



**HAL**  
open science

# On the complexity of minimizing the total calibration cost

Eric Angel, Evripidis Bampis, Vincent Chau, Vassilis Zissimopoulos

► **To cite this version:**

Eric Angel, Evripidis Bampis, Vincent Chau, Vassilis Zissimopoulos. On the complexity of minimizing the total calibration cost. 11th International Frontiers of Algorithmics Workshop (FAW 2017), Jun 2017, Chengdu, China. pp.1–12, 10.1007/978-3-319-59605-1\_1 . hal-01630651

**HAL Id: hal-01630651**

**<https://hal.science/hal-01630651v1>**

Submitted on 18 Aug 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# On the Complexity of Minimizing the Total Calibration Cost

Eric Angel<sup>1</sup>, Evripidis Bampis<sup>2</sup>, Vincent Chau<sup>3(✉)</sup>,  
and Vassilis Zissimopoulos<sup>4</sup>

<sup>1</sup> IBISC, Université d'Évry Val d'Essonne, Évry, France  
`angel@ibisc.fr`

<sup>2</sup> Sorbonne Universités, UPMC Univ. Paris 06, UMR 7606, LIP6, Paris, France  
`Evripidis.Bampis@lip6.fr`

<sup>3</sup> Shenzhen Institutes of Advanced Technology, Academy of Sciences,  
Shenzhen, China  
`vincentchau@siat.ac.cn`

<sup>4</sup> Department of Informatics and Telecommunications,  
National and Kapodistrian University of Athens, Athens, Greece  
`vassilis@di.uoa.gr`

**Abstract.** Bender et al. (SPAA 2013) proposed a theoretical framework for testing in contexts where safety mistakes must be avoided. Testing in such a context is made by machines that need to be often calibrated. Since calibrations have non negligible cost, it is important to study policies minimizing the calibration cost while performing all the necessary tests. We focus on the single-machine setting and we study the complexity status of different variants of the problem. First, we extend the model by considering that the jobs have arbitrary processing times and that the preemption of jobs is allowed. For this case, we propose an optimal polynomial time algorithm. Then, we study the case where there is many types of calibrations with different lengths and costs. We prove that the problem becomes NP-hard for arbitrary processing times even when the preemption of the jobs is allowed. Finally, we focus on the case of unit-time jobs and we show that a more general problem, where the recalibration of the machine is not instantaneous, can be solved in polynomial time.

## 1 Introduction

The scheduling problem of minimizing the number of calibrations has been recently introduced by Bender et al. in [2]. It is motivated by the Integrated Stockpile Evaluation (ISE) program [1] at Sandia National Laboratories for testing in contexts where safety mistakes may have serious consequences. Formally,

---

This work has been supported by the ALGONOW project of the THALES program and the Special Account for Research Grants of National and Kapodistrian University of Athens, by NSFC (Nos. 61433012, U1435215), and by Shenzhen basic research grant JCYJ20160229195940462.

the problem can be stated as follows: we are given a set  $\mathcal{J}$  of  $n$  jobs (tests), where each job  $j$  is characterized by its release date  $r_j$ , its deadline  $d_j$  and its processing time  $p_j$ . We are also given a (resp. a set of) testing machine(s) that must be calibrated on a regular basis. The calibration of a machine has a unit cost and it is instantaneous, i.e., a machine can be calibrated between the execution of two jobs that are processed in consecutive time-units. A machine stays calibrated for  $T$  time-units and a job can only be processed during an interval where the machine is calibrated. The goal is to find a feasible schedule performing all the tests (jobs) between their release dates and deadlines and minimizing the number of calibrations. Using the classical three-field notation in scheduling [4], the problem can be denoted as  $P | r_j, d_j, T | (\# \text{ calibrations})$ . Bender et al. [2] studied the case of *unit-time* jobs. They considered both the single-machine and multiple-machine problems. For the single-machine case, they showed that there is a polynomial-time algorithm, called the Lazy Binning algorithm that solves the problem optimally. For the multiple-machine case, they proposed a 2-approximation algorithm. However, the complexity status of the multiple-machine case with unit-time jobs remained open. Bender et al. [2] stated, “As a next step we hope to generalize our model to capture more aspects of the actual ISE problem. For example, machines may not be identical, and calibrations may require machine time. Moreover, some jobs may not have unit size”.

