

Exact algorithms for single-machine scheduling with time windows and precedence constraints

Davari M, Demeulemeester E, Leus R, Talla Nobibon F.



Exact algorithms for single-machine scheduling with time windows and precedence constraints

Morteza Davari · Erik Demeulemeester
Roel Leus · Fabrice Talla Nobibon

Abstract We study a single-machine scheduling problem that is a generalization of a number of problems for which computational procedures have already been published. Each job has a processing time, a release date, a due date, a deadline and a weight representing the penalty per unit-time delay beyond the due date. The goal is to schedule all jobs such that the total weighted tardiness penalty is minimized and both the precedence constraints as well as the time windows (implied by the release dates and the deadlines) are respected. We develop a branch-and-bound algorithm that solves the problem to optimality. Computational results show that our approach is effective in solving medium-sized instances, and that it compares favorably with existing methods for special cases of the problem.

Keywords single-machine scheduling · branch and bound · mixed-integer programming

M. Davari
Research Center for Operations Management, KU Leuven,
Belgium
E-mail: morteza.davari@kuleuven.be

E. Demeulemeester ✉
Research Center for Operations Management, KU Leuven
tel.: +32 16 32 69 72
fax: +32 16 32 66 24
E-mail: erik.demeulemeester@kuleuven.be

R. Leus
ORSTAT, KU Leuven
E-mail: roel.leus@kuleuven.be

F. Talla Nobibon
FedEx Express Europe, Middle East, Indian Subcontinent &
Africa
E-mail: ftallanobibon@fedex.com

1 Introduction

Scheduling problems arise in production planning [38], in balancing processes [37], in telecommunication [28] and more generally in all situations in which scarce resources are to be allocated to jobs over time [32]. Depending on the application, the corresponding scheduling problem can be such that each job must be processed within a given time window, where the lower bound (*release date* or *ready time*) of this time window represents the earliest start of the execution of the job and the upper bound (*deadline*) corresponds with the latest acceptable completion time, for instance the ultimate delivery time agreed upon with the customer [15, 31, 46]. For some of these applications, only release dates or only deadlines are considered [21, 30, 33, 40]. In practice, a job often also needs to be processed before or after other jobs, e.g., due to tool or fixture restrictions or for other case-dependent technological reasons, which leads to *precedence constraints* [25, 34, 42]. Finally, the contract with a client can also contain clauses that stipulate that penalties must be paid when the execution of a job is not completed before a reference date (*due date*) [1, 12, 20, 21, 39, 40].

In this article, we develop exact algorithms for a single-machine scheduling problem with total weighted tardiness (TWT) penalties. In the standard three-field notation introduced by Graham et al. [16], the problem that we tackle can be denoted as $1|r_j, \delta_j, prec|\sum w_j T_j$: the execution of each job is constrained to take place within a time window, and we assume the corresponding deadline to be greater than or equal to a due date, which is the reference for computing the tardiness of the job. The scheduling decisions are also subject to precedence constraints. In the following lines we briefly summarize the state of the art.

Abdul-Razaq et al. [2] survey different branch-and-bound (B&B) algorithms for $1||\sum w_j T_j$. A benchmark algorithm is the B&B procedure of Potts and Van Wassenhove [36]; an older reference is Held and Karp [18], who present a dynamic programming (DP) approach. Abdul-Razaq and Potts [1] introduce a DP-based approach to obtain tight lower bounds for the generalized version of the problem where the cost function is piecewise linear, and Ibaraki and Nakamura [20] extend this work and construct an exact method, called Successive Sublimation Dynamic Programming (SSDP), which solves instances with up to 50 jobs. Tanaka et al. [41] improve the SSDP of [20] and succeed in solving reasonably large instances (300 jobs) of $1||\sum w_j T_j$ within acceptable runtimes.

Single-machine scheduling for TWT with (possibly unequal) release dates ($1|r_j|\sum w_j T_j$) has also been studied by several authors. Akturk and Ozdemir [3, 4] and Jouglet et al. [21] develop B&B algorithms that solve small instances. Van den Akker et al. [45] propose a time-indexed formulation and a method based on column generation to solve this problem with identical processing times. Tanaka and Fujikuma [40] present an SSDP algorithm that can solve instances of $1|r_j|\sum w_j T_j$ with up to 100 jobs.

There are only few papers dealing with single-machine scheduling with deadlines and/or precedence constraints. Among these, we cite Posner [33] and Pan [30], who propose B&B algorithms for $1|\delta_j|\sum w_j C_j$, Pan and Shi [31], who develop a B&B algorithm to solve $1|r_j, \delta_j|\sum w_j C_j$, and Lawler [25] and Potts [34], who present exact B&B algorithms to solve $1|prec|\sum w_j C_j$. Tanaka and Sato [42] also propose an SSDP algorithm to solve a generalization of $1|prec|\sum w_j T_j$ (piecewise linear cost function). To the best of our knowledge, scheduling problems with release dates, deadlines and precedence constraints have not yet been studied in the literature. The goal of this paper is to fill this gap and to propose efficient B&B algorithms that solve all the foregoing subproblems within limited computation times.

The remainder of this paper is structured as follows. In Section 2 we provide some definitions and a formal problem statement, while Section 3 proposes two different integer programming formulations. In Section 4 we explain the branching strategies for our B&B algorithms, while the lower bounds and dominance rules are discussed in Section 5 and Section 6, respectively. Computational results are reported and discussed in Section 7. We provide a summary and conclusions in Section 8.

2 Problem description

The jobs to be scheduled are gathered in set $N = \{1, 2, \dots, n\}$. Job i is characterized by a processing time p_i , a release date r_i , a due date d_i , a deadline δ_i , and a weight w_i , which represents the cost per unit time of delay beyond d_i . Jobs can neither be processed before their release dates nor after their deadlines ($0 \leq r_i \leq d_i \leq \delta_i$). Precedence constraints are represented by a graph $G = (N', A)$, where $N' = N \cup \{0, n+1\}$, with 0 a dummy start job and $n+1$ a dummy end. Each arc $(i, j) \in A$ implies that job i must be executed before job j (job i is a predecessor of job j). We will assume that $G(N', A)$ is its own transitive reduction, that is, no transitive arcs are included in A . Let \mathcal{P}_i be the set of all predecessors of job i in A ($\mathcal{P}_i = \{k | (k, i) \in A\}$) and \mathcal{Q}_i the set of successors of job i ($\mathcal{Q}_i = \{k | (i, k) \in A\}$). We also define an associated graph $\hat{G} = (N', \hat{A})$ as the transitive closure of G . We assume that $\mathcal{P}_0 = \mathcal{Q}_{n+1} = \emptyset$, and that all jobs are successor of 0 and predecessor of $n+1$ in \hat{G} (apart from the job itself).

Throughout this paper, we use the term ‘sequencing’ to refer to ordering the jobs (establishing a permutation) whereas ‘scheduling’ means that start (or end) times are determined. We denote by π an arbitrary sequence of jobs, where π_k represents the job at the k^{th} position in that sequence. Let $\pi^{-1}(i)$ be the position of job i in π ; we only consider sequences for which $\pi^{-1}(i) < \pi^{-1}(j)$ for all $(i, j) \in A$. Value C_i is the completion time of job i . Each sequence π implies a schedule, as follows:

$$C_{\pi_i} = \begin{cases} \max\{r_{\pi_i}, C_{\pi_{i-1}}\} + p_{\pi_i} & \text{if } i > 1 \\ r_{\pi_i} + p_{\pi_i} & \text{if } i = 1 \end{cases}$$

Equivalently, the end of job i according to sequence π can also be written as $C_i(\pi)$. We denote by \mathcal{D} the set of all feasible permutations, where a permutation π is feasible ($\pi \in \mathcal{D}$) if and only if it generates a feasible schedule, which means that

$$r_{\pi_i} + p_{\pi_i} \leq C_{\pi_i} \leq \delta_{\pi_i} \quad \forall i \in N$$

Note that the set \mathcal{D} may be empty.


The weighted tardiness associated with the job at the i^{th} position in the sequence π is given by $W(\pi_i) = w_{\pi_i} (C_{\pi_i} - d_{\pi_i})^+$, where $x^+ = \max\{0, x\}$. A conceptual formulation of the problem P studied in this paper is the following:

$$P : \min_{\pi \in \mathcal{D}} \text{TWT}(\pi) = \sum_{i=1}^n W(\pi_i). \quad (1)$$

This problem is at least as hard as $1||\sum w_i T_i$, which is known to be strongly NP-hard [24, 26, 32]. A stronger

Job i	p_i	r_i	d_i	δ_i	w_i
Job 1	2	3	10	14	1
Job 2	3	4	11	13	2
Job 3	4	3	8	15	3
Job 4	2	2	6	9	1

Table 1: Job characteristics.

Fig. 1: Precedence graph $G(N', A)$.

result is that the mere verification of the existence of a feasible schedule that respects a set of ready times and deadlines is already NP-complete (problem SS1, page 236, [14]); we do not, however, incorporate the feasibility check as a formal part of the problem statement.

Example 1 Consider the following instance of P with $n = 4$ jobs. The processing times, release dates, due dates, deadlines and weights of the jobs are given in Table 1. The graph representing the precedence constraints is depicted in Figure 1, with arc set $A = \{(0, 1), (0, 4), (1, 2), (2, 3), (3, 5), (4, 5)\}$.

An optimal solution to this instance is $\pi = (4, 1, 2, 3)$, which leads to the schedule $C_1 = 6$, $C_2 = 9$, $C_3 = 13$ and $C_4 = 4$. The objective value is $w_4 \times 0 + w_1 \times 0 + w_2 \times 0 + w_3 \times (13 - 8) = 3 \times 5 = 15$.

3 Mathematical formulations

The conceptual formulation for P presented in the previous section is not linear, therefore it cannot be used by a standard (linear) mixed-integer programming (MIP) solver. In this section, we propose an Assignment Formulation (ASF) and a Time-Indexed Formulation (TIF) for the problem. These formulations are adaptations of those presented in [22, 39].

3.1 Assignment formulation

We use binary decision variables $x_{is} \in \{0, 1\}$ ($i \in N$, $s \in \{1, 2, \dots, n\}$), which identify the position of jobs in the sequence so that x_{is} is equal to 1 if job i is the s^{th} job processed and equal to 0 otherwise. In other words, $x_{is} = 1$ if and only if $\pi_s = i$. We also use additional continuous variables $T_i \geq 0$ representing the tardiness of job $i \in N$ and continuous variables $\tau_s \geq 0$ representing the machine idle time immediately before the

execution of the s^{th} job. The MIP formulation is given by:

$$\text{ASF : } \min \sum_{i=1}^n w_i T_i \quad (2)$$

subject to

$$\sum_{s=1}^n x_{is} = 1 \quad \forall i \in N \quad (3)$$

$$\sum_{i=1}^n x_{is} = 1 \quad \forall s \in \{1, 2, \dots, n\} \quad (4)$$

$$\sum_{s=1}^n x_{is} s \leq \sum_{t=1}^n x_{jt} t - 1 \quad \forall (i, j) \in A \quad (5)$$

$$\sum_{s=1}^n x_{is} s \leq \sum_{t=1}^n x_{jt} t - 1 \quad \forall (i, j) \in A \quad (6)$$

$$\tau_s \geq \sum_{i=1}^n x_{is} r_i - \sum_{t=1}^{s-1} \left(\sum_{i=1}^n (x_{it} p_i) + \tau_t \right) \quad \forall s \in N \quad (7)$$

$$\sum_{t=1}^s \tau_t + \sum_{t=1}^{s-1} \sum_{i=1}^n p_i x_{it} + \sum_{i=1}^n ((p_i - \delta_i) x_{is}) \leq 0 \quad \forall s \in N \quad (8)$$

$$T_i \geq \sum_{t=1}^s \tau_t + \sum_{t=1}^{s-1} \sum_{j=1}^n p_j x_{jt} + p_i - d_i - (1 - x_{is}) M_i \quad \forall i \in N, s \in N \quad (9)$$

$$x_{is} \in \{0, 1\}, \tau_s, T_i \geq 0 \quad \forall i \in N, s \in \{1, 2, \dots, n\} \quad (10)$$

The objective function (2) is a reformulation of (1). The set of constraints (3) ensures that all jobs are executed. Constraints (4) check that each position in the sequence is occupied by exactly one job. The set of constraints (6) enforces the precedence restrictions. The set of equations (7) computes the idle time of the machine between the jobs in positions $s - 1$ and s , and ensures that each job is not started before its release date. In this set of constraints, $\sum_{t=1}^{s-1} (\sum_{i=1}^n (x_{it} p_i) + \tau_t)$ equals the completion time of the $(s - 1)^{\text{th}}$ job. Constraints (8) ensure that each job is not completed after its deadline, where $\sum_{t=1}^s \tau_t + \sum_{t=1}^{s-1} \sum_{i=1}^n p_i x_{it}$ is the start time of the s^{th} job. Constraints (9) compute the correct value of the tardiness of job i , with $M_i = \delta_i - d_i$ the maximum tardiness of job i .

A variant of ASF is obtained by replacing the set of constraints (6) by the following:

$$\sum_{s=v}^n x_{is} + \sum_{s=1}^v x_{js} \leq 1 \quad \forall (i, j) \in A, \forall v \in \{1, \dots, n\}. \quad (11)$$

We refer to this alternative formulation as ASF'. We have the following result:

Lemma 1 ASF' is stronger than ASF .

Proof See Appendix. \square

The number of constraints in (11) is much higher than in (6). As a result, the additional computational effort needed to process this higher number of constraints might offset the improvement of a stronger bound, and we will empirically compare the performance of the two variants in Section 7.4.

3.2 Time-indexed formulation

Let T_S (respectively T_E) be a lower (respectively upper) bound on the time the execution of any job can be completed; we compute these values as $T_S = \min\{r_i + p_i | i \in N\}$ and $T_E = \max\{\delta_i | i \in N\}$. The time-indexed formulation uses binary decision variables $x_{it} \in \{0, 1\}$, for $i \in N$ and $T_S \leq t \leq T_E$, where $x_{it} = 1$ if job i is completed (exactly) at time t and $x_{it} = 0$ otherwise. We also introduce the set of parameters $T_{it} = (t - d_i)^+$, representing the tardiness of job i when it finishes at time t . The time-indexed formulation is given by:

$$\text{TIF : } \min \sum_{i=1}^n \sum_{t=r_i+p_i}^{\delta_i} w_i T_{it} x_{it} \quad (12)$$

subject to

$$\sum_{i=1}^n \sum_{s=\max\{t, r_i+p_i\}}^{\min\{\delta_i, t+p_i-1\}} x_{is} \leq 1 \quad \forall t, T_S \leq t \leq T_E \quad (13)$$

$$\sum_{t=r_i+p_i}^{\delta_i} x_{it} = 1 \quad \forall i \in N \quad (14)$$

$$\sum_{s=r_i+p_i}^{\delta_i} x_{is} s \leq \sum_{t=r_j+p_j}^{\delta_j} x_{jt} t - p_j \quad \forall (i, j) \in A \quad (15)$$

$$x_{it} \in \{0, 1\} \quad i \in N, r_i + p_i \leq t \leq \delta_i \quad (16)$$

The set of constraints (13) eliminates the parts of the solution space where the jobs overlap. The constraint set (14) ensures that all jobs are scheduled exactly once. We enforce precedence constraints in the formulation using the set of constraints (15).

