_j \geq \check{S}_j$ and $\hat{S}'_j \leq \hat{S}_j$ hold. The constraint requires an initialization step, which sorts the tasks according to the above defined order in $O(n \log n)$ time. Then, each propagation run takes $O(n)$ time. The algorithm first computes an upper bound of r_{max} (line 1). Then, it considers the tasks j in the above order, and maintains an earliest start time s of the task under consideration (lines 3–8). If j must be scheduled after r_{max}, then it must start after s (lines 4–6). Otherwise, if j cannot start after s, then it cannot be scheduled after r_{max} (lines 7 and 8). In the latter case, one of the following conditions hold: either j is not scheduled, in which case $\hat{S}'_j$ can be set to r_j; or j starts strictly before r_{max}. Overall, $\hat{S}'_j$ can be set to $\max(r_{max} - 1, r_j)$. Finally, the algorithm repeats the same type of inference for the latest finish times, considering the tasks in the reverse order (lines 10–15).

Fig. 4 Algorithm for propagating the dominance rule

```
PROCEDURE Propagate_Dominance()
1    LET rmax := max{j|1∈Dom(xj)} rj
2    s := 0
3    FORALL j ∈ {1, ..., n} IN INCREASING ORDER
4        IF xj is bound to 1 AND Šj ≥ rmax
5            Update Š'j := max(Šj, s)
6            s := Š'j + pj
7        ELSE IF s > Ŝj
8            Update Ŝ'j := max(rmax − 1, rj)
9    e := ∞
10   FORALL j ∈ {1, ..., n} IN DECREASING ORDER
11       IF xj is bound to 1 AND Šj ≥ rmax
12           Update Ĉ'j := min(Ĉj, e)
13           e := Ĉ'j − pj
14       ELSE IF e < Čj
15           Update Ŝ'j := max(rmax − 1, rj)
```

5.3 Bounds on the follower's cost

5.3.1 The upper bound

Our follower's subproblem is $1|r_j| \sum w_j^2 C_j$. Let $g(T_1)$ denote the follower's minimal cost when the leader accepts the task set T_1. It is obvious that if $T_1 \subseteq T_2$ then $g(T_1) \leq g(T_2)$. Therefore, the cost of any heuristic solution to the follower's problem with task set $T_{\max} = \{j \mid 1 \in \text{Dom}(x_j)\}$ can be used as $\hat{g}^{F \max}$. In our solver, we have implemented the constructive heuristic called *CPRTWT* with the *makeBetter* improvement step after the insertion of each task to the schedule, originally introduced in [20].

5.3.2 The lower bound

Computing $\check{g}^{L \min}$ requires obtaining a lower bound on a $\sum w_j^2 C_j$ problem subject to release times, deadlines, and optional activities. Our lower bound (LB) is based on the model of [29] for a similar problem, though, without optional activities:

$$\min_{z_j, C_j} \quad \sum_j w_j^2 C_j \tag{7}$$

subject to

$$\forall j \quad (r'_j + p_j)z_j \leq C_j \tag{8}$$

$$\forall j \quad C_j \leq d'_j z_j \tag{9}$$

$$\forall i, j \ (i \neq j) \quad (C_i \leq C_j − p_j z_j) \ \vee \ (C_j \leq C_i − p_i z_i) \tag{10}$$

$$\sum_j w_j^1 z_j \geq W \tag{11}$$

$$\forall j \quad z_j \in \text{Dom}(x_j) \tag{12}$$

In this formulation, variables z_j indicate if task j is processed ($z_j = 1$) or not ($z_j = 0$). Variables C_j denote the completion times. Constraints (8) and (9) specify the release times and deadlines of the task, and also ensure that C_j is 0 if j is not scheduled. Note that the time windows taken from the CP model can be used, since these are strengthened by propagation compared to the original values. Line (10) defines the unary resource constraint. Constraint (11) states that the total weight of the tasks selected for processing must be at least $W = \sum_j w_j^1 - UB + 1$ in order to achieve an improving solution for the leader. In a given search node, it can already be known for some tasks if they are already selected for processing by the leader or not, while it is still an open question for the rest (12). Note that $\text{Dom}(x_j) \subseteq \{0, 1\}$. Now, by dualizing (8) and (9) with multipliers a and b, we receive the following Lagrangian relaxation (LR):