Fineman and Sheridan [3] studied a first generalization of the problem by considering that the jobs have arbitrary processing times. They considered the multiple-machine case where the execution of a job is not allowed to be interrupted once it has been started. Since the feasibility problem is NP-hard, they considered a *resource-augmentation* [5] version of the problem. They were able to relate this version with the classical *machine-minimization* problem [8] in the following way: suppose there is an  $s$ -speed  $\alpha$ -approximation algorithm for the machine-minimization problem, then there is an  $O(\alpha)$ -machine  $s$ -speed  $O(\alpha)$ -approximation for the resource-augmentation version of the problem of minimizing the number of calibrations.

In this paper, we focus on the single-machine case without resource augmentation and we study the complexity status of different variants of the problem. In Sect. 2, we study the problem when the jobs may have arbitrary processing times and the preemption of the jobs is allowed: the processing of any job may be interrupted and resumed at a later time. We denote this variant of the problem as  $1 | r_j, d_j, pmtn, T | (\# \text{ calibrations})$ . Clearly, by using the optimal algorithm of Bender et al. for unit-time jobs, we can directly obtain a pseudopolynomial-time algorithm by just replacing every job by a set of unit-time jobs with cardinality equal to the processing time of the job. We propose a polynomial time algorithm for this variant of the problem. Then, in Sect. 3, we study the case of scheduling a set of jobs when  $K$  different types of calibrations are available. Each calibration type is associated with a length  $T_i$  and a cost  $f_i$ . The objective is to find a feasible schedule minimizing the total calibration cost. We show that the problem, denoted as  $1 | r_j, d_j, pmtn, \{T_1, \dots, T_K\} | \text{cost}(\text{calibrations})$  for arbitrary processing times is NP-hard, even when the preemption of the jobs is allowed.

Given the NP-hardness of the problem for arbitrary processing times, in Sect. 4, we study the case of unit-time jobs. We propose a polynomial time algorithm based on dynamic programming. We present the algorithm for a more general setting where each calibration takes  $\lambda$  units of time during which the machine cannot be used. We denote this variant as  $1 \mid r_j, d_j, p_j = 1, \lambda + \{T_1, \dots, T_K\} \mid \text{cost}(\text{calibrations})$ .

## 2 Arbitrary Processing Times and Preemption

We suppose here that the jobs have arbitrary processing times and that the preemption of the jobs is allowed. An obvious approach in order to obtain an optimal preemptive schedule is to divide each job  $j$  into  $p_j$  unit-time jobs with the same release date and deadline as job  $j$  and then apply the Lazy Binning (LB) algorithm of [2] that optimally solves the problem for instances with unit-time jobs. However, this idea leads to a pseudopolynomial-time algorithm. Here, we propose a more efficient way for solving the problem. Our method is based on the idea of Lazy Binning. Before introducing our algorithm, we briefly recall LB: at each iteration, a time  $t$  is fixed and the (remaining) jobs are scheduled, starting at time  $t+1$  using the Earliest Deadline First (EDF) policy<sup>1</sup>. If a feasible schedule exists (for the remaining jobs),  $t$  is updated to  $t+1$ , otherwise the next calibration is set to start at time  $t$  which is called the current *latest-starting-time* of the calibration. Then, the jobs that are scheduled during this calibration interval are removed and this process is iterated after updating  $t$  to  $t+T$ , where  $T$  is the calibration length. The polynomiality of the algorithm for unit-time jobs comes from the observation that the starting time of any calibration is at a distance of no more than  $n$  time-units before any deadline. In our case however, i.e. when the jobs have arbitrary processing times, a calibration may start at a distance of at most  $P = \sum_{j=1}^n p_j$  time-units before any deadline.

**Definition 1.** Let  $\Psi := \bigcup_i \{d_i - P, d_i - P + 1, \dots, d_i\}$  where  $P = \sum_{j=1}^n p_j$ .


**Proposition 1.** *There exists an optimal solution in which each calibration starts at a time in  $\Psi$ .*

*Proof.* Let  $\sigma$  be an optimal solution in which there is at least one calibration that does not start at a time in  $\Psi$ . We show how to transform the schedule  $\sigma$  into another optimal schedule that satisfies the statement of the proposition.

