A branch and bound method for the job-shop problem with sequence-dependent setup times

Christian Artigues · Dominique Feillet

Abstract This paper deals with the job-shop scheduling problem with sequence-dependent setup times. We propose a new method to solve the makespan minimization problem to optimality. The method is based on iterative solving via branch and bound decisional versions of the problem. At each node of the branch and bound tree, constraint propagation algorithms adapted to setup times are performed for domain filtering and feasibility check. Relaxations based on the traveling salesman problem with time windows are also solved to perform additional pruning. The traveling salesman problem is formulated as an elementary shortest path problem with resource constraints and solved through dynamic programming. This method allows to close previously unsolved benchmark instances of the literature and also provides new lower and upper bounds.

Keywords Job-shop scheduling · Sequence-dependent setup times · Branch and bound · Constraint propagation · Dynamic programming

1 Introduction

This work deals with the job-shop problem with sequence dependent setup times (SDST-JSP). The job-shop problem considers the scheduling of a set of jobs on distinct machines. This problem is widely investigated in the literature and many efficient approaches exist for its resolution (Blazewicz et al. 1996; Nowicki and Smutnicki 1996; Jain and Meeran 1999; Vaessens et al. 1996). The SDST-JSP is a variant problem where machines have to be re-configured between two consecutive operations. We address here the optimal solution of

C. Artigues (✉)
Université de Toulouse, LAAS–CNRS, 7 Avenue du Colonel Roche, 31077 Toulouse, France
e-mail: artigues@laas.fr

D. Feillet
LIA—Université d'Avignon, 339 Chemin Meinajariés, Agroparc, BP 1228, 84911 Avignon Cedex 9, France
e-mail: Dominique.Feillet@univ-avignon.fr

the SDST-JSP with a makespan minimization criterion. The method is based on the iterative solving of decisional versions of the problem, via branch and bound. Within branch and bound trees, constraint propagation algorithms are performed for domain filtering and feasibility check. Relaxations based on the traveling salesman problem with time window (TSPTW) are also solved to perform additional pruning. The traveling salesman problem is formulated as an elementary shortest path with resource constraints (Feillet et al. 2004) and solved through dynamic programming.

We present the considered problem in Sect. 2 and the traveling salesman problem with time windows relaxation in Sect. 3. The disjunctive graph representation is recalled in Sect. 4. A review of the relevant literature is given in Sect. 5. The sketch of the branch and bound algorithm is provided in Sect. 6 and its components are detailed in the subsequent sections. The heuristic used to compute feasible solutions is described in Sect. 7. Feasibility checks and domain filtering algorithms based on constraint propagation are described in Sect. 8. Section 9 gives the formulation of the TSPTW as an elementary shortest path problem with resource constraints and presents the dynamic programming algorithm used to solve it. In Sect. 10, the proposed branching scheme and the associated dominance rule are described. In Sect. 11, the results of the branch and bound method on the set of instances proposed by Brucker and Thiele (1996) are discussed. Concluding remarks are drawn in Sect. 12.

2 The job-shop problem with sequence dependent setup times

The SDST-JSP considers a set of $n \geq 1$ jobs $\mathcal{J} = \{J_i\}_{1 \leq i \leq n}$ and a set of $m \geq 1$ machines $\mathcal{M} = \{M_k\}_{1 \leq k \leq m}$. Each job $J_i \in \mathcal{J}$ is defined as a set of m operations $J_i = \{O_{ij}\}_{1 \leq j \leq m}$. Each operation O_{ij} has a non zero integer duration $p_{ij} \in \mathbb{N}$ and requires a single machine $m_{ij} \in \mathcal{M}$. The operations of a same job all require distinct machines and are subject to precedence constraints. Operation O_{ij} precedes operation $O_{i,j+1}$ for all $i = 1, \ldots, n$ and for all $j = 1, \ldots, m - 1$. The set of all operations is denoted $\mathcal{O} = \bigcup_{1 \leq i \leq n} J_i$ while $\mathcal{O}_k$ denotes the set of operations assigned to machine M_k.

A sequence dependent setup time, denoted s_{ijk}, is defined for each couple of distinct jobs (J_i, J_j) and for each machine M_k. An initial setup time s_{0ik} is defined for each job J_i and for each machine M_k. It is assumed that the triangular inequality holds: for each machine M_k and for each triplet of distinct jobs (J_i, J_j, J_x), we have $s_{ixk} \leq s_{ijk} + s_{jxk}$; for each machine M_k and for each couple of distinct jobs (J_i, J_j), we have $s_{0jk} \leq s_{0ik} + s_{ijk}$.

A schedule is a mapping $\mathcal{T}$ of operations to time periods where $\mathcal{T} = (t_{ij})_{O_{ij} \in \mathcal{O}}$ and $t_{ij} \in \mathbb{N}$. In this paper we consider the makespan objective for the SDST-JSP. This particular problem is denoted $J|s_{ij}|C_{\max}$ in the standard three-field scheduling notation and we will denote it (P) in the remaining of the paper.

$$(P) \quad \min \quad C_{\max} \tag{1}$$

subject to

$$C_{\max} \geq t_{ij} + p_{ij} \quad \forall O_{ij} \in \mathcal{O}, \tag{2}$$

$$t_{ij} + p_{ij} \leq t_{i,j+1} \quad \forall J_i \in \mathcal{J}, \forall j \in [1, m-1], \tag{3}$$

$$t_{ij} + p_{ij} + s_{ixk} \leq t_{xy} \quad \text{or}$$
$$t_{xy} + p_{xy} + s_{xik} \leq t_{ij} \quad \forall O_{ij}, O_{xy} \in \mathcal{O}_k, O_{ij} \neq O_{xy}, \tag{4}$$

$$t_{ij} \geq s_{0im_{ij}} \qquad \forall O_{ij} \in \mathcal{O}. \tag{5}$$

The objective (1) is the minimization of the makespan $C_{\max}$, the maximum completion time of all jobs (2). Constraints (3) are the job precedence constraints, stating that an operation $O_{i,j+1}$ cannot start before the end of its preceding operation O_{ij} in job J_i. It follows that the operations of a job form a chain. Constraints (4) are the machine constraints which state that two distinct operations O_{ij} and O_{xy} sharing the same machine k cannot be scheduled simultaneously and that the machine must be set up between two consecutive operations. Hence either O_{ij} cannot start before the end of O_{xy} plus the necessary setup time s_{xik}, or O_{xy} cannot start before the end of O_{ij} plus the necessary setup time s_{ixk}. Constraints (5) enforce the start time of any operation to occur after the initial setup time on its assigned machine.

An alternative representation of the setup times lies in defining a set $\mathcal{F}$ of $f \geq 1$ families such that each operation $O_{ij} \in \mathcal{O}$ is associated with a family $f_{ij} \in \mathcal{F}$. Now setup time s_{ixk} on machine M_k from an operation $O_{ij} \in \mathcal{O}_k$ of job J_i to an operation $O_{xy} \in \mathcal{O}_k$ of job J_x can be written $s_{f_{ij} f_{xy}}$. This representation can be interesting when the number of families is significantly smaller than the number of operations. Then, efficient preprocessing techniques can be carried out (see Sect. 8.1).

The job-shop scheduling problem with sequence dependent setup times is a NP-hard problem, as an extension of the standard job-shop scheduling problem, denoted $J||C_{\max}$, where all setup times are equal to zero. In this paper, the optimal solution of (P) is searched by iteratively solving the decisional version (FP) defined as follows. Let $T \geq 0$ be a tentative upper bound of the makespan. We consider problem $FP(T)$ as the problem of finding a schedule $\mathcal{T}$ such that $C_{\max} \leq T$ and constraints (2–5) are satisfied. Let UB denote the makespan of any feasible schedule. Let LB denote a lower bound of the optimal solution of (P). Let $C_{\max}^*$ denote the optimal solution of (P). We have:

$$C_{\max}^* = \min_{LB \leq T \leq UB} \{T | FP(T) \text{ has a solution}\}. \tag{6}$$

3 The traveling salesman problem with time windows relaxation

When $m = 1$, problem (P) is denoted $1|s_{ij}|C_{\max}$ and is equivalent to the traveling salesman problem (TSP) as stated by Conway et al. (1967). Since $m = 1$ we drop the machine and operation indices in the notations. We can define a depot represented by node 0 and n cities represented by nodes $1, \dots, n$. Then the distance matrix $(l_{ij})_{0 \leq i, j \leq n}$ is given by

$$l_{0i} = s_{0i} \quad \forall i \in [1, n], \tag{7}$$

$$l_{ij} = s_{ij} \quad \forall i, j \in [1, n], i \neq j, \tag{8}$$

$$l_{i0} = 0 \quad \forall i \in [0, n]. \tag{9}$$

Since the sum of the operation processing times is a constant, problem $1|s_{ij}|C_{\max}$ and the above defined TSP are equivalent. It follows that when $m > 1$, a lower bound LB_{TSP} of the optimal solution of (P) can be computed as follows by using $l^*(\text{TSP}_{\mathcal{O}_k})$, the length of the optimal tour of the traveling salesman relaxation considering only the operations on M_k.

$$LB_{\text{TSP}} = \max_{M_k \in \mathcal{M}} \left(l^*(\text{TSP}_{\mathcal{O}_k}) + \sum_{O_{ij} \in \mathcal{O}_k} p_{ij} \right). \tag{10}$$

This relaxation can be strengthened by computing time windows for all operations given an upper bound UB of the optimal makespan. Let r_{ij} denote the earliest start time of operation O_{ij}. Let d_{ij} denote the latest completion time of operation O_{ij}. For a given upper bound UB we can compute the following valid values for the time windows:

$$r_{i1} = s_{0im_{i1}} \quad \forall J_i \in \mathcal{J}, \tag{11}$$

$$r_{ij} = \max(s_{0im_{ij}}, r_{i(j-1)} + p_{i(j-1)}) \quad \forall J_i \in \mathcal{J}, \forall j \in [2, m], \tag{12}$$

$$d_{im} = UB \quad \forall J_i \in \mathcal{J}, \tag{13}$$

$$d_{ij} = d_{i(j+1)} - p_{i(j+1)} \quad \forall J_i \in \mathcal{J}, \forall j \in [1, m - 1]. \tag{14}$$

Any solution of (P) with a makespan not greater than UB must obviously verify $r_{ij} \leq t_{ij} \leq d_{ij} - p_{ij}$ for each operation $O_{ij} \in \mathcal{O}$. In Sect. 8, we will present methods to compute tighter valid time windows. Let us now consider the relaxation of (P) considering only operations assigned to machine k. We obtain a makespan minimization one machine problem with time windows and sequence-dependent setup time denoted $1|r_i, d_i, s_{ij}|C_{\max}$, the solution of which is a lower bound for (P). When setup times are all equal to 0, the problem is already NP-hard but it can be efficiently solved by the well-known Carlier's algorithm (Carlier 1982). Otherwise this problem is a traveling salesman problem with time windows (TSPTW) with depot 0, cities $1, \ldots, n$, distance matrix $(l_{ij})_{0 \leq i,j \leq n}$ given by (7–9), service times p_i for $1 \leq i \leq n$ and time windows $[r_i, d_i - p_i]$. More precisely the problem is a TSPTW variant where the objective function is the minimization of total travel times plus total waiting time (Focacci et al. 2002; Artigues et al. 2004; Balas et al. 2005). Let $C^*_{\max}(\text{TSPTW}_{\mathcal{O}_k})$ denote the optimal solution of the TSPTW relaxation considering only the operations on M_k. We obtain the lower bound LB_{TSPTW} as follows.

$$LB_{\text{TSPTW}} = \max_{M_k \in \mathcal{M}} (C^*_{\max}(\text{TSPTW}_{\mathcal{O}_k})). \tag{15}$$

Note that this lower bound depends on the upper bound UB and on the way the time windows are computed.