$$\min_{z_j, C_j} \quad \sum_j w_j^2 C_j + \sum_j \left[a_j((r_j' + p_j)z_j - C_j) + b_j(C_j - d_j'z_j) \right]$$

$$= \sum_j (w_j^2 - a_j + b_j)C_j + \sum_j \left[a_j(r_j' + p_j) - b_j d_j' \right]z_j \tag{13}$$

subject to

$$\forall i, j \ (i \neq j) \quad (C_i \leq C_j - p_j z_j) \ \vee \ (C_j \leq C_i - p_i z_i) \tag{14}$$

$$\sum_j w_j^1 z_j \geq W \tag{15}$$

$$\forall j \quad z_j \in \text{Dom}(x_j) \tag{16}$$

Next, we present how the LR problem can be solved to optimality for fixed non-negative Lagrangian multipliers a and b. We exploit that the first component of the objective function corresponds to $\sum_{\{j \mid z_j=1\}} w_j' C_j$ with $w_j' = w_j^2 - a_j + b_j$, while the second component does not contain completion time variables C_j. Therefore, for any fixed z, the optimal solution is a no-delay schedule containing the selected tasks in WSPT order.

Computing the LB in leaves of the search tree We have developed two different methods for computing the optimal solution of LR: one to be used in the leaves of the search tree, and another for internal search nodes. In leaves we exploit that there are no optional tasks, i.e., the variables z_j are fixed. In this case, $\check{g}^{L\min}$ equals the objective value of the WSPT schedule, which can be computed in $O(n \log n)$ time.

Computing the LB in internal search nodes In internal search nodes, the LR problem with a fixed choice of multipliers a and b can be solved by the following dynamic program (DP). As an initialization step, we sort the tasks j that may be scheduled ($1 \in \text{Dom}(x_j)$) by non-increasing w_j'/p_j, which corresponds to the WSPT order according to the modified weights w_j'. Tasks that cannot be scheduled ($1 \notin \text{Dom}(x_j)$) are completely ignored by the algorithm.

The DP fills in a three dimensional table whose cells are indexed by parameters k, v, and t. The content of each cell characterizes the optimal solution of a subproblem of LR, received by applying the following restrictions to LR:

- The task selected for processing are a subset of $\{1, ..., k\}$;

- The total leader's weight of the selected task must be equal to v;
- The schedule must end at time t.

The optimal schedule in cell (k, v, t) will be denoted by $\sigma(k, v, t)$, and its cost by $u(k, v, t)$. It is sufficient to store the cost $u(k, v, t)$ in the table, while we use $\sigma(k, v, t)$ only to show the correctness of the algorithm below. For combinations of parameters k, v, and t that do not lead to a feasible schedule, we consider $u(k, v, t) = \infty$.

The DP fills in the table layer-by-layer, using induction over the different values of k. The first layer ($k = 1$) contains two finite values only, which characterize the two trivial solutions: $u(1, w_1^1, p_1) = w_1^2 p_1$ for a schedule that contains task 1 only, and $u(1, 0, 0) = 0$ for the empty schedule.

Each of the values $u(k, v, t)$ in the subsequent layers, $k \geq 2$, can be derived in one of the following two ways. If task k is contained in schedule $\sigma(k, v, t)$, then k is the last task in this schedule due to the optimality of the WSPT order. Consequently, $\sigma(k, v, t)$ is the concatenation of the optimal schedule for the subproblem without task k, $\sigma(k - 1, v - w_k^1, t - p_k)$, and task k itself. Alternatively, if task k is not contained in $\sigma(k, v, t)$, then $\sigma(k, v, t)$ is identical to the schedule computed in the previous layer, $\sigma(k - 1, v, t)$. From these two candidate schedules, the one that leads to the lowest cost is selected. Hence, the value of $u(k, v, t)$ can be computed as:

$$u(k, v, t) = \min(u(k - 1, v - w_k^1, t - p_k) + tw_k^2, \ u(k - 1, v, t))$$

Note that a cell corresponds to a feasible solution of LR if and only if it has $v \geq W$, i.e., it satisfies constraint (15). All other constraints of LR are respected by all cells by definition. Hence, the optimal solution of LR can be retrieved from the last layer of the table, denoted as layer $k_{\max}$, as follows:

$$\check{g}^{L\,\min} = \min_{v,t}\{u(k_{\max}, v, t)|v \geq W\}$$

The DP runs in pseudo-polynomial time and space: its complexity is $O(nVP)$, where $V = \sum_j w_j^1$ and $P = \sum_j p_j$.