Let  $c_{i'}$  be the first calibration of  $\sigma$  that starts at time  $t' \notin \Psi$ . Let  $c_{i'}, \dots, c_i$  be the maximum set of consecutive calibrations such that when a calibration finishes another starts immediately. We denote by  $c_{i+1}$  the next calibration that is not adjacent to calibration  $c_i$ . We can push the set of calibrations  $c_{i'}, \dots, c_i$  to the right (we delay the calibrations) until:

---

<sup>1</sup> An EDF policy is a schedule in which at any time, the job with the smallest deadline among the available jobs is scheduled first.



**Fig. 1.** Illustration of Proposition 1. The first schedule is an optimal schedule. The second one is obtained after pushing the continuous block of calibrations  $c_{i'}, \dots, c_i$  to the right.

- either we reach the next calibration  $c_{i+1}$ ,
- or  $c_{i'}$  starts at a time in  $\Psi$  (Fig. 1).

Note that this transformation is always possible. Indeed, since  $c_{i'}$  starts at a time that is in a distance more than  $P$  from a deadline, it is always possible to push the scheduled jobs to the right. In particular, if there are no jobs scheduled when calibration  $c_{i'}$  starts, then there are no modifications for the execution of jobs. Otherwise, there is at least one job scheduled when calibration  $c_{i'}$  starts. Let  $a_1, \dots, a_e$  be the continuous block of jobs. Since the starting time of job  $a_1$  is at a distance (to the left hand side) more than  $P$  from a deadline, then all these jobs can be pushed to the right by one unit. This transformation is possible because no job of this block finishes at its deadline. Note that after this modification, jobs can be assigned to another calibration.

We can repeat the above transformation until we get a schedule satisfying the statement of the proposition.  $\square$

For jobs with arbitrary processing times when the preemption of the jobs is allowed, we propose the following algorithm whose idea is based on the Lazy Binning algorithm: we first compute the current latest-starting-time of the calibration such that no job misses its deadline (this avoids to consider every time in  $\Psi$ ). This calibration time depends on some deadline  $d_k$ . At each iteration, among the remaining jobs, we compute for every deadline the sum of the processing times of all these jobs (or of their remaining parts) having a smaller than or equal deadline and we subtract it from the current deadline. The current latest-starting-time of the calibration is obtained by choosing the smallest computed value. Once the calibration starting time is set, we schedule the remaining jobs in the EDF order until reaching  $d_k$  and we continue to schedule the available jobs until the calibration interval finishes. In the next step, we update the processing time of the jobs that have been processed. We repeat this computation until there is no processing time left. A formal description of the algorithm, that we call the Preemptive Lazy Binning (PLB) algorithm, is given below (Algorithm 1).

---

**Algorithm 1.** Preemptive Lazy Binning (PLB)

---

```
1: Jobs in  $\mathcal{J}$  are sorted in non-decreasing order of deadline
2: while  $\mathcal{J} \neq \emptyset$  do
3:    $t \leftarrow \max_{i \in \mathcal{J}} d_i$ ,  $k \leftarrow 0$ 
4:   for  $i \in \mathcal{J}$  do
5:     if  $t > d_i - \sum_{j \leq i, j \in \mathcal{J}} p_j$  then
6:        $t \leftarrow d_i - \sum_{j \leq i, j \in \mathcal{J}} p_j$ 
7:        $k \leftarrow i$ 
8:     end if
9:   end for
10:   $u \leftarrow t + \left\lceil \frac{d_k - t}{T} \right\rceil \times T$ 
11:  Calibrate the machine at time  $t, t + T, t + 2T, \dots, u - T$ 
12:  Schedule jobs  $\{j \leq k \mid j \in \mathcal{J}\}$  from  $t$  to  $d_k$  by applying the EDF policy and
  remove them from  $\mathcal{J}$ .
13:  Schedule fragment of jobs from  $k + 1, \dots, n$  in  $[d_k, u)$  in EDF order
14:  Let  $q_j$  for  $j = k + 1, \dots, n$  be the processed quantity in  $[d_k, u)$ 
15:  //Update processing time of jobs
16:  for  $i = k + 1, \dots, n$  do
17:     $p_i \leftarrow p_i - q_i$ 
18:    if  $p_i = 0$  then
19:       $\mathcal{J} \leftarrow \mathcal{J} \setminus i$ 
20:    end if
21:  end for
22: end while
```

---

We can prove the optimality of this algorithm using a similar analysis as the one for the Lazy Binning algorithm in [2].