Similarly as for the assignment formulation, we introduce an alternative formulation TIF' by replacing the set of constraints (15) by the following:

$$\sum_{s=t}^{\delta_i} x_{is} + \sum_{s=r_j+p_j}^{t-p_i} x_{js} \leq 1 \quad (17)$$

$$\forall (i, j) \in A; \forall t, \max\{r_i, r_j + p_j\} + p_i \leq t \leq \min\{\delta_i, \delta_j + p_i\}$$

Lemma 2 TIF' is stronger than TIF .




Fig. 2: The structure of a partial schedule.

Proof See [5, 9]. \square

As explained for the assignment formulation, the performance of the new formulation is not necessarily better. In fact, it can be much worse than TIF, since in a time-indexed formulation the number of additional constraints is quite large (pseudo-polynomial).

4 Branching strategies

In this section we discuss two different branching strategies for our B&B algorithm. The structure of the B&B search trees is as follows: each tree consists of a finite number of nodes and branches, and at each level of the tree we make a sequencing decision for one job. Each node thus corresponds with a selection $S_P \subset N$ containing the already scheduled jobs and a set of unscheduled jobs $U = N \setminus S_P$. Each node also has two feasible partial sequences σ_B and σ_E of the scheduled jobs (each $i \in S_P$ appears in exactly one of these two): σ_B (respectively σ_E) denotes the partial sequence of jobs scheduled from the beginning (respectively end) of the scheduling horizon; see Figure 2 for an illustration. All jobs that are not scheduled, belong to the set of unscheduled jobs $U = E_B \cup E_E \cup E_N$. E_B is subset of unscheduled jobs that are eligible to be scheduled immediately after the last job in σ_B , E_E is the subset of unscheduled jobs that are eligible to be scheduled immediately before the first job in σ_E and E_N is the subset of unscheduled jobs that are not in $E_B \cup E_E$.

The root node represents an empty schedule ($S_P = \sigma_B = \sigma_E = \emptyset$). Each node branches into a number of child nodes, which each correspond with the scheduling of one particular job, called the *decision job*, as early as possible after the last job in σ_B or as late as possible before the first job in σ_E . A branch is called a *forward branch* if it schedules a job after the last job in σ_B , and is called a *backward branch* if it schedules a job before the first job in σ_E . In our branching strategies, there will be either only forward branches or only backward branches emanating from each given node. We will say that a node is of type FB (respectively BB) if all its branches are forward (respectively backward) branches.

Although scheduling jobs backward (from the end of time horizon) often improves the tightness of lower bounds [36], it is not always possible in the presence

of non-equal release dates; see Section 4.2 for more details. Backward scheduling also does not always improve the performance of the B&B algorithm because the efficiency of other dominance rules may decrease when we schedule jobs backward. We propose two B&B algorithms, each applying one of the branching strategies: BB1 corresponds with branching strategy 1 where only FB nodes are used and BB2 corresponds with branching strategy 2 where both FB and BB are created. The bounding and the dominance properties discussed in the following sections are the same in both B&B algorithms.

Let $C_{\max}(\sigma)$ be the completion time of the last job in the sequence σ . Throughout the branching procedure, we maintain two vectors of *updated release dates*, namely $\hat{\mathbf{r}} = (\hat{r}_1, \dots, \hat{r}_n)$ and $\bar{\mathbf{r}} = (\bar{r}_1, \dots, \bar{r}_n)$, defined as follows:

$$\hat{r}_j = \max\{r_j, C_{\max}(\sigma_B)\}$$

$$\bar{r}_j = \max\left\{\hat{r}_j, \max_{k \in \mathcal{P}_j} \{\bar{r}_k + p_k\}\right\}.$$

Let $st(\pi)$ denote the start time of the first job according to sequence π . In line with the two vectors of updated release dates, we also introduce two vectors of *updated deadlines*, namely $\hat{\delta} = (\hat{\delta}_1, \dots, \hat{\delta}_n)$ and $\bar{\delta} = (\bar{\delta}_1, \dots, \bar{\delta}_n)$, which are recursively computed as follows:

$$\hat{\delta}_j = \min\{\delta_j, st(\sigma_E)\}$$

$$\bar{\delta}_j = \min\left\{\hat{\delta}_j, \min_{k \in \mathcal{Q}_j} \{\bar{\delta}_k - p_k\}\right\}.$$

We use these updated release dates and deadlines in computing lower bounds and dominance rules. $\bar{\delta}$ and $\bar{\mathbf{r}}$ are more restrictive than $\hat{\delta}$ and $\hat{\mathbf{r}}$ in each node of the search tree ($\bar{r}_j \geq \hat{r}_j$ and $\bar{\delta}_j \leq \hat{\delta}_j$). Although being restrictive often is positive, \hat{r}_j and $\hat{\delta}_j$ are occasionally preferred over \bar{r}_j and $\bar{\delta}_j$, specifically in parts of computations related to the dominance rules discussed in Section 6. Further explanations of these occasions are given in Section 6. There are many cases in which $\bar{r}_j = \hat{r}_j$ (respectively $\bar{\delta}_j = \hat{\delta}_j$) and either of the updated release dates (respectively deadlines) can be used. In these cases, we use \hat{r}_j (respectively $\hat{\delta}_j$) because less computations are needed.

4.1 Branching strategy 1

Branching strategy 1 only uses FB nodes. The search tree is explored depth-first such that among children of a node, those with larger out-degrees (number of transitive successors) of their decision jobs in \hat{G} are visited first. As a tie-breaking rule, among children with equal

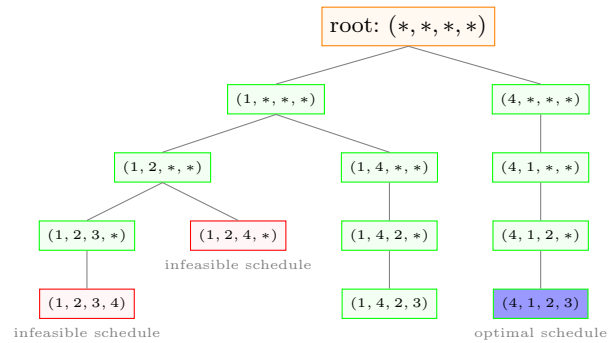


Fig. 3: Branching strategy 1 for Example 1 without dominance rules and without lower bounds.

out-degrees of their decision jobs, the node with lower index is visited first.

Figure 3 illustrates branching strategy 1 applied to Example 1; an asterisk ‘*’ indicates that the position has not been decided yet. Among the children of the root node, the node $(1, *, *, *)$ corresponds with the decision job (job 1) with the largest out-degree (namely 3). As a result, the node $(1, *, *, *)$ is visited first. The nodes $(2, *, *, *)$ and $(3, *, *, *)$ are not in the tree because they violate precedence constraints. Among the children of $(1, *, *, *)$, the node $(1, 2, *, *)$ is visited first because it has the decision job 2 with the largest out-degree. Among the children of $(1, 2, *, *)$, the node $(1, 2, 3, *)$ is visited first because its decision job has the largest out-degree and the smallest index. In Figure 3, green nodes are FB nodes; no BB nodes are present. Red nodes are considered infeasible because the completion of a job (namely job 4) occurs after its deadline. The node $(1, 4, 2, 3)$ corresponds with a feasible schedule, but it is not optimal: its objective value is greater than 15, which is attained by the optimal sequence $(4, 1, 2, 3)$.

4.2 Branching strategy 2

In branching strategy 2, we try to exploit the advantages of backward scheduling whenever possible, so the search tree consists of both FB and BB nodes. If the inequality $C_{\max}(\sigma_B) < r_{\max}(U) = \max_{j \in U} \{r_j\}$ holds, then the start times of the jobs in σ_E will depend on the order in which unscheduled jobs are processed. Therefore, if the inequality $C_{\max}(\sigma_B) < r_{\max}(U)$ holds, the corresponding node is of type FB. Otherwise, the completion time of the last job in σ_E can be computed regardless of the sequencing decisions for the jobs in U , and we have a BB node. The branching is *depth-first* for both FB and BB nodes. Among the children of a FB (respectively BB) node, those with higher (respectively

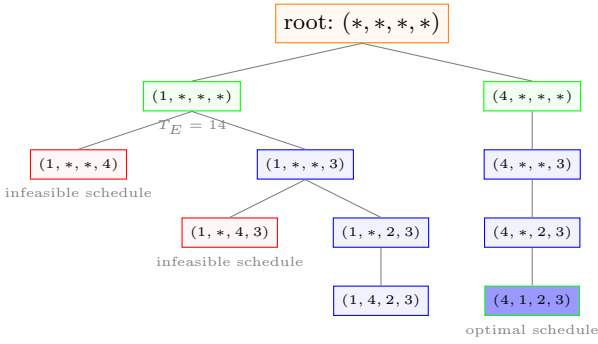


Fig. 4: Branching strategy 2 for Example 1 without dominance rules and without lower bounds.

lower) out-degrees of their decision jobs are visited first. As a tie-breaking rule, among children with equal out-degrees, the node with lower (respectively higher) index is visited first.

Figure 4 illustrates branching strategy 2 for Example 1; green nodes are of type FB and blue nodes are of type BB. The root node is FB because $C_{\max}(\emptyset) = 0 < 4 = r_{\max}(\{1, 2, 3, 4\})$. At the node labeled $(1, *, *, *)$, the completion time $C_{\max}(1, *, *, *) = 5$ of the decision job surpasses $r_{\max}(\{1, 2, 3, 4\}) = 4$, therefore the end of scheduling horizon is computed ($T_E = 5 + 3 + 4 + 2 = 14$) and the node is BB. The red nodes are infeasible because the completion time of job 4 falls after its deadline.

5 Lower bounding

In this section we describe the lower bounds that are implemented in our B&B algorithm. Section 5.1 first introduces a conceptual formulation for our problem, and in Section 5.2 we describe several lower bounds, which are mostly based on Lagrangian relaxation.

5.1 Another conceptual formulation

Let variable $C_j \geq 0$ denote the completion time of job $j \in N$ and let variable $T_j \geq 0$ represent the tardiness of job j . An alternative formulation of our problem is given by:

$$P : \min \sum_{j=1}^n w_j T_j$$

subject to

$$T_j \geq C_j - d_j \quad \forall j \in N \quad (18)$$

$$C_j \geq r_j + p_j \quad \forall j \in N \quad (19)$$

$$C_j \leq \delta_j \quad \forall j \in N \quad (20)$$

$$C_j \geq C_i + p_j \quad \forall (i, j) \in A \quad (21)$$

$$C_j \geq C_i + p_j \text{ or } C_i \geq C_j + p_i \quad \forall i, j \in N; i < j \quad (22)$$

$$T_j \geq 0 \quad \forall j \in N \quad (23)$$

$$C_j \geq 0 \quad \forall j \in N \quad (24)$$

In the above formulation, constraints (18) and (23) reflect the definition of job tardiness. Constraints (19) and (20) enforce time windows. Constraints (21) ensure that each job is scheduled after all its predecessors. Constraints (22) guarantee that jobs do not overlap. We will use this formulation in Section 5.2 for producing lower bounds.

To the best of our knowledge, a lower-bound procedure specifically for P has to date not been developed in the literature. Lower bounds proposed for $1|| \sum w_j T_j$, $1|prec| \sum w_j C_j$ and $1|r_j| \sum w_j C_j$, however, can also function as a lower bound for P; this is shown in the following theorems. These theorems are extensions of those presented in [3].

Let I be an instance of $1|\beta| \sum w_j T_j$. We construct an instance I' of $1|| \sum w_j T_j$ by removing all constraints implied by β and an instance I'' of $1|\beta| \sum w_j C_j$ by replacing all due dates with zeros. Let $\text{TWT}^*(I)$ be the optimal objective value of I . Given any valid lower bound $lb_{I'}$ on the optimal value of I' , we have:

Theorem 1 $lb_{I'} \leq \text{TWT}^*(I)$.

Proof Straightforward. \square

A job is called *early* if it finishes at or before its due date and is said to be *tardy* if it finishes after its due date. Let $C_j(S)$ be the completion time of job j in feasible solution S . For an optimal solution S^* to I , we partition N into two subsets: the set \mathcal{E} of early jobs and the set \mathcal{T} of tardy jobs. Let lb^E be a lower bound on the value $\sum_{j \in \mathcal{E}} w_j (d_j - C_j(S^*))$. Given any valid lower bound $\bar{lb}_{I''}$ on the optimal value of I'' , we have:

Theorem 2 $\bar{lb}_{I''} - \sum_j w_j d_j + lb^E \leq \text{TWT}^*(I)$.

Proof The following equality holds:

$$\begin{aligned} \text{TWT}^*(I) &= \sum_{j \in \mathcal{T}} w_j(C_j(S^*) - d_j) = \sum_{j \in N} w_j(C_j(S^*) - d_j) \\ &\quad - \sum_{j \in \mathcal{E}} w_j(C_j(S^*) - d_j) = \sum_{j \in N} w_j C_j(S^*) \\ &\quad - \sum_{j \in N} w_j d_j + \sum_{j \in \mathcal{E}} w_j(d_j - C_j(S^*)). \end{aligned}$$

Recall that \bar{lb}_J is a valid lower bound on the value $\sum_{j \in N} w_j C_j(S^*)$ and lb^E is a valid lower bound on the value $\sum_{j \in \mathcal{E}} w_j(d_j - C_j(S^*))$. \square

In the following, we remove several combinations of constraints in P to construct subproblems for which there exist polynomial-time-bounded algorithms for computing lower bounds. These bounds then directly lead to valid lower bounds for P via Theorem 1.

5.2 Lagrangian-relaxation-based bounds

In this section, we use Lagrangian relaxation for computing various lower bounds. Let P_0 be the subproblem of P in which constraints (19), (20) and (21) are removed. This problem is studied by Potts and Van Wassenhove [36] and is considered as our base problem. Let λ be a vector of Lagrangian multipliers. Potts and Van Wassenhove [36] obtain the following Lagrangian problem associated with P_0 :

$$\begin{aligned} \text{LR}_{P_0} : \quad L_0(\lambda) &= \min \sum_{j=1}^n (w_j - \lambda_j) T_j + \sum_{j=1}^n \lambda_j (C_j - d_j) \\ &\text{subject to constraints (22)–(24).} \end{aligned}$$

Parameter λ_j is the Lagrangian multiplier associated with job j ($0 \leq \lambda_j \leq w_j$). Potts and Van Wassenhove propose a polynomial-time algorithm to set the multipliers. Their algorithm yields a very good lower bound for P_0 ; they compute the optimal values of the multipliers in $O(n \log n)$ time, and for a given set of multipliers, the bound itself can be computed in linear time. Let λ_{PV} be the best Lagrangian multipliers computed by Potts and Van Wassenhove [36]; we refer to this lower bound as $\text{LB}_0 = L_0(\lambda_{PV})$. By Theorem 1, LB_0 is also a valid bound for P. Quite a number of aspects of the definition of P are completely ignored in LB_0 , however; in the following sections, we will examine a number of ways to strengthen LB_0 .

