



HAL
open science

Impact of imperfect monitoring on the optimal condition-based maintenance policy of a single-item system

Matthieu Roux, Yiping Fang, Anne Barros

► **To cite this version:**

Matthieu Roux, Yiping Fang, Anne Barros. Impact of imperfect monitoring on the optimal condition-based maintenance policy of a single-item system. 32nd European Safety and Reliability Conference, ESREL 2022, Aug 2022, Dublin, Ireland. pp.658-664, 10.3850/978-981-18-5183-4_R14-08-468-cd . hal-03762550

HAL Id: hal-03762550

<https://hal.science/hal-03762550v1>

Submitted on 7 Dec 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.

Impact of imperfect monitoring on the optimal condition-based maintenance policy of a single-item system

Matthieu Roux

LGI EA 2606, Univ. Paris-Saclay, CentraleSupélec, France.
E-mail: matthieu.roux@centralesupelec.fr

Yi-Ping Fang

LGI EA 2606, Univ. Paris-Saclay, CentraleSupélec, France.
E-mail: yiping.fang@centralesupelec.fr

Anne Barros

LGI EA 2606, Univ. Paris-Saclay, CentraleSupélec, France.
E-mail: anne.barros@centralesupelec.fr

In maintenance planning, condition-based maintenance (CBM) policies leverage the observation of the current condition of a degrading system to optimize the planning of future maintenance interventions. Researchers often make a strong assumption that remote sensors (potentially coupled with analytic technologies) perfectly capture the health condition of an industrial asset. For many reasons however (e.g., feasibility, cost), such assumption is inexact and the optimization model should take into account the inaccuracy of the monitoring system. This work focuses on a single-item system, continuously but imperfectly monitored, to illustrate how partially observable Markov decision processes (POMDPs) enable this inaccuracy to be considered within the maintenance optimization model. In this framework, the decision-maker minimizes the total discounted cost over an infinite time horizon by taking one maintenance decision at each time step. For the use case we consider, we choose to model three types of interventions: i) preventive maintenance (PM), (ii) corrective maintenance (CM), and iii) perfect inspection (I). A resource constraint is also added to the model to limit the availability of the repair crew. To solve the POMDP, we implement a point-based value iteration (PBVI) algorithm to compute the optimal CBM policy via approximate dynamic programming. In order to guide investment decisions into monitoring technologies, we develop a framework to quantify and compare the value of information provided by different monitoring systems with known average performances. Finally, we analyze how the monitoring quality impacts the structure of the optimal CBM policy.

Keywords: condition-based maintenance, imperfect monitoring, optimal maintenance planning, partially observable Markov decision process (POMDP), value of information, maintenance policy.

1. Introduction

In this paper, we intend to compute the optimal maintenance policy of a single-item system in a context of imperfect condition monitoring. Condition-based maintenance (CBM) strategies have received much attention recently in the literature (Alaswad (2017)). Such strategies rely on the continuous monitoring of the underlying operating system via the sensing and analytics of one or several measured physical parameters. This information collection is precious because it enables the decision-maker to adjust the maintenance interventions to the actual degradation state

of the item. If the system degrades faster than anticipated, then preventive maintenance should be scheduled sooner, and conversely if the system degrades slower than anticipated.

Nevertheless, this CBM framework requires most of the time to assume that the monitoring system is perfect, i.e., that the signal observed corresponds precisely to the degradation state of the system. This assumption may be reasonable in many applications, but is not always acceptable. The question is then to investigate CBM policies that would exploit the condition monitoring information while taking into account its potential

inaccuracy in the best possible way. According to Alaswad (2017) or de Jonge (2020), more research should be done to study the effects of imperfect condition monitoring on optimal dynamic CBM policies. Here, we focus on the 1-item system, but it is clear that more complex systems should also be investigated.

In the present work, we illustrate the efficiency of the *partially observable Markov decision processes* (POMDPs), first introduced by Sondik (1971), to solve such a maintenance planning problem. It is a rich and fecund framework when dealing with imperfect observations of an underlying system behaving as a Markov process. Papakonstantinou (2014a) and Papakonstantinou (2014b) proposed an interesting work summarizing the modeling framework of POMDPs, applied explicitly to maintenance problems, and illustrated on a real maintenance planning use case. However, we believe that too few studies have been conducted on the topic. More specifically, through this paper, we aim at 1) illustrating the convenience and efficiency of the POMDP framework for such a problem of maintenance planning optimization; 2) conducting a sensitivity analysis on the quality of the monitoring system to investigate how different levels of imperfection in the monitoring system impact the optimal maintenance policy; 3) illustrating in practice the concept of value of information (VoI) developed in Memarzadeh (2016) for the particular case of sequential decision-making. Since this work extends Roux (2022), we may sometimes be quick on the modeling details. The reader is nevertheless highly encouraged to look at this more detailed paper to find more precisions about the modeling or algorithmic details that could be missing in the present article.