Let us now consider the decisional version $FP(T)$ tackled in this paper. Since we have a tentative upper bound T and a feasibility problem, the subproblem reduces to the search feasible solutions for the above-defined TSPTW. Finding a machine k such that the associated TSPTW$_{\mathcal{O}_k}$ is infeasible would indeed mean than value T renders $FP(T)$ infeasible. This problem (denoted F-TSPTW in what follows) is known to be NP hard (Savelsberg 1985). In the next section, we will see how comparable relaxations are tackled by the previously proposed exact methods for the SDST-JSP and its variants. In this paper, we propose to solve the TSPTW relaxation of problem $FP(T)$ exactly by a dynamic programming algorithm (see Sect. 9), the time windows being sharpened by constraint propagation techniques (see Sect. 8). We use the solution of the relaxations to prune the search in the case unfeasibility is proven or to derive a feasible solution to (P) when a feasible TSPTW solution is found.

4 The disjunctive-graph representation

A useful tool for the solution of scheduling problems is the so-called disjunctive graph initially proposed for the standard job-shop problem (Roy and Sussman 1964). This graph provides an efficient representation of the decisions, while limiting the solution space. It is defined as follows.

Let $G = (X, U, E)$ be the disjunctive graph. The set of vertices X is made up of the set of operations plus two dummy vertices representing the beginning and the end of the schedule. Thus, X has $n \times m + 2$ vertices. A set of arcs U and a set of edges E are defined. Arcs in U represent precedence constraints between operations and are called conjunctive arcs. They are weighted with the processing time of the origin vertex of the arc. Edges in E represent disjunctive constraints between operations on a same machine and are called disjunctive arcs. Actually, disjunctive arcs can be interpreted as the union of two exclusive arcs with opposite directions.

By definition, a selection is a state of the disjunctive graph where a direction is chosen for some disjunctive arcs. A selection is said to be complete when every arc has a direction. A complete selection corresponds to a unique semi-active schedule (in which no operation can be left-shifted of one time unit) if the resulting graph is acyclic. Once they are directed, disjunctive arcs are weighted with the sum of the processing time of the origin vertex of the arc plus the setup time required between the origin and the destination vertices. Minimizing the makespan then reduces to the search of the longest path in the graph, the makespan being the length of such a path. Hence, (P) can be defined as the problem of finding a complete acyclic selection for which the longest path is minimum. This standpoint relies on the property that it is possible to consider only semi-active scheduling (when disjunctive and conjunctive constraints are satisfied), to find an optimal solution.

5 Literature review

Problem $J|s_{ij}|C_{\max}$ received only little attention in the literature, albeit being a natural extension of the job-shop problem with many applications in manufacturing (Allahverdi et al. 1999). In this section, we review the methods previously proposed to solve this problem and its close variants or extensions that we classify into three categories: priority-rule based (dispatching) methods, local search methods and exact methods.

5.1 Priority-rule based methods

The earliest proposed methods are based on simulation: each time a machine becomes available the dispatching rule selects a job already waiting for this machine, i.e. which as no predecessor or whose predecessor is completed before the machine gets available. Wilbrecht and Prescott (1969) were to our knowledge the first to study under this framework the influence of setup times on job shop performance. They propose a simple but efficient priority rule which aims at selecting the job yielding the smallest setup time. Kim and Bobrowski (1994) extend the preceding work to various performance measures and propose compound priority rules. Noivo and Ramalhinho-Lourenço (1998) propose also new priority-rules considering setup times and compare them with classical priority-rules. Ovacik and Uzsoy (1994a) propose a more sophisticated priority-rule based algorithm for problem $J|s_{ij}|L_{\max}$ with reentrant jobs (a job may use a machine more than once). Indeed the set of candidate operations considered when a machine becomes available is no more restricted to the already pending operations but extended to the operations having their preceding operation in process. This is justified intuitively because they may require a sufficiently large setup time to be taken in consideration. Theoretical aspects of the interest of extending this candidate set is linked to dominance properties of active schedules (see Sect. 7 and Artigues et al. 2005b). Ovacik and Uzsoy were to the best of our knowledge the first to use an exact solution of the TSPTW relaxation as a routine to guide the search. Indeed, among the proposed priority-rules, one amounts to select the first operation in the optimal sequence of the

one-machine subproblem restricted to the set of candidate operations. This subproblem is here denoted $1|r_i, s_{ij}|L_{\max}$ and its decisional variant amounts to F-TSPTW. This problem is possibly solved by a specific branch and bound method (Ovacik and Uzsoy 1994b) which has reasonable computational times only for less than 10 operations.

As a priority-rule heuristic embedded in their exact method that will be discussed later, Brucker and Thiele (1996) propose an extension of the Giffler-Thompson algorithm which defines an active schedule generation scheme in the standard job-shop problem. They propose a second heuristic based on maximal matchings on bipartite graph. Brucker and Thiele (1996) also generated a set of 15 SDST-JSP instances which we use in Sect. 11 to evaluate our method and compare our results with the state-of-the-art exact methods.

Artigues et al. (2005b) propose several new priority-rule based methods, improving the results obtained by Brucker and Thiele (1996) on the 10 largest BT instances. One of these methods being used in our exact procedure, we will give more details in Sect. 7 and explain why it has better dominance properties than all the previously-proposed priority-rule based methods.

5.2 Local search methods

Among local search methods, Choi and Korkmaz (1997) propose a mixed-integer programming formulation of the problem. They formally identify the problem as a multi-traveling salesman problem with side constraints (precedence constraints). They propose a heuristic based upon sequentially identifying a pair of operations that provides a minimum lower bound on the makespan of the associated two-job/m-machine problem with release times. They show their heuristic outperform the one proposed by Zhou and Egbelu (1989). Choi and Choi (2002) propose a local search method based on the iterative application of dispatching rules for a problem involving alternative operations in conjunction with setup-times.

The other local search methods proposed so-far for the SDST-JSP are generally based on the disjunctive graph representation. Based on this model, Candido et al. (1998) propose a genetic algorithm for a job-shop problem involving setup times among multiple practical characteristics such as multiresource operations, machine calendars and alternative routings. A simple tabu search heuristic has been tested on the BT instances by Buscaylet and Artigues (2003). This latter method has been applied by Artigues et al. (2005a) to improve the results on the 10 largest BT instances of the priority-rule based methods proposed by Artigues et al. (2005b).

A large number of local search methods are based on the shifting bottleneck heuristic. This efficient heuristic has been proposed by Adams et al. (1988) for the job shop problems. It has been successfully extended to tackle practical characteristics of actual job shops. Focusing only on extension to setup times, several variants of the shifting bottleneck heuristic have been proposed (Ovacik and Uzsoy 1992; Schutten 1995; Sun and Noble 1999; Mason et al. 2002; Balas et al. 2005). The shifting bottleneck heuristic can be defined as a large neighborhood search which solves at each iteration exactly or approximately the one-machine relaxation of the problem. Hence, most of these methods solve in a specific way the TSPTW relaxation and are therefore of interest for our study. However because of the complexity of the subproblem and because most of these works describe applications on large-size industrial problems, the sub-problem is mostly solved by heuristics (Ovacik and Uzsoy 1992; Schutten 1995; Sun and Noble 1999; Mason et al. 2002). The subproblem considered by Balas et al. (2005) corresponds to the TSPTW subproblem which provides the LB_{TSPTW} lower bound described in Sect. 3. To solve this subproblem, the authors transform the TSPTW in a TSP with precedence constraints having a special structure. This special structure allows to solve the TSP in linear time by dynamic programming (Balas 1996;

Balas and Simonetti 2001). However the TSPTW subproblem does not have in general the required restrictions and the methods is turned into an efficient heuristic. The shifting bottleneck heuristic proposed by Balas et al. (2005) obtains equivalent or better solutions than the multistart tabu search method of Artigues et al. (2005a) for all the BT instances.

5.3 Exact methods

To our knowledge, the only exact methods that have been developed so far for the SDST-JSP are the ones of Brucker and Thiele (1996), Focacci et al. (2000) and Artigues et al. (2004).

Brucker and Thiele (1996) propose a branch and bound method that extends the one proposed in Brucker et al. (1994b) for the standard JSP. The problem considered by Brucker and Thiele (1996) is the general shop problem with sequence-dependent setup time which includes the SDST-JSP as a special case. The method is based on the disjunctive graph representation and a node of the branch and bound tree corresponds to a partial selection. Lower bounds are computed at each node by applying several polynomially-solvable extensions of the Jackson's preemptive schedule relaxation (Carlier and Pinson 1989) to setup times. Branching is based on the analysis of blocks of operations on the critical path of a feasible schedule issued from the current node. The method was able to solve to optimality the 5 smallest BT instances and provided lower and upper bounds for the 10 largest instances.