**Proposition 2.** *The schedule returned by Algorithm PLB is a feasible schedule in which the starting time of each calibration is maximum.*

*Proof.* The condition in line 5 in Algorithm PLB ensures that we always obtain a feasible schedule. In fact, we compute the latest-starting-time at each step and this time is exactly the latest time of the first calibration.

By fixing a deadline  $d_i$ , we know that jobs that have a deadline earlier than  $d_i$  have to be scheduled before  $d_i$ , while the other jobs are scheduled after  $d_i$ . When we update  $t$  for every deadline  $d_i$  in the algorithm, we assume that there is no idle time between  $d_i - \sum_{j \leq i, j \in \mathcal{J}} p_j$  and  $d_i$ . Note that if  $d_i - \sum_{j \leq i, j \in \mathcal{J}} p_j < 0$ , then the schedule is not feasible. For the sake of contradiction, suppose that a feasible schedule exists in which some calibration is not started at a time computed by the algorithm. We will show that the starting time of this calibration is not maximum. Denote this time by  $t'$ . Since, the starting time of the calibration is not one of  $d_i - \sum_{j \leq i, j \in \mathcal{J}} p_j \forall i$ , then there is at least one unit of idle time between the starting time of the calibration and some deadline  $d_i$ . Hence, it is possible to delay all calibrations starting at  $t'$  or after, as well as the execution of the jobs inside these calibrations by keeping the EDF order. This can be done in a similar way as in the proof of Proposition 1.  $\square$

**Proposition 3.** *Algorithm PLB is optimal.*

*Proof.* It is sufficient to prove that Algorithm PLB returns the same schedule as Lazy Binning after splitting all jobs to unit-time jobs. We denote respectively *PLB* and *LB* the schedules returned by these algorithms.

Let  $t'$  be the first time at which the two schedules differ. The jobs executed before  $t'$  are the same in both schedules since the jobs are scheduled in the EDF order. Given that the schedules are the same before  $t'$ , the remaining jobs are the same after  $t'$ . Two cases may occur:

- a job is scheduled in  $[t', t' + 1)$  in *PLB* but not in *LB*. This means that the machine is not calibrated at this time slot in the schedule produced by *LB*. Since the calibrations are the same before  $t'$  in both schedules, then a calibration starting at  $t'$  is necessary in *PLB*. Thanks to Proposition 2, we have a contradiction to the fact that we were looking for the latest-starting-time of the calibration.
- a job is scheduled in  $[t', t' + 1)$  in *LB* but not in *PLB*. This means that there does not exist a feasible schedule starting at  $t' + 1$  with the remaining jobs. Hence, *PLB* is not feasible. This case cannot happen thanks to Proposition 2.  $\square$

**Proposition 4.** *Algorithm PLB has a time complexity in  $O(n^2)$ .*

*Proof.* We first sort jobs in the non-decreasing order of their deadlines in  $O(n \log n)$  time. At each step, we compute the first latest-starting-time of the calibration in  $O(n)$  time. Then the scheduling of jobs in the EDF order takes  $O(n)$  time. We need also to update the processing times of the jobs whose execution has been started. This can be done in  $O(n)$  time. At each step, we schedule at least one job. Hence, there are at most  $n$  steps.  $\square$

### 3 Arbitrary Processing Times, Preemption and Many Calibration Types

In this section, we consider a generalization of the model of Bender et al. in which there are more than one types of calibration. Every calibration type is associated with a length  $T_i$  and a cost  $f_i$ . We are also given a set of jobs, each one characterized by its processing time  $p_j$ , its release time  $r_j$  and its deadline  $d_j$ . Each job can be scheduled only when the machine is calibrated regardless of the calibration type. Our objective is to find a feasible preemptive schedule minimizing the total calibration cost. We prove that the problem is NP-hard.