Setting the Lagrangian multipliers The above methods result in an optimal solution of LR for any fixed non-negative multipliers a and b, assuming $w_j' = w_j^2 - a_j + b_j \geq 0$. To find the multipliers that provide the strongest LB, we embedded the above methods into a loop, and adjusted the multipliers after each cycle as follows. If task j violates its release time in the current optimal solution of LR, then its weight is decreased in order to move it later in the WSPT order. Namely, we set $w_j' = \frac{w_k' p_j}{p_k} - \varepsilon$, where k is the successor task of j in the schedule. Similarly, if j violates its deadline, then its weight is increased to $w_j' = \frac{w_k' p_j}{p_k} + \varepsilon$, where k is the predecessor of j. If task j respects both its release time and deadline then its weight is not changed. Note that tasks cannot violate their release time and deadline at the same time. The method was initialized with $a_j = b_j = 0$.

The best run times were achieved with the number of cycles set to 15 in leaves, and not using this method in internal search nodes (c.f. the experimental results for further details).

5.4 Lower bounds

The single level relaxation (SLR) of our problem is the $1|r_j| \sum w_j^1 U_j$ scheduling problem, where the optimization criterion $\sum w_j^1 U_j$ stands for the weighted number of late tasks. This problem is NP-complete. Nevertheless, various solution techniques and polynomial lower bounds are available from the literature. The current best algorithm for the SLR is the branch-and-bound of [27]. We have implemented the mixed-integer programming (MIP) formulation of the single level problem proposed in this paper, and solved its linear relaxation in each node of the search tree. The parameters r_j and d_j were updated in each node by the tighter time windows taken from the CP model.

The gap between the lower bound and the bilevel solution originates from two sources: solving the SLR instead of the bilevel problem, and the further linear relaxation of the SLR. In preliminary experiment we have found that over 75% of the gap is due to taking the SLR, and only 25% originates from solving the linear relaxation. Overall, this lower bound was not sufficiently tight to be used for pruning the search tree efficiently.

5.5 An enhanced propagator for C_{Opt}

A basic propagator for the follower's optimality constraint C_{Opt} can be built based on the generic scheme presented in Section 4.1. Below we present an enhanced algorithm that fully exploits the follower's lower and upper bounds during the exact solution of the follower's subproblem. This algorithm can be applied in bilevel problems where the leader's objective, f, does not depend on the follower's response.

The pseudo-code of the algorithm is shown in Fig. 5. In lines 3–5, the algorithm checks if the computed bounds on the follower's cost allow the existence of a solution. Afterwards, it solves the follower's subproblem *with* the leader's deadline constraints, resulting in solution Y^+ and cost g^+ (line 6). Then, the follower's subproblem is solved *without* the leader's deadline constraints, leading to cost g^* (line 10). Search can be aborted when a solution with $g^* < g^+$ is reached, because

```
PROCEDURE Propagate_C_Opt2()
1    IF X is not bound
2      RETURN
3    Compute ĝ^Fmax and ǧ^Lmin
4    IF ĝ^Fmax < ǧ^Lmin
5      Fail
6    LET Y^+ := arg min_Y'{g(X, Y') | C(X, Y') ∧ D(X, Y') ∧ g(X, Y') <= ĝ^Fmax}
7    IF Y^+ = 'no_solution'
8      Fail
9    LET g^+ := g(X, Y^+)
10   LET g^* := min_Y'{g(X, Y') | D(X, Y') ∧ g(X, Y') <= ĝ^Fmax}
11   IF g^+ > g^*
12     Fail
13   Instantiate Y ← Y^+
```

Fig. 5 Enhanced algorithm for propagating constraint C_{Opt}

this solution, as well as any potential improving solution, will lead to failure in lines 11 and 12 anyway. The solution Y^+ is optimal for the leader if and only if $g^+ = g^*$ (note that $g^+ \geq g^*$ always holds). Observe that the order of lines 6 and 10 is reversed w.r.t. the basic version of the propagator, which is advantageous because the problem faced in line 6 is tighter and generally easier-to-solve than the problem of line 10. This latter step exploits that f does not depend on the follower's response.