At each node, Brucker and Thiele (1996) apply filtering algorithms to compute the earliest start times r_{ij} of the operations, so-called heads, as well as the tails q_{ij} which are linked to the latest start times by the relation $q_{ij} = UB - d_{ij}$ and also to deduce implied precedence constraints that enrich the current selection. These filtering algorithms are based on the so-called immediate selections and edge-finding techniques that have been shown to be very efficient to solve the standard JSP (Carlier and Pinson 1989; Brucker et al. 1994a; Baptiste et al. 2001). These techniques can be defined as constraint propagation algorithms for the one-machine problem. To extend efficiently these algorithms to the presence of setup times, Brucker and Thiele propose several lower bound of the minimal setup time necessary between operations sharing the same machine.

Vilím and Barták (2002) proposed another version of the one-machine constraint propagation algorithms, in the context of a batching and scheduling problem with sequence-dependent setup times. In Sect. 8.2, we describe the simple version of the filtering algorithm we have implemented for the branch and bound method proposed in this paper, adapting the algorithms of Brucker and Thiele (1996), Vilím and Barták (2002) and Nuijten (1994) to the SDST-JSP.

Focacci et al. (2000) propose a branch-and bound method in a constraint programming framework for another variant of the SDST-JSP involving alternative resources for the operations. They also consider total setup time minimization in conjunction with makespan minimization. Their method combines the standard (with no setup) one-machine constraint propagation algorithms and reduced cost-based domain filtering techniques. The latter is based on the relaxation of the TSPTW into an assignment problem solvable in polynomial time. Solving the relaxation provides a lower bound and also a reduced cost matrix. Such reduced costs give information about the additional increase of the lower bound when the corresponding variable assignment is performed. Whenever the increase exceeds the upper bound then the value can be removed from the variable domain. It is worth to note that the reduced cost-based constraint propagation methods have also been successfully applied to solve TSPTW instances (Focacci et al. 2002). For the SDST-JSP, Focacci et al. (2000) do not report results for the makespan minimization only. However, when the objective is to minimize the makespan and the total setup time in a lexicographic way, they found all the

optimal makespans of the 5 smallest BT instances. They also report upper bounds for 2 large instances.

A preliminary version of the branch and bound procedure proposed in this paper was presented by Artigues et al. (2004). The main idea of the method is to solve exactly the F-TSPTW relaxation of problem (*FP*) to prune the search instead of solving the Jackson's preemptive schedule relaxation as in (Brucker and Thiele 1996) or the assignment problem relaxation as in Focacci et al. 2000. (*FP*) is formulated as a one-machine scheduling problem and solved by means of the commercial solver (ILOG scheduler). In addition, the TSPTW solutions found during the search are stored in a dictionary. Each time a TSPTW problem has to be solved, the dictionary is searched to avoid the subproblem resolution if a feasible solution is encountered. This method solved all the smallest BT instances to optimality and closed 2 open instances.

6 The branch and bound method

The branch and bound method we propose is based on the disjunctive graph representation. It aims at making further experiments in the direction initiated by Artigues et al. (2004). As in the method of Brucker and Thiele (1996), a node of the branch and bound tree corresponds to a partial selection $\mathcal{E}$ where $\mathcal{E}$ is a set of directed arcs obtained from a partial or complete orientation of disjunctive edges E. A node is also defined by operation time windows compatible with the current selection. Hence a node ν is denoted by a triple $(\mathcal{R}, \mathcal{D}, \mathcal{E})$ with $\mathcal{R} = (r_{ij})_{O_{ij} \in \mathcal{O}}$ and $\mathcal{D} = (d_{ij})_{O_{ij} \in \mathcal{O}}$. The branch and bound is given in Algorithm 1. It takes as input a tentative makespan T and returns a failure or a success and in the latter case a feasible schedule $\mathcal{T}$. The time windows are initialized according to T and (11–14) (step 1). The root node with an empty selection is created (step 2). A first constraint propagation algorithm (SHAVING) is applied to check the feasibility of the root node and to strengthen the time windows and insert arcs in the selection (step 3). The algorithm is described in Sect. 8.3. If no unfeasibility is detected by SHAVING, the branch and bound algorithm starts and a node stack is initialized with the root node (step 6). Steps 8–19 are repeated until the stack is empty or a maximal number of nodes is reached. For each node ν, a second (faster) constraint propagation algorithm (PROPAGATE) is applied (step 9). This algorithm is detailed in Sect. 8.2. If unfeasibility is not detected, PROPAGATE has possibly tightened the time windows and the selection of the current node, and a dynamic programming method SOLVETSPTW (described in Sect. 9) is used to solve the F-TSPTW relaxation (step 10). If SOLVETSPTW does not detect unfeasibility, a solution (represented by start time $\mathcal{T}'$) to the m TSPTW has been found. A heuristic is then applied (step 11). This heuristic, which searches a solution from the current node ν information and solution $\mathcal{T}'$ of the TSPTW, is described in Sect. 7. If solution $\mathcal{T}$ found by the heuristic has a makespan not greater than T, the problem is feasible and the branch and bound stops. Otherwise, we have to update the node stack by branching (step 16). The branching procedure is described in Sect. 10.

7 A priority-rule based serial heuristic

The function HEURISTIC returns a feasible schedule $\mathcal{T}$ taking as input the current node $\nu = (\mathcal{R}, \mathcal{D}, \mathcal{E})$ and a TSPTW solution $\mathcal{T}'$ possibly computed. Note that $\mathcal{T}'$ satisfies in any case the machine constraints but possibly violates one or more precedence constraints.

HEURISTIC first check whether setting $\mathcal{T} = \mathcal{R}$, $\mathcal{T} = \mathcal{D} - p$ or $\mathcal{T} = \mathcal{T}'$ yields a feasible schedule. In that case, this solution is returned. Otherwise we apply a priority-rule

---

Algorithm 1 BRANCHANDBOUND(T, $\mathcal{T}$): returns failure or success

---

1: initialize $\mathcal{R}^0$ and $\mathcal{D}^0$ with (11–14) and $UB = T$
2: $\nu^0 \leftarrow (\mathcal{R}^0, \mathcal{D}^0, \emptyset)$
3: **if** SHAVING(ν^0) = failure **then**
4: 　**Solution Found**: unfeasible
5: **else**
6: 　$BBQ \leftarrow \{\nu^0\}$
7: 　**while** BBQ is not empty **do**
8: 　　$\nu \leftarrow \text{pop}(BBQ)$
9: 　　**if** PROPAGATE(ν) $\neq$ failure **then**
10: 　　　**if** SOLVETSPTW($\mathcal{T}'$)(ν) $\neq$ failure **then**
11: 　　　　$\mathcal{T} \leftarrow$ HEURISTIC($\nu, \mathcal{T}'$)
12: 　　　　**if** $C_{\max}(\mathcal{T}) \leq T$ **then**
13: 　　　　　**return** success
14: 　　　　**else**
15: 　　　　　$BBQ \leftarrow$ BRANCH(ν, BBQ)
16: 　　　　**end if**
17: 　　　**end if**
18: 　　**end if**
19: 　**end while**
20: 　**return** failure
21: **end if**

---

based heuristic SERIALSGS which returns a feasible schedule taking a priority function $\pi : \mathcal{O} \to \mathbb{R}$.

SERIALSGS is applied 3 times by setting successively $\pi = \mathcal{R}$, $\pi = \mathcal{D} - p$, $\pi = \mathcal{T}'$ and the schedule with the lower makespan returned. We explain hereafter the principles and the interest of SERIALSGS.

None of the other priority-rule based heuristics described in Sect. 5.1 was explicitly designed for having the ability to generate (all) schedules in the classical active, semi-active and non-delay schedule sets. However, it is well-known that active and semi-active schedules are dominant for makespan generation. Artigues et al. (2005b) propose several schedule generation schemes aiming at generating semi-active, active and non-delay schedules and study the dominance properties of the reachable schedules. This analysis reveals that most of the priority-rule based heuristics proposed so far generate sets of schedules possibly excluding all optimal solutions. For instance, the extension of the Giffler and Thompson algorithm proposed by Brucker and Thiele (1996) does not generate dominant schedules. In this paper we use the serial schedule generation scheme (SGS) which generates only active schedules and for which the dominance property holds (Artigues et al. 2005b). This algorithm is a is an adaptation for the SDST-JSP of the serial SGS widely used for the resource-constrained project scheduling problem (Kolisch 1996). Algorithm 2 implements the serial schedule generation scheme with an $O(nm^2)$ time complexity.

Algorithm 2 SERIALSGS(π): returns a feasible schedule $\mathcal{T}$

1: $Q \leftarrow \{O_{i1} | i = 1, \ldots, n\}$
2: $S_k \leftarrow \emptyset, \forall k \in \mathcal{M}$
3: **while** Q is not empty **do**
4:     select $O_{i^* j^*} \leftarrow \arg\min_{O_{ij} \in Q} \pi(O_{ij})$ and let $k = m_{i^* j^*}$
5:     $t \leftarrow s_{0i^* k}$
6:     **for** each operations $O_{xy} \in S_k$ **do**
7:       **if** $t + p_{i^* j^*} + s_{i^* x k} \leq t_{xy}$ **then**
8:          insert $O_{i^* j^*}$ before O_{xy} in S_k
9:          **break**
10:       **else**
11:          $t \leftarrow \max(t, t_{xy} + p_{xy} + s_{xi^* k})$
12:       **end if**
13:     **end for**
14:     **if** $O_{i^* j^*} \notin S_k$ **then**
15:       append $O_{i^* j^*}$ to S_k
16:     **end if**
17:     $t_{i^* j^*} \leftarrow t$
18:     $Q \leftarrow Q \setminus \{O_{i^* j^*}\}$
19:     **if** $j^* < m$ **then**
20:       $Q \leftarrow Q \cup \{O_{i^*(j^*+1)}\}$
21:     **end if**
22: **end while**
23: **return** $\mathcal{T} = (t_{ij})_{O_{ij} \in \mathcal{O}}$

---

8 Feasibility tests and adjustments based on constraint propagation

8.1 Setup preprocessing

The feasibility tests we have implemented make use of extensions to setup times of well-known one-machine constraint propagation algorithms, such as immediate selections and edge-finding (Carlier and Pinson 1989; Brucker et al. 1994a; Baptiste et al. 2001). Such extensions have been proposed by Brucker and Thiele (1996) and Vilím and Barták (2002). Immediate selections aim at detecting new precedence constraints among operations assigned to the same machine, and consequently enrich the current selection by new conjunctive arcs. To perform these deductions the constraint propagation algorithms make use of the minimal duration without idle times of subsets of operations $\Omega \subseteq \mathcal{O}_k$, for each machine $M_k \in \mathcal{M}$. When there are no setup times, the minimal duration of a set Ω is denoted p_Ω and is trivially equal to $p_\Omega = \sum_{O_{ij} \in \Omega} p_{ij}$.