5.2.1 Retrieving precedence constraints

When $A \neq \emptyset$ then incorporating some or all of precedence constraints into the lower bound will improve






Fig. 5: This figure shows (a) an example graph G , (b) an associated VSP sub-graph G' and (c) G'' .

its quality. We partition arc set A as $A = A' \cup A''$, where $G' = (N, A')$ is a two-terminal *vertex serial-parallel* (VSP) graph and $G'' = (N, A'')$. Figure 5 depicts an example of this graph decomposition. For the precise definition of VSP graphs, we refer to Valdes et al. [43]. It should be noted that there exist two types of serial-parallel graphs: VSP graphs and *edge serial-parallel* (ESP) graphs. Valdes et al. [43] describe the link between these two types: a graph is VSP if and only if its so-called ‘line-digraph inverse’ is ESP. We split the set of constraints (21) into two subsets, as follows:

$$C_j \geq C_i + p_j \quad \forall (i, j) \in A' \quad (25)$$

$$C_j \geq C_i + p_j \quad \forall (i, j) \in A'' \quad (26)$$

We introduce P_1 , which is a generalization of P_0 where precedence constraints are retrieved by imposing constraints (25) and (26). We create the following associ-

ated Lagrangian problem:

$$\begin{aligned} \text{LR}_{P_1} : \quad L_1(\lambda, \mu) = & \min \sum_{j \in N} (w_j - \lambda_j) T_j \\ & + \sum_{j \in N} \lambda_j (C_j - d_j) + \sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} (C_j + p_k - C_k) \end{aligned}$$

subject to constraints (22)–(25).

Here $\lambda_j \geq 0$ is again the multiplier with job j and $\mu_{jk} \geq 0$ denotes the Lagrangian multiplier associated with the arc $(j, k) \in A$. We deliberately keep constraints (25) in the Lagrangian problem LR_{P_1} . The objective function can be rewritten as

$$\begin{aligned} & \sum_{j \in N} (w_j - \lambda_j) T_j + \sum_{j \in N} C_j \left(\lambda_j + \sum_{k \in Q_j} \mu_{jk} - \sum_{k \in P_j} \mu_{kj} \right) \\ & + \sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} p_k - \sum_{j \in N} \lambda_j d_j, \end{aligned}$$

so it can be seen that LR_{P_1} is a total-weighted-completion-times problem with serial-parallel precedence constraints, because all T_j will be set to zero. Lawler [25] proposes an algorithm that solves this problem in $O(n \log n)$ time provided that a *decomposition tree* is also given for the VSP graph G' . Valdes et al. [43] propose an $O(n + m)$ -time algorithm to construct a decomposition tree of a VSP graph, where m is the number of arcs in the graph. Calinescu et al. [8] show that any VSP graph (directed or undirected), including G' , has at most $2n - 3$ arcs. Therefore, for any given λ and μ , the problem LR_{P_1} is solvable in $O(n \log n)$ time. From the theory of Lagrangian relaxation (see Fisher [13]), for any choice of non-negative multipliers, $L_1(\lambda, \mu)$ provides a lower bound for P_1 . By Theorem 1, this lower bound is also valid for P . In Section 5.2.2, we explain how to choose appropriate values for λ and μ and Section 5.2.3 describes how to select a suitable VSP graph G' and how to construct a decomposition tree for G' .

5.2.2 Multiplier adjustment

We present a two-phase adjustment (TPA) procedure for the multipliers in $L_1(\lambda, \mu)$. Let λ_{TPA} and μ_{TPA} be Lagrangian multipliers adjusted by TPA; these lead to a new lower bound $\text{LB}_1 = L_1(\lambda_{\text{TPA}}, \mu_{\text{TPA}})$. The TPA procedure is heuristic, in the sense that it may not minimize L_1 in λ and μ .

In the first stage of TPA, we simply ignore precedence constraints altogether. For a feasible solution S ,

consider the function $g(\lambda, \mu, S)$ defined as follows:

$$\begin{aligned} g(\lambda, \mu, S) = & \sum_{j \in N} (w_j - \lambda_j) T_j + \sum_{j \in N} \lambda_j (C_j - d_j) \\ & + \sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} (C_j + p_k - C_k). \end{aligned}$$

We start with the Lagrangian problem $\hat{\text{LR}}_{P_1}$ where $\hat{L}_1(\lambda, \mu) = \min g(\lambda, \mu, S)$ subject to constraints (22)–(24), which is a relaxation of LR_{P_1} . We simply set all μ_{jk} to zero ($\mu = \mu_0 = (0, \dots, 0)$); with this choice, $\hat{L}_1(\lambda, \mu) = L_0(\lambda)$ and we set $\lambda_{\text{TPA}} = \lambda_{\text{PV}}$.

In the second stage of TPA, the multipliers μ_{jk} are adjusted assuming that $\lambda = \lambda_{\text{TPA}}$ is predefined and constant. This adjustment is an iterative heuristic; we adopt the *quick ascent direction* (QAD) algorithm proposed by van de Velde [44]. One iteration of TPA runs in $O(m + n \log n)$ time, where $m = |A|$. We have run a number of experiments to evaluate the improvement of the lower bound as a function of the number of iterations k_{max} . For a representative dataset, Table 2 shows that the average percentage deviation of LB_1 from LB_0 significantly increases in the first iterations, whereas after about five iterations the incremental improvement becomes rather limited; more information on the choices for k_{max} follows in Sections 5.2.3 and 7.2.

Theorem 3 $\text{LB}_0 \leq \text{LB}_1$.

Proof We argue that

$$\begin{aligned} \text{LB}_0 = L_0(\lambda_{\text{PV}}) = \hat{L}_1(\lambda_{\text{PV}}, \mu_0) & \leq \hat{L}_1(\lambda_{\text{TPA}}, \mu_{\text{TPA}}) \\ & \leq L_1(\lambda_{\text{TPA}}, \mu_{\text{TPA}}) = \text{LB}_1. \end{aligned}$$

The first inequality follows from Theorem 3 in [44], where it is shown that TPA generates a series of monotonically increasing lower bounds. The second inequality corresponds with Theorem 1. \square

5.2.3 Finding a VSP graph


LB_1 requires a decomposition of graph G into two subgraphs $G' = (N, A')$ and $G'' = (N, A'')$, such that $A' \cup A'' = A$ and G' is a VSP graph. The more arcs we can include in A' , the tighter the lower bound. In the following, we discuss procedures to find a VSP subgraph G' with maximum number of arcs; we refer to this problem as the *maximum VSP subgraph* (MVSP) problem.

Valdes et al. [43] state the following result:

Lemma 3 (From [43]) *A graph G is VSP if and only if its transitive closure does not contain the graph of Figure 6 as a subgraph.*

Table 2: The average percentage deviation of LB_1 from LB_0 .

n	$k_{\max} = 0$	$k_{\max} = 5$	$k_{\max} = 10$	$k_{\max} = 20$	$k_{\max} = 50$	$k_{\max} = 100$	$k_{\max} = \infty$
20	11.576	14.579	15.026	15.207	15.296	15.310	15.314
30	8.505	17.454	18.147	18.419	18.503	18.508	18.512
40	6.850	13.493	14.065	14.344	14.466	14.495	14.506

**Fig. 6:** The forbidden subgraph for VSP graphs.

Valdes et al. refer to the pattern in Figure 6 as the *forbidden subgraph*. Polynomial-time exact procedures exist for finding an ESP subgraph with maximum number of nodes (see [6], for instance), but to the best of our knowledge, no exact approach for MVSP has been proposed yet in literature. McMahan and Lim [27] suggest a heuristic traversal procedure to find and eliminate all forbidden subgraphs and, at the same time, construct a binary decomposition tree for the resulting VSP graph. Their procedure runs in $O(n + m)$ time. The number of arcs in a VSP graph is bounded by $2n - 3$ for an arbitrary non-VSP graph, but the maximum number of arcs for an arbitrary input graph is $O(n^2)$. We implement a slightly modified variant of the algorithm in [27] to compute G' ; we select arcs for removal so that the lower bound remains reasonably tight. Simultaneously, it constructs a decomposition tree for the obtained VSP graph. The time complexity of $O(n + m)$ is maintained.

The structure of our heuristic decomposition and arc-elimination procedure is described in the following lines. The procedure constructs a decomposition tree by exploiting *parallel* and *serial node reduction* [25]. Parallel reduction merges a job pair into one single job if both jobs have the same predecessor and successor set. In the decomposition tree, such jobs are linked by a P node, which means they can be processed in parallel (see Figure 7(b)). Serial reduction merges a job pair $\{i, j\}$ into one single job if arc $(i, j) \in A$, job i has only one successor and job j has only one predecessor. In the decomposition tree, such two jobs are linked by an S node, which means they cannot be processed in parallel (see Figure 7(d)). Whenever a forbidden subgraph is recognized, the procedure removes arcs such that the forbidden subgraph is resolved (removed) and the total number of removed arcs (including transitive and merged arcs) is approximately minimized (see Figures 7(b)–7(c)). Notice that some arcs may actually represent multiple merged arcs, so removing one arc in one

iteration might imply the removal of multiple arcs simultaneously in the original network G .

The proposed algorithm is run only once, in the root node of the search tree. In each other node of the search tree, graphs G' and G'' are constructed by removing from the corresponding graphs in the parent node the arcs associated with the scheduled jobs; the resulting graphs are then the input for computing LB_1 . Notice that for each child node, both graphs G' and G'' as well as the associated decomposition tree are constructed in $O(n)$ time.

To evaluate the impact of our arc elimination procedure on the quality of the bounds, we examine two variations of LB_1 , namely $LB_1(\text{VSP}) = L_1(\lambda_{\text{TPA}}, \mu_{\text{TPA}})$, where all forbidden graphs in G are resolved using the arc elimination procedure, and $LB_1(\text{NO}) = \hat{L}_1(\lambda_{\text{TPA}}, \mu_{\text{TPA}})$, in which we simply remove all arcs ($A' = \emptyset$ and $A'' = A$). Let k_{\max} be the maximum number of iterations for TPA, as explained in Section 5.2.2. Table 3 demonstrates the success of our proposed algorithm in tightening the bound. The distance between the bounds is decreasing with increasing k_{\max} , but in a B&B algorithm, a large value for k_{\max} becomes computationally prohibitive.

Theorem 4 $LB_1(\text{NO}) \leq LB_1(\text{VSP})$ for the same value of k_{\max} .

Proof $LB_1(\text{NO})$ is obtained by solving LR_{P_1} with $A' = \emptyset$ and $A'' = A$, so with the same multipliers the problem associated with $LB_1(\text{NO})$ is a relaxation of the problem associated with $LB_1(\text{VSP})$. The multipliers are updated with TPA in both cases, and will indeed be the same for a given k_{\max} , so the theorem holds. \square

5.2.4 Retrieving release dates

Bound LB_1 turns out not to be very tight when release dates are rather heterogeneous. Below, we examine two means to produce a stronger bound, namely block decomposition and job splitting.

Block decomposition We follow references [17, 31, 35] in setting up a decomposition of the job set into separate *blocks*: a *block* is a subset of jobs for which it is a dominant decision to schedule them together. We sort and renumber all jobs in non-decreasing order of their

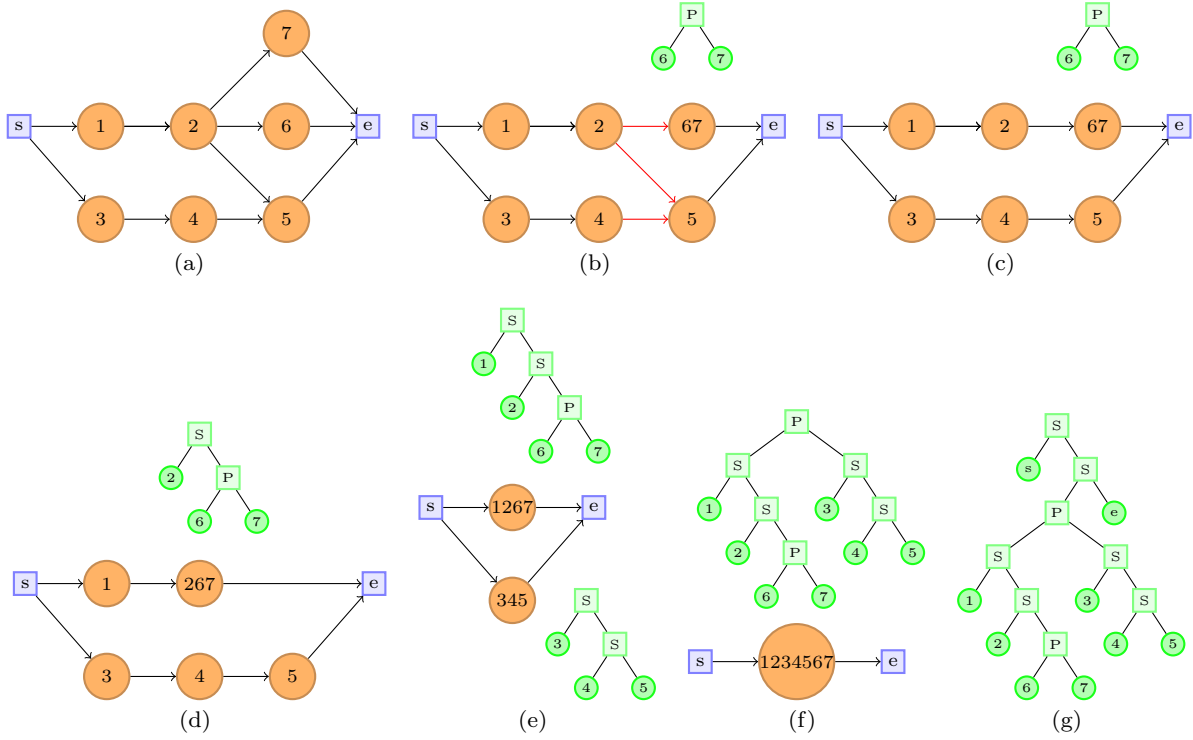


Fig. 7: Modified traversal algorithm applied to the input graph in (a).

Table 3: The average percentage deviation of LB_1 from LB_0 .

	$k_{\max} = 0$	$k_{\max} = 1$	$k_{\max} = 2$	$k_{\max} = 3$	$k_{\max} = 5$	$k_{\max} = 10$	$k_{\max} = 50$
$LB_1(\text{VSP})$	6.85	10.515	12.171	12.896	13.493	14.344	14.466
$LB_1(\text{NO})$	0	9.057	11.497	12.538	13.385	14.020	14.458

modified release dates \bar{r}_j ; as a tie-breaking criterion, we consider non-increasing order of w_j/p_j . The resulting *non-delay* sequence of jobs is given by $\sigma^r = (1, \dots, n)$, where a sequence is said to be ‘non-delay’ if the machine is never kept idle while some jobs are waiting to be processed [32]. Let $B_i = (u_i, \dots, v_i)$ be one block (in which jobs are sorted according to their new indices). The set $B = \{B_1, \dots, B_\kappa\}$ is a *valid* decomposition of the job set into κ blocks if the following conditions are satisfied:

1. $u_1 = 1$;
2. for each i, j with $1 < i \leq \kappa$ and $1 \leq j \leq n$, if $u_i = j$ then $v_{i-1} = j - 1$ and vice versa;
3. for each i, j with $1 \leq i \leq \kappa$ and $u_i \leq j \leq v_i$, we have $\bar{r}_{u_i} + \sum_{s=u_i}^{j-1} p_s \geq \bar{r}_j$.

