Department of Industrial Engineering & Innovation Sciences

Investigating the Application of Condition-Based Maintenance at the RNLN using Partially Observable Markov Decision Processes

Internship (1MSE15)

M.E. (Maud) Crooymans
MSc Operations Management and Logistics

Supervisors:

R. (Rob) Basten	(Eindhoven University of Technology)
W.W. (Wieger) Tiddens	(Royal Netherlands Navy)

Eindhoven, December 13, 2020

Contents

Summary

In this research, we investigate if a partially observable Markov decision process is a beneficial method to apply condition-based maintenance at the Royal Netherlands Navy. A total of eight guidelines are defined from interviews and literature to find suitable components. These guidelines have to be satisfied to model a system or component within the Royal Netherlands Navy as a partially observable Markov decision process. To apply the guidelines within the Royal Netherlands Navy, a flowchart is created that can be discussed with system-specific engineers. Hence, with a total of eight questions each corresponding to one of the guidelines, it can be decided if it is beneficial to use a partially observable Markov decision process for the system or component. The case study conducted on the diesel engine resulted in three components for which it is recommended to be modeled as a partially observable Markov decision process. These components are the air cooler, turbo and main bearing. A textual description of the model is given for each of these components. Finally, it is concluded that the partially observable Markov decision process would be a suitable method to apply condition-based maintenance at the Royal Netherlands Navy.

1 Introduction

Technical systems and physical assets are part of everyday operations in many different industries. These systems and assets typically deteriorate and are subject to breakdowns due to usage and age (Kim & Makis, 2013). Such breakdowns are seen as one of the biggest risks to daily business (LaRiviere et al., 2016). The high costs that are associated with unplanned breakdowns stimulate to optimize maintenance actions such as repairs or replacements, planned overhauls or corrective actions. However, besides the fact that these actions are very costly, these actions affect the flexibility, throughput time, and quality of the operations during the daily business (Tiddens, 2018). With this in mind, it appears to be important to plan maintenance actions before unexpected breakdowns occur but not too far in advance when sufficient remaining lifetime is left. Thus, this means that optimal maintenance actions should not be performed too late but also not too early.

In addition, technical systems and physical assets that are subject to breakdowns contain many different components. When talking about maintenance actions, one should consider that only some of these components have to be repaired, replaced or updated instead of the complete system or asset. In most of the literature, it is assumed that condition monitoring samples provide perfect information, meaning that the true state of the system is fully observable at time of condition monitoring (Kim & Makis, 2013). However, the actual state of such a system cannot always be determined. This could be, for example, due to the fact that it consists of multiple components. Consider, for example, an engine that consists of multiple components. It could be the case that there is sensor information available of only a part of the components which leads to incomplete information about the true state of the complete engine. Another reason for not knowing the true state could be that the sensor information is not readable when the engine is in use, or that the sensor does not provide reliable information due to deterioration of the sensor (Van Oosterom et al., 2017). In other words, there are multiple causes that could lead to uncertainty about the state of the technical system or physical asset. This uncertainty could, in turn, lead to non-optimal decisions on the timing of maintenance actions.

The Royal Netherlands Navy (RNLN) experiences such situations when optimizing her maintenance policies. At this moment, it is not obvious for the RNLN to decide what the optimal moment is to perform maintenance. This ambiguity occurs due to imperfect predictions or because the engineers do not always know the true state of the system or component within the system. Consequently, it is very difficult to make a choice between performing maintenance or doing nothing. In fact, it could be argued that in these situations it would be a better option to perform a more detailed inspection instead of making this two-sided decision. These kinds of situations can be modelled using partially observable Markov decision processes (POMDPs). Within this research it is investigated whether it is possible and beneficial to model specific systems at the RNLN using a POMDP. Furthermore, it is defined what guidelines are needed to select suitable components for the POMDP at the RNLN. Applying these guidelines defined at the diesel engine led to three components within the diesel engine for which POMDP would be possible and beneficial. This shows that POMDPs would be a suitable method to apply condition-based maintenance (CBM) at the RNLN.

In this report, we discuss the suitability of a POMDP at the RNLN. First, Section 1.1 contains a short introduction to the RNLN, the main tasks of the overarching Ministry of Defence, its departments and the contribution of the Data for Maintenance team. Next, Section 1.2 describes the problem; corresponding research questions are defined in Section

1.3. Chapter 2 discusses the maintenance policy at the RNLN. Subsequently, Chapter 3 provides some background information with regard to the POMDP model followed by a literature review on the applications within the maintenance field. Chapter 4 continues with the guidelines defined after which a small case study is performed and presented in Chapter 5. Finally, the report is concluded with a conclusion, and some recommendations for future research.

1.1 Short Introduction to the RNLN

The Netherlands Ministry of Defence is the military organization of the Kingdom of the Netherlands. The organization consists of armed forces, special organizational units, and supporting organizations (Ministerie van Defensie, 2020). The structure can be seen in Figure 1. The Minister of Defence is politically in charge of the Ministry of Defence. The armed forces are in the hands of the Commander of the Armed Forces and subdivided into the Royal Netherlands Navy (RNLN), the Royal Netherlands Army, the Royal Netherlands Air Forces, and the Royal Netherlands Military Police. Next to the armed forces, there are two other departments, namely the Joint Support Command and the Defence Materiel Organization (DMO). Where the Joint Support Command performs supporting tasks for the armed forces, DMO is a logistics supporting, administrative, and coordinating organization. Besides, DMO is responsible for purchasing new material. Altogether, the Ministry of Defence is responsible for the protection of The Netherlands and its overseas territories, and the human, financial and material resources available for this purpose (Ministerie van Defensie, 2020). Their main tasks can therefore be summarized as follows:

- Protection of the Dutch territory and that of its allies;
- Foster the (international) law order and stability;
- Offer support in case of disasters and crises.

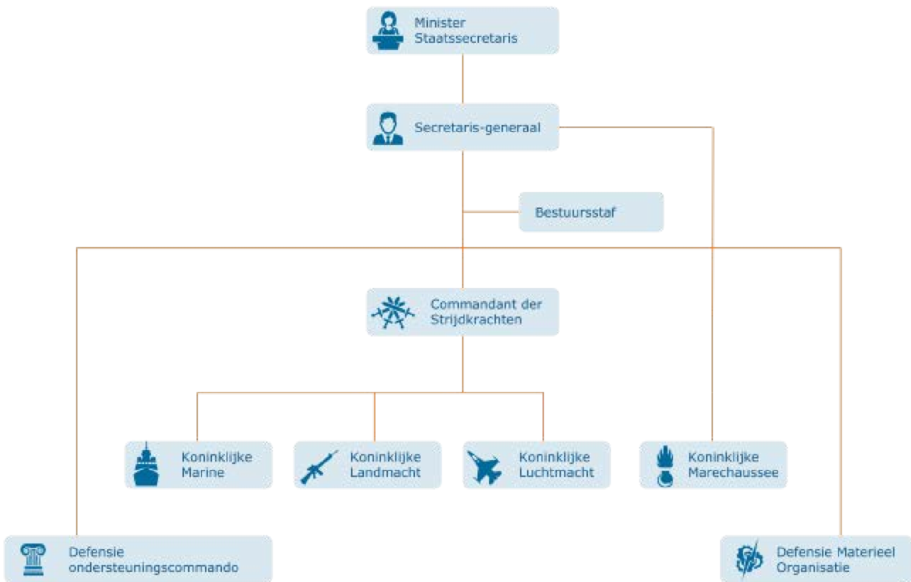

Figure 1: Organization of the Netherlands Ministry of Defence (Ministerie van Defensie, 2020)

The current research is performed at the RNLN. This armed force is the maritime division of the Ministry of Defence and operates from the sea. It has a total of 29 large and many smaller vessels in operation to ensure safety at sea by securing seaways and nodes (Ministerie van Defensie, 2020; Karremann, 2020). Lots of piracy, drug transport, and human trafficking are combatted. Besides the safety at sea, the RNLN also provides safety from sea meaning that RNLN-units can be used ashore from the sea to carry out or support land operations. Finally, there are several other important social contributions in The Netherlands itself. They, for example, fight terrorism, clean up unexploded explosives or support coastguards. Next to their main tasks, the Sail Plan of the RNLN indicates what needs to be realized before a specified year. Several goals are written and the aim is to achieve them on top of their common operations.