**Proposition 5.** *The problem of minimizing the calibration cost is NP-hard for jobs with arbitrary processing times and many types of calibration, even when the preemption is allowed.*

In order to prove the NP-hardness, we use a reduction from the UNBOUNDED SUBSET SUM problem (which is NP-hard) [6,7]. In an instance of the

UNBOUNDED SUBSET SUM problem, we are given a set of  $n$  items where each item  $j$  is associated to a value  $\kappa_j$ . We are also given a value  $V$ . We aim to find a subset of the items that sums to  $V$  under the assumption that an item may be used more than once.

*Proof.* Let  $\Pi$  be the preemptive scheduling problem of minimizing the total calibration cost for a set of  $n$  jobs that have arbitrary processing times in the presence of a set of  $K$  calibration types.

Given an instance of the UNBOUNDED SUBSET SUM problem, we construct an instance of problem  $\Pi$  as follows. For each item  $j$ , create a calibration length  $T_j = \kappa_j$  and of cost  $f_j = \kappa_j$ . Moreover, we create  $n$  jobs with positive arbitrary processing times such that  $\sum_i p_i = V$  with  $r_i = 0$  and  $d_i = V \forall i$ .

We claim that the instance of the UNBOUNDED SUBSET SUM problem is feasible if and only if there is a feasible schedule for problem  $\Pi$  of cost  $V$ .

Assume that the instance of the UNBOUNDED SUBSET SUM problem is feasible. Therefore, there exists a subset of items  $C'$  such that  $\sum_{j \in C'} \kappa_j = V$ . Note that the same item may appear several times. Then we can schedule all jobs, and calibrate the machine according to the items in  $C'$  in any arbitrary order. Since the calibrations allow all the jobs to be scheduled in  $[0, V)$ , then we get a feasible schedule of cost  $V$  for  $\Pi$ .


For the opposite direction of our claim, assume that there is a feasible schedule for problem  $\Pi$  of cost  $V$ . Let  $\mathcal{C}$  be the set of calibrations that have been used in the schedule. Then  $\sum_{j \in \mathcal{C}} T_j = V$ . Therefore, the items which correspond to the calibrations in  $\mathcal{C}$  form a feasible solution for the UNBOUNDED SUBSET SUM problem.  $\square$

## 4 Unit-Time Jobs, Many Calibration Types and Activation Length

Since the problem is NP-hard when many calibration types are considered even in the case where the calibrations are instantaneous, we focus in this section on the case where the jobs have unit processing times. We also assume that there is an activation length, that we denote by  $\lambda$ . This means that in this section, the calibrations are no more instantaneous, but each of them takes  $\lambda$  units of time during which no job can be processed. For feasibility reasons, we allow to recalibrate the machine at any time point, even when it is already calibrated. To see this, consider the instance given in Fig. 2. The machine has to be calibrated at time 0 and requires  $\lambda = 3$  units of time for being available for the execution of jobs. At time 3 the machine is ready to execute job 1 and it remains calibrated for  $T = 4$  time units. If we do not have the possibility to recalibrate an already calibrated machine then the earliest time at which we can start calibrating the machine is at time 7. This would lead to the impossibility of executing job 2. However, a recalibration at time 4 would lead to a feasible schedule.

It is easy to see that the introduction of the activation length into the model makes necessary the extension of the set of “important” dates that we have used





**Fig. 2.** An infeasible instance if we do not have the possibility to recalibrate at any time. We have a single machine, two unit-time jobs and a single type of calibration of length  $T = 4$ . The activation length, i.e. the time that is required in order for the calibration to be effective is  $\lambda = 3$ . Job 1 is released at time 3 and its deadline is 4. Job 2 is released at time 7 and its deadline is 8.

in Sect. 2 (Definition 1). Indeed, jobs can be scheduled at a distance bigger than  $n$  from a release date or a deadline. However, as we prove below, it is still possible to define a polynomial-size time-set.

In the worst case, we have to calibrate  $n$  times and schedule  $n$  jobs. Thus the calibration can start at a time at most  $n(\lambda + 1)$  time units before a deadline. Note that it is not necessary to consider every date in  $[d_i - n(\lambda + 1), d_i]$  for a fixed  $i$ . In the sequel, we suppose without loss of generality that jobs are sorted in non-decreasing order of their deadline,  $d_1 \leq d_2 \leq \dots \leq d_n$ .