Although the sequencing of the jobs within one block is actually still open, the sequencing of the blocks is pre-determined. Given a valid set of blocks B , we compute LB_1 for each block $B_i \in B$ separately. The value LB_2 is then the sum of the bounds per block; analogously

to [17, 31, 35], LB_2 can be shown to be a lower bound for P .

We define $LB_1^* = L_1(\lambda^*, \mu^*)$, where λ^* and μ^* are optimal choices for the Lagrangian multipliers for LB_1 , and LB_2^* , which is computed by adding the contribution $L_1(\lambda_{B_i}^*, \mu_{B_i}^*)$ for each block B_i , where $\lambda_{B_i}^*$ and $\mu_{B_i}^*$ are the optimal choices for the multipliers for block B_i .

Theorem 5 $LB_1^* \leq LB_2^*$.

Proof We introduce $g_{B_i}(\lambda, \mu, S)$ as follows:

$$g_{B_i}(\lambda, \mu, S) = \sum_{j \in B_i} (w_j - \lambda_j) T_j + \sum_{j \in B_i} \lambda_j (C_j - d_j) + \sum_{j \in B_i} \sum_{k \in Q_j} \mu_{jk} (C_j + p_k - C_k).$$

Let S_1^* be an optimal solution to LB_1^* and $S_2^* = (S_{B_1}^*, \dots, S_{B_\kappa}^*)$ an optimal solution to LB_2^* . The following re-

sult is derived.

$$\begin{aligned} \text{LB}_1^* &= g(\lambda^*, \mu^*, S_1^*) \leq g(\lambda^*, \mu^*, S_2^*) \\ &= \sum_{i=1}^{\kappa} g_{B_i}(\lambda^*, \mu^*, S_{B_i}^*) \leq \sum_{i=1}^{\kappa} g_{B_i}(\lambda_{B_i}^*, \mu_{B_i}^*, S_{B_i}^*) = \text{LB}_2^*. \end{aligned}$$

□

Although TPA might not find $\lambda_{B_i}^*$ and $\mu_{B_i}^*$ and the same result as Theorem 5 cannot be derived for LB_1 and LB_2 , experimental results show that LB_2 is on average by far tighter than LB_1 .

Job splitting It sometimes happens that the decomposition procedure fails to improve the bound (only one block is created and $\text{LB}_2 = \text{LB}_1$). Another approach is to explicitly re-introduce the release-date constraints (which have been removed previously). We define problem P_2 , which is a generalization of P_1 in which the release-date constraints (19) are included. The associated Lagrangian problem is:

$$\begin{aligned} \text{LR}_{P_2} : \quad L_2(\lambda, \mu) &= \min \sum_{j \in N} (w_j - \lambda_j) T_j \\ &+ \sum_{j \in N} C_j \left(\lambda_j + \sum_{k \in Q_j} \mu_{jk} - \sum_{k \in P_j} \mu_{kj} \right) \\ &+ \sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} p_k - \sum_{j \in N} \lambda_j d_j \end{aligned}$$

subject to constraints (19),(22)–(24).

Contrary to LR_{P_1} , we now remove the serial-parallel precedence constraints because they render the Lagrangian problem too difficult. Problem LR_{P_2} is a total-weighted-completion-times problem with release dates, where the phrase $\lambda_j + \sum_{k \in Q_j} \mu_{jk} - \sum_{k \in P_j} \mu_{kj}$ is the weight of job j and the phrase $\sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} p_k - \sum_{j \in N} \lambda_j d_j$ is a constant. This problem is known to be NP-hard [26], but a number of efficient polynomial algorithms, which are based on job splitting, have been proposed to compute tight lower bounds [7, 17, 29]. One of these algorithms is the SS procedure proposed by Belouadah et al. [7] which runs in $O(n \log n)$ time, and which we adopt here. Essentially, we again decompose the job set into a set of blocks B and compute $L_2(\lambda, \mu)$ for each block $B_i \in B$. The lower bound LB_2^{SSr} is again the sum of the contributions of the individual blocks. Experiments show that LB_2^{SSr} is typically tighter than LB_2 when the release dates are unequal. With equal release dates, on the other hand, normally $\text{LB}_2 \geq \text{LB}_2^{\text{SSr}}$ because LB_2 incorporates a part of the precedence graph. TPA is applied also here for multiplier updates.

5.2.5 Retrieving deadlines

We introduce P_2' , which is a generalization of P_1 where deadline constraints are retrieved by inclusion of the constraint set (20). The associated Lagrangian problem is:

$$\begin{aligned} \text{LR}_{P_2'} : \quad L_2'(\lambda, \mu) &= \min \sum_{j \in N} (w_j - \lambda_j) T_j \\ &+ \sum_{j \in N} C_j \left(\lambda_j + \sum_{k \in Q_j} \mu_{jk} - \sum_{k \in P_j} \mu_{kj} \right) \\ &+ \sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} p_k - \sum_{j \in N} \lambda_j d_j \end{aligned}$$

subject to constraints (20),(22)–(24).

$\text{LR}_{P_2'}$ is a total-weighted-completion-times problem with deadlines, with $\lambda_j + \sum_{k \in Q_j} \mu_{jk} - \sum_{k \in P_j} \mu_{kj}$ the weight of job j and $\sum_{j \in N} \sum_{k \in Q_j} \mu_{jk} p_k - \sum_{j \in N} \lambda_j d_j$ a constant. This problem is known to be NP-hard [26]. Posner [33] proposes a job-splitting lower bounding scheme for $\text{LR}_{P_2'}$ that uses $O(n \log n)$ time; the lower bound LB_2^{SSs} results from block decomposition and computation of $L_2'(\lambda, \mu)$ for each block. We again apply TPA for setting the multipliers.

5.2.6 Improvement by slack variables

Relaxed inequality constraints can be considered to be ‘nasty’ constraints because they decrease the quality of lower bounds. We follow Hoogeveen and van de Velde [19] in exploiting the advantages of slack variables to lessen the effect of such nasty constraints to improve the quality of the lower bounds.

We introduce two non-negative vectors of slack variables: vector $\mathbf{y} = (y_1, \dots, y_n)$ and vector $\mathbf{z} = (z_{11}, \dots, z_{1n}, \dots, z_{n1}, \dots, z_{nn})$. Consider the following sets of constraints:

$$T_j = C_j - d_j + y_j \quad \forall j \in N \quad (27)$$

$$C_j = C_i + p_j + z_{ij} \quad \forall (i, j) \in A \quad (28)$$

$$y_j, z_{ij} \geq 0 \quad \forall i, j \in N \quad (29)$$

Let problem P_3 be the variant of problem P_1 in which the sets of constraints (18) and (21) are replaced by the constraints (27)–(29). The Lagrangian problem associ-

ated with P_3 is:

$$\begin{aligned} \text{LR}_{P_3}: \quad L_3(\lambda, \mu) = & \min \sum_{j=1}^n (w_j - \lambda_j) T_j + \sum_{j=1}^n \lambda_j y_j + \\ & \sum_{j=1}^n \sum_{k \in \mathcal{Q}_j} \mu_{jk} z_{jk} + \sum_{j=1}^n C_j \left(\lambda_j + \sum_{k \in \mathcal{Q}_j} \mu_{jk} - \sum_{k \in \mathcal{P}_j} \mu_{kj} \right) \\ & + \sum_{j=1}^n \sum_{k \in \mathcal{Q}_j} \mu_{jk} p_k - \sum_{j=1}^n \lambda_j d_j \end{aligned}$$

subject to constraints (22)–(25) and (29).

The values of the variables T_j , y_j and z_{jk} are zero in any optimal solution to LR_{P_3} because for $i, j \in N$ the following inequalities hold: $0 \leq \lambda_j \leq w_j$ and $\mu_{jk} \geq 0$. In an optimal solution to P_3 , however, these values might not be zero. In fact, according to the set of constraints (27), unless $C_j = d_j$, either T_j or y_j is nonzero. Also, from constraints (28), z_{jk} may not be zero when job j has at least two successors or job k has at least two predecessors in G . We introduce three problems that each carry a part of the objective function of LR_{P_3} , one of which is LR_{P_1} and the other two are the following two slack-variable (SV) problems, where Y is the set of all \mathbf{y} -vectors corresponding to feasible solutions to P_3 and Z similarly contains all \mathbf{z} -vectors.

$$\begin{aligned} \text{P}_{\text{SV1}}: \quad \text{SV}_1(\lambda) = & \min \sum_{j=1}^n (w_j - \lambda_j) T_j + \sum_{j=1}^n \lambda_j y_j \\ & \text{subject to constraints (22),(23),(25) and } \mathbf{y} \in Y; \end{aligned}$$

$$\begin{aligned} \text{P}_{\text{SV2}}: \quad \text{SV}_2(\mu) = & \min \sum_{j=1}^n \sum_{k \in \mathcal{Q}_j} \mu_{jk} z_{jk} \\ & \text{subject to constraint } \mathbf{z} \in Z. \end{aligned}$$

Note that the term $\sum_{j=1}^n (w_j - \lambda_j) T_j$ appears in two of the problems, but it will be set to zero anyway in LR_{P_1} .

Hoogeveen and van de Velde [19] propose $O(n \log n)$ -time procedures to compute valid lower bounds for P_{SV1} and P_{SV2} . Let $\text{LB}_{\text{SV1}} \geq 0$ and $\text{LB}_{\text{SV2}} \geq 0$ be lower bounds for P_{SV1} and P_{SV2} , respectively. By adding LB_{SV1} and LB_{SV2} to LB_2 , a better lower bound LB_3 for P is obtained [19]. The same SV problems can also be constructed for $\text{LB}_2^{\text{SS}_r}$ and $\text{LB}_2^{\text{SS}_s}$ to lead to bounds $\text{LB}_3^{\text{SS}_r} = \text{LB}_2^{\text{SS}_r} + \text{LB}_{\text{SV1}} + \text{LB}_{\text{SV2}}$ and $\text{LB}_3^{\text{SS}_s} = \text{LB}_2^{\text{SS}_s} + \text{LB}_{\text{SV1}} + \text{LB}_{\text{SV2}}$. We have the following result:

Observation 1 $\text{LB}_2 \leq \text{LB}_3$, $\text{LB}_2^{\text{SS}_r} \leq \text{LB}_3^{\text{SS}_r}$ and $\text{LB}_2^{\text{SS}_s} \leq \text{LB}_3^{\text{SS}_s}$.

6 Dominance properties

Our search procedure also incorporates a number of dominance rules, which will be described in this section. We will use the following additional notation. Given two partial sequences $\pi = (\pi_1, \dots, \pi_k)$ and $\pi' = (\pi'_1, \dots, \pi'_{k'})$, we define a merge operator as follows: $\pi | \pi' = (\pi_1, \dots, \pi_k, \pi'_1, \dots, \pi'_{k'})$. If π' contains only one job j then we can also write $\pi | j = (\pi_1, \dots, \pi_k, j)$, and similarly if $\pi = (j)$ then $j | \pi' = (j, \pi'_1, \dots, \pi'_{k'})$.

6.1 General dominance rules

We use the lower bounds proposed in Section 5 to prune the search tree. Let $\text{LB}(U)$ represent any of the lower bounds described in Section 5, applied to the set U of unscheduled jobs, and let S_{best} be the currently best known feasible solution. The following dominance rule is then immediate:

Dominance rule 1 (DR₁) Consider a node associated with selection S_P . If

$$\text{TWT}(S_P) + \text{LB}(U) \geq \text{TWT}(S_{\text{best}}),$$

then the node associated with S_P can be fathomed.

As we already introduced in Section 4, a partial schedule can be denoted by either S_P or (σ_B, σ_E) . Multiple lower bounds can be used to fathom a node. The selection of lower bounds and the order in which they are computed, obviously influences the performance of the B&B algorithm. These issues are examined in Section 7.2.

The subset of *active schedules* is dominant for total weighted tardiness problems [10, 32]. A feasible schedule is called *active* if it is not possible to construct another schedule by changing the sequence of jobs such that at least one job is finishing earlier and no other job finishes later. The dominance of active schedules holds even when deadlines and precedence constraints are given.

Dominance rule 2 (DR₂) Consider a node associated with (σ_B, \emptyset) that is selected for forward branching, and let j be a job belonging to E_B . If $\bar{r}_j \geq \min_{k \in E_B} \{\bar{r}_k + p_k\}$, then the child node associated with the schedule $(\sigma_B | j, \emptyset)$ can be fathomed.

We also prune a branch whenever an obvious violation of the deadline constraints is detected. A partial schedule associated with a particular node is not always extended to a feasible schedule. Scheduling a job in one particular position may force other jobs to violate their deadline constraints, even though it does not violate its

own constraints. Let \mathcal{A} be an arbitrary subset of U and let $\Pi_{\mathcal{A}}$ be the set of all possible permutations of jobs in \mathcal{A} . The following theorem states when a job is scheduled in a ‘wrong position’, meaning that it will lead to a violation of deadline constraints.

Theorem 6 *Consider a partial schedule (σ_B, σ_E) . If there exists any non-empty subset $\mathcal{A} \subset U$ such that the inequality $\min_{\pi \in \Pi_{\mathcal{A}}} \{C_{\max}(\sigma_B | \pi)\} > \max_{j \in \mathcal{A}} \{\bar{\delta}_j\}$ holds, then the schedule (σ_B, σ_E) is not feasible.*

Proof If $\min_{\pi \in \Pi_{\mathcal{A}}} \{C_{\max}(\sigma_B | \pi)\} > \max_{j \in \mathcal{A}} \{\bar{\delta}_j\}$ then at least one job in \mathcal{A} cannot be scheduled before its deadline and the schedule (σ_B, σ_E) is not feasible. \square

The problem $\min_{\pi \in \Pi_{\mathcal{A}}} \{C_{\max}(\sigma_B | \pi)\}$, which equates with $1|r_j, \delta_j, prec|C_{\max}$, is NP-hard because the mere verification of the existence of a feasible schedule is already NP-complete. We remove deadlines and create a new problem whose optimal solution is computed in $O(n^2)$ [23]. For computational efficiency, we use a linear-time lower bound for this new problem. This lower bound is computed as follows: $\min_{j \in \mathcal{A} \cap E_B} \{\bar{r}_j\} + \sum_{j \in \mathcal{A}} p_j$.

Dominance rule 3 (DR₃) *The node associated with (σ_B, σ_E) can be eliminated if at least one of the following conditions is satisfied:*

1. if $\sigma_E = \emptyset$ and the condition of Theorem 6 is satisfied for the partial schedule (σ_B, \emptyset) ;
2. if $\sigma_E \neq \emptyset$ and $\max_{j \in U} \{\bar{\delta}_j\} < st(\sigma_E)$.