The RNLN, in turn, consists of the following divisions; Directie Operaties, Directie Materiële Instandhouding (DMI), Directie Personeel en Bedrijfsvoering, Commando der Zeemacht Caribisch gebied, and Kustwacht Caribisch gebied (see Figure 2) (Ministerie van Defensie, 2020). The primary processes of preparation and maintenance, the focus of this report, take place in DMI. DMI is the maintainer and disseminator of the RNLN to make deployment possible. In general, this department is responsible for the maintenance. Though, small maintenance tasks can be carried out by the users of the vessels making DMI responsible for large maintenance jobs only.

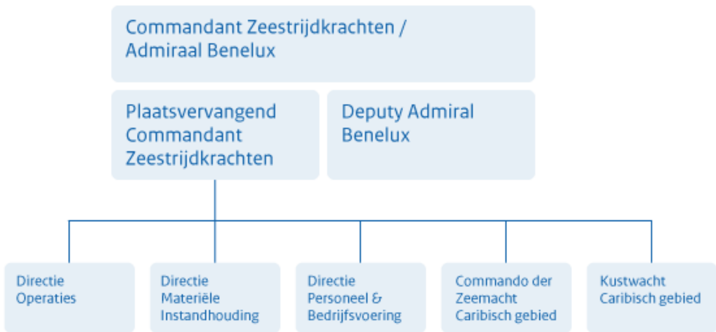

Figure 2: Organization of the RNLN (Ministerie van Defensie, 2020)

Within DMI there are multiple departments that are each responsible for a specific task regarding the maintenance. This research is conducted at one of these departments: Data for Maintenance (in Dutch: Data voor Onderhoud, DvO). The DvO team was founded in April 2019 and contributes to the vision of the RNLN as written in their Sail Plan 2030. DvO facilitates DMI in their transition towards smart maintenance using big data. This project consists of data infrastructure, data governance, data acquisition, asset management, and data analysis. The goal of the DvO team is to bring initiatives from initial development to practice realizing the strategic goals for the RNLN's future maintenance.

1.2 Research Motivation

One of the top priorities at the RNLN is optimizing its maintenance, as high costs are associated with it. Within the Sail Plan 2030 of the RNLN, many upcoming technologies and opportunities are mentioned. One of these opportunities is the use of big data to perform smart maintenance (Margés, 2019). At this moment, the RNLN is making significant steps forward in using physical models and data to diagnose and predict the failure of ship components. However, after making these kinds of predictions, it is of great importance to apply the right maintenance policy so the RNLN can take advantage of it.

In the current situation at the RNLN, a lot of maintenance is done based on calendar time or running hours. Occasionally the state of the system is observed whereafter a decision is made. In general, there are two main decisions that can be made when planning maintenance. The engineer could either choose that a preventive maintenance action should be conducted or decide to do nothing and wait for the next maintenance moment. Furthermore, in exceptional situations, the engineer can decide to perform an inspection in addition to the common choices. Subsequently, the engineer can decide based on this inspection which decision is optimal. The RNLN is wondering if it would be advantageous to perform intermediate and more in-depth inspections instead of taking the risk to wrongly decide on which action to take. Obviously, a lot of different aspects are involved in observing the state of the system, and subsequently, planning the maintenance actions. In many situations, it is not even possible to determine the true state of the system but only a belief (i.e. estimation) is made. In other words, the state of the system is only partially observable. As a consequence, additional and more in-depth inspections could be needed to find out this true state. These inspections are a third option in addition to the previous two choices the engineer could make. The question is though, when to perform such an inspection and when to decide to just perform maintenance?

In order to gain more control on when to perform such an inspection and when to just perform maintenance, the RNLN is interested in formally modeling these partially observable situations to support engineers in making their decisions. One way of doing this is by the use of a POMDP. This specific model is now investigated within the PrimaVera project in which the RNLN has a share. The "PrimaVera: Predictive maintenance for Very effective asset management" project develops new big data algorithms to better predict disruptions to infrastructure and production resources and thus to better plan maintenance (Ton et al., 2020). The vision of the PrimaVera project is to make the deployment of predictive maintenance easier and more effective. As the current research is also part of the PrimaVera project, the RNLN is interested in knowing if POMDP could be applied within their organization. Consequently, this research contributes to the investigation of POMDPs within the PrimaVera project.

POMDP is a specific form of a Markov decision process. However, in a POMDP it is assumed that the underlying state is not directly observable. The applicability of this model for the RNLN has not been investigated yet. Therefore, it is not clear if the model would be appropriate for some assets at the RNLN and thus if it could lead to an improved maintenance policy in specific situations. Though, when the model seems to be usable, the RNLN is also interested in knowing for which components it is possible to model the situation as a POMDP, and how to optimize the maintenance moment accordingly.

Altogether, the interest at the RNLN in knowing the applicability of POMDP leads to the following research objective for the organization:

> *Investigate whether a partially observable Markov decision process is a beneficial method to apply condition-based maintenance within the RNLN, and define guidelines to select suitable components for the partially observable Markov decision process to be able to optimize its maintenance planning accordingly.*

1.3 Research Questions

To make first steps towards realizing the objective defined in Section 1.2, a main question is stated. The presented goal for the RNLN requires investigation in the applicability of POMDP from which guidelines can be found that are necessary to model a system or component within the RNLN as POMDP, and finally optimizing the maintenance planning. However, within this research only the first two aspects are investigated. This leads to the following main question:

> *Would a partially observable Markov decision process be a suitable method for some components to apply condition-based maintenance within the RNLN?*

This main question is accompanied by several research questions to give further direction:

1. What is the current way of maintenance planning at the RNLN?
2. How could a POMDP contribute to the maintenance process when applying condition-based maintenance?
3. Which guidelines can be defined to model a system or component within the RNLN as (PO)MDP?
4. Can a component within the diesel engine be found for which these guidelines apply?
5. If we can find a specific component, how does the textual description of the model look like for this component?

First, we look at how the current way of maintenance planning looks like at the RNLN as the POMDP should apply to these situations (RQ1). Background information on maintenance policies is given followed by a description of the maintenance planning as executed at the RNLN. Subsequently, with the second research question (RQ2) it is explored if POMDP could be a benefitting method for maintenance at the RNLN. A more detailed description is given of the Markovian property, the Markov decision process, and the POMDP. Besides, a literature review is done on some of the current maintenance applications. Additionally, it is investigated which guidelines are needed to model a system or component as a POMDP answering the third research question (RQ3). These guidelines give the opportunity to find a component at the RNLN on which the POMDP can be applied (RQ4). Finally, the last research question (RQ5) gives a textual description of the POMDP models for the components found in RQ4 using the guidelines of RQ3. Thereafter, it is concluded if a POMDP would be a suitable model for some components to apply condition-based maintenance (CBM) at the RNLN and thus the main question is answered.

1.4 Method

The aim of this research is to investigate whether POMDP would be a beneficial method to apply condition-based maintenance at RNLN. First, definitions of the concepts used (e.g. maintenance policies, and POMDP) were explored and reported to get a clear understanding. Subsequently, via an interview, the current way of planning maintenance is discovered making sure RQ1 is answered. A literature review on maintenance applications of POMDP and several conversations with experts gave sufficient knowledge to answer RQ2 and RQ3. The guidelines of RQ3 are, in turn, confirmed by several researchers in the

PrimaVera project. Next, RQ4 is solicited from experts at the RNLN having much knowledge about the main engine using the flowchart designed. Finally, using the knowledge gathered from the expert about the components found we can answer RQ5.

To conclude, an extensive literature review and interviews with experts and engineers ensured that all research questions are answered. Accordingly, the main question was answered as a result of the research questions.