**Definition 2.** Let  $\Theta := \bigcup_i \{d_i - j\lambda - h, j = 0, \dots, n, h = 0, \dots, n\}$ .

**Proposition 6.** *There exists an optimal solution in which each calibration starts at a time in  $\Theta$ .*

*Proof.* We show how to transform an optimal schedule into another schedule satisfying the statement of the proposition without increasing the total calibration cost. Let  $c_j$  be the last calibration that does not start at a date in  $\Theta$ . We can shift this calibration to the right until:

- one job of this calibration finishes at its deadline and hence, it is no more possible to push this calibration to the right anymore. This means that there is no idle time between the starting time of this calibration and this deadline. Thus the starting time of this calibration is in  $\Theta$ .
- the current calibration meets another calibration. In this case, we continue to shift the current calibration to the right while this is possible. Perhaps, there will be an overlap between calibration intervals, but as we said before, we allow to recalibrate the machine at any time. If we cannot shift to the right anymore, either a job ends at its deadline (and we are in the first case), or there is no idle time between the current calibration and the next one. Since there is at most  $n$  jobs and the next calibration starts at a time  $d_i - j\lambda - h$  for some  $i, j, h$ , then the current calibration starts at a time  $d_i - (j+1)\lambda - (h+h')$  where  $h'$  is the number of jobs scheduled in the current calibration with  $h' + h \leq n$  and  $j \leq n - 1$ .  $\square$

Moreover the set of starting times of jobs has also to be extended by considering the activation length  $\lambda$ .


**Definition 3.** Let  $\Phi := \{t + a \mid t \in \Theta, a = 0, \dots, n\} \cup \bigcup_i \{r_i, r_i + 1, \dots, r_i + n\}$ .

As for the starting time of calibrations, the worst case happens when we have to recalibrate after the execution of every job.

**Proposition 7.** *There exists an optimal solution in which the starting times and completion times of jobs belong to  $\Phi$ .*

*Proof.* The first part of the proof comes from Proposition 1. Indeed, jobs can only be scheduled when the machine is calibrated. Let  $i$  be the first job that is not scheduled at a time in  $\Phi$  in an optimal solution. Thanks to Proposition 1, we know that a calibration occurs before a deadline. Job  $i$  belongs to some calibration that starts at time  $t \leq d_j$  for some other job  $j$ . By moving job  $i$  to the left, the cost of the schedule does not increase, since this job belongs to the same calibration. Two cases may occur:

- job  $i$  meets another job  $i'$  (Fig. 3(a)). In this case, we consider the continuous block of jobs  $i'', \dots, i', i$ . We assume that at least one job in this block is scheduled at its release date and job  $i$  is at a distance at most  $n$  of this release date (because there is at most  $n$  jobs). Otherwise, we can shift this block of jobs to the left by one time unit (Fig. 3(b)). Indeed, this shifting is possible because no job in  $\{i'', \dots, i'\}$  is executed at a starting time of a calibration (if it is the case, job  $i$  is in  $\Phi$  by definition). Since job  $i'$  was in  $\Phi$ , by moving this block, job  $i$  will be scheduled at a time in  $\Phi$ .
- job  $i$  meets its release date, thus its starting time is in  $\Phi$ . □



**Fig. 3.** Illustration of Proposition 7

**Definition 4.** Let  $S(j, u, v) = \{i \mid i \leq j \text{ and } u \leq r_i < v\}$ . We define  $F(j, u, v, t, k)$  as the minimum cost of a schedule of the jobs in  $S(j, u, v)$  such that:

- all these jobs are scheduled during the time-interval  $[u, v)$
- the last calibration of the machine is at time  $t$  for a length of  $\lambda + T_k$  (where the time-interval  $[t, t + \lambda)$  corresponds to the activation length)
- the first calibration is not before  $u$ .


We are now ready to give our dynamic programming algorithm. We examine two cases depending on whether  $r_j$  belongs to the interval  $[u, v)$ . Otherwise, there is two subcases: whether job  $j$  is scheduled in the last calibration or not.