6.2 Dominance rule based on two-job interchange

We describe a dominance rule based on job interchange. This dominance rule consists of two parts. The first part deals with the interchange of jobs in a FB node whereas the second part deals with the interchange of jobs in a BB node.

6.2.1 Interchanging jobs in a FB node

In an FB node, consider jobs $j, k \in E_B$ that are not identical (they differ in at least one of their parameters). We will always assume that $\hat{r}_k < \hat{r}_j + p_j$ and $\hat{r}_j < \hat{r}_k + p_k$, because otherwise Dominance rule 2 enforces the scheduling of the job with smaller \hat{r} before the job with larger \hat{r} ; note here that $\hat{r}_j = \bar{r}_j$ and $\hat{r}_k = \bar{r}_k$ because all predecessors of jobs j and k has already been scheduled and therefore the branching decisions cover the propagation of precedence constraints. We also assume that any successor of job k is also a successor of job j ($\mathcal{Q}_k \subset \mathcal{Q}_j$). Consider a node of the search tree in which job k is scheduled at or after the completion of the sequence σ_B . Suppose that the partial schedule




Fig. 8: Schedules S_1 and S'_1 .

associated to the current node can be extended to a feasible schedule S_1 in which job j is scheduled somewhere after job k . We define a set $\mathcal{B} = U \setminus \{j, k\}$ of jobs. We also construct a schedule S'_1 by interchanging jobs j and k while the order of jobs belonging to \mathcal{B} remains unchanged. Figure 8 illustrates schedules S_1 and S'_1 .

To prove that interchanging jobs j and k does not increase the total weighted tardiness, we argue that the gain of interchanging jobs j and k , which is computed as $TWT(S_1) - TWT(S'_1)$, is greater than or equal to zero, no matter when job j is scheduled. Let $st_j(S)$ denote the start time of job j in schedule S . Remember that $st(\pi)$ denotes the start time of a sequence π . Let τ_1 be the difference between the start time of job j in S_1 and the start time of k in S'_1 . If $st_k(S'_1)$ is less than $st_j(S_1)$ then τ_1 is negative, otherwise it is non-negative. By interchanging jobs j and k each job that belongs to set \mathcal{B} may be shifted either to the right or to the left. Let $\tau_2 \geq 0$ be the maximum shift to the right of the jobs belonging to set \mathcal{B} . Notice that if all jobs in \mathcal{B} are shifted to the left, then $\tau_2 = 0$. For each t as the start time of job j in S_1 , Jouglet et al. [21] define a function $\Gamma_{jk}(t, \tau_1, \tau_2)$ as follows:

$$\begin{aligned} \Gamma_{jk}(t, \tau_1, \tau_2) = & w_j \max\{0, t + p_j - d_j\} \\ & - w_k \max\{0, t + \tau_1 + p_k - d_k\} \\ & + w_k \max\{0, \hat{r}_k + p_k - d_k\} \\ & - w_j \max\{0, \hat{r}_j + p_j - d_j\} - \tau_2 \sum_{i \in \mathcal{B}} w_i. \end{aligned}$$

For the sub-problem of P where precedence and deadline constraints are removed, Jouglet et al. [21] show that $\Gamma_{jk}(t, \tau_1, \tau_2)$ is a lower bound for the gain of interchanging jobs j and k when $t = st_j(S_1)$. This result can be improved by adding the gain of shifting the jobs which are tardy in both schedules S_1 and S'_1 . We introduce the set \mathcal{B}' of jobs where each job $i \in \mathcal{B}'$ is certainly a tardy job in S'_1 . Let $\hat{\mathcal{P}}_i$ be the set of transitive predecessors of job i . The following set of jobs, which is a subset of \mathcal{B}' , is used in our implementations because the order based on which the jobs in \mathcal{B} are scheduled has not yet been defined and therefore computing \mathcal{B}' is

not possible:

$$\left\{ i \in \mathcal{B} \left| \hat{r}_j + p_j + \sum_{l \in (\mathcal{B} \cap \hat{\mathcal{P}}_i)} p_l + p_i \geq d_i \right. \right\}.$$

Let $\tau'_2 \geq 0$ be the minimum shift to the left of the jobs belonging to set \mathcal{B} . Note that at least one of the values τ'_2 and τ_2 equals zero. We define the function $\hat{\Gamma}_{jk}(t, \tau_1, \tau_2, \tau'_2)$ as follows:

$$\hat{\Gamma}_{jk}(t, \tau_1, \tau_2, \tau'_2) = \Gamma_{jk}(t, \tau_1, \tau_2) + \tau'_2 \sum_{i \in \mathcal{B}'} w_i.$$

The values τ_2 and τ'_2 cannot be negative. Therefore, we immediately infer $\Gamma_{jk}(t, \tau_1, \tau_2) \leq \hat{\Gamma}_{jk}(t, \tau_1, \tau_2, \tau'_2)$. We need the following result:

Theorem 7 $\hat{\Gamma}_{jk}(t, \tau_1, \tau_2, \tau'_2)$ is a valid lower bound for the gain of interchanging jobs j and k .

Proof If $\tau'_2 = 0$, then $\Gamma_{jk}(t, \tau_1, \tau_2) = \hat{\Gamma}_{jk}(t, \tau_1, \tau_2, \tau'_2)$ and the theorem holds based on [21]. If $\tau'_2 > 0$, all jobs in \mathcal{B} are shifted to the left at least τ'_2 units. Also, τ_2 equals zero because no job is shifted to the right. For all jobs $i \in \mathcal{B}'$ we have $C_i(S_1) \geq C_i(S'_1) \geq d_i$. Consequently, $\tau'_2 \sum_{i \in \mathcal{B}'} w_i \geq 0$ is a lower bound for the gain of rescheduling jobs in \mathcal{B} . The value $w_j \max\{0, t + p_j - d_j\} - w_j \max\{0, \hat{r}_j + p_j - d_j\}$ equals the gain of rescheduling job j and the value $w_k \max\{0, \hat{r}_k + p_k - d_k\} - w_k \max\{0, t + \tau_1 + p_k - d_k\}$ equals the gain of rescheduling job k . By adding lower bounds for rescheduling gains of all jobs in $U = \mathcal{B} \cup \{j, k\}$, a lower bound for the gain of interchanging jobs j and k is obtained. \square

In a general setting (problem P), however, job interchanges are not always feasible for every starting time t . We opt for verifying the feasibility of an interchange by ensuring that it does not cause any violation of deadlines and/or precedence constraints for all possible $t = st_j(S_1)$. Let Ψ be an upper bound for the completion time of the sequence S'_1 , computed as follows:

$$\Psi = \max \left\{ \max\{\hat{r}_j + p_j, \hat{r}_k\} + p_k, \max_{i \in \mathcal{B}} \{\hat{r}_i\} \right\} + \sum_{i \in \mathcal{B}} p_i.$$

The following theorem provides the conditions under which for every possible $t = st_j(S_1)$ interchanging jobs j and k is feasible.

Theorem 8 For each feasible schedule S_1 , an alternative feasible schedule S'_1 is created by interchanging jobs j and k , if the following conditions are satisfied:

1. $\bar{\delta}_j - p_j \leq \bar{\delta}_k - \tau_1 - p_k$ or $\Psi \leq \hat{\delta}_k$;
2. $\tau_2 = 0$ or $\Psi \leq \min_{i \in \mathcal{B}} \{\hat{\delta}_i\}$.

Proof We examine under which conditions the jobs belonging to the set $U = \mathcal{B} \cup \{j, k\}$ do not violate any of their deadlines and/or precedence constraints. Precedence constraints are not violated because jobs $j, k \in E_B$ are deliberately chosen such that $\mathcal{Q}_k \cap \mathcal{Q}_j = \mathcal{Q}_k$ and job j does not violate its deadline simply because $C_j(S'_1) \leq C_j(S_1) \leq \bar{\delta}_j$.

Condition 1 ensures that job k does not violate its deadline. We argue that $t = st_j(S_1) \leq \bar{\delta}_j - p_j$. If $\bar{\delta}_j - p_j \leq \bar{\delta}_k - \tau_1 - p_k$ holds, then we infer $C_k(S'_1) = t + \tau_1 + p_k \leq \bar{\delta}_k$. Also, if $\Psi \leq \hat{\delta}_k$, then all unscheduled jobs including j and k are completed at or before $\hat{\delta}_k$. Note that $\hat{\delta}_k$ is preferred over $\bar{\delta}_k$ because $\bar{\delta}_k \leq \hat{\delta}_k$, thus condition 1 is less violated, and the inequality $\Psi \leq \hat{\delta}_k$ also implies $C_k(S'_1) \leq \bar{\delta}_k$.

Condition 2 verifies that no job in \mathcal{B} violates its deadline. On the one hand, if $\tau_2 = 0$, then no job in \mathcal{B} is shifted to the right, which means $C_i(S'_1) \leq C_i(S_1) \leq \bar{\delta}_i$ for each job $i \in \mathcal{B}$. On the other hand, if $\tau_2 > 0$ and $\Psi \leq \min_{i \in \mathcal{B}} \{\hat{\delta}_i\}$, then for all jobs $i \in \mathcal{B}$ we conclude: $C_i(S'_1) \leq \Psi \leq \min_{i \in \mathcal{B}} \{\hat{\delta}_i\} \leq \hat{\delta}_i$. Again, $\hat{\delta}_i$ is preferred over $\bar{\delta}_i$ because of the same reasoning as for the preference of $\hat{\delta}_k$ over $\bar{\delta}_k$. \square

Jouglet et al. [21] prove that if $w_j \geq w_k$ then the value $\Gamma_{jk}(\max\{d_j - p_j, \hat{r}_k + p_k\}, \tau_1, \tau_2)$ is the minimum gain obtained by interchanging jobs j and k for the setting where deadlines and precedence constraints are removed. We derive a more general result using the following lemma.

Lemma 4 Let $f : t \rightarrow \alpha \max\{0, t - a\} - \beta \max\{0, t - b\} + C$ be a function defined on $[u, v]$ for $a, b, C \in \mathbb{R}$ and $\alpha, \beta, u, v \in \mathbb{R}^+$. The function f reaches a global minimum at value t^* computed as follows:

$$t^*(\alpha, \beta, a, b, u, v) =$$

$$\begin{cases} \min\{\bar{u}, v\} & \text{if } \alpha \geq \beta \\ u & \text{if } \alpha < \beta, b > a, \alpha(\bar{v} - \bar{u}) \geq \beta(\bar{v} - b) \\ v & \text{otherwise} \end{cases}$$

where $\bar{u} = \max\{u, a\}$ and $\bar{v} = \max\{v, b\}$.

Proof See Appendix. \square

Theorem 9 below follows from Theorem 7, Theorem 8 and Lemma 4, if we choose $\alpha = w_j$, $\beta = w_k$, $a = d_j - p_j$, $b = d_k - \tau_1 - p_k$, $u = \hat{r}_k + p_k$, $v = \delta_j - p_j$ and $C = w_k \max\{0, \hat{r}_k + p_k - d_k\} - w_j \max\{0, \hat{r}_j + p_j - d_j\} - \tau_2 \sum_{i \in \mathcal{B}} w_i + \tau'_2 \sum_{i \in \mathcal{B}'} w_i$. Let st_j^* be computed as follows:

$$st_j^* = t^*(w_j, w_k, d_j - p_j, d_k - \tau_1 - p_k, \hat{r}_k + p_k, \delta_j - p_j).$$

Theorem 9 $\Gamma_{jk}^*(\tau_1, \tau_2, \tau'_2) = \hat{\Gamma}_{jk}(st_j^*, \tau_1, \tau_2, \tau'_2)$ is the minimum gain obtained by interchanging jobs j and k , provided that for every possible $st_j(S_1)$ interchanging jobs j and k is feasible.

To compute $\Gamma_{jk}^*(\tau_1, \tau_2, \tau'_2)$, the values of τ_1 , τ_2 and τ'_2 must be known. We establish an exhaustive list of cases for which τ_1 , τ_2 and τ'_2 can be computed, which is summarized in Table 4. Given a particular case, the values τ_1 , τ_2 and τ'_2 are computed as follows:

$$\tau_1 = \begin{cases} 0 & \text{Cases 1,5} \\ \max_{i \in U} \{\hat{r}_i\} - \hat{r}_k - p_k & \text{Case 2} \\ \max\{\hat{r}_j + p_j, \max_{i \in \mathcal{B}} \{\hat{r}_i\}\} - \hat{r}_k - p_k & \text{Cases 3,4,6} \\ \hat{r}_j + p_j - \hat{r}_k - p_k & \text{Cases 7,8} \end{cases}$$

$$\tau_2 = \begin{cases} p_k - p_j & \text{Case 1} \\ \max_{i \in U} \{\hat{r}_i\} - \hat{r}_k - p_j & \text{Case 2} \\ 0 & \text{Cases 3,5,6} \\ \max\{\hat{r}_j + p_j, \max_{i \in \mathcal{B}} \{\hat{r}_i\}\} - \hat{r}_k - p_j & \text{Case 4} \\ \hat{r}_j - \hat{r}_k & \text{Case 7} \\ \hat{r}_j + p_j - \hat{r}_k - p_k & \text{Case 8} \end{cases}$$

$$\tau'_2 = \begin{cases} 0 & \text{Cases 1,2,4,5,7,8} \\ \hat{r}_k - \hat{r}_j & \text{Case 3} \\ \hat{r}_k + p_k - \max\{\hat{r}_j + p_j, \max_{i \in \mathcal{B}} \{\hat{r}_i\}\} & \text{Case 6} \end{cases}$$

Following the above results, the first part of Dominance rule 4 is derived.

Dominance rule 4 (DR₄; first part) *Given an FB node associated with (σ_B, \emptyset) , if there exist two non-identical jobs $j, k \in E_B$ with $\mathcal{Q}_k \cap \mathcal{Q}_j = \mathcal{Q}_k$ and the inequality $\Gamma_{jk}^*(\tau_1, \tau_2, \tau'_2) > 0$ holds, then $(\sigma_B|j, \emptyset)$ dominates $(\sigma_B|k, \emptyset)$.*

6.2.2 Interchanging jobs in a BB node

Let $j, k \in E_E$ where jobs j and k are not identical. We also assume that any unscheduled predecessor of job k is also a predecessor of job j . In other words, we have $\mathcal{P}_k \cap \mathcal{P}_j \cap U = \mathcal{P}_k \cap U$. Consider a BB node of the search tree with decision job k . The partial schedule associated with the current node can be extended to a feasible schedule S_2 in which job j is scheduled before job k but after all jobs in the sequence σ_B . The set \mathcal{B} is the set of all remaining unscheduled jobs where $\mathcal{B} = U \setminus \{j, k\}$. Let schedule S'_2 be constructed by interchanging jobs j and k while keeping the order based on which the jobs belonging to \mathcal{B} will be scheduled. Figure 9 illustrates schedules S_2 and S'_2 .



Fig. 9: Schedules S_2 and S'_2 .