6 Computational experiments

6.1 Experimentation of the proposed solution techniques

In this section we report computational results achieved on the sample scheduling problem. The solver was implemented in such a way that each inference technique presented in a separate section above could be switched on or off individually. Moreover, the bounds on the follower's cost (see Section 5.3) could be computed independently in the internal search nodes of the master problem using the DP, or in the leaves using the WSPT schedule. Preliminary experiments showed that all the presented techniques contribute to pruning the search tree, but it does not pay off in terms of search time to use the leader's lower bound or to compare the follower's bounds in internal search nodes. Therefore we decided to switch off these two components in the main experiments, even in solver version V1 that we will consider the complete solver in the sequel. Further versions, V2 and V3, were created by gradually switching of the introduced inference techniques. The last version, V4, used only the basic constraint model of Section 5.1 without any enhancements. The techniques used in the different versions are summarized in Table 1.

The solver was implemented in C++ using ILOG Solver and Scheduler, both for the bilevel master problem and the follower's subproblem solver. ILOG Cplex was used for the computation of the LP lower bound. The experiments were run on a 1.86 GHz Intel Xeon computer with 2 GB of RAM under Windows Server 2003. The time limit was set to 600 s per problem instance.

Problem instances have been generated similarly to the instances for the single level problem of minimizing the weighted number of late jobs in [9] and [27], with the only difference that we have also added the follower's weights w_j^2. The parameters of the generator are the number of tasks n, the range of release times, k_R (a larger value means a greater variance of the release times), and the tightness of the deadlines, k_D (the larger the value, the wider the time windows). Parameter n varied between 20 and 50 with increments of 5, while k_R and k_D were chosen from the set {1, 5, 10, 20}. Generating ten instances with all possible combinations of the three parameters resulted in 1,120 instances altogether. Processing times were generated

Table 1 Comparison of the tested versions of the solver

	V1	V2	V3	V4
Basic constraint model (Section 5.1)	+	+	+	+
Enhanced propagator for C_{Opt} (Section 5.5)	+	+	+	
Follower's dominance rules in master problem (Section 5.2)	+	+		
Bounds on the follower's cost, leaves only (Section 5.3)	+			

using $U[1, 100]$, release times from $U[0, k_R n]$, deadlines from $U[r_j + p_j, r_j + p_j + k_D n]$, while weights w_j^1 and w_j^2 from $U[1, 10]$, where $U[a, b]$ denotes the discrete uniform distribution over integers from the interval $[a, b]$.

The comparison of the results achieved with the four solver versions is displayed in Table 2, while Table 3 provides further statistics about runs of the propagator of C_{Opt} in the complete version, V1. In both tables, each row contains combined result for the instances with a given value of N and k_D. Column *Opt* displays the number of instances that could be solved to proven optimality out of 40, while column *Best* shows the number of instances on which the solver found the best solution known for the instance. *Time* contains the average computation time in seconds or 600 for instances where the time limit was hit. Column *Nodes* shows the average number of search nodes. The additional columns in Table 3 contain the number of times the propagator of C_{Opt} reached the different steps of computation, as well as the total time of these computation steps: calculating the follower's bounds (*Step (1)*), the follower's minimum cost when the leader's constraints are respected (*Step (2)*), and the minimum cost when the leader's constraints are ignored (*Step (3)*). These steps corresponds to lines 3–5, 6–9, and 10–12 of the pseudo-code in Fig. 5.

The results show that the stronger versions of the solver were able to solve instances with up to 20–25 tasks to optimality, whereas the naive version, V4, started to have difficulties even with some 20-task instances. On the whole, the complete version, V1, solved 6.6%, 9.9%, and 32.3% more instances to optimality than versions V2, V3, V4, respectively. The difference becomes slightly more significant as the problem size increases, and the comparison of average computation times brings roughly the same result. The two versions V1 and V2 (and the other two versions V3 and V4 likewise) generate the same number of search nodes for all instances that they could solve on time. This happened because the two versions differ only in the way of processing the leaves. For some larger problems it happened that the solvers did not find any solutions at all. This was the case, e.g., for parameters $N = 50$ and $k_D = 20$, where no feasible solution has been found by any solver version. Except for these cases, V1 always found the best solution among the four solver versions.