In particular, this work differentiates from Nguyen (2019) in several aspects. In that paper, the authors built a POMDP-based CBM policy to optimize inspection quality and maintenance decisions jointly. They compared their strategy with other CBM strategies through a numerical use case to show that wisely selecting the inspection quality at each time step could significantly improve the maintenance policy. The present paper,

however, imposes a fixed condition monitoring quality at the beginning of the optimization. Inaccurate observations are continuously collected by a remote condition monitoring system with a fixed performance. This imperfect monitoring information can be complemented, if necessary, by costly perfect inspections. We then use approximate dynamic programming to compute (to some numerical precision) the optimal policy of the maintenance problem modeled as a POMDP. We eventually compare the results for different levels of imperfect monitoring and analyze how it impacts the expected total maintenance cost and the structure of the optimal policy. To the best of our knowledge, such a sensitivity analysis has not been proposed in the literature and constitutes our main contribution.

The problem and modeling choices are described in Section 2. In Section 3, we briefly introduce the solving method that we have used for this article. Finally, we present our results in Section 4 where we analyze the impact of the quality of the monitoring system in terms of expected cost and structure of the optimal policy.

2. Problem description

2.1. Degradation model

We propose to optimize the maintenance interventions of a single-item system subjected to degradation and random failures. The degradation process is modeled by a Markov chain, where the different states of the system are enumerated in the finite set $\mathcal{S} = \{S_1, S_2, S_3, S_4, F\}$. We consider here four functioning states S_i and only one failed state F , where S_1 represents the *as-good-as-new* state, and S_4 is the most degraded functioning state.

The degradation process is defined by a transition matrix, noted P , describing the transition probabilities between two consecutive time steps. More specifically, we make the following assumption about the degradation process:

- i) The state of the system cannot spontaneously improve, meaning that transitions $F \rightarrow S_i$ and $S_i \rightarrow S_j$ with $j < i$

have a probability of zero.

$$\mathbb{P}(F \rightarrow S_i) = P[F, S_i] = 0, \quad \forall S_i$$

$$\mathbb{P}(S_i \rightarrow S_j) = P[S_i, S_j] = 0, \quad \forall i > j$$

- ii) The system can transition to failure from any functioning state, but the more degraded the state, the most likely a failure can occur at the next time step.

$$\mathbb{P}(S_i \rightarrow F) = P(S_i, F) > 0, \quad \forall S_i$$

$$P(S_i, F) < P(S_j, F), \quad \forall i < j$$

2.2. Maintenance actions

To maintain our system, we decide to model elementary maintenance interventions. However, it is sufficiently generic to accurately represent what is at stake in a maintenance planning problem. This model could easily be adapted with slight modifications (e.g., different levels of imperfect repairs), but the core ideas and insights would remain valid. We recall that, in the POMDP framework, the decision-maker should select exactly one of the following actions at each time step.

- **NA** (no action): No action is taken, meaning that no maintenance intervention will be performed on the system at the current time step;
- **PM** (preventive maintenance): Preventively replace a functioning item with a new one; we then consider that the new item is in *as-good-as-new* state S_1 ;
- **CM** (corrective maintenance): Replace a failed item with a new one; we then consider that the new item is *as-good-as-new* state S_1 ;
- **I** (perfect inspection): Inspect the system to have access to its true degradation state $s \in \mathcal{S}$; if desired, based on the inspection result, the decision-maker may decide to conduct a PM in the same time step following the inspection. This action may be helpful to the maintenance operator to complement the inaccurate information collected by the continuous monitoring system.

A cost is assigned to conducting each of these actions. Actions **NA** costs 0. Generally, **CMs** are much more expensive than **PMs**, explaining why it is crucial to wisely schedule **PMs** before the failure of the item. We note those costs $cost_{CM}$ and $cost_{PM}$, with $cost_{PM} < cost_{CM}$. Moreover, we consider a use case where inspections **I** are relatively cheap compared to other interventions but not free (typically, it can be costly to stop the system for inspection, or costly to pay the maintenance crew if the inspection is not straightforward); we note this cost $cost_I$. Eventually, to ensure that the decision-maker is willing to replace any failed item as soon as possible, an opportunity cost $cost_{OP}$ will be incurred for every time step when the system remains failed.

