# Multi-Machine Scheduling Lower Bounds Using Decision Diagrams☆

Pim van den Bogaerdt*, Mathijs de Weerdt

*Delft University of Technology*
*Faculty of Electrical Engineering, Mathematics and Computer Science*
*Van Mourik Broekmanweg 6, 2628 XE Delft, The Netherlands*

## Abstract

We consider parallel multi-machine scheduling with due times, where a partition of jobs is given where jobs in the same partition have a common release time, possibly precedence constraints, and cannot overlap. A formulation of decision diagrams for this problem greatly improves upon a more natural extension of the state-of-the-art for single-machine scheduling, and can provide decent lower bounds, outperforming existing solvers given the same short runtime limit, for problem instances with large time scales and tight due times.

---

☆This work is based on the first author's Master's thesis [1].
*Corresponding author.
*Email address:* `P.vandenBogaerdt@tudelft.nl` (Pim van den Bogaerdt)

## 1. Introduction

Multi-machine scheduling is a fundamental scheduling problem modeling a range of important real-life problems from scheduling tasks on a CPU to optimizing manufacturing processes [2, Ch.1]. Such problems have shown to be computationally very challenging and are an active area of research [3]. Recently, single-machine scheduling problems have been solved more quickly by lower bounds produced by relaxed decision diagrams [4, 5]. We show that this can be generalized to finding lower bounds for parallel multi-machine scheduling instances. Decision diagrams provide a new class of bounds, based on a discrete relaxation of the problem, which may be useful e.g. in combining with other bounds, such as additive bounding [6].

We are given $n$ *jobs* that are to be scheduled without preemption on $m \leq n$ *machines*. Let $\mathcal{J}$ denote the set of all jobs. Each job $j \in \mathcal{J}$ has a *processing time* and a *due time*; let $\mathbf{p}$ (resp. $\mathbf{d}$) denote the $n$-vector of processing (resp. due) times.

We assume the jobs are partitioned into $k$ subsets (partitions), and that the jobs in a given subset may not overlap in time ("partition constraints"). Let $\mathcal{P}$ be the set of partitions. We allow each partition $i \in \mathcal{P}$ to have a *release time*, meaning its jobs cannot start before that time. Let $\mathbf{r}$ be the $k$-vector of release times, and let $P(j)$ denote the partition of job $j$. We also allow *precedence constraints* between jobs of the same partition. The precedence constraint $i \to j$ specifies that job $i$ has to end before $j$ starts. Let $\omega(j) = \{i \mid i \to j\}$.

A *schedule* assigns each job a starting time and a machine such that each machine processes at most one job at a time. In a given schedule $S$, a job $j$ has *starting time* $s(j)$, and $m(j)$ denotes the machine $j$ is scheduled on. The *processing time of a machine $i$ is* $\max_{m(j)=i} (s(j) + p_j)$.

The objective function defines the cost of a schedule. We assume objective functions are non-decreasing in the job start times. This is a reasonable assumption: we typically want jobs to complete as soon as possible. We also assume the objective function can be written as a sum over jobs, i.e., $z(S) = \sum_j z_j(s(j))$. For example, the objective used in the experiments in this paper is the *total tardiness*, defined by $t(S) = \sum_j \max\{0, s(j) + p_j - d_j\}$. *Previous results.* Relaxations of Mixed Integer Programming formulations (MIPs) for multi-machine scheduling problems (e.g. [7]) can provide lower

2

bounds on the optimum. In particular, Baptiste et al. [8] focus on lower bounds for multi-machine scheduling by giving MIPs and a reduction to a flow problem, among other approaches. For single-machine scheduling, very good lower bounds can be found using decision diagrams [5, 4]. In spite of this promising result, there is no previous work on the use of decision diagrams for finding lower bounds for multi-machine scheduling. Finding lower bounds enables one to appreciate the quality of a schedule, and save time by not searching for a better schedule if the lower bound is equal to the optimum.

*Our results.* We present and prove correctness of a decision diagram (DD) formulation for a multi-machine scheduling problem which greatly improves upon a more natural extension of the state of the art for single-machine scheduling [5]. We use a known reduction of an optimal schedule to a permutation of jobs, and propose a decision diagram relaxation based on this representation. We furthermore show that this provides decent lower bounds, outperforming existing solvers given the same short runtime limit, for problem instances with large time scales and tight due times.

## 2. Background on Decision Diagrams

Many problems can be seen as a minimization problem involving variables $x_1, \ldots, x_n$ with finite domain, including the scheduling problem we consider. A *decision diagram* (DD) for an instance $I$ of such a problem is a directed graph with the following properties (see, e.g., [9] for an elaborate introduction into decision diagrams for optimization). The nodes are partitioned into $n+1$ *layers* $L_1, \ldots, L_{n+1}$, and each node $v$ contains a *state* $S$. For a node $v$ in layer $L_i$, the choices to $x_i$ we can make in state $S$ are defined by the set of *feasible choices* $F(S)$, and each choice $x_i \in F(S)$ induces an outgoing edge of $v$ with cost $c(S, x_i)$ to a node in layer $L_{i+1}$ with state $\varphi(S, x_i)$. Layer $L_1$ (resp. $L_{n+1}$) consists of a single *root node* (resp. *leaf node*).