When there are setup times, the minimal duration of $\Omega \subseteq \mathcal{O}_k$, is equal to $p_\Omega + s_\Omega$ where s_Ω is the solution of a modified TSP_Ω where the distance from the depot to each node is null. Hence, given a set Ω, computing s_Ω is NP-hard.

When the setup times are computed with operations families, operations of a same family can be gathered together and sequenced consecutively. Hence, the TSP might equivalently be solved with a single operation of each family $f \in \mathcal{F}_\Omega = \bigcup_{O_{ij} \in \Omega} \{f_{ij}\}$. If $|\mathcal{F}_\Omega|$ is sufficiently lower than $|\Omega|$ then the computational time of the corresponding TSP may become tractable. In the following we note equivalently $s_{\mathcal{F}_\Omega}$ and s_Ω.

The filtering algorithm needs more precisely, for each family subset $\mathcal{F}' \subseteq \mathcal{F}_{\mathcal{O}_k}$ and for each family $f \in \mathcal{F}'$, the values denoted $s_{f \to \mathcal{F}'}$ and $s_{\mathcal{F}' \to f}$. These values correspond to the minimal setup times to schedule a set of operations yielding family set $\mathcal{F}'$ if f is the family of the first (last) operation of the set, respectively. Note that we have $s_{\mathcal{F}'} = \min_{f \in \mathcal{F}'} s_{f \to \mathcal{F}'} = \min_{f \in \mathcal{F}'} s_{\mathcal{F}' \to f}$.

Brucker and Thiele (1996) use several TSP lower bounds to estimate $s_{f \to \mathcal{F}'}$ and $s_{\mathcal{F}' \to f}$ and use also the optimal value obtained by complete enumeration, with an $O(|\mathcal{F}'|!)$ time complexity.

In this paper we follow the method of Vilím and Barták (2002) who propose to precompute values $s_{f \to \mathcal{F}'}$ and $s_{\mathcal{F}' \to f}$ once before starting the search. Each family set $\mathcal{F}'$ is coded by a binary representation of the corresponding permutation so that $s_{f \to \mathcal{F}'}$ and $s_{\mathcal{F}' \to f}$ are obtainable in $O(1)$ once the precomputation is made. Values $s_{f \to \mathcal{F}'}$ for each subset $\mathcal{F}' \subseteq \mathcal{F}_{\mathcal{O}_k}$ and each family $f \in \mathcal{F}'$ can be computed in $O(|\mathcal{F}_{\mathcal{O}_k}|^2 2^{|\mathcal{F}_{\mathcal{O}_k}|})$ by the following recursion:

$$s_{f \to \{f\}} = 0 \quad \forall F_f \subseteq \mathcal{F}_{\mathcal{O}_k}, \tag{16}$$

$$s_{f \to \mathcal{F}' \cup \{f\}} = \min_{g \in \mathcal{F}'} \{s_{fg} + s_{g \to \mathcal{F}'}\} \quad \forall \mathcal{F}' \subseteq \mathcal{F}_{\mathcal{O}_k}, \forall f \in \mathcal{F}_{\mathcal{O}_k} \setminus \mathcal{F}'. \tag{17}$$

Values $s_{\mathcal{F}' \to f}$ for each subset $\mathcal{F}' \subseteq \mathcal{F}_{\mathcal{O}_k}$ and each family $f \in \mathcal{F}'$ can be computed in a symmetric fashion.

We also introduce the following notation linked to a set of operations Ω useful in what follows.

$$r_\Omega = \min_{O_{ij} \in \Omega} r_{ij}, \tag{18}$$

$$d_\Omega = \max_{O_{ij} \in \Omega} d_{ij}. \tag{19}$$

8.2 Local constraint propagation

The function PROPAGATE($\mathcal{R}, \mathcal{D}, \mathcal{E}$) of our branch and bound procedure aims at computing adjustments on earliest start times $\mathcal{R}$ and on latest completion times $\mathcal{D}$ compatible with the current selection $\mathcal{E}$, returning a failure when unfeasibility is detected.

It calls iteratively functions UPDATEEARLIESTSTART, UPDATELATESTCOMPLETION, IMMEDIATESELECTION and EDGEFINDING, each performing specific deductions. The sequence of calls is repeated until a failure occurs or no more adjustments are performed. Algorithms UPDATEEARLIESTSTART and UPDATELATESTCOMPLETION are adapted from Brucker and Thiele (1996). They update r_{ij} and d_{ij} for all operations $O_{ij} \in \mathcal{O}$ considering only the precedence relations given by the job precedence constraints and current selection $\mathcal{E}$. The earliest start times can be updated with the following rules:

$$r_{ij} \geq s_{0im_{ij}} \quad \forall O_{ij} \in \mathcal{O}, \tag{20}$$

$$r_{ij} \geq r_{i(j-1)} + p_{i(j-1)} \quad \forall O_{ij} \in \mathcal{O}, j > 1, \tag{21}$$

$$r_{ij} \geq r_\Omega + p_\Omega + s_{\mathcal{F}_{\Omega \cup \{f_{ij}\}} \to f_{ij}} \quad \forall O_{ij} \in \mathcal{O}, \forall \Omega \subseteq \{O_{xy} | (O_{xy}, O_{ij}) \in \mathcal{E}\}. \tag{22}$$

The latest completion times can be updated with the following symmetric rules:

$$d_{ij} \leq T, \tag{23}$$

$$d_{ij} \leq d_{i(j+1)} - p_{i(j+1)} \quad \forall O_{ij} \in \mathcal{O}, j < m, \tag{24}$$

$$d_{ij} \leq d_\Omega - p_\Omega - s_{f_{ij} \to \mathcal{F}_{\Omega \cup \{f_{ij}\}}} \quad \forall O_{ij} \in \mathcal{O}, \forall \Omega \subseteq \{O_{xy}|(O_{xy}, O_{ij}) \in \mathcal{E}\}. \qquad (25)$$

Algorithm 3 describes UPDATEEARLIESTSTART. Its time complexity is $O(n^2 m \log n)$. Brucker and Thiele (1996) prove that the algorithm provides the largest adjustment according to rule (22) provided that the triangular inequality is verified for setup times. Note that the complexity of our algorithm is lower than the one obtained by Brucker and Thiele (1996) thanks to the setup preprocessing. Since the procedures are applied iteratively the consistency check performed in step 12 detects any cycle in $G(\mathcal{E}) = (X, U \cup \mathcal{E})$. Algorithm UPDATELATESTCOMPLETION is symmetric.

Algorithm 3 UPDATEEARLIESTSTART($\mathcal{R}, \mathcal{D}, \mathcal{E}$): returns failure or success

1: **for** $O_{ij} \in \mathcal{O}$ **do**
2:     $r_{ij} \leftarrow \max(r_{ij}, s_{0im_{ij}})$
3:     **if** $j > 1$ **then**
4:         $r_{ij} \leftarrow \max(r_{ij}, r_{i(j-1)} + p_{i(j-1)})$
5:     **end if**
6:     $\Omega \leftarrow \{O_{xy}|(O_{xy}, O_{ij}) \in \mathcal{E}\}$
7:     sort operations of Ω according to increasing earliest start times
8:     **for** $O_{xy} \in \Omega$ **do**
9:         $r_{ij} \leftarrow \max(r_{ij}, p_\Omega + s_{\mathcal{F}_{\Omega \cup \{f_{ij}\} \to f_{ij}}})$
10:        $\Omega \leftarrow \Omega \setminus \{O_{xy}\}$
11:    **end for**
12:    **if** $r_{ij} + p_{ij} > d_{ij}$ **then**
13:        **return** failure
14:    **end if**
15: **end for**
16: **return** success

Algorithm IMMEDIATESELECTION is a simple constraint propagation algorithm for the disjunctive resource constraints whose objective is to detect new precedence constraints between two operations assigned to the same machine. It is based on the rules: $\forall M_k \in \mathcal{M}, \forall O_{ij}, O_{x,y} \in \mathcal{O}_k, O_{ij} \neq O_{x,y}$:

$$r_{ij} + p_{ij} + p_{xy} + s_{ixk} > d_{xy} \implies (O_{xy}, O_{ij}) \in \mathcal{E}, \qquad (26)$$

$$(O_{xy}, O_{ij}) \in \mathcal{E} \implies r_{ij} \geq r_{xy} + p_{xy} + s_{xik} \quad \text{and} \quad d_{xy} \leq d_{ij} - p_{ij} - s_{xik}. \qquad (27)$$

It is implemented trivially in $\mathcal{O}(n^2 m)$. IMMEDIATESELECTION returns a failure whenever any update on the time window of an operation O_{ij} yields $r_{ij} + p_{ij} > d_{ij}$.

Last, algorithm EDGEFINDING aims at performing additional immediate selections and time windows adjustments considering an operation O_{ij} assigned to a machine M_k and a set $\Omega \subseteq \mathcal{O}_k \setminus \{O_{ij}\}$. More precisely it is based on the following so-called primal and dual rules. The primal rule states that $\forall M_k \in \mathcal{M}, \forall O_{ij} \in \mathcal{O}_k, \forall \Omega \subseteq \mathcal{O}_k \setminus \{O_{ij}\}$:

$$r_{\Omega \cup \{o_{ij}\}} + p_{\Omega \cup \{o_{ij}\}} + s_{\mathcal{F}_{\Omega \cup \{o_{ij}\}}} > d_\Omega \implies \forall O_{xy} \in \Omega, (O_{xy}, O_{ij}) \in \mathcal{E}. \qquad (28)$$

The dual rule states that $\forall M_k \in \mathcal{M}, \forall O_{ij} \in \mathcal{O}_k, \forall \Omega \subseteq \mathcal{O}_k \setminus \{O_{ij}\}$

$$d_{\Omega \cup \{o_{ij}\}} - p_{\Omega \cup \{o_{ij}\}} - s_{\mathcal{F}_{\Omega \cup \{o_{ij}\}}} < r_\Omega \implies \forall O_{xy} \in \Omega, (O_{ij}, O_{xy}) \in \mathcal{E}. \qquad (29)$$

Roughly both primal and dual rules detect that operation O_{ij} is not insertable inside set Ω. The primal rule detects in addition that O_{ij} has to be scheduled after all the operations of Ω on machine M_k whereas the dual rule detects that O_{ij} has to be scheduled before all the operations of Ω. Rules (22) and (25) can be combined with the primal and dual rules, respectively, to perform the time windows adjustments brought by the new arcs added in $\mathcal{E}$.

