A Discrete Time Markov Decision Process to support the scheduling of re-manufacturing activities

Alessio Angius[1], Massimo Manzini[1] and Marcello Urgo[1]

Politecnico di Milano, Mechanical Dept., Italy
alessio.angius, massimo.manzini, marcello.urgo@polimi.it

Keywords: Markov Decision Process, re-manufacturing, scheduling.

1 Introduction and problem statement

Rotor blades are one of the most expensive components in gas turbines for power generation due to the specific materials used and the complex manufacturing process needed. For this reason, turbine blades are one of the component whose re-manufacturing is economically viable during the maintenance of gas turbines. Nevertheless, rework activities are subject to a considerable degree of uncertainty due to the unpredictable degree of damage affecting the blades. The wear of the rotor blades could usually occur in term of lack of material or the presence of cracks whose depth is difficult to estimate in advance. The repair process consists in the removal of the hard coating and of the damaged parts, and the addition of the missing material through additive manufacturing processes. Then, the blades have to undergo a material removal process to obtain the final desired shape. This material removal process is executed through Electrical Discharge Machining technology (EDM), operating to lots of turbine blades belonging to the same stage of the gas turbine and, thus, sharing the same geometrical features. To be able to process a lot of blades, the EDM machine has to undergo a set-up to mount the right electrode. The duration of the processing of a lot also entails a certain degree of uncertainty. Some of the blades, in fact, during the rework process, results having damages that are not possible to repair and, thus, have to be discarded. For this reason, the number of blades to be manufactured in a lot cannot be known in advance. Once the lot of blades have been manufactured in the EDM shop, it undergoes further process steps, it is integrated with new blades to complement the missing ones and then shipped to the customer's premises to be made available to the turbine, thus defining a due date to be respected.

In this paper, we will focus on the EDM shop where both new and repaired blades have to be processed. New blades follow the standard manufacturing process and the associated production plans. Repaired blades arrive as soon as the repair process has been completed and compete for the same resources, i.e., an EDM machine. Thus, a proper approach is needed to schedule the processing of both the classes of blades. We model the presence of multiple EDM machines where an already defined production plan sequences the lot of new blades to be processed while a set of lots of repaired blades is known to be about to require the same machines to be processed. The scheduling approach considers the need of a set-up to be able to process a new lot of blades and aims at minimizing the tardiness of both lots of new and repaired blades.

2 The model

The model under investigation considers K machines processing two classes of production lots: *production* and *repair*. The main difference between the two classes is that production lots are immediately available and can be scheduled on machines in advance.

On the contrary, the arrival of repair lots is uncertain, thus, any decision about their processing is delayed to the moment they become available. The production of a lot is never preempted, hence, a machine must finish the production of a lot before processing a new one. A set-up is needed to move from the production of a lot to a new one.

The model considers a time interval $[0, T]$ where P production lots are produced by knowing in advance that M repair lots will require to be processed. The scheduling of the production lots is defined and conveniently described by a vector $S = |s_1, ...s_k|$ where lots between 1 and s_1 will be produced in sequence on the first machine, lots between $s_1 + 1$ and s_2 will be produced by the second machine, and so on. Each lot is associated to a due date $d_{c,i}$ and size $l_{c,i}$ where c indicates the class and i the index of the lot.

The model assumes that the M repair lots arrive together into the system with the same due date and only one arrival is possible in $[0, T]$. The state of the system is defined by a vector $|k_1, \ldots, k_K, l_{p,1}, \ldots, l_{p,P}, l_{m,1}, \ldots, l_{m_M}|$ with $K + P + M$ entries, where $k_i \in \{R, S\}$ (*Running* and *Set-up*) refers to the state of each machine; $l_{p,i} \in [1, K] \cup \{D\}$ represents the state of the ith production lot that can be assigned to a machine or completed (*Done*). Similarly, $l_{m,i} \in [1, K] \cup \{D, NA, A\}$ represents the ith repair lot, where the two additional states NA and A (*Not Arrived* and *Arrived*) are necessary to discriminate if the lot has arrived or not. A repair lot can be assigned to a machine only if it has been arrived.

The initial state of the system consider all the machines as running and working the firsts lots assigned in the schedule S, e.g., with $K = 3$ we will have $l_{p,1} = 1$, $l_{p,s_1+1} = 2$ and $l_{p,s_2+1} = 3$, and all the repair lots marked as NA. The model divides the time in units and assumes a syncronous system, hence, in each time unit more than one change can occur in the system. Each transition is the consequence of several events due to the the change of state of the machines and the arrival of repair lots. The arrival of repair lots changes their state from NA to A. Instead, the change of state of a machine k from R to S leads to the completion of a lot. This transition will move the system in a state where the machine will be in the state S and the corresponding lot will change state in D. In this case, if exist a $l_{m,j} = k$, $1 \leq j \leq M$, then the machine was processing a repair lot, otherwise the machine was processing the first production lot not in the state D, by following the sequence in S. Vice-versa, a machine can return to the state R from the state S. Whenever this transition is performed with both production and repair lots available, a decision must be taken, i.e., the machine has to decide if it has to follow the schedule or choose a repair lot. If machine k decides to process a repair lot, then the chosen lot will move from the state A to the state k.