2.3. POMDP model

The time is discretized (typically with a granularity of one day), and one maintenance decision (or *action*) should be taken at each time step. The purpose is to find the optimal maintenance strategy, noted π^* , that minimizes the discounted sum of maintenance costs over an infinite time horizon. Such a time horizon is convenient because by doing so, we are ensured to find a stationary policy. It is then necessary to discount the maintenance costs to avoid obtaining an infinite objective function. However, the main difficulty comes from the imperfection in the monitoring system. Because of that, the policy cannot directly use the degradation state $s \in \mathcal{S}$ of the system to decide which action to recommend.

In fact, in a POMDP framework, maintenance decisions should be based, not on the true degradation state s of the system, which is unknown, but on what is called the *belief* $b \in \mathcal{B}$ about the state of the system. As the condition monitoring is imperfect, our knowledge about the system is encapsulated in this *belief*, which is a probability distribution over \mathcal{S} describing our uncertainty about the true degradation state. Formally, we note $b \in \mathcal{B} = \{b \in [0, 1]^{|\mathcal{S}|} : \sum_{s \in \mathcal{S}} b[s] = 1\}$.

The optimization problem can then be formulated. It is a sequential decision-making problem, for which at each time step t , exactly one action a_t should be chosen among $\mathcal{A} = \{\mathbf{NA}, \mathbf{PM}, \mathbf{CM}, \mathbf{I}\}$.

It consists in searching for the optimal policy π^* minimizing the expected sum of discounted maintenance costs:

$$V_{\pi^*}(b) := \min_{\pi \in \Pi} V_{\pi}(b) = \mathbb{E} \left[\sum_{t=0}^{+\infty} \gamma^t c(s_t, a_t) \right] \quad (1)$$

with $a_t := \pi(b_t)$ action taken at time t

b_t : belief at time t ($b_0 = b$)

$0 < \gamma < 1$: discount factor

s_t : degradation state at time t

s_0 : initial state of the system, distributed as b

$c(s, a)$: cost resulting from being in state s and taking action a

The synthetic notation $c(s, a)$ is widely used in MDPs and POMDPs models. Here, for example, we would have $c(S_i, \text{PM}) = \text{cost}_{PM}$, or $c(F, \text{CM}) = \text{cost}_{CM}$.

2.4. Imperfect monitoring matrix

The core notion in our study is the concept of what we call the *monitoring matrix*. Indeed, the condition monitoring may be imperfect, but it can still output valuable data. Following traditional POMDP models, like in Papakonstantinou (2014a) and Papakonstantinou (2014b), we suppose that our monitoring system has a finite set of outputs, noted \mathcal{O} . In that context, collecting an (imperfect) observation at time t means receiving a certain value $o_t \in \mathcal{O}$. The potential inaccuracy of this observation means that, from o_t , there is no way to conclude with absolute certainty the true degradation state s of the system. However, we assume that states s and observations o are statistically correlated, and this dependence should be exploited to update the belief b_t when observing the output o_t .

More specifically, in order to optimize the maintenance strategy, we assume that we know the conditional probability $Q(s, o) = \mathbb{P}(o|s)$ of receiving the observation o given that the system is in state s . The matrix Q is called the *monitoring matrix*, and the rest of our paper aims at illustrating to what extent this matrix influences the optimal policy.

2.5. Modeling details

We are aware that we were a bit quick on the modeling part, so we strongly encourage the curious reader to look at Roux (2022), where the same problem is exposed with much more modeling details. In particular, we would like to be precise that the model we implement is slightly more complex since it comprises the two additional constraints mentioned in Roux (2022), but which were not detailed here for simplicity.

3. Solving method

Solving a POMDP is a difficult task in general, more complicated than a classical Markov decision process (MDP). This complexity is due to the belief state space \mathcal{B} , which is continuous and forces us to consider approximate dynamic programming methods to compute approximate solutions (see Hauskrecht (2000)). The current state-of-the-art POMDP solver is SARSOP, for Successive Approximations of the Reachable Space under Optimal Policies, and details can be found in Kurniawati (2008).

We, however, choose to implement a slightly different approximate solving method, yet very much inspired by SARSOP. Like SARSOP, our method is a *point-based value iteration* (PBVI) algorithm, which is an adaptation of the value iteration algorithm for classical MDPs. The value function V is computed on specific points $b_{grid} \in \hat{\mathcal{B}}_{dyn}$ of a dynamic grid, and the whole value function is approximated by interpolation for points $b \notin \hat{\mathcal{B}}_{dyn}$ that are not grid points. The method is explained more in detail in Roux (2022). Still, such a method has a very interesting property as it provides a lower bound on the value of the problem. That bound will be valuable for estimating our approximate solutions' quality.