Smaller values of k_D made the problems easier to solve for all versions, because then the leader had a smaller choice of task sets to accept, and those sets are identified relatively efficiently without the follower's optimality condition, too. This is made apparent especially by the low number of calls to the propagator of C_{Opt} with $k_D = 1$ (column *Step (1)* in Table 3). The results also depend on k_R (small k_R makes them easier to solve), but much less than on k_D or N.

The analysis of the runs of the enhanced propagator in Table 3 shows that follower's bounds computation inferred the infeasibility of the leaf in 67% of the cases. The exact CP solver had to be called with the leader's constraints in the remaining 33% of the runs (*Step (2)*), and without the leader's constraint in only 0.5% of the runs (*Step (3)*). On the one hand, this low percentage is an excellent result, since the last step of the algorithm is the most time consuming. On the other hand, it also shows that at least 99.5% of the leaves did not contain a solution that is both feasible for the leader and optimal for the follower. Hence, future research should address the efficient propagation of the follower's optimality constraint C_{Opt} also in the internal search nodes. Overall, 89% of the total computation time was spent in the propagator of C_{Opt} in the leaves (23%, 62%, and 4% in steps (1), (2), and (3), respectively).

Table 2 Comparison of four different versions of the solver

N	k_D	V1				V2				V3				V4			
		Opt	Best	Nodes	Time	Opt	Best	Nodes	Time	Opt	Best	Nodes	Time	Opt	Best	Nodes	Time
20	1	40	40	43	0.02	40	40	43	0.02	40	40	43	0.02	40	40	43	0.02
	5	40	40	240	0.05	40	40	240	0.14	40	40	309	0.18	40	40	309	0.29
	10	40	40	503	0.16	40	40	503	0.43	40	40	688	0.52	40	40	688	4.02
	20	40	40	4,385	6.53	40	40	4,385	20.06	40	40	11,542	24.51	26	29	7,681	270.53
25	1	40	40	97	0.02	40	40	97	0.03	40	40	97	0.03	40	40	97	0.03
	5	40	40	1,670	0.43	40	40	1,670	2.38	40	40	1,942	2.56	40	40	1,942	5.16
	10	40	40	6,974	4.98	40	40	6,974	20.24	40	40	10,568	24.26	35	38	8,607	142.98
	20	29	37	35,507	230.72	21	24	18,151	307.42	21	24	71,127	330.87	6	9	13,132	546.85
30	1	40	40	284	0.05	40	40	284	0.10	40	40	285	0.10	40	40	285	0.11
	5	38	39	30,690	33.10	38	38	11,624	41.45	38	38	13,328	42.23	37	37	7,276	70.56
	10	37	40	44,316	83.90	31	33	38,476	177.01	30	33	85,007	203.11	11	15	25,908	502.53
	20	13	24	133,946	442.32	12	18	65,989	453.52	9	12	228,712	521.80	0	3	11,701	600.00
35	1	40	40	701	0.12	40	40	701	0.24	40	40	706	0.24	40	40	706	0.27
	5	38	38	58,298	43.34	38	38	38,762	72.39	36	36	56,082	95.30	28	30	21,472	255.52
	10	26	35	229,963	312.59	16	18	119,492	396.13	15	17	251,134	419.90	3	5	21,570	575.97
	20	5	8	178,155	561.73	2	4	70,353	581.79	1	3	200,364	593.90	0	0	1,612	600.00
40	1	39	39	11,571	15.21	39	39	3,792	15.29	39	39	3,814	15.28	39	39	3,436	15.31
	5	34	36	276,036	154.51	30	31	156,687	246.98	28	30	222,036	286.20	15	16	57,854	467.86
	10	12	22	433,949	512.41	10	15	316,517	535.60	2	12	403,749	598.32	0	1	8,168	600.00
	20	0	3	152,912	600.00	0	1	52,706	600.00	0	0	207,786	600.00	0	0	310	600.00
45	1	39	39	13,618	16.22	39	39	6,799	18.59	39	39	6,637	18.65	39	39	6,371	19.34
	5	21	32	500,967	357.07	16	21	232,003	395.01	14	19	270,896	419.36	4	6	22,899	555.90
	10	2	15	445,267	586.29	0	9	238,465	600.00	0	4	260,543	600.00	0	0	4,566	600.00
	20	0	2	127,072	600.00	0	2	35,495	600.00	0	1	160,533	600.00	0	0	497	600.00
50	1	39	39	26,054	19.24	39	39	22,030	27.36	39	39	22,283	28.01	39	39	21,934	32.47
	5	13	27	683,993	484.13	8	15	373,070	518.58	7	13	412,258	549.67	1	2	18,848	586.40
	10	0	7	239,020	600.00	0	3	108,710	600.00	0	1	231,011	600.00	0	0	1,348	600.00
	20	0	0	58,843	600.00	0	0	19,428	600.00	0	0	156,308	600.00	0	0	108	600.00
∑		745	842			699	747			678	720			563	588		