For each t as the start time of job j in S_2 , we define a function $\Delta_{jk}(t)$ as follows:

$$\Delta_{jk}(t) = w_j \max\{0, t + p_j - d_j\} - w_k \max\{0, t + p_k - d_k\} + w_k \max\{0, st(\sigma_E) - d_k\} - w_j \max\{0, st(\sigma_E) - d_j\} - \max\{0, p_k - p_j\} \sum_{i \in \mathcal{B}} w_i.$$

In a BB node, for each t as the start time of job j , $\Delta_{jk}(t)$ is a lower bound of the gain of interchanging jobs k and j , if the conditions of Theorem 10 are satisfied. Theorem 10 provides the conditions on which for every possible $t = st_j(S_1)$ interchanging jobs j and k is feasible.

Theorem 10 *For each feasible schedule S_2 , a feasible schedule S'_2 can be created by interchanging jobs j and k , if the following conditions are satisfied:*

1. $st(\sigma_E) \leq \hat{\delta}_j$;
2. $p_k - p_j \leq 0$ or $st(\sigma_E) - p_j \leq \min_{i \in \mathcal{B}} \hat{\delta}_i$.

Proof Similar to the proof of Theorem 8. \square

Theorem 11 follows from Theorem 10 and Lemma 4, if we choose $\alpha = w_j$, $\beta = w_k$, $a = d_j - p_j$, $b = d_k - p_k$, $u = C_{\max}(\sigma_B)$, $v = st(\sigma_E) - p_k - p_j$ and $C = w_k \max\{0, st(\sigma_E) - d_k\} - w_j \max\{0, st(\sigma_E) - d_j\} - \max\{0, p_k - p_j\} \sum_{i \in \mathcal{B}} w_i$. Let $st_j^{*'}$ be computed as follows:

$$st_j^{*'} = t^*(w_j, w_k, d_j - p_j, d_k - p_k, C_{\max}(\sigma_B) + \sum_{i \in \mathcal{P}_j \cap U} p_i, st(\sigma_E) - p_k - p_j).$$

Theorem 11 $\Delta_{jk}^* = \Delta_{jk}(st_j^{*'})$ is the minimum gain obtained by interchanging jobs j and k , provided that for every possible $t = st_j(S_1)$ interchanging jobs j and k is feasible.

Following the above results, the second part of Dominance rule 4 is derived.

Dominance rule 4 (DR₄; second part) *Given a BB node associated with (σ_B, σ_E) , if there exist two non-identical jobs $j, k \in E_E$ with $\mathcal{P}_k \cap \mathcal{P}_j \cap U = \mathcal{P}_k \cap U$ and $\Delta_{jk}^* > 0$, then $(\sigma_B, j|\sigma_E)$ dominates $(\sigma_B, k|\sigma_E)$.*

Case	$(\hat{r}_j + p_j - \hat{r}_k - p_k)$	$(p_j - p_k)$	$(\max_{i \in U} \{\hat{r}_i\} - \hat{r}_k - p_k)$	$(\hat{r}_j - \hat{r}_k)$	$(\max_{i \in U} \{\hat{r}_i\} - \hat{r}_k - p_j)$
1	≤ 0	< 0	≥ 0	-	-
2	≤ 0	< 0	< 0	≤ 0	> 0
3	≤ 0	< 0	< 0	≤ 0	≤ 0
4	≤ 0	< 0	< 0	> 0	-
5	≤ 0	≥ 0	≥ 0	-	-
6	≤ 0	≥ 0	< 0	-	-
7	> 0	< 0	-	-	-
8	> 0	≥ 0	-	-	-

Table 4: Interchange cases.

Fig. 10: Schedule S''_1 .

6.3 Dominance rule based on job insertion

We describe a dominance rule based on job insertion. This dominance rule, similar to the dominance rule based on job interchange, consists of two parts. The first part deals with the insertion of a job in an FB node whereas the second part deals with the insertion of a job in a BB node.

6.3.1 Inserting a job in an FB node

In a FB node, let $j, k \in E_B$ where jobs j and k are not identical. Again we assume that $\hat{r}_k < \hat{r}_j + p_j$ and $\hat{r}_j < \hat{r}_k + p_k$, otherwise Dominance rule 2 enforces scheduling the job with smaller \hat{r} before the job with larger \hat{r} (remind that $\hat{r}_j = \bar{r}_j$ and $\hat{r}_k = \bar{r}_k$ because all predecessors of jobs j and k have already been scheduled and therefore the branching decisions cover precedence constraints propagation). Consider an FB node of the search tree in which job k is scheduled after the jobs in sequence σ_B . Assume that the partial schedule associated with the current node can be extended to the feasible schedule S_1 depicted in Figure 8. We construct a schedule S''_1 by inserting the job j before job k while keeping the order of jobs belonging to \mathcal{B} . Figure 10 illustrates the construction of the schedule S''_1 .

Let τ_3 be the maximum shift to the right of the jobs belonging to \mathcal{B} , which is computed as follows:

$$\tau_3 = \max \left\{ 0, \hat{r}_j + p_j + p_k - \max \left\{ \hat{r}_k + p_k, \min_{i \in \mathcal{B}} \{\bar{r}_i\} \right\} \right\}.$$

For each t as the start time of job j in schedule S_1 , we define a function $\Gamma'_{jk}(t, \tau_3)$ as follows:

$$\begin{aligned} \Gamma'_{jk}(t, \tau_3) &= w_j \max\{0, t + p_j - d_j\} \\ &\quad - w_k \max\{0, \hat{r}_j + p_j + p_k - d_k\} \\ &\quad + w_k \max\{0, \hat{r}_k + p_k - d_k\} \\ &\quad - w_j \max\{0, \hat{r}_j + p_j - d_j\} - \tau_3 \sum_{i \in OJ} w_i. \end{aligned}$$

Job insertion, similar to job interchange, is not always feasible for every starting time t of job j . We verify feasibility of an insertion by ensuring that it does not cause any deadline and/or precedence-constraint violation for all possible $t = st_j(S_1)$. Let Ψ' be an upper bound for the completion time of the sequence S'_1 , computed as follows:

$$\Psi' = \max \left\{ \hat{r}_j + p_j + p_k, \max_{i \in \mathcal{B}} \{\hat{r}_i\} \right\} + \sum_{i \in \mathcal{B}} p_i.$$

The following theorem provides the conditions under which for every possible $t = st_j(S_1)$ inserting job j before job k is feasible.

Theorem 12 For each feasible schedule S_1 , another feasible schedule S''_1 can be created by inserting job j before job k if the following conditions hold:

1. $\hat{r}_j + p_j + p_k \leq \hat{\delta}_k$;
2. $\tau_3 = 0$ or $\Psi' \leq \min_{i \in \mathcal{B}} \{\hat{\delta}_i\}$.

Proof Similar to the proof of Theorem 8. \square

Theorem 13 below follows from Theorem 12.

Theorem 13 $\Gamma'_{jk}(\tau_3) = \Gamma'_{jk}(\hat{r}_k + p_k, \tau_3) = \Gamma_{jk}(\hat{r}_k + p_k, \hat{r}_j + p_j - \hat{r}_k - p_k, \tau_3)$ is the minimum gain obtained by inserting job j before job k provided that for every possible $t = st_j(S_1)$ inserting job j before job k is feasible.

Following the above results, the first part of Dominance rule 5 is derived.

Dominance rule 5 (DR₅; first part) Consider an FB node associated with (σ_B, \emptyset) . If there exist two non-identical jobs $j, k \in E_B$ for which the inequality $\Gamma'_{jk}(\tau_3) > 0$ holds, then $(\sigma_B|j, \emptyset)$ dominates $(\sigma_B|k, \emptyset)$.




Fig. 11: Schedule S_2'' .

6.3.2 Inserting a job in a BB node

In a BB node, let $j, k \in E_E$ where jobs j and k are not identical. Consider a node of the search tree in which job k is scheduled before sequence σ_E . Assume that the partial schedule associated with the current node can be extended to the feasible schedule S_2 depicted in Figure 9. We also construct a schedule S_2'' by inserting the job j to be scheduled after job k but before the jobs in the sequence σ_E and by keeping the order of jobs belonging to \mathcal{B} . Figure 11 illustrates schedule S_2'' .

For each t , which is the start time of job j in schedule S_2 , we define the function $\Delta'_{jk}(t)$ as follows:

$$\begin{aligned} \Delta'_{jk}(t) = & w_j \max\{0, t + p_j - d_j\} \\ & - w_k \max\{0, st(\sigma_E) - p_j - d_k\} \\ & + w_k \max\{0, st(\sigma_E) - d_k\} \\ & - w_j \max\{0, st(\sigma_E) - d_j\}. \end{aligned}$$

Similarly to the previous results, for each feasible schedule S_2 , a feasible schedule S_2'' is constructed by inserting jobs j after job k , if $st(\sigma_E) \leq \hat{\delta}_j$. The following theorem is obtained:

Theorem 14 $\Delta'_{jk}^* = \Delta'_{jk}(C_{\max}(\sigma_B) + \sum_{i \in \mathcal{P}_j \cap U} p_i)$ is the minimum gain obtained by inserting job j after job k provided that $st(\sigma_E) \leq \hat{\delta}_j$.

Following the above results, the second part of Dominance rule 5 is derived.

Dominance rule 5 (DR₅; second part) Consider a BB node associated with (σ_B, σ_E) . If there exist two non-identical jobs $j, k \in E_E$ for which the inequality $\Delta'_{jk}^* > 0$ holds, then $(\sigma_B, j|\sigma_E)$ dominates $(\sigma_B, k|\sigma_E)$.

6.4 Dominance rules on scheduled jobs

The *dominance theorem of dynamic programming* (see Jouglet et al. [21]) is another existing theorem that can be used to eliminate nodes in the search tree. It compares two partial sequences that contain identical subsets of jobs and eliminates the one having the larger total weighted tardiness. When total weighted tardiness values are the same, then only one of the sequences is kept. Let us consider two feasible partial sequences σ_1 and σ_2 (σ_2 is a feasible permutation of σ_1) of k jobs,

where $k < n$. Let \mathcal{C} be the set of jobs in either σ_1 or σ_2 . We are going to decide whether it is advantageous to replace σ_2 by σ_1 in all (partial) schedules in which σ_2 orders the last k jobs. The set of scheduled jobs and the set of unscheduled jobs are identical for both σ_1 and σ_2 . Sequence σ_1 is as good as sequence σ_2 if it fulfills one of the following conditions:

1. $C_{\max}(\sigma_1) \leq C_{\max}(\sigma_2)$ and $\text{TWT}(\sigma_1) \leq \text{TWT}(\sigma_2)$;
2. $C_{\max}(\sigma_1) > C_{\max}(\sigma_2)$ and the following inequality also holds:

$$\text{TWT}(\sigma_1) + (\min_{i \in U} \{\bar{r}_i^{\sigma_1}\} - \min_{i \in U} \{\bar{r}_i^{\sigma_2}\}) \sum_{i \in U} w_i \leq \text{TWT}(\sigma_2),$$

where $\bar{r}_i^{\sigma_1}$ is the updated release date associated with the sequence σ_1 and $\bar{r}_i^{\sigma_2}$ is the updated release date associated with the sequence σ_2 .

Jouglet et al. [21] determine the sequences that can be replaced by a dominant permutation. They find that sequence σ_1 dominates sequence σ_2 if the following two conditions hold:

1. sequence σ_1 is as good as sequence σ_2 ;
2. sequence σ_2 is not as good as σ_1 or σ_1 has lexicographically smaller release dates than σ_2 .

Note that the second condition enforces a tie-breaking rule where a lexicographical number associated to each sequence is computed and among those sequences that are equivalent, the one with lower lexicographic number is selected. To avoid conflicts with Dominance rule 2, jobs are renumbered in non-decreasing order of their release dates r_j .

Dominance rule 6 (DR₆) If there exists a better feasible permutation of σ_B and/or a better feasible permutation of σ_E , then the node (σ_B, σ_E) is fathomed.

If $\sigma_E = \emptyset$ and there is a better feasible permutation of σ_B , then the dominance is proven similarly to Theorem 13.6 in [21]. If $\sigma_E \neq \emptyset$, then all jobs belonging to the set U will be scheduled between $C_{\max}(\sigma_B)$ and $st(\sigma_E) = C_{\max}(\sigma_B) + \sum_{j \in U} p_j$. Therefore, all permutations of σ_E start at time $st(\sigma_E)$ and if there exists at least one better feasible permutation of σ_E , then fathoming the node associated with (σ_B, σ_E) does not eliminate the optimal solution.

Dominance rule 6 where only permutations of the last k jobs are considered, is referred to as DR₆^k. Computing DR₆ⁿ amounts to enumerating all $O(n!)$ feasible solutions, which would yield an optimal solution but is computationally prohibitive. In our B&B algorithm, we therefore choose $k < n$. There is a trade-off between the computational effort needed to compute DR₆^k and the improvement achieved by eliminating dominated nodes.

Based on initial experiments, we observe that the algorithms perform worse when $k > 4$ and we therefore only implement DR_6^2 , DR_6^3 and DR_6^4 .

7 Computational results

All algorithms have been implemented in VC++ 2010, while CPLEX 12.3 is used to solve the MIP formulations. All computational results were obtained on a laptop Dell Latitude with 2.6 GHz Core(TM) i7-3720QM processor, 8GB of RAM, running under Windows 7.

7.1 Instance generation

To the best of our knowledge, there are no benchmark sets of instances of problem P available, and so we have generated our own instances with $|N| = n = 10, 20, 30, 40$ and 50 jobs. Two different instance sets are generated: Ins^S and Ins^L . The former set contains instances with small processing times, the latter holds instances with large processing times: the values p_i ($1 \leq i \leq n$) are sampled from the uniform distribution $U[1, \alpha]$, where $\alpha = 10$ for Ins^S and $\alpha = 100$ for Ins^L . Release dates r_i are drawn from $U[0, \tau P]$, where $P = \sum_{i \in N} p_i$ and $\tau \in \{0.0, 0.5, 1.0\}$. Due dates d_i are generated from $U[r_i + p_i, r_i + p_i + \rho P]$ with $\rho \in \{0.05, 0.25, 0.50\}$ and weights w_i stem from $U[1, 10]$. Deadlines are chosen from $U[d_i, d_i + \phi P]$ with $\phi \in \{1.00, 1.25, 1.50\}$. The precedence graph is constructed using RanGen [11] with $OS \in \{0.00, 0.25, 0.50, 0.75\}$, where OS is the *order strength* of the graph (a measure for the density of the graph). For each combination of $(\alpha, n, \tau, \rho, \phi, OS)$, four instances are generated; the total number of instances is thus $2 \times 5 \times 3 \times 3 \times 3 \times 4 \times 4 = 4320$. In all our experiments, the time limit is set to 1200 seconds. If an instance is not solved to guaranteed optimality, it is labeled ‘unsolved’ for the procedure. Throughout this section, we report the averages computed only over the solved instances.

7.2 Lower bounds

We compare the quality of the lower bounds for the subset of instances with large processing times and $n = 30$. We set $k_{\max} = 10$ for all lower bounds. The detailed results of this comparison are reported in Table 5.