4. Numerical experiment

4.1. Case study

This section applies our POMDP maintenance model to an illustrative use case. Numerical values do not come from real industrial use cases, but orders of magnitude are directly inspired from Yildirim (2017) and the wind turbine industry. The

problem may seem relatively poor, with a limited number of actions for example, but the purpose here is really to illustrate the potential offered by the POMDP framework and start deriving some insights about the quality of the monitoring. Our main contribution is to study the impact of different monitoring matrices on the optimal maintenance policy obtained. For concision purposes, the cost parameters, transition matrix P , and discount factor γ used in this case study are the same as in Roux (2022) and are therefore not explicated here.

To study the sensitivity to the quality of the monitoring, we generate four monitoring matrices, noted $Q_1, Q_2, Q_3,$ and Q_4 , described below. To keep the model simple and easily interpretable, we suppose there is one observation o_s for each degradation state $s \in \mathcal{S}$ (e.g., an observation o_1 indicates that the system is *probably* in state S_1). Hence intuitively, the closer to one the diagonal elements, the more accurate the monitoring.

- Low monitoring performance

$$Q_1 = \begin{bmatrix} 0.7 & 0.2 & 0.09 & 0.01 \\ 0.25 & 0.4 & 0.3 & 0.05 \\ 0.1 & 0.32 & 0.43 & 0.15 \\ 0.01 & 0.05 & 0.3 & 0.64 \end{bmatrix}$$

- Low+ monitoring performance

$$Q_2 = \begin{bmatrix} 0.8 & 0.13 & 0.06 & 0.01 \\ 0.17 & 0.6 & 0.2 & 0.03 \\ 0.07 & 0.21 & 0.62 & 0.1 \\ 0.01 & 0.03 & 0.2 & 0.76 \end{bmatrix}$$

- Good- monitoring performance

$$Q_3 = \begin{bmatrix} 0.85 & 0.1 & 0.045 & 0.005 \\ 0.125 & 0.7 & 0.15 & 0.025 \\ 0.05 & 0.16 & 0.715 & 0.0125 \\ 0.005 & 0.025 & 0.15 & 0.82 \end{bmatrix}$$

- Good monitoring performance

$$Q_4 = \begin{bmatrix} 0.925 & 0.05 & 0.022 & 0.003 \\ 0.062 & 0.85 & 0.075 & 0.13 \\ 0.025 & 0.08 & 0.858 & 0.037 \\ 0.002 & 0.013 & 0.075 & 0.91 \end{bmatrix}$$

4.2. Methodology

To evaluate the proposed POMDP framework and solving method, we proceed in two phases:

- optimization phase: for each *monitoring matrix* Q , we compute an approximate policy π^Q ;
- simulation phase: then, we simulate the execution of this policy multiple times on a time horizon of 15 000 time steps.

We recall that because the solving method is approximate, we are not able to compute the optimal policy, but only an approximation π^Q (which is dependent on Q). Moreover, thanks to the property of the PBVI-interpolation algorithm, we can get a lower bound on the value of the problem. It is helpful in practice since it provides an excellent way to check the policy obtained at the end of the procedure. Ideally, we would like the average total cost computed via Monte-Carlo simulations to be sufficiently close to the lower bound, indicating that we have found a good quality approximate policy.

The tests are performed on an office laptop, Intel Core i7, with 32 Go RAM. The solutions obtained below required, to be computed, one hour of optimization phase and one hour of simulation phase for each *monitoring matrix*.

4.3. Value of information (VoI)




Fig. 1. Average total discounted cost (simulation) vs. monitoring quality

From Figure 1, we can clearly see how much

the expected total maintenance cost is dependent on the quality of the monitoring. It plays an important role and, consequently, should not be neglected in the modeling of CBM strategies. Obviously, it was expected that the better the monitoring system, the lower the maintenance cost. However, we believe it constitutes a natural and interesting extension to the analysis proposed by Memarzadeh (2016), which suggests a way to evaluate the value of information (VoI) of a flow of information in a sequential decision-making framework. If we were to quantify the VoI brought by each monitoring system, this could immediately be obtained by reading the gap with the "no monitoring" case. On the opposite, by looking at the "perfect monitoring" case, we can see the maximum value one should be ready to pay for a monitoring system because no condition monitoring system can have a higher VoI than the "perfect monitoring" one.