Table 3 Detailed results achieved with the complete version, V1, of the solver

N	k_D	Opt	Best	Nodes	Time	Step (1)		Step (2)		Step (3)	
						Runs	Time	Runs	Time	Runs	Time
20	1	40	40	43	0.02	6	0.00	3	0.00	3	0.00
	5	40	40	240	0.05	51	0.00	14	0.00	4	0.00
	10	40	40	503	0.16	104	0.00	38	0.00	4	0.00
	20	40	40	4,385	6.53	2,687	1.10	1,291	3.93	9	0.23
25	1	40	40	97	0.02	6	0.00	3	0.00	3	0.00
	5	40	40	1,670	0.43	868	0.20	37	0.00	4	0.00
	10	40	40	6,974	4.98	2,914	2.15	1,074	1.03	10	0.35
	20	29	37	35,507	230.72	19,492	20.13	13,263	205.25	19	2.33
30	1	40	40	284	0.05	20	0.00	6	0.00	4	0.00
	5	38	39	30,690	33.10	24,704	24.03	8,891	5.38	163	0.53
	10	37	40	44,316	83.90	11,786	15.35	6,130	57.75	25	2.88
	20	13	24	133,946	442.32	75,825	62.48	21,176	350.13	13	11.35
35	1	40	40	701	0.12	35	0.00	19	0.00	5	0.00
	5	38	38	58,298	43.34	28,561	27.18	9,223	6.90	75	0.50
	10	26	35	229,963	312.59	63,536	75.03	31,118	172.68	400	27.60
	20	5	8	178,155	561.73	102,183	111.03	16,730	411.23	6	12.95
40	1	39	39	11,571	15.21	9,525	10.63	4,860	3.70	5	0.00
	5	34	36	276,036	154.51	36,983	54.15	16,676	31.55	692	25.30
	10	12	22	433,949	512.41	82,944	134.20	38,362	264.23	541	28.60
	20	0	3	152,912	600.00	88,775	106.93	11,535	468.90	5	0.00
45	1	39	39	13,618	16.22	8,641	11.48	3,391	3.03	17	0.05
	5	21	32	500,967	357.07	69,785	128.40	38,407	93.08	1,338	40.50
	10	2	15	445,267	586.29	104,149	136.25	38,077	335.18	121	20.63
	20	0	2	127,072	600.00	66,531	137.60	4,918	436.50	0	0.00
50	1	39	39	26,054	19.24	7,199	11.95	2,245	2.75	170	0.83
	5	13	27	683,993	484.13	80,435	150.78	34,328	141.30	723	36.55
	10	0	7	239,020	600.00	76,213	121.50	27,487	399.50	56	20.63
	20	0	0	58,843	600.00	34,352	101.43	3,939	482.95	0	0.00
Avg.						35,654	51.57	11,901	138.46	158	8.28

6.2 Results with translation to QCSP

To verify our contributions, we have implemented our scheduling problem in Qe-Code 2.0 using two different encodings to QCOP+. Our first QCOP+ model is based on the rewriting presented in Section 3.4, with one pair of existential and universal quantifiers. This model proved to be rather inefficient and memory consuming, since all possible decisions of the follower had to be enumerated explicitly in the computed strategy as possible values for the universally quantified variables. For this reason, instances with at most four tasks were trackable only, which is nearly an order of magnitude smaller than the instances we were interested in.