The probability driving the system transitions are assumed distributed according to *phase-type distributions* (PH). It is determined by a vector α, which gives the initial probabilities of the transient states and a matrix A containing the intensities of the transitions among the transient state (see (Horváth, A. 2002)). The rates toward the absorption state are collected in a firing vector $f = -A\mathbf{1}$ where $\mathbf{1}$ is a vector of ones having the same size of the matrix A. This class of distributions is able to approximate any general distribution on the positive axis with a pre-determined accuracy whilst the overall process preserves the Markovian property. This allows us to model the distribution of the lot completion time in a statistically sound manner (as described in (Angius *et. al.* 2018)). In the following, the time that machine k requires to complete the lot i of class c follows a PH distribution represented by $(\alpha_{k,c,i}, A_{k,c,i})$ and a firing vector $f_{k,c,i}$. Similarly, the set-up times and repair arrival are distributed according to a PH distribution represented by (α_s, A_s) and (α_m, A_m).

Since transitions from a state z to a state z' are always combinations of events that involve the machines and the arrival of repair lots, any transition probability is the result of a Kronecker product of the form $B_A(z, z') \otimes_{k=1}^{K} B_k(z, z')$. The function $B_A(z, z')$ is equal to A_m if the repair lots are not arrived in both z and z', it is equal to the firing vector if the repair lots arrived in z', and it is equal to 1 otherwise. Instead, the function $B_k(z', z')$ describes the dynamic of machine k in state z. If the machine is processing a lot, this function behaves as $B_A(z, z')$ by using the corresponding values $(\alpha_{k,c,i}, A_{k,c,i})$. Otherwise, if the machine is performing the set-up, the function $B_A(z, z')$ takes values from (α_s, A_s) and a behaviour similar to the previous case, but starting the execution of a new lot after the completion of the set-up. This is done by multiplying the firing vector by the initial vector of the corresponding lot to be processed. The firing of the set-up coincides with a decision every time a machine has to choose between a production and a repair lot. Because of the presence of non-deterministic decisions, the underlying process is a Discrete Time Markov Decision Process (DTMDP) which is fully characterized by a matrix P containing all the dynamics that do not depend on decisions and a set of matrices $D_1, \ldots, D_V$ describing the different strategies in selecting the next lot to be processed in each machine.

3 Problem definition

Given the DTMDP described in Section 2, the aim of this work is to analyze the tardiness of each lot as a function of a scheduler. We define a scheduler $\mathcal{W} = |w_1, \ldots, w_T|$ as a sequence of T entries $w_i \in [1, \ldots, V]$ that determine which decision matrix has to be used in each $t \in [1, T]$. Given a scheduler $\mathcal{W}$, the distribution $\pi(t)$ of the DTMDP evolves on time according to the formula $\pi(t) = \pi(t-1)(P + D_{w_{(t-1)}})$ starting from an initial vector having all the probability mass in the initial state. Let us denote the random variable describing the completion of the ith production lot as $X_{p,i}$, computed as follows:

$$Pr\{X_{p,i} \geq d_{p,i}\} = 1 - \left(\sum_{t=1}^{d_{p,i}} \pi(t-1) F^{(l_{p,i} <> D)} (P + D_{w_{(t-1)}}) F^{(l_{p,i} == D)} \right) \times \mathbf{1}$$

where $F^{<cond>}$ is a filtering matrix whose entries are equal to one on the diagonal only if the state satisfies the boolean condition $< cond >$. Filtering matrices exploit basic linear algebra to select only the transitions of interest from the matrix used to catch the moment in which a production lot completes its execution. For this reason, the filtering matrix on the *lhs* selects only the source state in which lot i is still under processing ($l_{p,i} <> D$), while the matrix on the *rhs* selects the destination states in which the lot i is completed ($l_{p,i} == D$). This guarantees that the process performs only those transitions that lead to the completion of the considered lot. Instead, the summation over t and the multiplication by $\mathbf{1}$ are used to evaluate all the time units until the due date, and to generate a scalar number from the distribution vector.