The average gap for LB_1 is less than or equal to that for LB_0 , especially when the precedence graph is dense; for $OS = 0$, on the other hand, there are no precedence constraints and LB_0 and LB_1 are essentially the same. A similar observation can be made for LB_1

and LB_2 , where the gap for LB_2 is noticeably smaller than that for LB_1 when release dates are imposed, while in the case $\tau = 0$, only one block is created and the lower bounds LB_1 and LB_2 coincide. The average gap for LB_3 is indeed smaller than that for LB_2 , as was to be expected according to Observation 1.

Although we have no theoretical result that would indicate a better performance of $LB_2^{SS_r}$ in comparison with LB_2 , the average gap for $LB_2^{SS_r}$ is less than that for LB_2 in case of non-zero release dates. When release dates are zero, however, the gap for $LB_2^{SS_r}$ is larger than or equal to that for LB_2 . In fact, when release dates are zero, only one block is created and constraints (19) can be removed from LR_{P_2} , and thus LR_{P_2} is a relaxation of LR_{P_1} . $LB_2^{SS_s}$ performs better than LB_2 and $LB_2^{SS_r}$ for most of the instances. The gap for $LB_3^{SS_r}$ is less than that for $LB_2^{SS_r}$ and a similar observation holds for $LB_3^{SS_s}$ versus $LB_2^{SS_s}$, which again confirms the result in Observation 1.

In our final implementation, we will not compute all the bounds for all the nodes because this consumes too much effort. We start with computing LB_0 and LB_{SV_1} for the unscheduled jobs. Let S_{best} be the best feasible schedule found. If the node is fathomed by DR_1 , then we backtrack; otherwise if $TWT(S_P) + \lceil LB_0 + LB_{SV_1} \rceil \times 1.4 < TWT(S_{best})$ then we do not compute the remaining lower bounds and continue branching. If the latter equality does not hold, then we anticipate that with a better bound we might still be able to fathom the node, and we compute LB_3 and/or $LB_3^{SS_s}$. For all lower bounds we choose $k_{\max} = 0$ if $OS < 0.5$ and $k_{\max} = 1$ otherwise. Preliminary experiments indicate that the extra computational effort to compute $LB_2^{SS_r}$ and $LB_3^{SS_r}$ typically outweighs the improvement in the bound, and so we decide not to use these two values.

7.3 Dominance rules

In each node of the B&B algorithm, dominance rules are tested. Based on some preliminary experiments, we find that applying the rules in the following order performs well, and we will therefore follow this order throughout the algorithm:

$DR_2, DR_3, DR_6^2, DR_4, DR_5, DR_6^3, DR_6^4, DR_1$.

In order to evaluate the effectiveness of the rules, we examine a number of scenarios with respect to the selection of the implemented bounds; the list of scenarios is given in Table 6. Scenario 1 includes the simplest combination of dominance rules, namely DR_2 and DR_3 . From Scenario 2 to Scenario 5, extra rules are gradually added. In Scenario 6, all dominance rules are active

Table 5: Average percentage gap from optimal value.

		LB ₀	LB ₁	LB ₂	LB ₂ ^{SS_r}	LB ₂ ^{SS_s}	LB ₃	LB ₃ ^{SS_r}	LB ₃ ^{SS_s}	LB _{Best}
<i>OS</i>	0.00	50.505	50.505	44.369	43.313	36.857	43.142	42.086	35.630	35.372
	0.25	67.776	63.465	53.444	52.702	52.300	51.955	51.013	50.568	49.834
	0.50	71.461	66.108	52.378	51.890	52.021	51.216	50.727	50.767	50.352
	0.75	77.055	69.836	50.769	50.520	50.565	49.430	49.182	49.041	48.863
τ	0.00	38.141	29.169	29.169	29.182	24.784	28.152	28.165	23.667	23.667
	0.50	76.712	73.157	59.554	58.724	57.699	57.876	57.046	55.922	55.656
	1.00	85.704	85.539	62.494	61.255	61.655	61.117	59.878	60.233	59.309
All	-	67.161	62.911	50.638	49.950	48.268	49.275	48.586	46.824	46.425

except DR₁, and in Scenario 7, only DR₄ is inactive. Scenario 8 and Scenario 9 similarly include all dominance rules except DR₅ and DR₆, respectively. Finally, in Scenario 10, all dominance rules are active.

For each of these implementations, we report the average CPU times and the average number of nodes explored in the search tree in Table 7; the results pertain to the instances of Ins^L with $n = 10, 20, 30$. We observe that both the CPU time as well the number of nodes in Scenario 2 are significantly lower than in Scenario 1; this confirms the value of DR₆ for two successive jobs. Scenarios 3 and 4 show the effect of DR₄ and DR₅. In Scenario 3, DR₄ improves the performance of both algorithms whereas in Scenario 4, DR₅ has a beneficial effect only for BB2. Scenarios 5 and 6 reflect the impact of DR₆ for three and four jobs, respectively. It can be seen that permutation of three and four successive jobs does not improve the results in small instances. For larger instances, however, the improvement is distinct. Comparing Scenario 6 to Scenario 10, we see that inclusion of DR₁ has a strong beneficial effect on both algorithms; the effect is strongest in BB2 because tighter bounds can be computed by scheduling backward. From Table 7, we learn that apart from DR₂, which is always crucial in total tardiness scheduling problems, the most important dominance rule is DR₆: deactivating this rule triggers a huge increase in the average number of nodes and the average CPU times; incorporating DR₄ also has a marked effect (compare Scenarios 7 and 10). Among all dominance rules tested, DR₅ is the least important; removing DR₅ slightly increases the node count and the runtimes in BB2. In BB1, removing DR₅ even decreases the number of nodes and the runtimes; it turns out that for $n > 30$, however, the effect of DR₅ is also (slightly) beneficial for BB1, and so we decide to adopt Scenario 10 as the final setting in which the experiments in the following sections will be run.

As a side note, we observe that for all the foregoing dominance rules, after the root node, omitting the precedence constraints implied by sets \mathcal{Q}_j and \mathcal{P}_j from the updates of \bar{r}_j and $\bar{\delta}_j$ has only little effect. We will therefore not include these precedence constraints into

the updated release dates and deadlines and thus avoid the additional computational overhead.

7.4 Branch-and-bound algorithms

In this section we discuss the performance of our B&B algorithms. We compare the performance of BB1 and BB2 with the MIP formulations discussed in Section 3 in Table 8; the columns labeled ‘#’ indicate the number of instances solved (out of 432).

Based on Table 8, we conclude that the time-indexed formulations are far better than the assignment formulations when processing times are small. For large processing times, the performance of ASF is slightly better than TIF. Although ASF’ and TIF’ are tighter than their counterparts with aggregate precedence constraints, the extra computational effort needed to process the larger models increases CPU times in both TIF’ and ASF’. The B&B algorithms BB1 and BB2 both clearly outperform the MIP formulations regardless of the size of the processing times. On average, BB1 performs better than BB2, although this does not hold for all parameter settings (more details follow below). The algorithms fail to solve a few instances with 40 jobs and around 5% of the instances with 50 jobs. We will indicate below that all these unsolved instances belong to a specific class; it is worth mentioning that the difficult instances are not the same for the two B&B algorithms.

The number of precedence constraints obviously affects the performance of the algorithms. On the one hand, by adding precedence constraints, the set of feasible sequences shrinks; on the other hand, the lower bounds also become less tight. The net result of these two effects is a priori not predictable. For instance classes without release dates and deadlines ($r_j = 0$ and $\delta_j = \infty$), the quality of the lower bound is very good when $OS = 0$, therefore the effect of a weaker bound due to higher OS will be more pronounced than when release dates and deadlines are also imposed.

To identify the classes of difficult instances, we focus on case $n = 50$. Table 10 shows the outcomes of the

Table 6: The list of scenarios.

Scenario	DR ₂	DR ₃	DR ₆ ²	DR ₄	DR ₅	DR ₆ ³	DR ₆ ⁴	DR ₁
1	✓	✓	-	-	-	-	-	-
2	✓	✓	✓	-	-	-	-	-
3	✓	✓	✓	✓	-	-	-	-
4	✓	✓	✓	✓	✓	-	-	-
5	✓	✓	✓	✓	✓	✓	-	-
6	✓	✓	✓	✓	✓	✓	✓	-
7	✓	✓	✓	-	✓	✓	✓	✓
8	✓	✓	✓	✓	-	✓	✓	✓
9	✓	✓	-	✓	✓	-	-	✓
10	✓	✓	✓	✓	✓	✓	✓	✓

Table 7: The effect of the dominance rules.

Method	Scenario	$n = 10$		$n = 20$		$n = 30$	
		CPU	Nodes	CPU	Nodes	CPU	Nodes
BB1	1	0.002	342.252	-	-	-	-
	2	0.001	55.046	0.033	15191.410	-	-
	3	0.001	34.743	0.028	5636.516	5.538 ⁽¹⁾	6845380.421
	4	0.001	33.081	0.026	5563.919	5.924 ⁽¹⁾	6726021.116
	5	0.001	27.688	0.021	1735.356	0.469	396619.023
	6	0.001	27.088	0.025	1248.784	0.302	168552.146
	7	0.001	37.407	0.012	2114.178	1.483	1177730.269
	8	0.001	25.634	0.010	598.218	0.080	21240.525
	9	0.001	111.405	0.667	655033.273	-	-
	10	0.001	24.729	0.011	596.940	0.091	22575.081
BB2	1	0.002	149.144	-	-	-	-
	2	0.001	48.016	0.022	9936.183	-	-
	3	0.001	29.951	0.026	4353.984	3.202	4719358.465
	4	0.001	26.741	0.020	4063.060	2.459	2735455.618
	5	0.001	24.222	0.022	1561.509	0.647	557498.373
	6	0.001	23.829	0.022	1228.444	0.459	293902.986
	7	0.001	36.725	0.012	1649.051	0.499	356457.007
	8	0.001	23.981	0.009	592.192	0.123	50681.481
	9	0.001	53.252	0.106	154295.438	-	-
	10	0.001	21.340	0.009	523.567	0.087	31303.123

The number in brackets indicates the number of instances that are not solved within the time limit of 1200 seconds. ‘-’ means that the implementation fails to solve most of the instances within the time limit.

experiments for each combination of τ , ρ and OS . According to this table, the most time-consuming class of instances is the one where release dates are neither loose nor tight ($\tau = 0.50$), due dates are loose ($\rho = 0.50$) and the set of precedence constraints is empty ($OS = 0$). No clear pattern can be distinguished for the algorithmic performance as a function of the tightness of the deadlines, so these results are excluded from the table. We distinguish subsets of instances: the set of all instances is $\text{Ins} = \text{Ins}^S \cup \text{Ins}^L$. The subset of Ins with $\tau = 0$ is denoted by $\text{Ins}^{\tau=0}$. Similarly, we introduce $\text{Ins}^{\tau=0.5}$, $\text{Ins}^{\tau=1}$, $\text{Ins}^{\rho=0.05}$, $\text{Ins}^{\rho=0.25}$ and $\text{Ins}^{\rho=0.50}$. Each of these six subsets contains 1440 instances. The unsolved instances are distributed differently for the two algorithms. For BB1, $\text{Ins}^{\tau=0}$ contains three unsolved instances, $\text{Ins}^{\tau=0.5}$ contains 32 unsolved instances and $\text{Ins}^{\tau=1}$ contains three such instances. For BB2, 36 unsolved instances are in $\text{Ins}^{\tau=0.5}$ and only two unsolved instances are in $\text{Ins}^{\tau=1}$. Clearly, BB2 performs better

than BB1 when release dates are equal (zero) and BB1 performs faster than BB2 when release dates are non-equal. In both B&B algorithms, the unsolved instances are mostly in $\text{Ins}^{\rho=0.50}$ rather than in $\text{Ins}^{\rho=0.25}$ and $\text{Ins}^{\rho=0.05}$. The runtimes of both algorithms are increasing with ρ , but BB2 performs better than BB1 when due dates are tight (ρ is small) and performs worse than BB1 when due dates are loose (ρ is large).

7.5 Branch and bound versus SSDP

We compare the performance of our B&B algorithms with the SSDP algorithm proposed by Tanaka and Sato [42]. Since the procedure in [42] was not developed for instances with time windows, we remove all deadlines from the instances in $\text{Ins}^{\tau=0}$; we refer to the new set as $\text{Ins}^{\tau=0, \phi=\infty}$. Table 11 shows the computational results for our B&B algorithms and for the SSDP algorithm

Table 8: Overall results.

α	Method	n									
		10		20		30		40		50	
		CPU	#	CPU	#	CPU	#	CPU	#	CPU	#
10	ASF	0.81	432	–	–	–	–	–	–	–	–
	ASF'	0.80	432	–	–	–	–	–	–	–	
	TIF	0.43	432	2.02	432	53.47	429	–	–	–	–
	TIF'	0.64	432	2.97	432	88.17	420	–	–	–	–
	BB1	0.00	432	0.01	432	0.05	432	3.43	432	28.15	418
	BB2	0.00	432	0.00	432	0.05	432	3.30	432	32.46	407
100	ASF	0.92	432	–	–	–	–	–	–	–	–
	ASF'	0.95	432	–	–	–	–	–	–	–	–
	TIF	6.54	432	–	–	–	–	–	–	–	–
	TIF'	21.783	432	–	–	–	–	–	–	–	–
	BB1	0.00	432	0.01	432	0.09	432	6.65	430	24.92	408
	BB2	0.00	432	0.01	432	0.09	432	5.29	430	29.35	413

Table 9: Number of jobs versus density of the precedence graph.

Method	n	OS							
		0		0.25		0.50		0.75	
		CPU	#	CPU	#	CPU	#	CPU	#
BB1	30	0.20	216	0.04	216	0.02	216	0.01	216
	40	18.89	214	1.23	216	0.12	216	0.02	216
	50	59.92	179	49.58	210	2.42	216	0.05	216
BB2	30	0.13	216	0.10	216	0.03	216	0.01	216
	40	12.72	214	4.07	216	0.42	216	0.02	216
	50	38.94	189	74.43	202	13.71	214	0.18	216

Table 10: Effect of release-date tightness, due-date tightness and density of the precedence graph for the instances with 50 jobs.

Method	τ	ρ	OS							
			0		0.25		0.50		0.75	
			CPU	#	CPU	#	CPU	#	CPU	#
BB1	0	0.05	1.90	24	2.54	24	0.68	24	0.03	24
		0.25	45.57	24	81.43	24	2.04	24	0.06	24
		0.50	227.66	21	146.48	24	10.08	24	0.06	24
	0.5	0.05	34.70	22	15.67	24	1.84	24	0.04	24
		0.25	74.67	16	37.84	24	1.56	24	0.05	24
		0.50	567.96	3	150.36	23	4.88	24	0.08	24
	1	0.05	5.56	24	0.48	24	0.28	24	0.04	24
		0.25	15.86	23	1.74	24	0.18	24	0.03	24
		0.50	29.35	22	14.63	24	0.26	24	0.02	24
BB2	0	0.05	0.67	24	3.22	24	0.73	24	0.16	24
		0.25	0.64	24	6.99	24	0.57	24	0.02	24
		0.50	1.29	24	7.48	24	0.40	24	0.02	24
	0.5	0.05	6.53	24	132.87	23	18.80	24	0.11	24
		0.25	195.90	22	240.35	22	27.47	24	0.16	24
		0.50	0	0	392.75	13	77.23	22	0.98	24
	1	0.05	12.26	24	1.27	24	0.36	24	0.07	24
		0.25	67.47	24	1.90	24	1.99	24	0.05	24
		0.50	39.87	22	45.24	24	1.15	24	0.05	24

(which was run on the same computer). The SSDP algorithm solves instances in very short runtimes when there are no precedence constraints. SSDP performs worse, however, when the precedence graph is dense, while the B&B algorithms will tend to perform better exactly in this case. To conclude this comparison, we underline the fact that our algorithms have been devel-

oped to solve the more general setting in which time windows are also imposed, whereas the instance set examined here does not contain such time windows.