1.5 Scope

The scope of this research is limited due to the short amount of time. Each ship of the RNLN consists of multiple thousands of components. A component is defined as a small, self-contained part of a larger entity. The components dealt with in this research are considered line replaceable units (LRUs) within the RNLN. A system, on the other hand, consists of multiple components. Examining all the components of a ship will take a lot of time. Therefore, it is decided to focus on the diesel engine of the ship meaning that the guidelines are only checked for the components within this system. The diesel engine is one of the critical components of a warship. In addition, DvO currently develops CBM methods for the diesel engine making them interested in modeling techniques that could be applied. Besides, it is decided not to optimize the maintenance planning and optimal moment mathematically. We make some first steps towards the objective of the RNLN by investigating if POMDP would be beneficial. The result of the research consists of a study of the applicability of POMDP at the RNLN and a set of guidelines to identify which components are suitable for the model. In other words, the conclusion includes if it is beneficial to model specific systems or components at the RNLN as POMDP and which guidelines are required for a component to be modeled as POMDP.

2 Current Maintenance Optimization at the RNLN

This chapter provides background information about different maintenance policies. Besides, the current maintenance policy at the Royal Netherlands Navy (RNLN) is briefly described. Concluding this chapter, the first research question (RQ1) is answered.

2.1 Literature on Maintenance Policies

Keeping capital assets up and running is not only very costly but also of critical importance (Arts, 2017). Where regular production operations are investigated by demand from the customer, maintenance operations are by the need for maintenance of equipment. When a component or equipment must be replaced or maintained is determined by the used maintenance policy. Figure 3 shows an overview of the maintenance policies for assets. A distinction is made between modificative, preventive and breakdown corrective maintenance (Arts, 2017) which can also be seen as aggressive, proactive and reactive maintenance, respectively (Tiddens, 2018).

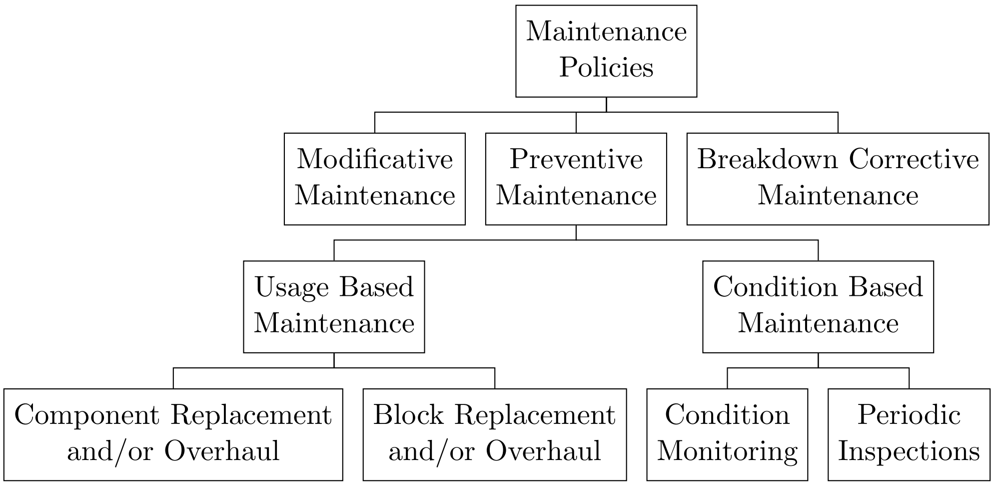

Figure 3: Maintenance policies adapted from Arts (2017)

Modificative maintenance interchanges a component with a technically more advanced component to make the equipment perform better (i.e. the design is improved) (Arts, 2017). Breakdown corrective maintenance is applied when the component has failed, while under preventive maintenance the aim is to replace components before a failure occurs (Arts, 2017). Since repair or replacement actions affect flexibility, throughput time, and quality of the operations and logistics (Tiddens, 2018), it is important to plan maintenance before a failure occurs but not far before the end of the component's service life. Therefore, preventive maintenance is a preferred policy. Moreover, the current research focuses on this preferred maintenance policy. Preventive maintenance can be further divided into usage-based maintenance, which means a component is maintained based on calendar time or running hours, and condition-based maintenance (CBM), see Figure 3. CBM determines the optimal maintenance moment based on condition monitoring information, such as power consumption, temperature or vibration (Karabağ et al., 2020). In this way, maintenance is only conducted based on the actual condition of the asset. Another way to find the condition of the component is by executing periodic inspections. According to Karabağ et al. (2020), CBM should help to reduce costs, increase the reliability of the system and maximize components' useful life. Besides, conducting maintenance only based on the actual condition of the component has proven to be more successful in preventing unexpected failures (Tiddens, 2018).

According to Mobley (2002), all preventive maintenance policies assume that components or systems will degrade within a time frame that matches their classification, and that this degradation depends on the usage. Therefore, predictive maintenance would be an interesting addition on preventive maintenance. Predictive maintenance can also be seen as a condition-driven preventive maintenance policy (Mobley, 2002). It refers to a policy that triggers maintenance activities by predictions of failures (Tiddens, 2018). With these predictions, the condition of the system can be estimated, and maintenance actions could be scheduled upfront based on condition that will lead to a more optimal planning. As a consequence, it can be stated that the maintenance planning depends on prediction of the condition of the asset.

2.2 Maintenance at the RNLN

To describe the maintenance optimization at the RNLN, an interview was held with an Operational Systems Engineer at DMI. In addition to this interview, multiple conversations are held with employees at the RNLN to supplement the information. Currently, the maintenance at the RNLN is usually planned preventively based on calendar time or running hours. Tasks, which vary in difficulty, are planned for an upcoming period leading to a planning horizon for each vessel. In addition, it is sometimes needed to perform corrective maintenance actions. Hence, within the maintenance at the RNLN we make a distinction in two areas, namely the periods in which maintenance is subdivided and the intensity of the maintenance tasks. These areas are described in the following paragraphs, thereafter the method used by the RNLN to plan their maintenance tasks is explained.

First, the periods in which maintenance is subdivided. Roughly spoken, a distinction is made between BO and AM. BO is the abbreviation for Benoemd Onderhoud (i.e., appointed maintenance) which takes place once every five years. During a BO, vessels are usually dry-docked at the base in Den Helder for approximately a year. Within this year large-scale maintenance is performed that cannot take place in the water but also other planned jobs are executed. During the four years between these BO periods, AM periods are planned. AM is called Assisted Maintenance and represents additional maintenance tasks that have to be conducted with the assistance of DMI. These periods are planned twice a year and last about four weeks. Note that there is some overlap in maintenance jobs that can be conducted in both AM and BO periods. However, during an AM period, it is not preferred to dry-dock a vessel.

Second, the intensity of each maintenance task varies. Maintenance planned during the different periods is usually based on a system or component-specific PO (i.e., periodic or scheduled maintenance). For a PO there are various tasks that can be performed that are planned based on labor level. At the RNLN there are, for example, daily tasks, tasks based on running hours, or yearly tasks. The intensity of these maintenance tasks differs and can be roughly divided into three levels. First, there is Organic Level Maintenance (OLM), or maintenance by the unit itself. In other words, this is what the maintainer performs on board of the vessel. Second, there is maintenance that is assisted by DMI. This level of tasks is called Intermediate Level Maintenance (ILM). ILM tasks are often performed during AM periods. However, in some cases, they can be postponed or brought forward to a BO period. Lastly, there is Depot Level Maintenance (DLM) that is performed during a BO period as the vessel then usually has to be dry-docked for these tasks. Note that it is not mandatory to dry-dock the vessel for a DLM task.

In addition to classifying the maintenance tasks, all maintenance tasks defined have to be scheduled. Currently, the RNLN makes use of a maintenance planning that is generated by SAP, the used enterprise resource planning (ERP) software. SAP generates notifications for the actions that must be carried out. So, for the daily, monthly or weekly tasks everything is processed in SAP. These tasks are preventive, thus often usage- and time-based because the system or component is maintained based on calendar time or running hours. Depending on the level of the maintenance task it can either be performed directly by the crew on board or with the assistance of DMI. More advanced or comprehensive tasks are planned within a BO or AM period. Tasks are visible beforehand in SAP and can, subsequently, be scheduled. This means that before the BO or AM period it is already clear what maintenance is performed. For example, if a specific task must be performed in 3 months, and the AM period is before that time, you will schedule the task earlier.