A root-leaf path corresponds to an assignment to $x_1, \ldots, x_n$. The definitions of $S, F, c, \varphi$ are problem-specific and also depend on the type of DD. We discuss two types of DDs in this paper.

An *exact DD* is a DD such that the cost of a root-leaf path equals the objective value for the corresponding assignment, and the set of root-leaf paths corresponds exactly to the set of feasible solutions. The shortest root-leaf path in an exact DD corresponds with an optimal solution of $I$.

A problem with exact DDs is their possibly exponential size, and we therefore consider another type of DD: relaxed DDs. These provide *lower bounds* on the optimum. Given an integer $w \geq 1$, a *relaxed DD of width* $\leq w$ is a DD with the following additional properties. The cost of a root-leaf path is *at most as high* as the objective value for the corresponding assignment to $x_1, \ldots, x_n$; the set of root-leaf paths corresponds to a set of solutions that is a *superset* of the feasible solutions; and each layer contains at most $w$ nodes.

A relaxed DD can be built layer-by-layer. After building each layer, multiple nodes can be *merged* into one to prevent exceeding the allowed width $w$. This algorithm is called *top-down compilation*. We define the merged nodes by means of a merge operator $\oplus$ that transforms two states into one merged state. We assume the operator is associative, so that we can unambiguously allow it to have multiple states as input.

A theorem by Hooker [5, Theorem 1] describes when a merge operator indeed yields a relaxed DD (called *validity*). The theorem is fundamental for proving the correctness of our DD formulation (in Section 3). We present this theorem in a slightly more formal way, using the following definition of a state relaxation relation.

**Definition 1.** *A* state relaxation relation *is a binary operator $\preceq$ that operates on states with the following properties:*

- *(R1) $\preceq$ is reflexive and transitive.*

- *(R2) If $S' \preceq S$, then $F(S') \supseteq F(S)$.*

- *(R3) If $S' \preceq S$, then for $j \in F(S), c(S', j) \leq c(S, j)$.*

- *(R4) If $S' \preceq S$, then for $j \in F(S), \varphi(S', j) \preceq \varphi(S, j)$.*

**Theorem 1** (Hooker [5]). *Given a state relaxation relation $\preceq$, a merge operator $\oplus$ is valid if $S \oplus S' \preceq S, S'$ for all states $S, S'$.*

To compute the shortest root-leaf path in a DD, we keep track of the partial objective in each node. We can use this value also as a heuristic to choose what nodes to merge [5]: nodes with a high partial objective are not likely to lie on the shortest path. Thus, merging these may maintain the same lower bound on the shortest root-leaf path.

## 3. Multi-Machine Scheduling Using DDs

The DD formulation for multi-machine scheduling we introduce in this paper uses an efficient search space, state description, and merge operator. We explain these elements and prove their correctness in this section.

### 3.1. Minimal Schedules

To reduce the search space, we use the concept of *minimal* schedules and that it is sufficient to consider permutations of jobs with non-decreasing start times [10, 11, 12, 13, 14].

We say a schedule is *minimal* if each job is scheduled 1) as early as possible, and 2) on the machine with earliest completion time (ignoring this and later jobs; breaking ties by the smallest machine index).

The set of minimal schedules contains an optimal one: any optimal schedule can be made minimal without affecting optimality by iteratively scheduling a job earlier or by swapping the machine assignments of a subset of jobs.

Each minimal schedule can be represented as a permutation of the jobs by using the Smallest Processing Time First algorithm (SPTF) to transform a permutation into a schedule: schedule each job on the machine with smallest processing time so far (breaking ties by machine index), taking into account the release time of its partition. When scheduling a job, update the release time of its partition to the end time. SPTF takes into account precedence constraints if for each constraint $i \rightarrow j$, job $i$ occurs before $j$ in the permutation.

To see that a minimal schedule can be generated by SPTF, consider the permutation obtained by sorting the jobs on start time, breaking ties by machine index. Executing SPTF on this permutation gives back the same schedule because SPTF chooses the machine and time corresponding to the definition of a minimal schedule [14].

This idea also works with the precedence and partition constraints. Namely, the permutation obtained adheres to the precedence constraints and so the schedule by SPTF also does. Further, the partition constraints may be considered as a "dynamic release time": as we schedule a job in a partition, the release time of the unscheduled jobs of this partition are changed to the end time of that job. Hence, when scheduling a job, the partition constraints reduce to release time constraints, and so do not pose a complication compared to [14].