The combination of the primal and dual rule gives the following feasibility check, also called consistency rule. It generates a failure whenever the following condition holds. $\forall M_k \in \mathcal{M}, \forall \Omega \subseteq \mathcal{O}_k$

$$r_\Omega + p_\Omega + s_{\mathcal{F}_\Omega} > d_\Omega. \tag{30}$$

Vilím and Barták (2002) propose a $O(|\mathcal{F}|n^2)$ algorithm to perform all edge finding deductions and an other $O(n \log n)$ algorithm to perform all deductions of the consistency rule. In this paper we propose an extension of Nuijten's edge finding algorithm (Nuijten 1994) for which all consistency checks and adjustments are performed jointly. This extension, which runs in $O(|\mathcal{F}|n^2)$, is detailed in Algorithm 4 for the primal part.

We recall the principles of Nuijten's algorithm and explain the changes we have made to take account of setup times, in the spirit of the method proposed by Vilím and Barták (2002).

The main loop enumerate all possible operation families f. The second level loop makes all the updates relative to all relevant sets Ω such that $d_\Omega = d_{ij}$. The first third-level loop (steps 6–17), plays two different roles. Through the update of C, it performs the consistency checks (30). Since setup times verify the triangular inequality, set $\Omega(O_{xy}, O_{ij}) = \{O_{vw} \in \mathcal{O}_k | d_\Omega \le d_{ij}, r_\Omega \ge r_{xy}\}$ is dominant for rule (30) over any other set Ω' such that $r_{\Omega'} = r_{xy}$ and $r_{\Omega'} = d_{ij}$. The set of sets Ω generated at step 9 includes all sets $\Omega(O_{xy}, O_{ij})$. Hence C provides the largest possible left side value of inequality (30). The second role is to prepare, through the computation of D, the updates of the earliest start times of family f operations due to rules (28) and (22). Hence at step 16, c_{xy} is the largest earliest start time update of an operation O_{xy} if O_{xy} has to be scheduled after all previously generated relevant sets Ω, such that $r_\Omega \ge r_{xy}$. Note that this update is only valid here if $f_{xy} = f$. The second third-level loop (steps 6–17) makes the adjustments detected by rules (28) and (22). The same relevant sets Ω are generated in the reverse order, starting from the last generated one. At step 21, H corresponds to the largest earliest completion time of the all previously generated sets Ω such that $r_\Omega \le r_{xy}$ provided that a type f operation will be inserted in the set, considering only the setup time of this operation (the left side of the inequality of rule (28) minus the operation duration which is here still unknown). Any operation O_{xy} such that $d_{xy} > d_{ij}$ is candidate for being updated by rule (28). Step 25 performs the test of rule (28) for such an operation O_{xy} and current set Ω, verifying $r_\Omega \ge r_{xy}$. Note that the current set is stronger than all the sets that remains to be generated for operation O_{xy} and rule (28), since the setup times verify the triangular inequality and $P = p_\Omega$ is maximal. Step 26 updates the earliest start time if the test is positive using value c_{xy} (see definition above). Step 28 performs the test of rule (28) for operation O_{xy} and the strongest set previously generated, and verifying $r_\Omega \le r_{xy}$ (see computation of H above). Step 29 performs the update of the earliest start time using value D (see definition above), the largest earliest completion time among all sets Ω such that $d_\Omega = d_{ij}$, (because no operation of any of these sets can start after O_{xy}).

8.3 Global constraint propagation

At the root node of the branch and bound tree, we perform a global constraint propagation algorithm, so-called SHAVING, which can be seen as a one level breadth first search from

---

Algorithm 4 PRIMALEDGEFINDING($k, \mathcal{R}, \mathcal{D}$): returns failure or success

---

1: **for** $f \in \mathcal{F}_{\mathcal{O}_k}$ **do**
2:    **for** $O_{ij} \in \mathcal{O}_k$ according to increasing earliest start times **do**
3:       $P \leftarrow 0$
4:       $C \leftarrow 0$
5:       $\Omega \leftarrow \emptyset$
6:       **for** $O_{xy} \in \mathcal{O}_k$ according to decreasing earliest start times **do**
7:          **if** $d_{xy} \leq d_{ij}$ **then**
8:             $P \leftarrow P + p_{xy}$
9:             $\Omega \leftarrow \Omega \cup \{O_{xy}\}$
10:            $C \leftarrow \max(C, r_{xy} + P + s_{\mathcal{F}_\Omega})$
11:            $D \leftarrow \max(D, r_{xy} + P + s_{\mathcal{F}_{\Omega \cup \{f\}} \to f})$
12:            **if** $C > d_{ij}$ **then**
13:               **return** failure
14:            **end if**
15:          **end if**
16:          $c_{xy} \leftarrow D$
17:       **end for**
18:       $H \leftarrow 0$
19:       **for** $O_{xy} \in \mathcal{O}_k$ according to increasing earliest start times **do**
20:          **if** $d_{xy} \leq d_{ij}$ **then**
21:             $H \leftarrow \max(H, r_{xy} + P + s_{\mathcal{F}_{\Omega \cup \{f\}}})$
22:             $P \leftarrow P - p_{xy}$
23:             $\Omega \leftarrow \Omega \setminus \{O_{xy}\}$
24:          **else if** $f_{xy} = f$ **then**
25:             **if** $r_{xy} + P + p_{xy} + s_{\mathcal{F}_{\Omega \cup \{f\}}} > d_{ij}$ **then**
26:               $r_{xy} \leftarrow \max(r_{xy}, c_{xy})$
27:             **end if**
28:             **if** $H + p_{xy} > d_{ij}$ **then**
29:               $r_{xy} \leftarrow \max(r_{xy}, D)$
30:             **end if**
31:          **end if**
32:       **end for**
33:    **end for**
34: **end for**
35: **return** success

---

the root node where the possible start times of all operations inside their time windows are tried, yielding possibly $\sum_{O_{ij} \in \mathcal{O}} (d_{ij} - p_{ij} - r_{ij})$ tries. More precisely for each operation O_{ij}, a tentative to reduce its time window to $[r_{ij}, r_{ij}]$ is done and function PROPAGATE is called. If a failure occurs then r_{ij} is set to $r_{ij} + 1$ and the process iterates until a global failure occurs (the time window becomes empty) or there are no more adjustments. If no global failure occurs, another operation is considered. The same process is carried out for the latest start time. At the end of the SHAVING process, either a global failure is detected, or the time windows are "shaved". We refer to (Peridy 1996; Martin and Shmoys 1996; Torres and Lopez 2000; Demassey et al. 2005) for more efficient implementations and other variants of the shaving technique.

9 Feasibility tests based on TSPTW solutions

The TSPTW can be defined as a particular elementary shortest path problem with resource constraints (ESPPRC) (Feillet et al. 2004). The ESPPRC considers a network, a origin node, a destination node and arcs valued by a cost c_{uv} and consuming resources. The number of resources is denoted Q. Traversing an arc (u, v) consumes an amount l_{uv}^q of each resource q, with $1 \leq q \leq Q$. Values l_{uv}^q are assumed to satisfy the triangle inequality for each resource. Each node u of the network is associated with an interval $[a_u^q, b_u^q]$ such that the consumption of a resource q along a path from the origin node to u is constrained to belong to the interval. More precisely, if the elementary path uses arc (u, v), the consumption W_v^q of resource q between the origin node and v has to satisfy $W_v^q \geq \max(a_v^q, W_u^q + l_{uv}^q)$ and $W_v^q \leq b_v^q$. The objective is to find an elementary path of minimal cost from the origin to the destination node while satisfying these resource constraints.

F-TSPTW relaxation of (*FP*) can be represented as an ESPPRC. Let us consider a machine k and the set of operations $\mathcal{O}_k$. A node is introduced for each operation of $\mathcal{O}_k$. An arc (u, v) is introduced between each pair of operations $(u = O_{ij}, v = O_{xy})$, with a cost $c_{uv} = -M$, where M is an arbitrary large constant. A single resource ($Q = 1$) is defined, with $l_{uv}^1 = s_{ixk} + p_{xy}$. The resource window for node $u = O_{ij}$ is set as $[a_u^1, b_u^1] = [r_{ij} + p_{ij}, d_{ij}]$. Two additional nodes are introduced for the origin and the destination. An arc with a resource consumption s_{0ik} is added between the origin and every operation node $u = O_{ij}$, while an arc with a resource consumption 0 is added between every operation node and the destination. Note that the triangular inequality is respected since we have $s_{ijk} \leq s_{izk} + s_{zjk} \implies s_{ijk} + p_j \leq s_{izk} + p_z + s_{zjk} + p_j$, for all distinct $i \in [0, n]$, $j, z \in [1, n]$. Costs are defined such that the optimal ESPPRC solution visits all nodes if possible. Solving the ESPPRC thus permits to solve the F-TSPTW, as expected.

We solve the problem by using the dynamic algorithm proposed by Feillet et al. (2004) This algorithm follows the classical Bellman's algorithm. The principle is to associate a label with each possible partial path and to extend these labels checking the resource constraint, until the best feasible paths are obtained. Dummy resources (indicating the reachability of nodes) are introduced to preserve path elementarity. Dominance rules are used to compare labels and remove some of them. The algorithm is adapted to avoid extending a label for which a non-visited node is unreachable.

To speed up the resolution process, we use a lower bound computation for each generated label. Let t denote the time associated to a label corresponding to a partial path from the origin node to a node $v = O_{ij}$, i.e. $t = W_v^q$ for this label. Let $\overline{\mathcal{O}}_k \subset \mathcal{O}_k$ denote the set of operations on machine k still unvisited by the partial path. We can use consistency rule (30) adapted to set $\Omega \cup \{O_{ij}\}$. Then, it is not useful to extend the current label whenever $\exists \Omega \subseteq \overline{\mathcal{O}}_k$, $t + p_{ij} + p_\Omega + s_{f_{ij} \to \Omega \cup \{f_{ij}\}} > d_\Omega$. Since the number of generated labels can be very large we keep the complexity of the lower bound linear by considering only $\Omega = \overline{\mathcal{O}}_k$.