All in all, the current maintenance optimization within the RNLN is predominantly usage- and time-based. Preventive maintenance tasks are planned based on calendar time or running hours registered in SAP. Where OLM tasks can be performed at almost any time, ILM and DLM are usually scheduled during AM and BO periods respectively. Nevertheless, as the number of sensors available increases, the RNLN sees potential in using the condition of an asset rather than fixed time epochs to plan maintenance actions. There already are some examples where the condition of an asset serves as input for knowing the state and accordingly decide on the maintenance action. According to Marc Los there are, for example, four monthly vibration measurements on all sorts of systems. However, the RNLN would like to use more CBM to enrich their passive system of periodic maintenance.

3 Description of the Partially Observable Markov Decision Process

In this chapter, we answer RQ2 by explaining the modeling technique that we investigate: the partially observable Markov decision process (POMDP). First, the basics of the model, i.e. the Markov property, are described after which the Markov decision process (MDP) is elaborated. Finally, the extension is added to this MDP leading to the POMDP. The chapter is concluded by a literature review on applications of POMDP in maintenance.

3.1 Markov Process

The evolvement of degradation of a large number of components can be modeled according to a Markov chain approach (Giorgio et al., 2011; Lin et al., 2017). In fact, in Markov chains, the basics of general stochastic simulation processes can be found. A Markov chain is one of the types of Markov processes. Markov processes are random processes in which only the present state influences the next future state (Ibe, 2013). It describes a sequence of possible events which future events are, given the present state of the process, independent of the past. This property is also referred to as the Markov property and illustrates the fact that the process is memory-less (Ibe, 2013). There are four basic types of Markov processes, namely discrete-time Markov chain, continuous-time Markov chain, discrete-time Markov process and continuous-time Markov process, see Table 1 (Ibe, 2013). Both the discrete-time Markov chain and the continuous-time Markov chain are characterized by their discrete-state space, where the discrete-time Markov process and the continuous-time Markov process have a continuous state-space.

Table 1: Markov processes adapted from Ibe (2013)

		State Space	
		Discrete	Continuous
Time	Discrete	Discrete-time Markov chain	Discrete-time Markov process
	Continuous	Continuous-time Markov chain	Continuous-time Markov process

Here, we focus on the discrete-state space and discrete-time process where each Markov process has a set of states which are part of the state-space (Maltby et al., 2020). Each state has its own definition depending on the type of process that is described. In the application of maintenance, there are three states that are often referred to; healthy, warning state and failure. However, depending on the component or system that is modelled these definitions and number of states can differ. The stochastic process makes movements between discrete states at times that can be fixed or random (Ibe, 2013). These movements are called transitions and occur with a certain probability. The system enters a state, sojourns an amount of time in that state, and then moves another state where it spends another amount of time.

Formally, a Markov chain is characterized by a set of states $S = \{s_1, s_2, \ldots, s_n\}$. The process starts in one of these states and, subsequently, moves from one state to another. A movement from state s_i to s_j occurs with transition probability p_{ij}. The process can remain in its current state, this occurs with probability p_{ii}. The transition probabilities satisfy the following conditions (Ibe, 2013):

1. $0 \leq p_{ij} \leq 1$
2. $\sum_j p_{ij} = 1$, $i = 1, 2, \ldots, n$, following from the fact that states are mutually exclusive.

The $n \times n$ transition matrix P contains all transition probabilities:

$$P = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1n} \\ p_{21} & p_{22} & \cdots & p_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ p_{n1} & p_{n2} & \cdots & p_{nn} \end{bmatrix}$$

Both the first and the second condition of the transition probabilities can be verified within this transition matrix P as the row sums should be equal to 1. In order to visualize such a problem a transition diagram can be made. Take for example the following problem Q:

$$Q = \begin{bmatrix} 0.60 & 0.40 \\ 0.35 & 0.65 \end{bmatrix}$$

Translating this problem to a transition diagram will give the representation as shown in Figure 4. The circles represent the states defined which are only two in this problem, where the arcs show the transition probabilities between the states.

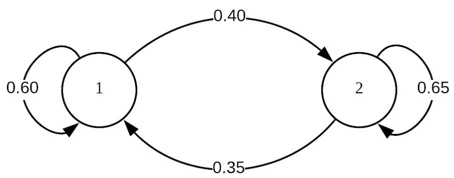

Figure 4: Transition diagram for problem Q

To determine the probability that the system will be in state j after exactly n transitions, given that it is presently in state i, the conditional probabilities $p_{ij}(n)$ are needed. These conditional probabilities can be obtained by multiplying the transition matrix by itself for n times (Ibe, 2013). When multiplying the transition probability matrix by itself many times, it is found that the values remain constant and all probabilities of one column converge to the same value. These limit values are called the limiting-state probabilities and can be found using the following prerequisite:

$$\pi_i \sum_{j \neq i} p_{ij} = \sum_{j \neq i} \pi_j p_{ji}, \quad i \in S$$

where $\sum_j p_{ij} = 1$, $i = 1, 2, \ldots, n$.

The resulting set of equations, also called the balance equations, balance the probability of leaving and entering a state in equilibrium. Solving them will give the limiting-state probability for each of the states defined. For problem Q these values are equal to $\pi_1 = 7/15$ and $\pi_2 = 8/15$ meaning that the probability that the process will be in state 1 is $7/15$ and in state 2 is $8/15$.

As presented in Table 1, the discrete-state space makes a distinction between discrete-time Markov chains and continuous-time Markov chains. The difference between these two is that in discrete-time Markov chain there is a "jump" to a new state at fixed time epochs but in continuous-time Markov chain the "jump" to a new state may occur at any time $t > 0$ (Ibe, 2013). The amounts of time the process spends in the current state before making a transition into a different state are independent and exponentially distributed random variables. For the continuous-time Markov chain a subscription of time (t) is added to the notation. It can be solved using the Chapman-Kolmogorov equation, see Ibe (2013) for more details on the continues-time Markov chain.

3.2 Markov Decision Process

A Markov decision process (MDP) models decision making in discrete and sequential environments (Littman, 2001). A decision-maker, or agent, inhabits an environment which changes state randomly in response to action choices made by the decision-maker. In other words, in contrast to the earlier described Markov processes, MDPs include a decision-maker who is required to make a sequence of decisions, also called actions, over time, see Figure 5. These decisions, however, have uncertain outcomes (i.e. random behavior). Each action taken by the decision-maker can either lead to a reward or provoke a cost (Ibe, 2013). The combination of the current state in which the system is located and the decision the agent makes, affects the transition probabilities of the future state. The overall objective of the decision-maker is to minimize the costs and thus maximize the long-term reward.

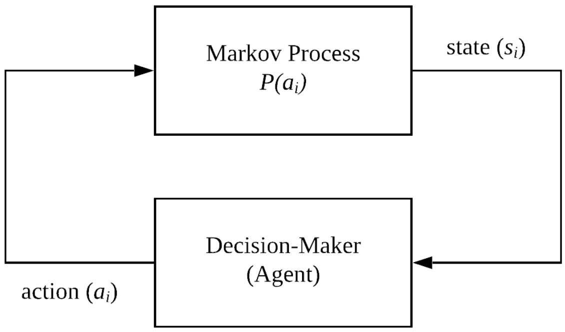

Figure 5: Illustration of MDP adapted from Krishnamurthy (2016)

Like the Markov chain, the MDP model consist of a state space $S = \{s_1, s_2, \ldots, s_n\}$ between which it can make a transition. The movements occur with transition probabilities p_{ij}. However, these transition probabilities depend on the action that is chosen by the decision-maker. In Figure 5, a decision-maker knows the current state is s_i and will choose an action a_i at each decision point in time. This action is part of the action space $A = \{a_1, a_2, \ldots, a_n\}$, and will result in a transition to another state occurring according to the transitions probabilities in matrix $P(a_i)$. The probability we transition to state s_j if we choose action a_i in state s_i is denoted by $P(s_j|s_i, a_i)$. Each time epoch (i.e. decision moment) the decision-maker uses all the information available until that time to choose an action using policy μ_i (Krishnamurthy, 2016). A policy μ_i is a mapping from the information set to the action set. After transitioning from state s_i to s_j due to action a_i, there is an immediate reward or cost incurred by the decision-maker denoted by $R_{(a_i)}(s_i, s_j)$. After the state evolves, the decision-maker updates the available information and the process will be repeated at the next time epoch.