Moreover, by construction, the jobs in the permutation obtained this way are sorted on non-decreasing start times, so we only need to consider such permutations [13 cited by 14], also when computing a lower bound on the optimum.

*3.2. Exact DD Formulation*

To define a DD, we give definitions for $S$, $F$, $c$, and $\varphi$ (as explained in Section 2). Our exact DD formulation for multi-machine scheduling represents precisely all permutations with non-decreasing start times.

The state $S(v)$ of a node $v$ is a tuple $(V, U, \mathbf{f}, \mathbf{t}, \mathbf{f^u}, \mathbf{t^u}, g)$. Below we introduce the terms in this tuple step-by-step, using a bar over the variables to indicate the value after a transition. First, the pair $(V, \mathbf{f})$ already provides sufficient information to build an exact DD, but without the partition and precedence constraints. Like [5], we define $V$ as the set of jobs that appear on the path from the root to $v$. However, instead of a single number $f$, we now have a *vector* $\mathbf{f} = (f_1, \ldots, f_m)$ of length $m$ that contains the total processing time *of each machine*. The root has $V = \varnothing$ and $\mathbf{f} = \mathbf{0}$. Since we consider identical machines, the ordering of elements in $\mathbf{f}$ is irrelevant. It is efficient for the merge operator of the relaxed DD to ensure that $\mathbf{f}$ is sorted in increasing order.

When we choose job $j$ as the next item in the permutation, the new state $\varphi(S, j)$ contains

$$\overline{V} = V \cup \{j\}.$$

To take the partition constraints into account, the state contains a $k$-vector $\mathbf{t}$, which represents when a job can start because the previous jobs of the partition have completed. The root has $\mathbf{t} = \mathbf{r}$ (release times). If we choose job $j$, resulting in completion time $x$, the new state $\bar{\mathbf{t}}$ is updated with

$$\overline{t_{P(j)}} = x.$$

For $1 \leq i \leq k, i \neq P(j)$, we set $\bar{t}_i = t_i$. In words, as we schedule job $j$, other jobs of its partition can only start after job $j$ ends. We can now define the transition for $\mathbf{f}$: job $j$ is scheduled at the first machine in $\mathbf{f}$ (like SPTF: recall that $\mathbf{f}$ is sorted), and can only start after time $t_{P(j)}$. That is,

$$\bar{\mathbf{f}} = \mathrm{sort}(\, (\max\{f_1, t_{P(j)}\} + p_j, f_2, \ldots, f_m)\,).$$

The cost of this transaction is computed as

$$c(S, j) = z_j(\max\{f_1, t_{P(j)}\}).$$

6

The set $U$ is included in the state to incorporate precedence constraints (similar to [4]). It represents which jobs occur on *some* path from the root to $v$, as opposed to $V$, which contains jobs that occur on *every* path from the root to $v$. The root has $U = \varnothing$. We only include a job $j$ in the feasible set $F(S)$ if $\omega(j) \subseteq U$. When we make a job choice $j$, we add $j$ to the set $U$ in the new node, just like we do for $V$:

$$\overline{U} = U \cup \{j\}.$$

So in an exact DD, $U = V$, but these sets may be different in a relaxed DD.

The last step of our exact DD formulation is related to only considering permutations with non-decreasing start times (to make the DD concise). For this, we include a number $g$ representing the starting time of the (last) incoming job, i.e.,

$$\bar{g} = \max\{f_1, t_{P(j)}\},$$

and an $m$-vector $\mathbf{f^u}$ and $k$-vector $\mathbf{t^u}$. The vectors $\mathbf{f^u}$ and $\mathbf{t^u}$ again represent the finishing times of the machines and partitions, respectively. In our exact DD formulation, $\mathbf{f^u} = \mathbf{f}$ and $\mathbf{t^u} = \mathbf{t}$, but in a relaxed DD these represent the respective upper bounds. The root has $\mathbf{f^u} = \mathbf{0}, \mathbf{t^u} = \mathbf{r}, g = 0$.

We may ignore any (next) job $j$ which would be scheduled at time $\max\{f_1^u, t_{P(j)}^u\} < g$, because it is sufficient to consider permutations with non-decreasing start times (see Subsection 3.1). Also, we do not schedule jobs that have already been scheduled ($V$) and jobs for which a preceding job $i \in \omega(j)$ has not been completed. The definition of the feasible set is thus

$$F(S) = \mathcal{J} - (V \cup \{j \mid \omega(j) \not\subseteq U\} \cup \{j \mid \max\{f_1^u, t_{P(j)}^u\} < g\}).$$

The transition for $\mathbf{f^u}$ uses these vectors as input also:

$$\overline{\mathbf{f^u}} = \text{sort}(\,(\max\{f_1^u, t_{P(j)}^u\} + p_j, f_2^u, \ldots, f_m^u)\,).$$

Similarly, we define $\overline{t_{P(j)}^u} = \max\{f_1^u, t_{P(j)}^u\} + p_j$ and $\overline{t_i^u} = t_i^u$ for $i \neq P(j)$.