10 Branching scheme and dominance rule

Several branching schemes have been proposed for the standard job-shop problem (Jain and Meeran 1999). For the SDST-JSP, Brucker and Thiele (1996) consider blocks of operations on the critical path of a feasible schedule issued from the current node. Branching is based on the property that to improve the current feasible solution, at least one operation has to be scheduled either before all the other ones in its block (if it not already the first one) or after all the other ones (if it not already the last one in its block). Hence from a current selection,

the child nodes are generated by numerating these different possibilities and by making the induced selection updates. Focacci et al. (2000) use a sequence-based branching scheme by memorizing which operation has been scheduled last on each machine. The branching scheme consist in selecting an unscheduled operation and perform a binary branching for this operation. For the left branch, the operation is scheduled next in the sequence of its machine. For the right branch the operation is scheduled as a successor but not next of the last sequenced operation.

In this paper we use a mixed chronological/active branching scheme based on the current partial selection. Let $E_{ij} \subseteq E$ denote the set of undirected disjunctive edges connected to O_{ij} at the current node $v = (\mathcal{R}, \mathcal{D}, \mathcal{E})$. When the selection $\mathcal{E}$ is complete, we have $E_{ij} = \emptyset$, for each operation $O_{ij} \in \mathcal{O}$.

Let $O_{i^*j^*}$ denote the operation such that $E_{i^*j^*} \neq \emptyset$ and

$$r_{i^*j^*} = \min_{O_{ij} \in \mathcal{O}, E_{ij} \neq \emptyset} r_{ij}. \tag{31}$$

Let $\mathcal{C}$ denote the set of operations "conflicting" with $O_{i^*j^*}$, i.e. such that:

$$\mathcal{C} = \{O_{ij} \in \mathcal{O} | E_{ij} \neq \emptyset, m_{ij} = m_{i^*j^*}, r_{ij} < r_{i^*j^*} + p_{i^*j^*} + s_{i^*im_{ij}}\}. \tag{32}$$

The branching scheme we propose aims generates $|\mathcal{C}|$ nodes $\{v_{ij}\}_{O_{ij} \in \mathcal{C}}$ from the current node v where $v_{ij} = (\mathcal{R}, \mathcal{D}, \mathcal{E} \cup \{(O_{ij}, O_{xy})\}_{O_{xy} \in \mathcal{C} \setminus O_{ij}})$.

It has to be underlined that such a conflict set definition may result in non active schedules since the conflict set is defined only with necessary conditions. Indeed, there may be some operations in $\mathcal{C}$ whose earliest start time is not fixed and may increase in the descendence of the current node. This cannot be the case for operation $O_{i^*j^*}$ because it has the smallest earliest start time among operations linked to undirected disjunctive edges. However considering an operation $O_{ij} \neq O_{i^*j^*}$ of set $\mathcal{C}$ then a child node of v_{ij} for which r_{ij} becomes greater than or equal to $r_{i^*j^*} + p_{i^*j^*} + s_{i^*im_{ij}}$ can be erased. This can be simply achieved by setting $d_{ij} = r_{i^*j^*} + p_{i^*j^*} + s_{i^*im_{ij}} + p_{ij} - 1$.

11 Computational results

In this section, we present computational experiments conducted to evaluate the quality of our approach. For this purpose, we use the benchmark instances from Brucker and Thiele (1996). These instances are issued from the classical Lawrence instances (Lawrence 1984) devoted to the Job Shop Problem, introducing setup times. Each instance is characterized by a number of machines, a number of jobs to be scheduled and a number of setup types for the operations. These three parameters define a triplet with the format (machines × jobs × types). There are 15 instances with sequence dependent setup times (named t2-ps01 to t2-ps15). Instances t2-ps01 to t2-ps05 are of type $5 \times 10 \times 5$ (small instances). Instances t2-ps06 to t2-ps10 are of type $5 \times 15 \times 5$ (medium instances). Instances t2-ps11 to t2-ps15 are of type $5 \times 20 \times 10$ instances (large instances). Brucker and Thiele (1996) remarked that the corresponding Lawrence instances without setup times (LA01 to LA15) were all easily solved by their branch and bound method. On the opposite, only the small sequence-dependent setup time instances (t2-ps01 to t2-ps05) were solved to optimality.

Since the branch and bound method is designed to solve problem $FP(T)$ we apply it different times to find the optimal solution of (P) according to principle (6). Due to the successive runs, we set a limit NBB to the number of nodes of the branch and bound described in Sect. 6 and we also apply the following refinements:

(i) To keep the time spent at each branch and bound node reasonable, we set a limit *NDP* on the number of iterations (label extensions) of the dynamic programming algorithm used to solve the TSPTW relaxations (see Sect. 9). When this limit is reached, nothing can be deduced from SOLVETSPTW and branching is necessary.

(ii) For diversification purposes, we also use a randomized version of the heuristic SERIALSGS described in Sect. 7. This randomized version selects at step 4 an operation different than the one with the minimal priority according to a random factor. At each node, SERIALSGS is called 6 times with the priority values given in Sect. 7 but for both the deterministic and randomized versions.

(iii) The heuristic keeps track of the best found solution, even if its makespan is greater than the current threshold *T*. Furthermore, each time the heuristic improves the best solution found so far, the randomized version of SERIALSGS is called *NH* times with the current priority value, for intensification purposes.

Then the following strategies have been implemented to solve (P):

Strategy 1: A binary search is performed between an initial trivial lower bound (set to the longest job duration) and an initial upper bound set of one call of the serial SGS with the earliest start priority rule. Parameters are $NBB = 500000$, $NDP = 1500$ and $NH = 50000$.

Strategy 2: An increasing linear search is performed, starting from the same initial trivial lower bound as for strategy 1 with parameters $NBB = 20 \times 10^6$, $NDP = 3000$, $NH = 50000$.

The algorithms are coded in C++ and the tests are carried out on a PC with AMD64 architecture under Linux. In Table 1, we give the results of strategy 1. For each instance, we give the best found lower (LB1) and upper (UB1) bounds, the total number of branch and bound nodes (#nodes), the total CPU time in seconds (CPU), the initial lower bound value (LB0), the initial upperbound value (UB0), the number of calls of BRANCHANDBOUND by the binary search (IT) and the number of times a TSPTW relaxation could not be solved because the iteration limit *NDP* has been reached (TO).

As seen in Table 1, Strategy 1 solves to optimality 8 problems out of 15: all the 5 small instances, 3 out of 5 medium instances and none of the large instances. The previously unsolved problem t2-ps06 is closed with an optimal value of 1009, in 6157 seconds. The large CPU time values for the medium and large instances are due to the very large initial gap (since we did not take the best known lower and upper bound as initial bounds, but rather trivial ones) and to the maximal number of nodes which is reached for all values between the best found lower and upper bounds. All the best known lower bounds (see also Table 4) of the previously unsolved instances are improved. Two best known upper bounds are improved for instances t2-ps06 and t2-ps14 (see also Table 5).

In Table 2, we give the results of strategy 2 giving the same information as for strategy 1, but omitting the number of iterations and the number of time outs for the dynamic programming method. Note the number of iterations is obtained here by subtracting the initial lower bound from the obtained lower bound. For the computation of lower bounds, strategy 2 obtains better results than strategy 1 (see also Table 4). Indeed, strategy 2 solves all the small instances and 4 out of 5 medium instances. The previously unsolved problem t2-ps08 is closed with an optimal value of 963 in 349923 seconds. None of the large instances can be solved to optimality but all lower bounds are significantly improved, at the expense of very large computational times. Except for instance t2-ps08, the upper bounds found by strategy 2 are not as good as the ones obtained by strategy 1 (see also Table 5). This is not surprising since the linear search is performed increasingly from the lower bound. Consequently, the heuristics are never guided by a feasible makespan value, except for the optimal one.

Besides the results shown in Tables 1 and 2, we notice that solving the TSPTW problems at each node appears crucial for the efficiency of the method. To take a single example,

Table 1 Best lower and upper bounds obtained by strategy 1

Problem	LB1	UB1	#nodes	CPU	LB0	UB0	#IT	#TO
t2-ps01	**798***	**798***	173697	400.6	433	844	8	747
t2-ps02	**784***	**784***	24380	105.1	434	992	7	2810
t2-ps03	**749***	**749***	150885	352.6	359	946	9	8323
t2-ps04	**730***	**730***	2414	16.3	399	921	7	215
t2-ps05	**691***	**691***	3345	16.2	390	733	5	182
t2-ps06	<u>**1009**</u>*	**1009***	1143776	6156.2	433	1120	8	150768
t2-ps07	**970***	**970***	1533045	10012.6	416	1129	10	336768
t2-ps08	<u>946</u>	982	2583663	22716.5	399	1066	10	507641
t2-ps09	<u>1049</u>	1061	2210868	26754.5	412	1174	9	490508
t2-ps10	**1018***	1047	1499685	6391.1	463	1187	9	20895
t2-ps11	<u>1373</u>	1494	4512218	39489.9	483	1719	13	671506
t2-ps12	<u>1219</u>	1381	3072119	26678.8	498	1425	10	411984
t2-ps13	<u>1317</u>	1457	5414824	50336.8	462	1531	14	864255
t2-ps14	<u>1429</u>	**1483**	3583123	37236.1	463	1549	10	509888
t2-ps15	<u>1392</u>	1661	4342505	46590	431	1749	13	696812

Value in bold: best known result on the instance is reached
Value underlined: previously best known result on the instance is improved
*Tight bound

removing the TSPTW resolution (i.e. only the constraint propagation algorithms are applied) for the t2-ps05 instance increases the number of nodes up to 389992 nodes and the CPU time to 346.71 whereas the problem is solved in 3345 nodes and 16.2 seconds if the TSPTW relaxation is used. Martin and Shmoys (1996) notice that for the standard job-shop problem, solving exactly the one-machine problem at each node with the Carlier's algorithms brings only a little improvement compared to using edge finding and immediate selections. This seems to be no more the case when sequence-dependent setup times are introduced.

Table 3 provides the results of the strategy 1-based exact method on the corresponding job-shop instances without setup times (Lawrence 1984). All of these instances are easily solved by our method, which confirms the difficulty of the setup times constraints, as stated by Brucker and Thiele (1996).