Table 11: The comparison of the performance of B&B and SSDP for $\text{Ins}^{\tau=0, \phi=\infty}$.

α	n	Method	OS							
			0		0.25		0.50		0.75	
			CPU	#	CPU	#	CPU	#	CPU	#
10	40	BB1	3.45	36	165.38	32	0.18	36	0.01	36
		BB2	0.04	36	0.70	36	0.06	36	0.00	36
		SSDP	0.05	36	2.08	36	1.41	36	0.33	36
	50	BB1	19.70	34	165.78	32	8.48	36	0.07	36
		BB2	0.21	36	34.40	36	0.62	36	0.05	36
		SSDP	0.11	36	6.24	36	3.39	36	1.66	36
100	40	BB1	1.68	36	2.03	36	0.13	36	0.03	36
		BB2	0.11	36	0.45	36	0.07	36	0.03	36
		SSDP	0.10	36	7.58	36	4.05	36	0.40	36
	50	BB1	7.62	36	137.14	34	2.52	36	0.06	36
		BB2	0.60	36	21.81	36	0.81	36	0.06	36
		SSDP	0.23	36	26.44	36	17.38	36	1.55	36

8 Summary and conclusion

In this article, we have developed exact algorithms for the single-machine scheduling problem with total weighted tardiness penalties. We work with a rather general problem statement, in that both precedence constraints as well as time windows (release dates and deadlines) are part of the input; this generalizes quite a number of problems for which computational procedures have already been published. We develop a branch-and-bound algorithm that solves the problem to guaranteed optimality. Computational results show that our approach is effective in solving medium-sized instances, and that it compares favorably with two straightforward linear formulations. Our procedure was also compared with an existing method (SSDP) for a special case of the problem without time windows. The SSDP algorithm requires only very low runtimes in the absence of precedence constraints, but it performs worse when the precedence graph is dense, which is exactly the easiest setting for our B&B algorithms.

Acknowledgments

We are very grateful to professor Shunji Tanaka from Kyoto University for providing us with the code of the SSDP algorithm that was used in Section 7.5.

Appendix

Proof of Lemma 1

Lemma 1. *ASF' is stronger than ASF.*

Proof Consider the set of constraints (11). For each $(i, j) \in A$, the following inequalities hold:

$$\begin{aligned}
 x_{i1} + \cdots + x_{in} &\leq 1 - x_{j1} = x_{j2} + \cdots + x_{jn} \\
 x_{i2} + \cdots + x_{in} &\leq 1 - x_{j1} - x_{j2} = x_{j3} + \cdots + x_{jn} \\
 &\vdots \\
 x_{i(n-1)} + x_{in} &\leq 1 - x_{j1} - \cdots - x_{j(n-1)} = x_{jn} \\
 x_{in} &\leq 1 - x_{j1} - \cdots - x_{jn} = x_{j1} + \cdots + x_{jn} - 1
 \end{aligned}$$

By adding the above inequalities, we obtain

$$\begin{aligned}
 x_{i1} + 2x_{i2} + 3x_{i3} + \cdots + nx_{in} &\leq \\
 x_{j1} + 2x_{j2} + 3x_{j3} + \cdots + nx_{jn} &- 1.
 \end{aligned}$$

This is exactly the associated constraint in the set of constraints (6). As a result, the solution space of the LP relaxation of ASF' is included in that of ASF. To show that the inclusion is strict, consider the following fractional values for the decision variables corresponding with a couple $(i, j) \in A$: $x_{i1} = x_{i5} = 0.5$ and $x_{j4} = 1$. These values can be seen to respect the weak but not the strong formulation. \square

Proof of Lemma 4

Lemma 4. *Let $f : t \rightarrow \alpha \max\{0, t - a\} - \beta \max\{0, t - b\} + C$ be a function defined on $[u, v]$ for $a, b, C \in \mathbb{R}$ and $\alpha, \beta, u, v \in \mathbb{R}^+$. The function f reaches a global minimum at value t^* computed as follows:*

$$\begin{aligned}
 t^*(\alpha, \beta, a, b, u, v) &= \\
 \begin{cases} \min\{\bar{u}, v\} & \text{if } \alpha \geq \beta \\ u & \text{if } \alpha < \beta, b > a, \alpha(\bar{v} - \bar{u}) \geq \beta(\bar{v} - b) \\ v & \text{otherwise} \end{cases}
 \end{aligned}$$

where $\bar{u} = \max\{u, a\}$ and $\bar{v} = \max\{v, b\}$.




Fig. 12: Four possible cases for the parameter combinations in the proof of Lemma 4.

Proof Let f have a global minimum at value t^* . Depending on the values of the parameters α, β, a and b , the function f behaves differently. We discuss four possible cases for the parameter combinations to prove this lemma (see also Figure 12). In the two first cases, we assume that $\alpha \geq \beta$. Case (a): in this case, $a \leq b$, and then f is constant on interval $[u, a]$ and is increasing on interval $[a, v]$, as shown in Figure 12(a). Case (b): $a > b$, f is constant on interval $[u, b]$, decreasing on interval $[b, a]$ and increasing on interval $[a, v]$, in line with Figure 12(b). The following results are valid for these two cases: 1- If $u \leq a \leq v$ then $t^* = a$. 2- If $a < u$, $t^* = u$ because f is always increasing on interval $[u, v]$. 3- If $a > v$, $t^* = v$ because f is always decreasing on interval $[u, v]$. We conclude that $t^* = \min\{\max\{a, u\}, v\}$ for the first two cases.

In the next two cases, we assume that $\alpha < \beta$. Case (c): $a < b$, f is constant for $[u, b]$, increasing for $[b, a]$ and decreasing for $[a, v]$, as shown in Figure 12(c). In this case, t^* equals either u or v . On the one hand, if $\alpha(b - \max\{a, u\}) \geq (\beta - \alpha)(\max\{v, b\} - b) \Rightarrow \alpha(\bar{v} - \bar{u}) \geq \beta(\bar{v} - b)$ then $f(v) \geq f(u)$ is inferred and $t^* = u$ is concluded. On the other hand, if $\alpha(\bar{v} - \bar{u}) < \beta(\bar{v} - b)$ then $t^* = v$ is concluded. Case (d): $a \geq b$, f is constant for $[u, b]$ and decreasing for $[b, v]$; see Figure 12(d). We find that t^* equals v for this case. \square

References

1. T.S. Abdul-Razaq and C.N. Potts. Dynamic programming state-space relaxation for single-machine scheduling. *Journal of the Operational Research Society*, 39(2):141–152, 1988.
2. T.S. Abdul-Razaq, C.N. Potts, and L.N. Van Wassenhove. A survey of algorithms for the single machine total weighted tardiness scheduling problem. *Discrete Applied Mathematics*, 26:235–253, 1990.
3. M.S. Akturk and D. Ozdemir. An exact approach to minimizing total weighted tardiness with release dates. *IIE Transactions*, 32:1091–1101, 2000.
4. M.S. Akturk and D. Ozdemir. A new dominance rule to minimize total weighted tardiness with unequal release dates. *European Journal of Operational Research*, 135(2):394–412, 2001.
5. C. Artigues, S. Demasse, and E. Neron. *Resource-Constrained Project Scheduling: Models, Algorithms, Extensions and Applications*. ISTE, 2007.
6. W. Bein, J. Kamburowski, and M. Stallmann. Optimal reduction of two-terminal directed acyclic graphs. *SIAM Journal on Computing*, 21(6):1112–1129, 1992.
7. H. Belouadah, M.E. Posner, and C.N. Potts. Scheduling with release dates on a single machine to minimize total weighted completion time. *Discrete Applied Mathematics*, 36(3):213–231, 1992.
8. G. Calinescu, C. Fernandes, H. Kaul, and A. Zelikovsky. Maximum series-parallel subgraph. *Algorithmica*, 63(1-2):137–157, 2012.
9. N. Christofides, R. Alvarez-Valdes, and J.M. Tamarit. Project scheduling with resource constraints: A branch and bound approach. *European Journal of Operational Research*, 29:262–273, 1987.

10. R.W. Conway, W.C. Maxwell, and L.W. Miller. *Theory of Scheduling*. Addison Wesley, Reading, MA, 1967.
11. E. Demeulemeester, M. Vanhoucke, and W. Herroelen. Rangen: A random network generator for activity-on-the-node networks. *Journal of Scheduling*, 6:13–34, 2003.
12. M.E. Dyer and L.A. Wolsey. Formulating the single machine sequencing problem with release dates as a mixed integer program. *Discrete Applied Mathematics*, 26(23):255–270, 1990.
13. M.L. Fisher. The Lagrangian relaxation method for solving integer programming problems. *Management Science*, 27(1):1–18, 1981.
14. M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
15. V. Gordon, E. Potapneva, and F. Werner. Single machine scheduling with deadlines, release and due dates. *Optimization*, 42(3):219–244, 1997.
16. R.L. Graham, E.L. Lawler, J.K. Lenstra, and A.H.G. Rinnooy Kan. Optimization and approximation in deterministic sequencing and scheduling: A survey. *Annals of Discrete Mathematics*, 5:287–326, 1979.
17. A.M.A. Hariri and C.N. Potts. An algorithm for single machine sequencing with release dates to minimize total weighted completion time. *Discrete Applied Mathematics*, 5(1):99–109, 1983.
18. M. Held and R.M. Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10(1):196–210, 1962.
19. J.A. Hoogeveen and S.L. van de Velde. Stronger Lagrangian bounds by use of slack variables: Applications to machine scheduling problems. *Mathematical Programming*, 70:173–190, 1995.
20. T. Ibaraki and Y. Nakamura. A dynamic programming method for single machine scheduling. *European Journal of Operational Research*, 76(1):72–82, 1994.
21. A. Jouglet, P. Baptiste, and J. Carlier. *Handbook of Scheduling: Algorithms, Models and Performance Analysis*, chapter 13, Branch and Bound Algorithms for Total Weighted Tardiness. CRC Press, Boca Raton, FL, USA, 2004.
22. A.B. Keha, K. Khowala, and J.W. Fowler. Mixed integer programming formulations for single machine scheduling problems. *Computers & Industrial Engineering*, 56:357–367, 2009.
23. E.L. Lawler. Optimal sequencing of a single machine subject to precedence constraints. *Management Science*, 19(5):544–546, 1973.
24. E.L. Lawler. A pseudopolynomial algorithm for sequencing jobs to minimize total tardiness. *Annals of Discrete Mathematics*, 1:331–342, 1977.
25. E.L. Lawler. Sequencing jobs to minimize total weighted completion time subject to precedence constraints. *Algorithmic Aspects of Combinatorics*, 2:75–90, 1978.
26. J.K. Lenstra, A.H.G. Rinnooy Kan, and P. Brucker. Complexity of machine scheduling problems. In *Studies in Integer Programming*, volume 1 of *Annals of Discrete Mathematics*, pages 343–362. Elsevier, 1977.
27. G.B. McMahon and C.J. Lim. The two-machine flow shop problem with arbitrary precedence relations. *European Journal of Operational Research*, 64(2):249–257, 1993.
28. G. Nemeth, I. Lovrek, and V. Sinkovic. Scheduling problems in parallel systems for telecommunications. *Computing*, 58:199–223, 1997.
29. R. Nessah and I. Kacem. Branch-and-bound method for minimizing the weighted completion time scheduling problem on a single machine with release dates. *Computers & Operations Research*, 39(3):471–478, 2012.
30. Y. Pan. An improved branch and bound algorithm for single machine scheduling with deadlines to minimize total weighted completion time. *Operations Research Letters*, 31(6):492–496, 2003.
31. Y. Pan and L. Shi. Dual constrained single machine sequencing to minimize total weighted completion time. *IEEE Transactions on Automation Science and Engineering*, 2(4):344–357, 2005.
32. M.L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer, 2008.
33. M.E. Posner. Minimizing weighted completion times with deadlines. *Operations Research*, 33:562–574, 1985.
34. C.N. Potts. A lagrangean based branch and bound algorithm for single machine sequencing with precedence constraints to minimize total weighted completion time. *Management Science*, 31(10):1300–1311, 1985.
35. C.N. Potts and L.N. Van Wassenhove. An algorithm for single machine sequencing with deadlines to minimize total weighted completion time. *European Journal of Operational Research*, 12(4):379–387, 1983.
36. C.N. Potts and L.N. Van Wassenhove. A branch and bound algorithm for the total weighted tardiness problem. *Operations Research*, 33(2):363–377, 1985.
37. B.A. Shirazi, K.M. Kavi, and A.R. Hurson, editors. *Scheduling and Load Balancing in Parallel and Distributed Systems*. IEEE Computer Society Press, 1995.
38. D.R. Sule. *Production Planning and Industrial Scheduling: Examples, Case Studies and Applications*. CRC Press, 2007.
39. F. Talla Nobibon and R. Leus. Exact algorithms for a generalization of the order acceptance and scheduling problem in a single-machine environment. *Computers & Operations Research*, 38(1):367–378, 2011.
40. S. Tanaka and S. Fujikuma. A dynamic-programming-based exact algorithm for general single-machine scheduling with machine idle time. *Journal of Scheduling*, 15:347–361, 2012.
41. S. Tanaka, S. Fujikuma, and M. Araki. An exact algorithm for single-machine scheduling without machine idle time. *Journal of Scheduling*, 12:575–593, 2009.
42. S. Tanaka and S. Sato. An exact algorithm for the precedence-constrained single-machine scheduling problem. *European Journal of Operational Research*, 229(2):345–352, 2013.
43. J. Valdes, R. Tarjan, and E. Lawler. The recognition of series parallel digraphs. *SIAM Journal on Computing*, 11(2):298–313, 1982.
44. S.L. van de Velde. Dual decomposition of a single-machine scheduling problem. *Mathematical Programming*, 69(1-3):413–428, 1995.
45. J. van den Akker, G. Diepen, and J. Hoogeveen. Minimizing total weighted tardiness on a single machine with release dates and equal-length jobs. *Journal of Scheduling*, 13:561–576, 2010.
46. J. Xu and D. Parnas. Scheduling processes with release times, deadlines, precedence and exclusion relations. *IEEE Transaction on Software Engineering*, 16(3):360–369, 1990.

FACULTY OF ECONOMICS AND BUSINESS
Naamsestraat 69 bus 3500
3000 LEUVEN, BELGIË
tel. + 32 16 32 66 12
fax + 32 16 32 67 91
info@econ.kuleuven.be
www.econ.kuleuven.be