Summarizing, the root node has state

$$S(r) = (\varnothing, \varnothing, \mathbf{0}, \mathbf{r}, \mathbf{0}, \mathbf{r}, 0),$$

and edges are defined by

$$\varphi((V, U, \mathbf{f}, \mathbf{t}, \mathbf{f^u}, \mathbf{t^u}, g), j) = (\overline{V}, \overline{U}, \bar{\mathbf{f}}, \bar{\mathbf{t}}, \overline{\mathbf{f^u}}, \overline{\mathbf{t^u}}, \bar{g}).$$

7

*3.3. Relaxed DD Formulation*

A relaxed DD formulation is implied by a state relaxation relation $\preceq$ and merge operator $\oplus$.

**Definition 2.** *The* state relaxation relation $\preceq$ *between two states holds if each of the following relations hold pairwise on the elements in the state tuples:* $\subseteq, \supseteq, \leq, \leq, \geq, \geq, \leq$. *On vectors,* $\leq$ *and* $\geq$ *denote pairwise comparisons.*

**Definition 3.** *The* merge operator $\oplus$ *on two states applies the following operators on the elements in the state tuples:* $\cap, \cup, \min, \min, \max, \max, \min$. *The* $\min$ *(*max*) operator on vectors takes pairwise minima (maxima).*

There is now a difference between $U$ and $V$. Intuitively, to obtain a valid relaxation, we need to underestimate $V$ (to ensure we do not disregard any schedule) but overestimate $U$ (to ensure we are not too strict regarding the precedence constraints). This, and the need for $\mathbf{f^u}$ and $\mathbf{t^u}$, becomes clear in the proof of Lemma 2.

To show that $\preceq$ is a state relaxation relation, we use the following lemma.

**Lemma 1.** *Let* $\mathbf{a} \leq \mathbf{b}$ *be sorted vectors of length* $m$, *and let* $a'_1, b'_1 \in \mathbb{R}$. *If* $a'_1 \leq b'_1$, *then* $\mathrm{sort}(\,(a'_1, a_2, a_3, \ldots, a_m)\,) \leq \mathrm{sort}(\,(b'_1, b_2, b_3, \ldots, b_m)\,)$.

The proof of this lemma is straightforward but technical; it can be found in the first author's Master's thesis [1, Section 3.3].

**Lemma 2.** *The relation* $\preceq$ *is a state relaxation operator (cf. Definition 1).*

*Proof.* We show each of the properties (R1)-(R4) in order. Let $S, S'$ be states on the same layer $i$ such that $S' \preceq S$. (R1) follows directly from the reflexivity and transitivity of $\subseteq, \supseteq, \leq, \geq$ (pairwise and regular).

For (R2), let $j$ be a job in $F(S)$. In other words, assume the following about $j$. First, $j \notin V$. Second, $\omega(j) \subseteq U$. Third, $\max\{f^u_1, t^u_{P(j)}\} \geq g$. We have to show that $j$ is in $F(S')$, that is, we have to show three similar predicates.

The predicate $j \notin V'$ follows from $V' \subseteq V$. Next, $\omega(j) \subseteq U'$ follows from $U' \supseteq U$. For the inequality, note that $\mathbf{f^{u'}} \geq \mathbf{f^u}, \mathbf{t^{u'}} \geq \mathbf{t^u}, g' \leq g$, so indeed $\max\{f^{u'}_1, t^{u'}_{P(j)}\} \geq \max\{f^u_1, t^u_{P(j)}\} \geq g \geq g'$.

For (R3), let $j$ be a feasible choice in $S$. We have to show $c(S', j) \leq c(S, j)$. The cost $c(S', j)$ is $z_j(\max\{f'_1, t'_{P(j)}\})$, which is indeed at most the cost of

making the choice in $S$, $z_j(\max\{f_1, t_{P(j)}\})$, since $\mathbf{f}' \leq \mathbf{f}$ and $\mathbf{t}' \leq \mathbf{t}$ and the $z_j$ are non-decreasing.

For (R4), consider a job choice $j$ that is feasible in $S$. Let $\overline{S'}$ and $\overline{S}$ denote the state after making this choice in $S'$ and $S$, respectively. We have to show $\overline{S'} \preceq \overline{S}$. We do this by showing the predicates in the definition of $\preceq$.

$\overline{V'} \subseteq \overline{V}$ and $\overline{U'} \supseteq \overline{U}$ follow directly from $V' \subseteq V, U' \supseteq U$.