Table 4 gives the result of several lower bounds on the 15 SDST-JSP instances, to evaluate the different components of our method and to make comparisons with other lower bounds. Column LB0 recalls the value of the trivial lower bound based on longest path computations in the precedence graph. Column SHAVING gives the lower bound obtained if only the shaving process is applied, i.e. the greatest value for which shaving can prove infeasibility. The lower bound value and the CPU times are given. Column TSPTW gives the lower bound obtained by applying only shaving followed by the resolution of the TSPTW relaxation. This is done by setting the branch and bound node limit to 1 and by setting no limit to the number of label extensions of the dynamic programming algorithm. The lower bound value and the maximal experienced number of iterations of the dynamic programming algorithm are provided. Column LB1 and LB2 recall the results of the best lower bounds by the two strategies we have implemented. We provide the results obtained by the root node lower

Table 2 Best lower and upper bounds obtained by strategy 2

Problem	LB2	UB2	#nodes	CPU	LB0	UB0
t2-ps01	**798***	**798***	1	56.7	433	844
t2-ps02	**784***	**784***	9498	242.3	434	992
t2-ps03	**749***	**749***	181090	699.3	359	946
t2-ps04	**730***	**730***	3594	251.6	399	921
t2-ps05	**691***	**691***	2918	58.2	390	733
t2-ps06	**1009**_*	**1009**_*	256520	1797.6	433	1120
t2-ps07	**970***	**970***	71106	781.8	416	1129
t2-ps08	**963**_*	**963**_*	33491952	349923	399	1066
t2-ps09	**1051**	1061	20105686	169582	412	1174
t2-ps10	**1018**_*	**1018**_*	227	35.1	463	1187
t2-ps11	**1395**	1617	57981046	916833	483	1719
t2-ps12	**1242**	1424	69387629	914086	498	1425
t2-ps13	**1342**	1457	58651163	895059	462	1531
t2-ps14	**1432**	1499	19999854	306899	463	1549
t2-ps15	**1406**	1671	52337255	792196	431	1749

Value in bold: best known result on the instance is reached
Value underlined: previously best known result on the instance is improved
*Tight bound

bound of the branch and bound method proposed by Brucker and Thiele (1996) (BT96), the root node and the best lower bound of the branch and bound method proposed in Artigues et al. (2004) (ABF04). The root node lower bound of Artigues et al. (2004) corresponds to the resolution of the TSPTW relaxation without tightening the time window by the shaving process. The best lower bound of (Artigues et al. 2004) was obtained by branch and bound.

The following ranking can be expected and is verified:
LB(BT96) < LB root(ABF04), LB(SHAVING) < LB(TSPTW) < LB1 < LB2. Note, however, that since there is a number of dynamic programming iterations limited to 1500 in our implementation of strategy 1, the TSPTW bound is not obtained without branching in our methods. One can observe that truncating the dynamic programming was necessary, due to an important number of iterations (label extensions) for the medium instances. The TSPTW bound could not even be obtained in a reasonable amount of time for the large instances. The shaving based-bound is tight 2 times and is generally better than the TSPTW bound without shaving (LB root ABF04) and than the Brucker and Thiele (1996) bound (LB root BT96), which indicates the good results of the constraint propagation algorithms and confirms the power of the shaving technique already experienced for the standard job-shop problem (Peridy 1996; Martin and Shmoys 1996). The computational time of the shaving algorithm is rather high but one could considerably accelerate it by performing a bisection search on the time window as described by Martin and Shmoys (1996). As a main result of our approach, strategy 2 obtains all the best known lower bounds on all instances, outperforming all other methods.

Table 5 gives the results of the state-of-the-art method in terms of the best obtained feasible solutions. BT96, FNL00, ALA05, ABF04, ABF05, BSV05 stand for the solutions found

Table 3 Results of strategy 1 on the instances with no setup times

Problem	Opt	#nodes	CPU	#IT	#TO
LA01 (t2-ps01)	**666***	1	7.5	7	1
LA02 (t2-ps02)	**655***	22150	70.1	8	752
LA03 (t2-ps03)	**597***	2	8.5	8	1
LA04 (t2-ps04)	**590***	42	13	8	8
LA05 (t2-ps05)	**593***	1	2.2	8	1
LA06 (t2-ps06)	**926***	1	4.6	6	1
LA07 (t2-ps07)	**890***	107	17.2	7	74
LA08 (t2-ps08)	**863***	1	27.42	8	1
LA09 (t2-ps09)	**951***	1	39.4	9	1
LA10 (t2-ps10)	**958***	1	32.2	8	1
LA11 (t2-ps11)	**1222***	1	14.6	9	1
LA12 (t2-ps12)	**1039***	15	112.6	9	15
LA13 (t2-ps13)	**1150***	19	29.33	8	19
LA14 (t2-ps14)	**1292***	1	207.1	8	1
LA15 (t2-ps15)	**1207***	1080	92.3	10	366

by (Brucker and Thiele 1996; Focacci et al. 2000; Artigues et al. 2005b; Artigues et al. 2004; Artigues et al. 2005a; Balas et al. 2005), respectively. The results show that the heuristic solutions found by our algorithm are not so far from the best known solutions mostly obtained by the shifting bottleneck heuristic of Balas et al. (2005) with reasonable computational requirements. We improve upon the results of Balas et al. (2005) for small instance t2-ps05, for 3 medium instances out of 5 (t2-ps06, t2-ps07 and t2-ps08) and for 1 large instance out of 5 (t2-ps14). The results of Balas et al. (2005) are better than ours for instance t2-ps09 and for all large instances except t2-ps14. As already underlined, the solutions found by strategy 1 on large instances are better than the one obtained by strategy 2. Note the initial heuristic (UB0), which carries out a single run of the serial SGS with the earliest-start priority-rule, obtains better results than the truncated branch-and-bound methods of Brucker and Thiele (1996) and Focacci et al. (2000) on 1 medium instance and on 3 large instances. It follows, as already mentioned by Artigues et al. (2005b), that the latter methods could be highly strengthened by integrating the serial SGS in the set of heuristics they use.

Last, Table 6 gives the evolution the gap between the best known lower and upper bound across all concerned studies since the paper of Brucker and Thiele (1996) as well as the number of improved lower bounds, upper bounds and closed instances obtained by each previous work. The results show that the present study has significantly reduced the gap on all unsolved instances. The gap is reduced to less than 1% on the unsolved medium instance and to less that 8% on all the large instances. Solving all the medium sized instances will be probably achieved in a near future. However the large instances remain challenging.

12 Concluding remarks

We have proposed a new exact method to solve the job-shop problem with sequence-dependent setup times which significantly reduces all gaps for the unsolved problem in-

Table 4 Lower bound comparison

Problem	LB0	SHAVING		TSPTW		LB1	LB2	LB root	LB root	LB
		LB	(CPU)	LB	(DPIT)			BT96	ABF04	ABF04
t2-ps01	433	781	(9.5)	796	(10847)	**798***	**798***	756	796	**798***
t2-ps02	434	745	(7.3)	745	(13667)	**784***	**784***	705	715	**784***
t2-ps03	359	710	(6.7)	710	(9079)	**749***	**749***	658	678	**749***
t2-ps04	399	707	(9.7)	707	(10945)	**730***	**730***	627	647	**730***
t2-ps05	390	687	(4.2)	690	(9049)	**691***	**691***	653	671	**691***
t2-ps06	433	1006	(52.2)	1006	(239221)	**1009***	**1009***	986	996	996
t2-ps07	416	**970***	(20.0)	**970***	(927515)	**970***	**970***	940	927	**970***
t2-ps08	399	923	(14.7)	930	(1130061)	946	**963***	913	923	923
t2-ps09	412	1011	(16.1)	1012	(858939)	1049	**1051**	1001	1012	1037
t2-ps10	463	**1018***	(19.3)	**1018***	(788711)	**1018***	**1018***	1008	**1018***	**1018***
t2-ps11	483	1360	(83.5)	–	–	1373	**1395**	1322	NA	NA
t2-ps12	498	1175	(145.8)	–	–	1219	**1242**	1139	1159	1159
t2-ps13	462	1280	(68.6)	–	–	1317	**1342**	1250	1250	1250
t2-ps14	463	1412	(181.8)	–	–	1429	**1432**	1402	NA	NA
t2-ps15	431	1357	(182.4)	–	–	1392	**1406**	1307	NA	NA

Value in bold: best known result on the instance is reached

Value underlined: previously best known result on the instance is improved

*Tight bound

Table 5 Upper bound comparison

Problem	UB0	BT96	FNL00	ALA05	ABF04	ABF05	BSV05	UB1	UB2
t2-ps01	844	**798***	**798***	818	**798***	**798***	**798***	**798***	**798***
t2-ps02	992	**784***	**784***	829	**784***	**784***	**784***	**784***	**784***
t2-ps03	946	**749***	**749***	782	**749***	771	**749***	**749***	**749***
t2-ps04	921	**730***	**730***	745	**730***	743	**730***	**730***	**730***
t2-ps05	733	**691***	**691***	704	**691***	693	693	**691***	**691***
t2-ps06	1120	1056	NA	1026	1026	1026	1018	<u>**1009**</u>*	<u>**1009**</u>*
t2-ps07	1129	1087	NA	1033	**970***	1022	1003	**970***	**970***
t2-ps08	1066	1096	NA	1002	1002	994	975	982	<u>**963**</u>*
t2-ps09	1174	1119	NA	**1060**	**1060**	**1060**	**1060**	1061	1061
t2-ps10	1187	1058	NA	1036	**1018***	**1018***	**1018***	1047	**1018***
t2-ps11	1719	1658	NA	1478	NA	1509	**1470**	1494	1617
t2-ps12	1425	1528	1448	1319	1319	1305	**1305**	1381	1424
t2-ps13	1531	1549	1658	1439	1439	1439	**1439**	1457	1457
t2-ps14	1549	1592	NA	1492	NA	1492	1485	<u>**1483**</u>	1499
t2-ps15	1749	1744	NA	1559	NA	1556	**1527**	1661	1671

Value in bold: best known result on the instance is reached
Value underlined: previously best known result on the instance is improved
*Tight bound

Table 6 Chronological evolution of the gap $(1 - LB/UB) \times 100$ for the SDST-JSP instances

Problem	BT96	FNL00	ALA05	ABF04	ABF05	BSV05	Our results
t2-ps01	0						
t2-ps02	0						
t2-ps03	0						
t2-ps04	0						
t2-ps05	0						
t2-ps07	13.52	–	9.00	0			
t2-ps10	4.73	–	2.70	0			
t2-ps06	6.62	–	3.90	2.92	–	2.16	0
t2-ps08	16.70	–	8.88	7.88	7.14	5.33	0
t2-ps09	10.54	–	5.56	2.17	–	–	0.85
t2-ps11	20.27	–	10.55	–	–	10.06	5.10
t2-ps12	25.46	21.33	13.64	12.13	11.18	–	4.83
t2-ps13	19.30	–	13.13	–	–	–	6.74
t2-ps14	11.93	–	6.03	–	–	5.59	3.44
t2-ps15	25.05	–	16.16	–	16.00	14.40	7.92
LB impr	15	0	0	6	0	0	8
UB impr	15	1	10	2	3	5	3
closed	5	0	0	2	0	0	2

stances proposed by Brucker and Thiele (1996). The results are obtained thanks to a cooperation between constraint propagation techniques extended to setup times and a truncated dynamic programming-based resolution of TSPTW relaxations. The feasible solutions are obtained by a serial schedule generation scheme with a priority-rule based on the solutions of the TSPTW relaxations.