The goal in a MDP is to find a good policy for the decision-maker. The sequence of policies that the decision-maker uses from time 0 to $T-1$ can be denoted by:

$$\boldsymbol{\mu} = (\mu_0, \mu_1, \ldots, \mu_{T-1})$$

To evaluate the performance of the sequence of policies we need to specify a performance criterion or objective function $J_\mu(s_i)$. Subsequently, the objective is to choose a sequence of policies μ that minimizes the cumulative costs for every initial state (Krishnamurthy, 2016):

$$\boldsymbol{\mu}^* = \underset{\boldsymbol{\mu}}{\operatorname{argmin}} J_\mu\left(s_i\right)$$

Alternatively, one can also find an optimal sequence of policies while maximizing the cumulative rewards. Note that a decision made at time t affects the costs incurred in the future, the choice of actions in the future, and the state's value in the future Krishnamurthy (2016). The optimal policy $\boldsymbol{\mu}^*$ can be found via value-iteration or the Bellman equation, see Tijms (1986) and Krishnamurthy (2016) respectively.

Figure 6 shows an example of a transition diagram of a simple MDP. The circles represent three states from which two actions can be chosen as illustrated by the squares. It can be seen that the transition probabilities are dependent on the combination of current state and action chosen. Besides, each combination occurs at a specific cost or reward represented by the orange numbers on the arrows. It is assumed that the decision-maker knowns in which state he or she is located. Therefore, the MDP as currently described is a "fully observed" model meaning that at each time epoch the agent observes the true state s_i (Krishnamurthy, 2016), and accordingly decides based on the actions and corresponding rewards.

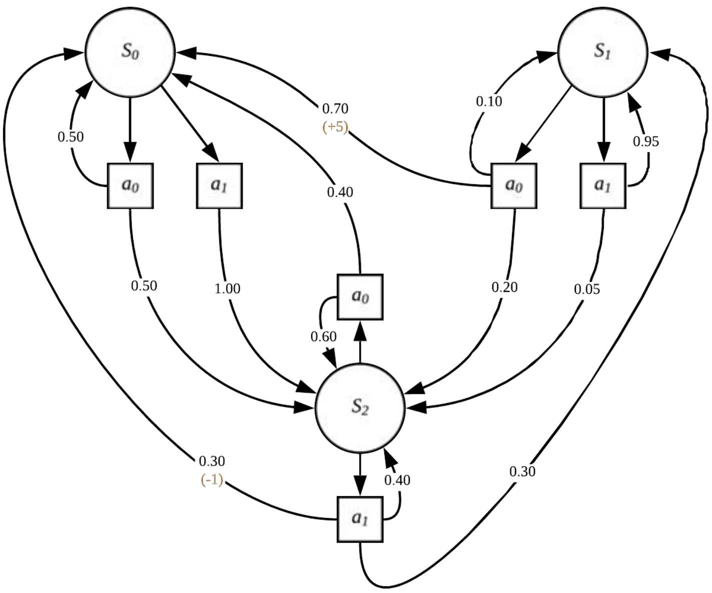

Figure 6: Example of transition diagram for MDP

However, in real world situations it is not always clear in which state the decision maker is located. In the current context this is referred to as partially observable. These kinds of partially observable situations can be modelled as well using the Markov properties described before. In fact, MDP provides a starting point for partially observable Markov decisions processes (POMDPs), which will be elaborated in the following section.

3.3 Partially Observable Markov Decision Process

In contrast to the MDP described in the previous section, a partially observable Markov decision process (POMDP) is not fully observable. The process being modelled in a POMDP is assumed to be a Markov process of which the states are unobservable. According to (Krishnamurthy, 2016) this process is also called a hidden Markov model (HMM) when the agent is disregarded. Within a HMM, noisy observations of the underlying state of the Markov chain can be estimated using the "HMM filter". The HMM filter makes an estimation of the posterior distribution of the state, referred to as the belief state. A POMDP can be seen as a *controlled* HMM (Krishnamurthy, 2016) because a decision-maker is involved. Besides, a POMDP is an adapted version of the previous described MDP. The HMM filter applied in a POMDP ensures an estimation of the posterior of the state of the underlying process for this model also referred to as belief state. In a POMDP, the decision-maker uses the belief state to choose the next action. Figure 7 shows a schematic representation of the POMDP. The decision-maker who currently beliefs to be in state s_i, chooses an action a_i based on this belief. This action causes a state transition from s_i to s_j of which a general measurement is visible. This general measurement (e.g. a noisy sensor) will in turn ensure an observation o_i of which the "HMM filter" can make a new estimation of the belief state.

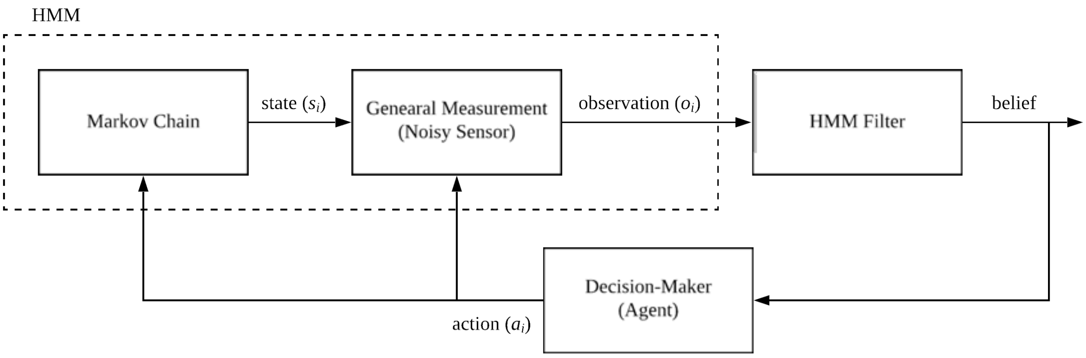

Figure 7: Illustration of POMDP adapted from Krishnamurthy (2016)

When considering a POMDP, the component should be suitable for condition-based maintenance (CBM). As explained before, CBM is a maintenance policy whereby systems or components are preventively maintained according to their perceived condition (Van Staden & Boute, 2020). A system or component might be suitable for CBM if there is freedom in scheduling inspections and maintenance operations based on the perceived condition. This free scheduling is, for example, possible for a car. You can bring your car to the garage whenever you like. Often you will decide to go to the garage depending on the current performance of your car illustrating the usefulness of a CBM policy. Another choice you can make is to preventively perform maintenance at home or you decide to do nothing and keep driving. These three options illustrate the most basic actions in the action structure of POMDP. The decision to do nothing will lead to no new information where inspection will reveal the true condition and thus leads to new information. Based on this new information one can choose to perform maintenance (e.g. repairing or replacing a component). One could also make a decision based on the belief state only. However, this belief state formed is probabilistically related to the internal state and can be observed wrong with a certain probability leading to, for example, performing maintenance too early or too late. It is usually assumed that the maintained or new component is as-good-as-new. Note that for both performing an inspection or performing preventive maintenance, there are costs associated. In the context of the Royal Netherlands Navy (RNLN) the same actions are of interest. Sometimes it would be more cost efficient to perform an inspection before directly performing maintenance. Though, some inspections could be difficult and costly to execute not leading to any advantages.