For $\overline{\mathbf{f}'} \leq \overline{\mathbf{f}}$, consider the definition of $\varphi$:

$$\overline{\mathbf{f}} = \mathrm{sort}(\,(\max\{f_1, t_{P(j)}\} + p_j, f_2, \ldots, f_m)\,)$$
$$\overline{\mathbf{f}'} = \mathrm{sort}(\,(\max\{f_1', t_{P(j)}'\} + p_j, f_2', \ldots, f_m')\,)$$

We have $\max\{f_1', t_{P(j)}'\} + p_j \leq \max\{f_1, t_{P(j)}\} + p_j$ and we may thus apply Lemma 1 to conclude $\overline{\mathbf{f}'} \leq \overline{\mathbf{f}}$.

For $\overline{\mathbf{t}'} \leq \overline{\mathbf{t}}$, note that from $f_1' \leq f_1$ and $t_{P(j)}' \leq t_{P(j)}$, we get $\overline{t_{P(j)}'} \leq \overline{t_{P(j)}}$. For $1 \leq i \leq k, i \neq P(j)$, we have equality in $\overline{t_i'} \leq \overline{t_i}$.

For $\overline{\mathbf{f^{u'}}} \geq \overline{\mathbf{f^u}}$, consider the equivalent predicate $\overline{\mathbf{f^u}} \leq \overline{\mathbf{f^{u'}}}$, which we can prove similarly to $\overline{\mathbf{f}'} \leq \overline{\mathbf{f}}$. The inequality $\overline{\mathbf{t^{u'}}} \geq \overline{\mathbf{t^u}}$ can be shown similarly to $\overline{\mathbf{t}'} \leq \overline{\mathbf{t}}$. Lastly, we have $\overline{g'} = \max\{f_1', t_{P(j)}'\} \leq \max\{f_1, t_{P(j)}\} = \overline{g}$. $\qquad\square$

We now consider the merge operator $\oplus$. This merge operator is associative, as required. We now prove the validity of this merge operator.

**Theorem 2.** *The merge operator $\oplus$ is valid.*

*Proof.* Let $S, S'$ be states. Using Theorem 1 we only have to show that $S \oplus S' \preceq S, S'$.

First, $V \cap V' \subseteq V, V'$ and $U \cup U' \supseteq U, U'$ are trivial. Next, note that in the merged state $S \oplus S'$, the vectors $\min(\mathbf{f}, \mathbf{f}')$ and $\max(\mathbf{f^u}, \mathbf{f^{u'}})$ are automatically sorted because $\mathbf{f}, \mathbf{f}', \mathbf{f^u}, \mathbf{f^{u'}}$ are sorted. We have $\min(\mathbf{f}, \mathbf{f}') \leq \mathbf{f}, \mathbf{f}'$ because we take pairwise minima and do pairwise comparisons. Similarly, $\min(\mathbf{t}, \mathbf{t}') \leq \mathbf{t}, \mathbf{t}'$ as well as $\max(\mathbf{f^u}, \mathbf{f^{u'}}) \geq \mathbf{f^u}, \mathbf{f^{u'}}$ and $\max(\mathbf{t^u}, \mathbf{t^{u'}}) \geq \mathbf{t^u}, \mathbf{t^{u'}}$. Lastly, $\min\{g, g'\} \leq g, g'$ holds. $\qquad\square$

The nodes with the highest partial objective are merged, breaking ties by (high) *slack* of the last scheduled job. The slack is the difference between due time and completion time.

Figure 1: Average gap over time for our DD formulation and the more natural extension we introduced based on the single-machine formulation by Hooker [5]. [single column]



## 4. Experiments

We evaluate the relaxed DD formulation using the total tardiness objective function. Our implementation is written in C# (x64, .NET 4.7). Whenever we compare against an upper bound, we use the best upper bound found by any of the solvers. Also, in our implementation, we used parallelization for building and merging: each thread builds and merges part of each layer. Recall that we merge nodes with highest partial objective.

In all experiments, we supplied all solvers with an initial solution provided by an adaptation of [10] (dealing with partition and precedence constraints). This algorithm quickly gives a very good upper bound.

### 4.1. Comparison to More Natural Formulation

We investigate to what extent our formulation improves upon the following more natural extension of the single-machine scheduling formulation by [5]: instead of using permutations, use a list of $n$ (job, machine) pairs. Each job is scheduled on the machine it is paired to, and the order of jobs on each machine is given by the order of the list. Thus, the choices do not define just the job but also the machine. The costs, feasible sets and transition functions are defined accordingly. In this formulation the vector $\mathbf{f}$ is not sorted, and we do not apply the permutation optimization.

For the relaxation of both this and our new DD formulation, a width needs to be selected. We construct DDs iteratively with increasing width as follows: first we use the adaptation of [10] to obtain a schedule (upper bound).

10

Then we compute lower bounds for widths $2^{12}$ and $2^{15}$, and extrapolate these linearly on the logarithm to determine the width for which the lower bound is likely equal to the upper bound. We then repeat this process with this extrapolated width, or a width increased by a factor of $2^3$, whichever is smaller.