The cooperation between constraint propagation and TSPTW relaxations could be further increased, for instance by solving smaller TSPTW corresponding to special subsets of operations sequenced on the same machine or conversely by designing time windows update methods based on the analysis of the TSPTW. The feasible solutions of the TSPTW could also be maintained by local search algorithms adapting to the branching decisions. In the near future, we will focus on improvements that can be brought to speed up the solving process of the dynamic programming algorithm. A better resolution of the TSPTW could be the key to solve larger instances.

References

Adams, J., Balas, E., & Zawack, D. (1988). The shifting bottleneck procedure for job-shop scheduling. *Management Science*, *34*, 391–401.

Allahverdi, A., Gupta, J. N. D., & Aldowaisan, T. (1999). A review of scheduling research involving setup considerations. *Omega*, *27*, 219–239.

Artigues, C., Belmokhtar, S., & Feillet, D. (2004). A new exact solution algorithm for the job shop problem with sequence-dependent setup times. In J. C. Régin & M. Rueher (Eds.), *1st international conference on integration of AI and OR techniques in constraint programming for combinatorial optimization problems: Vol. 3011. Lecture notes in computer science* (pp. 37–49). Berlin: Springer.

Artigues, C., Buscaylet, F., & Feillet, D. (2005a). Lower and upper bound for the job shop scheduling problem with sequence-dependent setup times. In *Proceedings of the second multidisciplinary international conference on scheduling: theory and applications (MISTA'2005)*. New York.

Artigues, C., Lopez, P., & Ayache, P. D. (2005b). Schedule generation schemes and priority rules for the job-shop problem with sequence-dependent setup times: Dominance properties and computational analysis. *Annals of Operations Research*, *138*(1), 21–52.

Balas, E. (1996). *New classes of efficiently solvable generalized traveling salesman problems* (Technical Report #MSRR-615). Graduate School of Industrial Administration, Carnegie Mellon University.

Balas, E., & Simonetti, N. (2001). Linear time dynamic programming algorithms for new classes of restricted TSPs. *INFORMS, Journal on Computing*, *13*, 56–75.

Balas, E., Simonetti, N., & Vazacopoulos, A. (2005). Job shop scheduling with setup-times, deadlines and precedence constraints. In *Proceedings of the second multidisciplinary international conference on scheduling: theory and applications (MISTA'2005)*. New York.

Baptiste, P., Le Pape, C., & Nuijten, W. (2001). In *Constraint-based scheduling: Vol. 39. International series in operations research & management science*. Berlin: Springer.

Blazewicz, J., Domschke, W., & Pesch, E. (1996). The job scheduling problem: Conventional and new solution techniques. *European Journal of Operational Research*, *93*(1), 1–33.

Brucker, P., & Thiele, O. (1996). A branch and bound method for the general-shop problem with sequence-dependent setup times. *Operations Research Spektrum*, *18*, 145–161.

Brucker, P., Jurisch, P., & Krämer, A. (1994a). The job-shop problem and immediate selection. *Annals of Operations Research*, *50*, 73–114.

Brucker, P., Jurisch, P., & Sievers, B. (1994b). A fast branch and bound algorithm for the job-shop scheduling problem. *Discrete Applied Mathematics*, *49*, 107–127.

Buscaylet, F., & Artigues, C. (2003). A fast tabu search method for the job-shop problem with sequence-dependent setup times. In *Metaheuristic international conference MIC'2003*.

Candido, M. A. B., Khator, S. K., & Barcias, R. M. (1998). A genetic algorithm based procedure for more realistic job shop scheduling problems. *International Journal of Production Research*, *36*(12), 3437–3457.

Carlier, J. (1982). The one machine sequencing problem. *European Journal of Operational Research*, *11*, 42–47.

Carlier, J., & Pinson, E. (1989). An algorithm for solving the job-shop problem. *Management Science*, *35*, 164–176.

Choi, I.-C., & Korkmaz, O. (1997). Job shop scheduling with separable sequence-dependent setups. *Annals of Operations Research*, *70*, 155–170.

Choi, I.-N., & Choi, D.-S. (2002). A local search algorithm for job-shop scheduling problems with alternative operations and sequence-dependent setups. *Computers and Industrial Engineering*, *42*, 43–58.

Conway, R. W., Maxwell, W. L., & Miller, L. W. (1967). *Theory of scheduling*. Reading: Addison–Wesley.

Demassey, S., Artigues, C., & Michelon, P. (2005). Constraint propagation-based cutting planes: an application to the resource-constrained project scheduling problem. *INFORMS Journal on Computing*, *17*(1), 52–65.

Feillet, D., Dejax, P., Gendreau, M., & Gueguen, C. (2004). An exact algorithm for the elementary shortest path problem with resource constraints: Application to some vehicle routing problems. *Networks*, *44*(3), 216–229.

Focacci, F., Laborie, P., & Nuijten, W. (2000). Solving scheduling problems with setup times and alternative resources. In *Fifth international conference on artificial intelligence planning and scheduling* (pp. 92–101).

Focacci, F., Lodi, A., & Milano, M. (2002). A hybrid exact algorithm for the tsptw. *INFORMS, Journal on Computing, 14*, 403–417.

Jain, A. S., & Meeran, S. (1999). Deterministic job-shop scheduling: Past, present and future. *European Journal of Operational Research*, *113*(2), 390–434.

Kim, S. C., & Bobrowski, P. M. (1994). Impact of sequence-dependent setup time on job shop scheduling performance. *International Journal of Production Research*, *32*(7), 1503–1520.

Kolisch, R. (1996). Serial and parallel resource-constrained project scheduling methods revisited: theory and computation. *European Journal of Operational Research*, *90*, 320–333.

Lawrence, S. (1984). *Resource constrained project scheduling: an experimental investigation of heuristic scheduling techniques (supplement)* (Technical report). Graduate School of Industrial Administration, Carnegie Mellon University.

Martin, P., & Shmoys, D. B. (1996). A new approach to computing optimal schedules for the job-shop scheduling problem. In W. H. Cunningham, S. T. McCormick & M. Queyranne (Eds.), *Proceedings of the 5th international conference on integer programming and combinatorial optimization IPCO'96* (pp. 389–403). Vancouver, British Columbia, Canada.

Mason, S. J., Fowler, J. W., & Matthew Carlyle, W. (2002). A modified shifting bottleneck heuristic for minimizing total weighted tardiness in complex job shops. *Journal of Scheduling*, *5*(3), 247–262.

Noivo, J. A., & Ramalhinho-Lourenço, H. (1998). *Solving two production scheduling problems with sequence-dependent set-up times* (Technical Report No. 138). Department of Economics and Business, Universitat Pompeu Fabra, Barcelona.

Nowicki, E., & Smutnicki, C. (1996). A fast taboo search algorithm for the job shop problem. *Management Science*, *42*, 797–813.

Nuijten, W. P. M. (1994). *Time and resource constrained scheduling: a constraint satisfaction approach*. Eindhoven University of Technology: Ph.D. thesis.

Ovacik, I. M., & Uzsoy, R. (1992). A shifting bottleneck algorithm for scheduling semiconductor testing operations. *Journal of Electronics Manufacturing*, *2*, 119–134.

Ovacik, I. M., & Uzsoy, R. (1994a). Exploiting shop floor status information to schedule complex job shop. *Journal of manufacturing systems*, *13*, 73–84.

Ovacik, I. M., & Uzsoy, R. (1994b). Rolling horizon algorithms for a single machine dynamic scheduling problem with sequence-dependent setup times. *International Journal of Production Research*, *32*(6), 1243–1263.

Peridy, L. (1996). *Le problème de job-shop: arbitrages et adjustments*. Université de Technologie de Compiègne: Ph.D. thesis.

Roy, B., & Sussman, B. (1964). *Les problèmes d'ordonnancement avec contraintes disjonctives* (Technical Report Note DS No. 9bis). SEMA, Paris.

Savelsberg, M. W. P. (1985). Local search in routing problems with time windows. *Annals of Operations Research*, *4*, 285–305.

Schutten, J. M. J. (1995). *Practical job shop scheduling* (Technical Report LPOM-95-12). Laboratory of Production and Operations Management, Department of Mechanical Engineering, University of Twente, The Netherlands.

Sun, X., & Noble, J. S. (1999). A modified shifting bottleneck approach to job shop scheduling with sequence dependent setups. *Journal of Manufacturing Systems*, *18*(6), 416–430.

Torres, P., & Lopez, P. (2000). Overview and possible extensions of shaving techniques for job-shop problems. In *Proceedings of the workshop on integration of AI and OR techniques in constraint programming for combinatorial optimization problems, CPAIOR'00* (pp. 181–186). Paderborn, Germany.

Vaessens, R. J. M., Aarts, E. H. L., & Lenstra, J. K. (1996). Job shop scheduling by local search. *INFORMS Journal on Computing*, *8*, 302–317.

Vilím, P., & Barták, R. (2002). Filtering algorithms for batch processing with sequence dependent setup times. In M. Ghallab, J. Hertzberg & P. Traverso (Eds.), *Proceedings of the sixth international conference on artificial intelligence planning and scheduling (AIPS 2002)* (pp. 312–320). Menlo Park: AAAI Press.

Wilbrecht, J. K., & Prescott, W. B. (1969). The influence of setup time on job shop performance. *Management Science*, *16*(4), B274–B280.

Zhou, C., & Egbelu, P. G. (1989). Scheduling in manufacturing shop with sequence-dependent setups. *Robotics and Computer Integrated Manufacturing*, *5*, 73–81.