Overall, within the MDP it is assumed that the complete true deterioration state of the system being monitored is visible. However, this may not always be possible due to factors such as outside interference, inherent limitations in sensor quality or misspecifications in the deterioration model (Van Staden & Boute, 2020). Consequently, the exact condition of each component cannot be monitored. Instead, a system level failure or defect can be observed leading to a partial representation of the real state and thus asks for a POMDP.

Formally, at each decision epoch t, the system is in a specific state $s_i \in S$ where $S = \{s_1, s_2, \ldots, s_n\}$ is a set of states. The decision-maker chooses an action $a_i \in A$ which causes the system to transition to state $s_j \in S$ with probability $P(s_j|s_i, a_i)$. Subsequently, the decision-maker receives an observation $o \in \Omega$ where $\Omega = \{o_1, o_2, \ldots, o_k\}$ is a set of observations. This observation state is related to the state of the system with probability $O(o_i|s_j, a_i)$. We might assume that for every belief state the decision-maker forms from the observation, he or she has a certain percentage chance of giving it a wrong qualification, and a certain percent chance of getting the qualification correct. These percentages are depicted in the observation matrix. Finally, the decision maker will receive a reward or incur a cost $R_{(a_i)}(s_i, s_j)$ as described for the MDP. This process repeats. Again, like the MDP, the goal for the decision-maker is to choose a sequence of policies that minimizes the cumulative costs for every initial state, see Krishnamurthy (2016) for a detailed and mathematical elaboration.

3.4 Maintenance Applications of POMDP

To apply POMDP in the maintenance context, some relevant papers on POMDP applied for maintenance are reviewed in this section. First, we briefly discuss the importance of CBM. This concept is followed by an introduction to MDPs as they provide a starting point for POMDPs. Subsequently, the relevant papers on POMDP are elaborated with respect to the most common actions, the effect of access to different information levels, interaction with other organizational operations, applying the technique to large state spaces, and an application on deteriorating sensors instead of deteriorating systems or components.

CBM seems an important understanding within the literature on POMDP. The literature on POMDP goes hand in hand with CBM. This maintenance policy receives growing attention by researchers from diverse disciplines in multiple industries (Karabağ et al., 2020). CBM determines the optimal maintenance moment based on condition monitoring. It should help to reduce costs, increase the reliability of the system and maximize components' useful life (Karabağ et al., 2020). Additionally, Tiddens (2018) states that conducting maintenance only based on the actual condition proved to contribute to the reduction of needed corrective maintenance actions. His work, in fact, showed a three-stage funnel-based selection method to identify systems or components for which preventive maintenance could be useful. Within this three-stage funnel it is found that components with a high average downtime and low fault frequency are commonly suitable to perform CBM on.

Many works in the literature on CBM models are devoted to maintenance optimization for single-component systems that deteriorate according to a Markov chain over a finite set of degradation levels (Van Oosterom et al., 2017). Derman (1963) presents one of the earliest works on scheduling replacements to minimize long-run average costs. He models the problem where corrective maintenance is more costly than preventive replacement as a Markov decision process (MDP). In following research, many variations and extensions to this classic problem have been investigated. From all these variations and extensions, the majority uses Markovian degradation of the concerning system or component modelled. The importance of availability of degradation information is demonstrated by Lin et al. (2017). They consider a spare parts stock point that serves an installed base of machines. The degradation behavior of the critical components of the machines are described by a Markov process. The condition of these parts is monitored perfectly meaning the problem is modeled as a discrete-time Markov decision problem (MDP). The research showed that their model is an efficient and effective heuristic policy.

Markov decision processes are commonly used for maintenance optimization because most of the maintenance problems deal with sequential decision making (Van Staden & Boute, 2020). However, it often comes forward that it is not always technically or economically possible to install sensors dedicated to each component (Van Staden & Boute, 2020). When the true state of the system cannot be fully observed, a POMDP model is applied. One of the earliest works on POMDPs is from Ohnishi et al. (1986). They investigate a problem where there are three actions. These actions are close to those possible for the RNLN and thus equal to continue system operation with incomplete information, perform a perfect inspection at a cost, or replace the system. Ohnishi et al. (1986) showed that there exists an optimal inspection and replacement policy in the class of monotonic four-region policies. Only by assuming an additional condition with respect to the probabilistic property of the monitoring mechanism the results could be extended interestingly. The assumption of three actions is also observed by Morato et al. (2019) and Cassandra (1997). Morato et al. (2019) consider a model were examples of actions are "do-nothing" or "repair". On top of that, it is possible to gather observations by means of inspection or monitoring but one can also decide not to inspect. It was found that within the context of maintenance planning less inspections might be necessary if information is gathered by a general measurement. This is an important finding as inspections could be complex and expensive. According to Cassandra (1997) POMDP is applicable to problems of which the production and inspection observations are probabilistically related to the internal states. Besides, a POMDP can account for both state transition and observation uncertainty (Cassandra, 1997). Their work surveys several areas where POMDP can be applied. They find that sending out or performing inspection is time and cost intensive in almost every application. Dismantling a machine, for example, to determine the full internal state requires an expenditure of time and personnel.

There are of course multiple levels of inspection meaning that a perfect inspection will expose more information than an imperfect inspection. The effect of having deviating levels of information was researched by Van Staden & Boute (2020). They consider the effect of access to different levels of deterioration data quality. Within the model analyzed the decision-maker can either choose to perform maintenance, to pay for external sensor information, or to do nothing and continue with internal sensor information only. Van Staden & Boute (2020) concluded that the benefit in reduced cost may not be significant enough to justify investment in a more complex maintenance policy because using multiple external sensors reduce average long-run costs but also ensures that the policy will become more complex.

In practice maintenance decisions are often related to other organizational operational decisions. Karabağ et al. (2020) investigate when to intervene for maintenance and which spare parts to bring along. In their work the critical components of the system considered are all subject to deterioration during the time that the system is operating. Markov chain models are employed to represent the component degradation. Their work proposes a POMDP formulation to solve the joint problem of maintenance timing and spare parts selection. There is a general sensor providing partial information about the condition of the system. With a numerical experiment it is shown that using the optimal policy instead of the corrective and preventive policies leads to cost decreases. Besides, Karabağ et al. (2020) show that having full information leads to cost decrease compared to having partial information on the components' deterioration levels. Nevertheless, having full information is found to be more valuable for cheaper, less reliable components compared with more expensive, more reliable components. Celen & Djurdjanovic (2020), in turn, propose a method based on a POMDP to deal with interactions between maintenance and production operations in flexible manufacturing systems in which equipment conditions are not perfectly observable. Their model was benchmarked against traditional CBM policies and thus showed that, regardless of uncertainties in knowing the true state, jointly making maintenance and production sequencing decisions outperforms.

The research conducted by Papakonstantinou & Shinozuka (2014) showed that the superior attributes of POMDPs enable optimum decisions based on the belief state. Their work is focused on modeling and solving the problem of finding optimal policies for the maintenance and management of aging structures. The POMDP framework proposed has a large state space that can adequately and sufficiently describe real-life problems. Papakonstantinou & Shinozuka (2014) prove that despite that it is hard to solve, a POMDP framework is an excellent choice for decision making and asset management under uncertainty. This statement applies especially for large models with thousands of states.

Van Oosterom et al. (2017) consider another way of applying POMDPs. Within their model it is not the system deteriorating but the sensor's costless observations of the binary state of the system become less informative over time. A full inspection provides the opportunity to perfectly determine the state of the system and to replace the sensor. The problem of adaptively scheduling full inspections and sensor replacements is formulated using a POMDP showing the broad applications of the modeling technique.

4 Guidelines for POMDP at the RNLN

In this chapter, the guidelines for applying partially observable Markov decision processes (POMDPs) at the Royal Netherlands Navy (RNLN) are elaborated, answering the third research question (RQ3). Section 4.1 presents the guidelines that were drawn up in response to the literature review and interviews with partners of the PrimaVera project. The experts on POMDP with who the guidelines were discussed and verified are Ayse Sena Eruguz, assistant professor at the Vrije Universiteit Amsterdam, Ragnar Eggertsson, PhD candidate at Eindhoven University of Technology, Rob Basten, associate professor at Eindhoven University of Technology, and Thom Bading, PhD candidate at Radboud University. Furthermore, from the guidelines a flowchart is created to be able to discuss it with experts on the technical systems of the vessels, see Section 4.2.