Regarding the instances, we used the parameters and distributions based on those used by [5] for single-machine scheduling. The processing times $\mathbf{p}$ were integers drawn from $\{10, \ldots, 16\}$, the release times $\mathbf{r}$ were drawn from $\{0, \ldots, 5n\}$, and the due times $\mathbf{d}$ were drawn from $\{r_{P(j)} + (p_j + x)/m \mid x = 4 \cdot 10, \ldots, 4 \cdot 16\}$. We added the divisor $m$ to model the fact that with $m$ machines, we can roughly finish $m$ times as quickly. The jobs were partitioned as evenly as possible, and each job had a random partition assigned. We set $n = 10, m = 2, k = 3$.

For half of the instances, we set precedence constraints as follows. Consider a partition with $p$ jobs $j_i, j_{i+1}, \ldots, j_{i+p-1}$. We added precedence constraints $j_i \rightarrow j_{i+1}, j_{i+2} \rightarrow j_{i+3}, \ldots$, ignoring the last job if $p$ is odd. The reason we used these precedence constraints is that adding more may cause the search space to be reduced to a rather large extent.

We used 1500 of such instances. This experiment was performed on an Intel Core i7-3537U (2.00 GHz, 4 threads), 12 GB RAM, Windows 10, with a maximum width of $2^{22}$. Figure 1 shows the decrease in gap (using the upper bound by [10]) for both our "real" DD formulation and the more natural formulation. We observe that with more time, and thus DDs with incrementally larger widths, the average gaps for both formulations decreases significantly. Moreover, we see that our formulation is clearly superior. For the more natural formulation, the median time for a non-zero lower bound was about 0.35 seconds, and for only 465 instances it gave a final lower bound equal to the upper bound. To compare, the new formulation we propose in this paper gave a lower bound equal to the upper bound for 1473 instances, and the median time to solve these instances to optimality was about 0.1 seconds.

We conclude that extending [5] to multi-machine scheduling is not trivial, and that the relaxation we introduced in this paper is essential to obtain decent performance.

### 4.2. Comparison to CP and MIP

We also compare our DD formulation to a constraint program (CP) and two mixed integer programs (MIPs). This experiment was run on an Intel

Figure 2: Constraint Program.

$$\text{Minimize} \quad \sum_j \max\{0, \text{StartOf}(I_j) + p_j - d_j\}$$

$$
\begin{aligned}
\text{s.t.} \quad & I_{j,M} \in \text{Intervals}([r_{P(j)}, \infty), p_j) && (j \in \mathcal{J}, \ M \in \mathcal{M}) \\
& I_j \in \{I_{j,M} \mid 1 \le M \le m\} && (j \in \mathcal{J}) \\
& \text{NoOverlap}(\{I_{j,M} \mid 1 \le j \le n\}) && (M \in \mathcal{M}) \\
& \text{NoOverlap}(\{I_j \mid P(j) = i\}) && (i \in \mathcal{P}) \\
& \text{EndBeforeStart}(i, j) && (i, j \in \mathcal{J}, \ i \to j)
\end{aligned}
$$

Figure 3: Extra constraints for the MIP by [7], with $i, j \in \mathcal{J}$ with $i < j$, and the $Z_{i,j}$ variables are binary.

$$
\begin{aligned}
& C_j - C_i + L(1 - Z_{i,j}) \ge p_i && (P(i) = P(j)) \\
& C_i - C_j + L Z_{i,j} \ge p_j && (P(i) = P(j)) \\
& C_j \ge r_{P(j)} + p_j && (\text{replaces (6)}) \\
& Z_{i,j} = 1 && (i \to j) \\
& Z_{i,j} = 0 && (j \to i)
\end{aligned}
$$

Xeon E5-1620 v2 (3.70 GHz, 8 threads), 8 GB RAM, Windows 7.

The CP (Figure 2) was solved by IBM ILOG CP Optimizer 12.8 and is based on one of the samples provided with the solver. We enabled presolving and used the "restart" method, which gives decent performance as well as intermediate lower bounds.

The MIPs were solved by both IBM ILOG CPLEX 12.8 and Gurobi 8.0.1 with default parameters; for each instance, the best result is reported. The first MIP is adapted from [7], which has few variables. See Figure 3. The second MIP is based on a time-indexed one given by [8]. It has a binary variable for each (job, time) pair, where "time" is any time step a job could be scheduled. For this MIP, we used the makespan of a random schedule as horizon $H$. See Figure 4. We used this MIP itself as well as its LP relaxation (as suggested in [8]), both with and without presolving, again reporting only the best result. We start measuring time after having built the model.