Moreover, we also represent the lower bound given by the dynamic programming value function on the graph. In all of our six cases, we observe that the gap between the analytical lower bound and the Monte-Carlo average cost is very small. This reinforces our confidence in the result and provides an additional argument in favor of using this framework to quantify the VoI (as defined in Memarzadeh (2016)) of an imperfect monitoring system.

For example, we believe that being able to test easily and efficiently different monitoring scenarios can be a substantial advantage when comparing different investments in sensors and analytic technologies with different costs and performances. The decision-maker would only have to compute the VoI of each potential *monitoring matrix* Q and compare them to select the most profitable one.

4.4. Structure of the optimal maintenance policy

From Figure 2, we can see how the quality of the monitoring influences the structure of the optimal policy. The following observations can be made:

- 1) the better the monitoring system, the lower the number of **CMs**; this means




Fig. 2. Average # interventions (simulation) vs. monitoring quality

that with more accurate monitoring, we can prevent some unexpected failures, which are very expensive;

- 2) under a certain minimum quality of monitoring, **PMs** are never performed alone but are always checked before by an inspection **I**; such inspections are used to make sure **PMs** are really worth it;
- 3) the more inaccurate the monitoring, the more simple inspections we make; this means that with low monitoring quality, we tend to perform many inspections, but a lot of them are not followed by a **PM** because the system is in fact not so degraded.

Through this simple illustration, we also intend to highlight that maintenance policies are very dependent on the quality of the monitoring system. This is why it seems crucial to us that monitoring uncertainties are correctly modeled and taken into account when they exist. It also raises some questions because this high sensitivity of the maintenance policy to the monitoring performance can also be dangerous when some uncertainties about the *monitoring matrix* subsist. Therefore, we believe there is more space for additional research in that direction.

5. Conclusion

In this paper, we modeled and practically applied a classical sequential maintenance problem using the POMDP framework. Our case study illustrated that such a modeling framework is very convenient and efficient, and we encourage more researchers and practitioners in the domain to pay attention to this tool. In particular, we illustrated on a practical example the fecund concept of value of information applied to this quite particular case of sequential monitoring observations. We believe that this work opens two research directions in the coming future. First, it could be interesting to look at strategies that could be adopted when we have some uncertainty about the monitoring performance. Second, it seems that this framework should be extended to multi-items systems. However, such a POMDP approach will quickly become intractable. To the best of our knowledge, satisfying heuristic approaches are still to be investigated to scale up to larger and more complex systems.

Acknowledgement

This work is funded by the Chair on Risk and Resilience of Complex Systems (CentraleSupélec, EDF, Orange, SNCF).

References

- Alaswad, S., and Xiang, Y. (2017). A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability engineering & system safety*, 157, 54–63.
- de Jonge, B., & Scarf, P. A. (2020). A review on maintenance optimization. *European journal of operational research*, 285(3), 805–824.
- Sondik, E. J. (1971). The optimal control of partially observable Markov processes. *Stanford University*.
- Papakonstantinou, K. G., and Shinozuka, M. (2014a). Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part I: Theory. *Reliability Engineering & System Safety*, 130, 202–213.
- Papakonstantinou, K. G., and Shinozuka, M. (2014b). Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation. *Reliability Engineering & System Safety*, 130, 214–224.
- Memarzadeh, M., and Pozzi, M. (2016). Value of information in sequential decision making: Component inspection, permanent monitoring and system-level scheduling. *Reliability Engineering & System Safety*, 154, 137–151.
- Yildirim, M., Gebraeel, N. Z., and Sun, X. A. (2017). Integrated predictive analytics and optimization for opportunistic maintenance and operations in wind farms. *IEEE Transactions on power systems*, 32(6), 4319–4328.
- Roux, M., Fang, Y.-P., and Barros, A. (2022). Maintenance planning under imperfect monitoring: an efficient POMDP model using interpolated value function. In *The 18th IFAC Workshop on Control Applications of Optimization: Gif-sur-Yvette, France July 18th-July 22nd: Proceedings*. To be published.
- Hauskrecht, M. (2000). Value-function approximations for partially observable Markov decision processes. *Journal of artificial intelligence research*, 13, 33–94.
- Kurniawati, H., Hsu, D., & Lee, W. S. (2008). SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*, 2008.
- Nguyen, K. T., Do, P., Huynh, K. T., Brenguer, C., & Grall, A. (2019). Joint optimization of monitoring quality and replacement decisions in condition-based maintenance. In *Reliability Engineering & System Safety*, 189, 177–195.