4.1 Guidelines for POMDP at the RNLN

Table 2 gives an overview of the eight guidelines for POMDP including those that are relevant within the context of the RNLN but not formally needed. The guidelines are explained one by one in the remainder of this section.

Table 2: Overview of the eight guidelines for POMDP at the RNLN

	Guideline	**Formally needed?**
1	It has to be useful to perform condition-based maintenance on the asset.	Yes
2	There has to be a degradation process/pattern over time for the asset.	Yes
3	The degradation process or pattern over time for the asset should follow a Markov process.	Yes
4	The exact condition and/or all underlying failure mechanisms of the asset cannot be completely observed.	Yes
5	There has to be a measurement (e.g. sensor or recurring measurement) that provides a general impression/observation of the asset's current performance/condition.	Yes
6	The general observations have to be probabilistically related to the exact condition of the internal states of the asset.	Yes
7	A more in-depth inspection has to reveal the complete condition of the asset.	No, for RNLN
8	It has to be beneficial (i.e. in terms of cost) to perform a more in-depth inspection compared to preventive maintenance.	No, for RNLN

GL1: It has to be useful to perform condition-based maintenance on the asset. Within the work of Tiddens (2018) it was written how to better understand the use and adoption of predictive maintenance and based on these observations a tool was developed to support the practical application of predictive maintenance (i.e. CBM). His three-stage funnel-based selection method identifies for which components condition-based maintenance (CBM) would be suitable to apply. This approach checks whether or not CBM would be beneficial based on the average downtime and the fault frequency. The components that have a high average downtime and low fault frequency are suitable for CBM Tiddens (2018). Besides, it should be possible to freely schedule inspections and maintenance action to perform CBM. This proposes the first important guideline for applying a POMDP, namely that it has to be useful to perform condition-based maintenance on

the asset (GL1). When the asset does not satisfy this first guideline, POMDP has no added value for the specific system or component as we want to take actions based on the condition. In fact, it is not useful to perform a CBM policy for the asset at all. Note that to check for this guideline the three-stage funnel-based selection method of Tiddens (2018) can be used.

GL2: There has to be a degradation process/pattern over time for the asset.
Subsequently, all literature speaks of the condition of the asset on which maintenance actions should be decided. However, to be able to use this condition we should be able to distinguish different states. The condition of the asset should deteriorate over time otherwise no maintenance actions are needed. In other words, components have to be subject to deterioration during the time that the system is operating (Karabağ et al., 2020). Therefore, the second guideline states that there has to be a degradation process or pattern over time for the asset meaning that it has to be possible to define separable states for the condition of the system or component.

GL3: The degradation process or pattern over time for the asset should follow a Markov process.
The need for separable states as referred to in GL2 also points to the mathematical requirement that the underlying degradation process of the asset has to follow a Markov model. The next future state should only depend on the present state (Ibe, 2013), and the action chosen by the decision-maker. The previous sequence of actions and transitions should not influence what occurs in the future. The degradation process between states should evolve according to a Markov chain with a finite state space, and there should be at most one state transition per period (Lin et al., 2017). Therefore, the third guideline states that the degradation process or pattern over time for the asset should follow a Markov process (GL3). In other words, the separable states for the condition of the system or component should comply to a Markov model. However, it should be stated that if the process or pattern seems history-dependent it does not necessarily mean that the system is not Markovian (Powell, 2014). Instead, the state variables of he process are not properly modeled meaning that the so-called history-dependent problems have an incomplete state variable. Nevertheless, in the current case it is assumed that when it is hard to properly model the state variables, the process is too complex, and thus modeling the asset as Markov process is not desirable in practice.

GL4: The exact condition and/or all underlying failure mechanisms of the asset cannot be completely observed.
As mentioned before, CBM determines the optimal maintenance moment based on condition monitoring information (Karabağ et al., 2020). There can be different levels of availability of information but also different types of observations such as power consumption, temperature or vibration. To model a Markov decision problem, it should be assumed that the condition is monitored perfectly at the beginning of each time epoch (Lin et al., 2017). Though, the assumption of having full information on the condition of one component should not be true for the current situation since then we deal with a MDP instead of a POMDP in which we are interested (Cassandra, 2020). An example of having incomplete information is that there could be multiple failure mechanisms of which it is not clear which one is at issue. It could also be the case that we just observe an overhead measurement that only gives a general impression of the condition of the asset. However, in all cases we should have some control over the state transitions meaning that a decision-maker has to be involved (Cassandra, 2020). This control distinctions the decision process from a Markov chain model. Hence, the exact condition and/or all

underlying failure mechanisms of the asset cannot be completely observed is included as fourth guideline (GL4).

GL5: There has to be a measurement (e.g. sensor or recurring measurement) that provides a general impression/observation of the asset's current performance/condition.
Within a POMDP, the decision-maker chooses an action which causes the asset to transition from state to state with a certain probability. The decision-maker should be able to receive an observation to form a believe about the new state the system is in. The fourth guideline concerns the fact that there has to be an overhead measurement that provides a general observation of the current performance or condition of the asset. POMDP can account for observation uncertainty (Cassandra, 1997); though, there should be a general observation to decide which action to choose (Krishnamurthy, 2016). Note that formally spoken, it is not needed to observe a condition. Instead, a believe state can be developed based on time. Though, in that case the problem boils down to making use of the failure distribution and thus CBM is no longer applicable. Therefore, it is stated in the fifth guideline that there has to be a measurement (e.g. sensor or recurring measurement) that provides a general impression/observation of the asset's current performance or condition (GL5). In the context of the RNLN an overhead measurement could either be a sensor value or a weekly simple inspection on the outside of the system or component. In turn, further in-depth inspections could reveal the true state.

GL6: The general observations have to be probabilistically related to the exact condition of the internal states of the asset.
Next, after an observation state is defined, we have to relate this state to one of the internal states of the asset with a certain probability. Cassandra (1997) showed that POMDP is applicable to problems of which the production and inspection observations are probabilistically related to the internal states. The overhead measurement required should give a general observation that is related to the internal states that were defined on the degradation process. The probabilities combined with these relations depend on the observation matrix. So, for example, we might assume that with a chance of 80% we have the qualification correct but with a 10% change the inspector has qualified a state lower or higher. With the observation being probabilistically related, the decision-maker can form the belief state and, subsequently, decide which action to choose. The sixth guideline is therefore as follows; the general observations have to be probabilistically related to the exact condition of the internal states of the asset (GL6).

GL7: A more in-depth inspection has to reveal the complete condition of the asset; and GL8: It has to be beneficial (i.e. in terms of cost) to perform a more in-depth inspection compared to preventive maintenance.
In the context of the RNLN, it was experienced that the uncertainty about the state of the technical system or physical asset could lead to non-optimal decisions on the timing of the maintenance actions. The engineer's current choice is often between doing nothing or performing preventive maintenance. As suggested, they are interested in knowing whether performing a more detailed inspection instead of making the two-sided decision would make sense. Therefore, the previous guidelines are of interest as POMDP should be suitable at first. However, to make sure POMDP could be beneficial for the RNLN it was decided to include two more guidelines concerning the carrying out of inspections to find the complete condition of the asset. It was already argued that this might improve the maintenance decision because instead of performing maintenance based on the general impression it can be chosen to obtain more information. Karabağ et al. (2020) found that having full

information is valuable and could decrease costs in the long run. The seventh guideline included will therefore be that a more in-depth inspection should be able to reveal the complete condition of the asset (GL7). It should be noted that there are different levels of inspection for which costs deviate. In proportion, the cost of an inspection should not exceed the cost of carrying out preventive maintenance. This means that when we require an inspection, this inspection has to be economically interesting compared to just performing preventive maintenance. Consequently, the eighth guideline will be that it has to beneficial (i.e. in terms of cost) to perform a more in-depth inspection compared to preventive maintenance (GL8).