Figure 4: Extra constraints for the MIP by [8]. We define $g(t, j) = \max\{0, t - p_j + 1\}$ and $\omega^{-1}(j) = \{i \mid j \to i\}$.

$$\sum_{\substack{j \\ P(j)=i}} \sum_{t'=g(t,j)}^{t} x_{j,t'} \leq 1 \qquad (i \in \mathcal{P},\ 0 \leq t < H)$$

$$\sum_{i \in \omega^{-1}(j)} \sum_{t'=0}^{t} x_{i,t'} \ \leq\ (t+1) \cdot |\omega^{-1}(j)| \cdot (1 - x_{j,t})$$

$$(j \in \mathcal{J},\ 0 \leq t < H)$$

We created instances based on the TF/RDD model [15, 16] adapted to multiple machines. Given TF, RDD $\in [0..1]$, the due times were drawn uniformly from the integers between $\lfloor r_j + \sum_j p_j(1 - \text{TF} - \text{RDD}/2)/m \rfloor$ and $\lfloor r_j + \sum_j p_j(1 - \text{TF} + \text{RDD}/2)/m \rfloor$ inclusive, with negative values removed. Thus, TF (RDD) controls the expectation (variance) of the due times. The processing times were drawn from $\{75, \ldots, 150\}$. These are rather large because for much smaller times the MIP by [8] outperforms all other formulations. The release times are drawn from $\{0, \ldots, 111\}$. We set $n = 4k, m = k - 2$, let each partition contain 4 jobs, and set precedence constraints for half of the instances like before.

We fixed a width of $w = 1,000$ and gave each solver a time limit of 15 seconds. Figure 5 shows the performance for these instances after 50ms (our formulation was done within this time for all but two instances). We see that DD performs well if the due times are tight on average (large TF) and have a small variance (small RDD). The reason is that such due times cause relatively high tardiness, so that the merge heuristic is well-informed. In contrast, DD do not work well for TF = 0.2.

If more time is available, the CP, LP and MIP solvers give (much) better bounds. For example, after about 350ms, the average gap of CP and the MIP by [7] are better than DD for TF = 0.8, RDD = 0.2, $k = 7$, except if we fix $m = 2$. After about 3 seconds, the LP relaxation of [8] gives average gaps close to zero for TF = 0.8, and after about 5.5 seconds, its averages are all less than those of DD.

However, for TF = 0.8, RDD = 0.2, and $m = 2$, DD outperforms all

13

Figure 5: Average gap with standard deviation after 50ms for TF/RDD instances as a function of $k$, with $n = 4k, m = k - 2$. The number of instances per plot varies between 510 and 551. Bal. (Bap.) refers to the MIP based on [7] ([8]). The error bars are shifted based on the solver so as not to let them overlap. [single column]
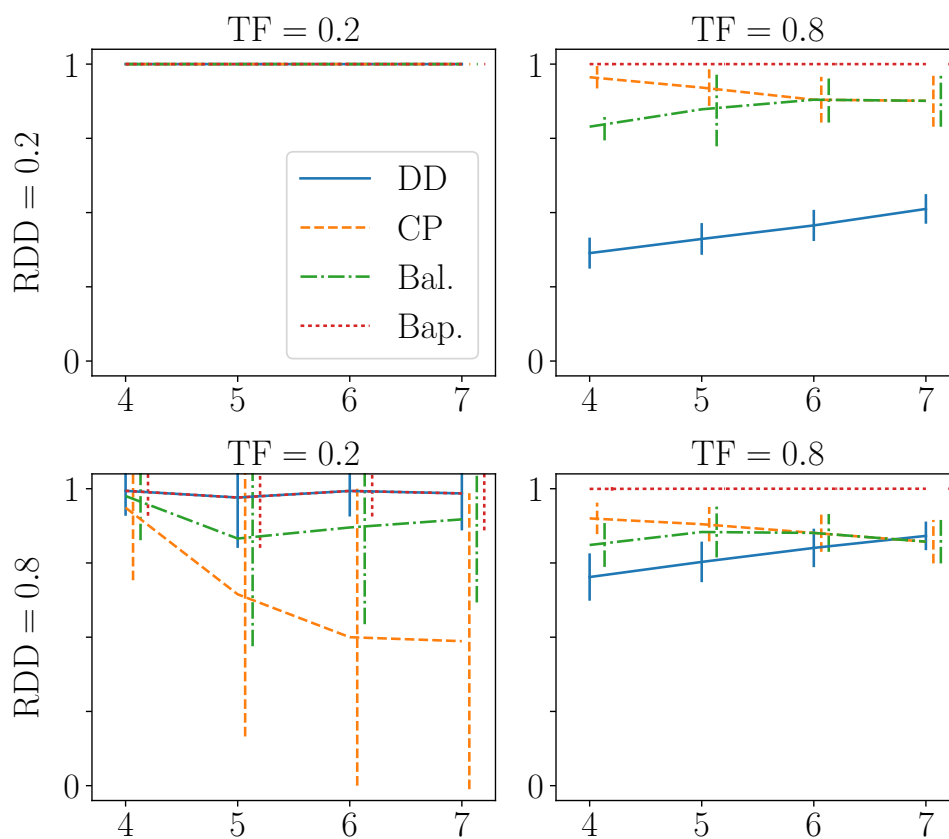
other models. Only after running about 20 to 100 times longer (1s to 5s), the model based on [8] eventually gives smaller average gaps than DD does. Moreover, this formulation uses many variables and constraints, and it takes time and memory to build these (which we did not take into account for the runtime). In particular, if the time scale is larger than what we used, this model may not be practical, while the runtime and memory of our DD formulation virtually do not depend on the time scale.