We have also implemented an alternative QCOP+ model with two existential quantifiers, and two different optimization criteria in the different quantifier scopes. The outer scope corresponds to the leader's choice, while the inner scope encodes the follower's subproblem. Note however that modeling tricks were required to overcome two shortcomings of the QCOP+ formalism. First, since QeCode does not

allow to state a constraint set C different from D, we had to embed a measure of violations of C into f. Furthermore, to ensure that a solution according to the optimistic assumption is computed, we added $\varepsilon \cdot f$ as tiebreaker to g. Using this model, we were able to solve instances with at most five tasks of our scheduling problem.

The above results show that even the naive version, V4, of our bilevel solver outperforms the approach of translation to QCOP+. This occurs because the techniques proposed in this paper exploit the bilevel problem structure, which is typically not present in general QCSP problems. Hence, these results provide justification for research on specialized solution techniques for discrete bilevel problems.

7 Conclusions

This paper introduced novel CP-based modeling and solution techniques for discrete bilevel optimization problems. Since bilevel problems are computationally difficult—they are often outside NP—, techniques that improve the efficiency of the solver are of key importance. Hence, we have presented how classical techniques of operations research, such as dominance rules or lower bounds, can be applied to bilevel problems. New algorithms for propagating the follower's optimality constraint and computing bounds on the follower's cost were proposed. These techniques were illustrated on a bilevel scheduling problem and evaluated in computational experiments.

We think that an interesting direction for future research is the development of new inference techniques for discrete bilevel problems. Depending on the specific problem, these can include the filtering of the leader's variable domains based on inference from the follower's optimality condition, or the re-use of the follower's response computed in earlier visited leaves.

References

1. Baptiste, P., Le Pape, C., & Nuijten, W. (2001). *Constraint-based scheduling*. Kluwer.
2. Bard, J. F. (1983). Coordination of a multidivisional organization through two levels of management. *Omega, 11*, 457–468.
3. Benedetti, M., Lallouet, A., & Vautard, J. (2007). QCSP made practical by virtue of restricted quantification. In *International joint conference on artificial intelligence* (pp. 38–43).
4. Benedetti, M., Lallouet, A., & Vautard, J. (2008). Modeling adversary scheduling with QCSP+. In *Proc. of the 2008 ACM symposium on applied computing* (pp. 151–155).
5. Benedetti, M., Lallouet, A., & Vautard, J. (2008). Quantified constraint optimization. In *CP2008, principles and practice of constraint programming. LNCS* (Vol. 5202, pp. 463–477). Springer.
6. Brown, G. G., Carlyle, W. M., Royset, J., & Wood, R. K. (2005). On the complexity of delaying an adversary's project. In *Operations research/computer science interfaces* (Vol. 29, Chapter 1). Springer.
7. Brown, K. N., Little, J., Creed, P. J., & Freuder, E. C. (2004). Adversarial constraint satisfaction by game-tree search. In *Proc. of ECAI 2004* (pp. 151–155).
8. Clark, P. A., & Westerberg, A. (1990). Bilevel programming for steady-state chemical process design. I: Fundamentals and algorithms. *Computers and Chemical Engineering, 14*(1), 87–97.
9. Dauzère-Pérès, S., & Sevaux, M. (2003). Using Lagrangean relaxation to minimize the weighted number of late jobs on a single machine. *Naval Research Logistics, 50*(3), 273–288.