The computation of the time for the completion of a repair lot is slightly more complicated because the arrival is stochastic and, as a consequence, the due date is shifted on time. Thus, the calculations involve also the isolation of the moment in which the lot arrives into the system. By denoting the completion of the ith repair lot with $X_{m,i}$, we have that:

$$\begin{aligned} Pr\{X_{m,i} \geq d_{m,i}\} = 1 - \Big(& \textstyle\sum_{t=1}^{T} \pi(t-1) F^{(l_{m,i} <> NA)} (P + D_{w_{(t-1)}}) F^{(l_{m,i} == A)} \times \\ & \textstyle\prod_{t'=t+1}^{d_{m,i}-1} (P + D_{w_{(t-1)}}) \times F^{((l_{m,i} <> NA) \wedge (l_{m,i} <> A))} (P + D_{w_{(t-1)}}) F^{(l_{m,i} == A)} \Big) \times \mathbf{1} \end{aligned}$$

The first term isolates only those transitions starting from a state in which the repair lot is not arrived ($l_{m,i} <> NA$) and ending in a state where it is ($l_{m,i} == A$), while the second term carries on the process for $d_{m,i} - 1$ time units. The third term is used to catch only the moment in which the lot processing is completed ($(l_{m,i} <> NA) \wedge (l_{m,i} <> A)$) as already done for the production lots.

4 Numerical example

In this section we show the importance of the problem under investigation by means of a numerical example. We performed an experiment by assuming a system having $K = 3$ machines that has to produce 7 production lots and is waiting for 4 repair lots. The scheduling is such that the first three lots are scheduled at the first machine, the fourth and the fifth are scheduled at the second machine and the remaining lots are scheduled at the third machine. The sizes of production lots are equal to $|30, 20, 10, 30, 20, 30, 20|$ whereas repair lots are all composed of 30 parts each. Each part requires on average 1 time unit (TU) for being produced whereas set-ups require 1.25 TU on average. The probability that repair lots arrive in the next time unit is 0.5. In order to underline the strong impact that different policies have on the tardiness of each lot, we defined two matrices, D_1 and D_2, that represent the two extreme cases. We defined matrix D_1 in such a way that it always selects the next production lot. On the contrary, matrix D_2 gives always precedence to repair lots. We performed an experiment by four different schedulers: the first scheduler uses constantly matrix D_1; the second always uses matrix D_2; finally, the other two schedulers select randomly D_1 or D_2. We expect that the tardiness of production lots will be minimized by the first scheduler and maximized by the second. Vice versa, the second scheduler will minimize the tardiness of the repair lots and maximizes the one of the production lots. The third and fourth schedulers are expected to provide results in between the two extremes. Figure 1 shows the probability of the tardiness of the third production lot and the third repair lot as function of the due date. It is possible to notice that the results confirms the expectations. In fact, the probability to complete the third production lot on time is maximized by the first scheduler and minimized by the second. On the contrary, the second scheduler provides the best results for the repair lot whereas it is detrimental for the production lot. Furthermore, as expected, the trajectories referring to the random schedulers can be found between the trajectories of generated by the first and second scheduler.

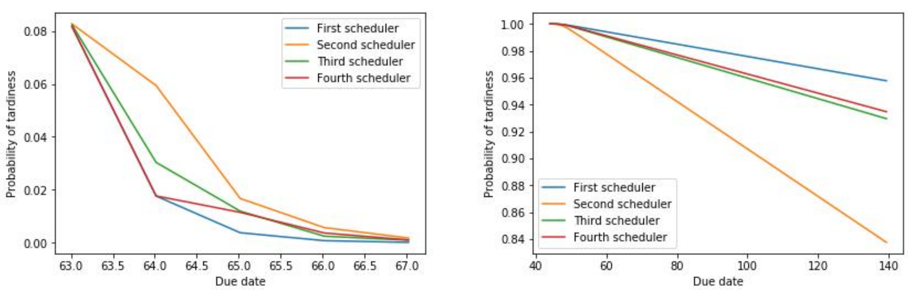

Fig. 1. Probability of the tardiness of the third production lot and the third repair lot as function of the due date.

5 Conclusive Remarks and Future Works

The paper presents a DTMDP that provides the tools for optimizing the scheduling of lots whose arrival into the system is uncertain and cannot be planned in advance. We tested two different scheduling strategies affecting the tardiness function in different ways and validating the model. Future works will regard the identification of optimal scheduling policies.

References

Horváth, A., 2011, "Approximating non-Markovian Behavior by Markovian Models", *PhD Thesis, Department of Telecommunications, Budapest University of Technology and Economics.*

Angius, A. and Colledani, M. and Yemane, A., 2018, "Impact of condition based maintenance policies on the service level of multi-stage manufacturing systems", *Control Engineering Practice*, Vol. 76, pp. 65-78.