**Proposition 8.** *One has  $F(j, u, v, t, k) = F'$*

$$F' := \begin{cases} F(j-1, u, v, t, k) & \text{if } r_j \notin [u, v) \\ \min_{\substack{u' \in \Phi, r_j \leq u' < t + T_k + \lambda \\ t' + \lambda \leq u' < v \\ u' < d_j}} F(j-1, u, u', t, k) & \text{last calibration} \\ \min_{\substack{u' \in \Phi, r_j \leq u' < t' + T_{k'} + \lambda \\ t' \in \Theta, t' + \lambda \leq u' < v \\ u' < d_j \\ 1 \leq k' \leq K}} \left\{ \begin{array}{l} F(j-1, u, u', t', k') \\ + F(j-1, u'+1, v, t, k) \end{array} \right\} & \text{otherwise} \end{cases}$$

with  $F(0, u, v, t, k) := f_k, \forall t + \lambda \leq v \ \& \ t \geq u$ .  
 $F(0, u, v, t, k) := +\infty$  otherwise.

The objective function for our problem is  $\min_{t \in \Theta, 1 \leq k \leq K} F(n, \min_i r_i, \max_i d_i, t, k)$  (Fig. 4).



**Fig. 4.** Illustration of Proposition 8

*Proof.* When  $r_j \notin [u, v)$ , we have necessarily  $F(j, u, v, t, k) = F(j-1, u, v, t, k)$ . In the following, we suppose that  $r_j \in [u, v)$  which includes two cases. The first one is when job  $j$  is scheduled in the last calibration.

**We first prove that  $F(j, u, v, t, k) \leq F'$ .**

We consider a schedule  $S_1$  that realizes  $F(j-1, u, u', t', k')$  and a schedule  $S_2$  that realizes  $F(j-1, u'+1, v, t, k)$ . We build a schedule as follows: from time  $u$  to time  $u'$  use  $S_1$ , then execute job  $j$  in  $[u', u'+1)$ , and finally from  $u'+1$  to time  $v$  use  $S_2$ . Moreover, it contains all jobs in  $\{i \mid i \leq j \text{ and } u \leq r_i < v\}$ . Since the first calibration in  $S_2$  does not begin before  $u'+1$ , then we have a feasible schedule.

So  $F(j, u, v, t, k) \leq F'$ .

**We now prove that  $F(j, u, v, t, k) \geq F'$ .**

Since  $j \in \{i \mid i \leq j \text{ and } u \leq r_i < v\}$ , job  $j$  is scheduled in all schedules that realize  $F(j, u, v, t, k)$ .

Among such schedules, let  $\mathcal{X}$  denote the schedule of  $F(j, u, v, t, k)$  in which the starting time of job  $j$  is maximal. We claim that all jobs in  $\{i \leq j, u \leq r_i < v\}$  that are released before  $u'$  are completed at  $u'$ . If it is not the case, we could swap the execution of such a job with job  $j$ , getting in this way a feasible schedule with the same cost as before. Formally, let  $i$  be a job with  $\{i \leq j, u \leq r_i < u'\}$  that is scheduled after  $u' + 1$ . We can swap the execution of job  $i$  with job  $j$ , the resulting schedule is feasible since job  $j$  has larger deadline than job  $i$ , and job  $i$  is released before  $u'$ . This will contradict the fact that the starting time of job  $j$  is maximal.

We consider a schedule  $S_1$  that realizes  $F(j-1, u, u', t', k')$  and a schedule  $S_2$  that realizes  $F(j-1, u'+1, v, t, k)$ . Then, the restriction of  $S_1$  in the schedule  $\mathcal{X}$  to  $[u, u']$  will be a schedule that meets all constraints related to  $F(j-1, u, u', t', k')$ . Hence its cost is greater than  $F(j-1, u, u', t', k')$ . Similarly, the restriction of  $S_2$  in the schedule  $\mathcal{X}$  to  $[u'+1, v]$  is a schedule that meets all constraints related to  $F(j-1, u'+1, v, t, k)$ .