## 5. Discussion and Future Work

When aiming for better bounds, the memory usage in our formulation may become a bottleneck. Each node in our formulation uses $O(m + n)$ memory for its state, so that the total memory usage is $O(nw(m + n))$. We found empirically that one has to double $w$ to get a constant increase in the lower bound, so the memory is exponential in the quality. Although under specific circumstances, an exponential DD is inevitable with our model (see Appendix A in [1]), a more sophisticated merge heuristic may reduce $w$ significantly and is thus an interesting direction for further research.

Given that the DD formulation proposed in this paper is so effective if a small amount of runtime is available, it furthermore is very promising to embed it in algorithms that have worked with DDs successfully in the past, such as branch and bound [9, Chapter 6], logic-based Benders decomposition [17], constraint programming filtering techniques [4], and Lagrangian methods [18].

Finally, we may tackle other multi-machine scheduling problems using similar DDs, such as for unrelated machines.

## Acknowledgements

## References

[1] P. van den Bogaerdt, Multi-Machine Scheduling Lower Bounds Using Decision Diagrams, Master's thesis, Delft University of Technology, The Netherlands, 2018.

[2] M. L. Pinedo, Scheduling: theory, algorithms, and systems, Springer, 4 edn., 2016.

[3] M. O. Adamu, A. O. Adewumi, A survey of single machine scheduling to minimize weighted number of tardy jobs, Journal of Industrial & Management Optimization 10 (1) (2014) 219–241.

[4] A. A. Cire, W.-J. van Hoeve, Multivalued decision diagrams for sequencing problems, Operations Research 61 (6) (2013) 1411–1428.

[5] J. N. Hooker, Job Sequencing Bounds from Decision Diagrams, in: Proceedings of the International Conference on Principles and Practice of Constraint Programming, Springer, 565–578, 2017.

[6] M. Slusky, Integrating Relaxations for Combinatorial Optimization, Ph.D. thesis, Carnegie Mellon University, 2015.

[7] N. Balakrishnan, J. J. Kanet, V. Sridharan, Early/tardy scheduling with sequence dependent setups on uniform parallel machines, Computers & Operations Research 26 (2) (1999) 127–141.

[8] P. Baptiste, A. Jouglet, D. Savourey, Lower bounds for parallel machine scheduling problems, International Journal of Operational Research 3 (6) (2008) 643–664.

[9] D. Bergman, A. A. Cire, W.-J. van Hoeve, J. Hooker, Decision diagrams for optimization, Springer, 2016.

[10] R. Rodrigues, A. Pessoa, E. Uchoa, M. Poggi de Aragão, Heuristic algorithm for the parallel machine total weighted tardiness scheduling problem, Relatório de Pesquisa em Engenharia de Produçao 8 (10) (2008) 1–11.

[11] B. Simons, Multiprocessor scheduling of unit-time jobs with arbitrary release times and deadlines, SIAM Journal on Computing 12 (2) (1983) 294–299.

[12] J. M. J. Schutten, List scheduling revisited, Operations Research Letters 18 (4) (1996) 167–170.

[13] F. Yalaoui, C. Chu, New exact method to solve the $Pm/r_j/\sum C_j$ schedule problem, International Journal of Production Economics 100 (1) (2006) 168–179.

[14] A. Jouglet, D. Savourey, Dominance rules for the parallel machine total weighted tardiness scheduling problem with release dates, Computers & Operations Research 38 (9) (2011) 1259–1266.

[15] J. E. Beasely, Weighted Tardiness (OR library), URL `http://people.brunel.ac.uk/~mastjjb/jeb/orlib/wtinfo.html`, originally described in [16], [Online accessed 26-Feb-2018], 2018.

[16] J. E. Beasley, OR-Library: distributing test problems by electronic mail, Journal of the Operational Research Society 41 (11) (1990) 1069–1072.

[17] A. A. Ciré, J. N. Hooker, The Separation Problem for Binary Decision Diagrams, in: Proceedings of the International Symposium on Artificial Intelligence and Mathematics, 2014.

[18] D. Bergman, A. A. Cire, W.-J. van Hoeve, Lagrangian bounds from decision diagrams, Constraints 20 (3) (2015) 346–361.