10. Dempe, S. (2001). *Discrete bilevel optimization problems*. Technical report, Universität Leipzig.
11. Dempe, S. (2002). *Foundations of bilevel programming*. Kluwer.
12. Fanghänel, D., & Dempe, S. (2009). Bilevel programming with discrete lower level problems. *Optimization: A Journal of Mathematical Programming and Operations Research, 58*(8), 1029–1047.
13. Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability—A guide to the theory of NP-completeness*. Freeman.
14. Gecode (2009). *Generic constraint development environment, version 3.1*. http://www.gecode.org/.
15. Gent, I., Nightingale, P., & Stergiou, K. (2005). QCSP-Solve: A solver for quantified constraint satisfaction problems. In *Proceedings of IJCAI-2005* (pp. 138–143).
16. Gent, I. P., Nightingale, P., Rowley, A., & Stergiou, K. (2008). Solving quantified constraint satisfaction problems. *Artificial Intelligence, 172*, 738–771.
17. Graham, R. L., Lawler, E. L., Lenstra, J. K., & Rinnooy Kan, A. H. G. (1979). Optimization and approximation in deterministic sequencing and scheduling: A survey. *Annals of Operations Research, 5*, 287–326.
18. Hooker, J. N., & Ottosson, G. (2003). Logic-based benders decomposition. *Mathematical Programming, 96*, 33–60.
19. Jouglet, A., Baptiste, P., & Carlier, J. (2004). Branch-and-bound algorithms for total weighted tardiness. In *Handbook of scheduling: Algorithms, models, and performance analysis* (Chapter 13). Chapman & Hall/CRC.
20. Jouglet, A., Savourey, D., Carlier, J., & Baptiste, P. (2008). Dominance-based heuristics for one-machine total cost scheduling problems. *European Journal of Operational Research, 184*, 879–899.
21. Karlof, J. K., & Wang, W. (1996). Bilevel programming applied to the flow shop scheduling problem. *Computers and Operations Research, 23*(5), 443–451.
22. Kis, T., & Kovács, A. (2009). On bilevel machine scheduling problems. *OR Spektrum*, in print. doi:10.1007/s00291-010-0219-y.
23. Labbé, M., Marcotte, P., & Savard, G. (1998). A bilevel model of taxation and its application to optimal highway pricing. *Management Science, 44*(12), 1608–1622.
24. Lenstra, J. K., Rinnooy Kan, A. H. G., & Brucker, P. (1977). Complexity of machine scheduling problems. *Annals of Discrete Mathematics, 1*, 343–362.
25. Lukač, Z., Šorić, K., & Vojvodić Rosenzweig, V. (2008). Production planning problem with sequence dependent setups as a bilevel programming problem. *European Journal of Operational Research, 187*, 1504–1512.
26. Marcotte, P., Mercier, A., Savard, G., & Verter, V. (2009). Toll policies for mitigating hazardous materials transport risk. *Transportation Science, 43*(2), 228–243.
27. M'Hallah, R., & Bulfin, R. L. (2007). Minimizing the weighted number of tardy jobs on a single machine with release dates. *European Journal of Operational Research, 176*, 727–744.
28. Nightingale, P. (2009). Non-binary quantified CSP: Algorithms and modelling. *Constraints, 14*(4), 539–581.
29. Pan, Y., & Shi, L. (2005). Dual constrained single machine sequencing to minimize total weighted completion time. *IEEE Transactions on Automation Science and Engineering, 2*(4), 344–357.
30. Prékopa, A. (1995). *Stochastic programming*. Kluwer.
31. QeCode (2009). An open QCSP solver, version 2.0. www.univ-orleans.fr/lifo/members/vautard/qecode.
32. Rossi, F., van Beek, P., & Walsh, T. (Eds.) (2006). *Handbook of constraint programming*. Elsevier.
33. Sherali, H. D., Soyster, A. L., & Murphy, F. H. (1983). Stackelberg-Nash-Cournot equilibria: Characterizations and computations. *Operations Research, 31*, 253–276.
34. Stackelberg, H. (1934). *Marktform and gleichgewicht*. Vienna: Julius Springer.
35. Tsoukalas, A., Wiesemann, W., & Rustem, B. (2009). Global optimisation of pessimistic bi-level problems. In P. M. Pardalos, & T. F. Coleman (Eds.), *Lectures on global optimization* (pp. 215–243). American Mathematical Society.
36. Verger, G., & Bessière, C. (2006). Blocksolve: a bottom-up approach for solving quantified CSPs. In *CP2006, principles and practice of constraint programming. LNCS* (Vol. 4204, pp. 635–649). Springer.
37. von Stengel, B. (2007). Equilibrium computation for two-player games in strategic and extensive form. In N. Nisan, T. Roughgarden, É. Tardos, & V. V. Vazirani (Eds.), *Algorithmic game theory* (Chapter 3). Cambridge University Press.