Finally,  $F(j, u, v, t, k) \geq F'$ .  $\square$

**Proposition 9.** *The problem of minimizing the total calibration cost with arbitrary calibration lengths, activation length and unit-time jobs can be solved in time  $O(n^{16}K^2)$ .*

*Proof.* This problem can be solved with the dynamic program in Proposition 8. Recall that the objective function is  $\min_{t \in \Theta, 1 \leq k \leq K} F(n, \min_i r_i, \max_i d_i, t, k)$ . The size of both sets  $\Theta$  and  $\Phi$  is  $O(n^3)$ . Indeed, by rewriting the set  $\Phi$ , we have

$$\begin{aligned} \Phi &= \bigcup_i \{r_i, r_i + 1, \dots, r_i + n\} \cup \{t + a \mid t \in \Theta, a = 0, \dots, n\} \\ &= \bigcup_i \{r_i, r_i + 1, \dots, r_i + n\} \bigcup_i \left\{ \begin{array}{l} d_i - j\lambda - k + a, j = 0, \dots, n \\ k = 0, \dots, n, a = 0, \dots, n \end{array} \right\} \\ &= \bigcup_i \{r_i, r_i + 1, \dots, r_i + n\} \bigcup_i \{d_i - j\lambda + k, j = 0, \dots, n, k = -n, \dots, n\} \end{aligned}$$

The size of the table is  $O(n^{10}K)$ . When each value of the table is fixed, the minimization is over the values  $u', t'$  and  $k'$ , so the time complexity is  $O(n^6K)$ . Therefore the overall complexity time is  $O(n^{16}K^2)$ .  $\square$

Note that when there is no feasible schedule, the objective function  $\min_{t \in \Theta, 1 \leq k \leq K} F(n, \min_i r_i, \max_i d_i, t, k)$  will return  $+\infty$ .

## 5 Conclusion

We considered different extensions of the model introduced by Bender et al. in [2]. We proved that the problem of minimizing the total calibration-cost on a

single machine can be solved in polynomial time for the case of jobs with arbitrary processing times when the preemption is allowed. Then we proved that the problem becomes NP-hard for arbitrary processing times when there are many calibration types, even if the preemption of jobs is authorized. Finally, we considered the case with many calibration types, where the calibrations are not instantaneous but take machine time, and we proved that the problem can be solved in polynomial time using dynamic programming for unit-time jobs. An interesting question is whether it is possible to find a lower time-complexity algorithm for solving this version of the problem, either optimally, or in approximation. Of course, it would be of great interest to study the case where more than one machines are available. Recall that the complexity of the simple variant studied by Bender et al. remains unknown for the multiple machines problem.

## References

1. New integrated stockpile evaluation program to better ensure weapons stockpile safety, security, reliability (2006). <http://www.sandia.gov/LabNews/060331.html>
2. Bender, M.A., Bunde, D.P., Leung, V.J., McCauley, S., Phillips, C.A.: Efficient scheduling to minimize calibrations. In: 25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2013, pp. 280–287. ACM (2013). <http://doi.acm.org/10.1145/2486159.2486193>
3. Fineman, J.T., Sheridan, B.: Scheduling non-unit jobs to minimize calibrations. In: Blelloch, G.E., Agrawal, K. (eds.) Proceedings of the 27th ACM on Symposium on Parallelism in Algorithms and Architectures, SPAA 2015, Portland, OR, USA, 13–15 June 2015, pp. 161–170. ACM (2015). <http://doi.acm.org/10.1145/2755573.2755605>
4. Graham, R.L., Lawler, E.L., Lenstra, J.K., Kan, A.: Optimization and approximation in deterministic sequencing and scheduling: a survey. *Ann. Discret. Math.* **5**, 287–326 (1979)
5. Kalyanasundaram, B., Pruhs, K.: Speed is as powerful as clairvoyance. *J. ACM* **47**(4), 617–643 (2000). <http://doi.acm.org/10.1145/347476.347479>
6. Lai, T.: Worst-case analysis of greedy algorithms for the unbounded knapsack, subset-sum and partition problems. *Oper. Res. Lett.* **14**(4), 215–220 (1993). [http://dx.doi.org/10.1016/0167-6377\(93\)90072-O](http://dx.doi.org/10.1016/0167-6377(93)90072-O)
7. Lueker, G.: Two NP-complete problems in nonnegative integer programming. Report No. 178, Computer Science Laboratory, Princeton University (1975)
8. Phillips, C.A., Stein, C., Torng, E., Wein, J.: Optimal time-critical scheduling via resource augmentation. *Algorithmica* **32**(2), 163–200 (2002). <http://dx.doi.org/10.1007/s00453-001-0068-9>