Summarizing, a total of seven guidelines are defined. If the asset satisfies the first five guidelines, it can be concluded that POMDP is a suitable and recommended model. When all seven guidelines are satisfied POMDP is concluded to be beneficial for the RNLN as well because then the POMDP becomes useful for practice. Furthermore, to conclude that a CBM policy would be interesting, at least GL1 should be satisfied. In the case that GL1 is satisfied, but one of the guidelines from GL2 up to and including GL6 is not, it can be concluded that CBM is useful, but the model will not be a POMDP. GL7 and GL8 were added for completeness in the context of practitioners as the RNLN but not formally required to model a component or system as POMDP. Note that when GL7 and GL8 hold, GL1 has to hold as well, but not the other way around. With these guidelines the third research question (RQ3) is answered from where the next step is to find a possible component for which the guidelines apply (RQ4).

4.2 Flowchart for POMDP at the RNLN

To find suitable assets for POMDP using the guidelines defined, expect knowledge is needed. It was decided to translate the guidelines into a flowchart. All guidelines are translated into one corresponding question. The flowchart can be discussed with the engineers of the RNLN, see Figure 8. Using the total of eight questions that are related to the guidelines, it can be decided if the system or component discussed is found to be beneficial. In other words, the flowchart is used as input for finding a component to which the guidelines defined apply. The next chapter shows the application through a case study in which the flowchart is applied to the diesel engine.

The first question of the flowchart relates to GL1 and reads as follows:

1. Would condition-based maintenance be a useful policy to apply?

As explained before, this question will be answered using the approach of Tiddens (2018) meaning that this question will not be asked to the RNLN's engineer. His approach consists of a three-stage funnel-based selection method. The first step of the funnel has to reduce the number of suitable systems or components by a traditional filtering on failure frequency and impact on the company. The second and third step filter potential showstoppers and investigate the technical and economic feasibility of developing a CBM model for the selected systems or components. Tiddens (2018) uses the four-quadrant chart based on the work of Labib (2004), Lee et al. (2009) and Tinga et al. (2017) to find the candidates for which CBM would be promising. These candidates found usually have a high average downtime and low fault frequency. See Tiddens (2018) for the exact understanding of his three-stage funnel-based selection method. The components qualified in the predictive maintenance quadrant are found to be useful for CBM.

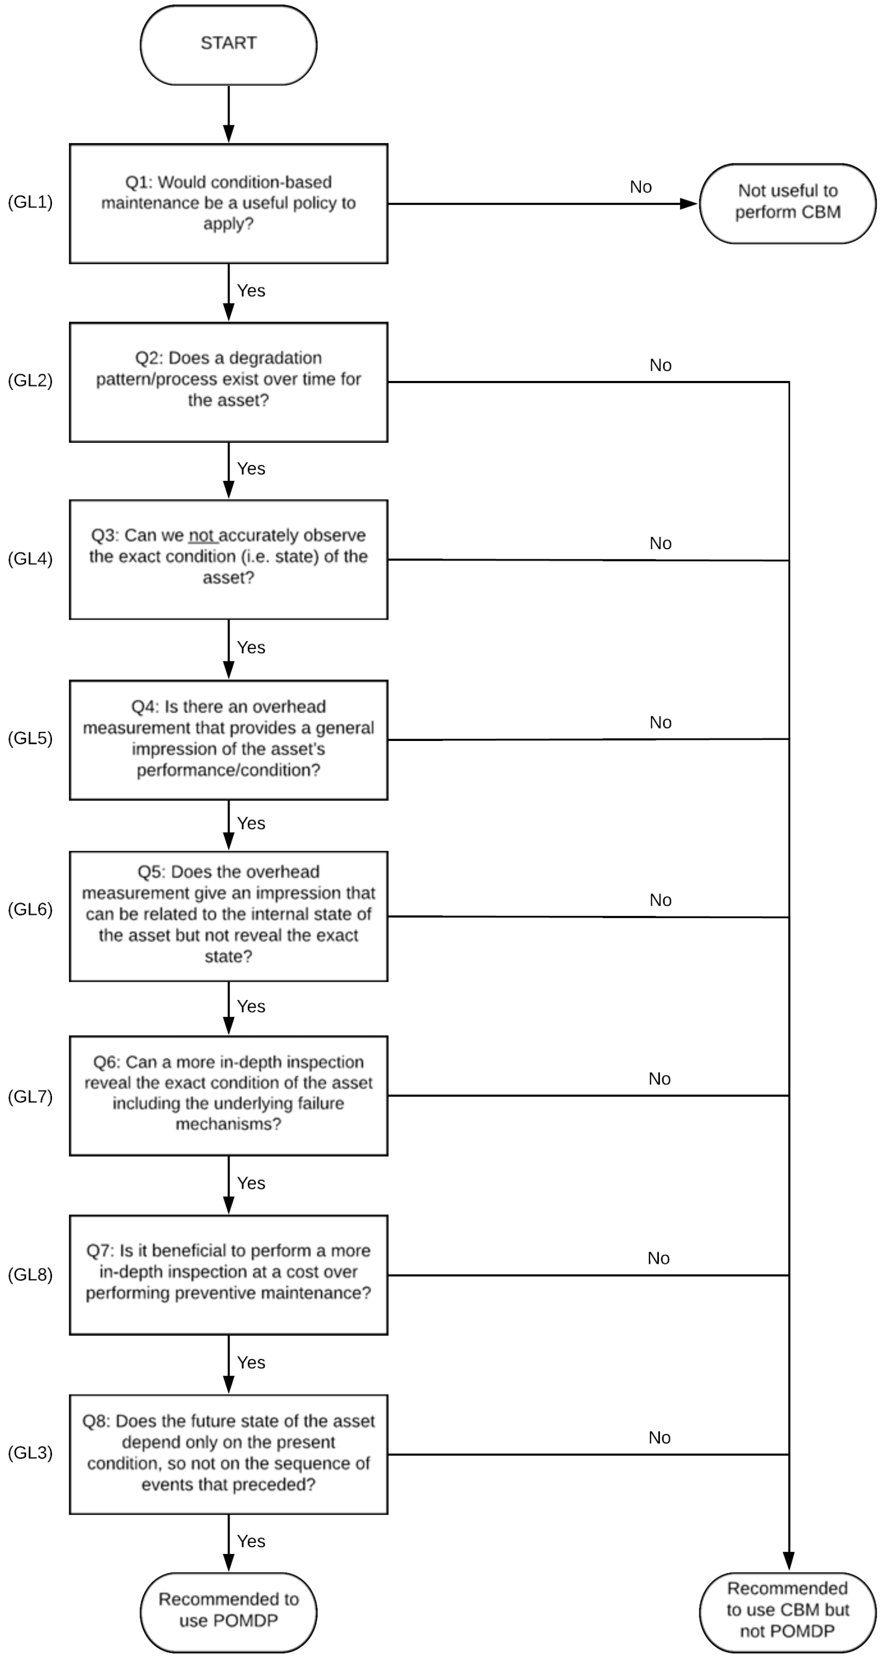

Figure 8: Flowchart based on guidelines for POMDP

The second and third guideline are related with the following two questions:

2. Does a degradation pattern/process exist over time for the asset?

8. Does the future state of the asset depend only on the present condition, so not on the sequence of events that preceded?

First, GL2 refers to the presence of a degradation process where GL3 is about the Markov property. As this latter concept is assumed to be hard to understand for an engineer it is decided to place this question at the end of the flowchart, therefore this question is numbered with 8. Note that the engineer agreeing with this eighth question does not lead to evidence of the degradation process being a Markov process. Though, a confirmation on this property is assumed to be satisfactory in the current research. Moreover, as explained in Section 4.1 it does not mean that a process that seems to be history-dependent is not Markovian. Though, in the current case it is assumed that it is not desirable to model these processes as Markov process leading to exclusion of the asset based on the eighth question.

Subsequently, GL4 and GL5 are translated to the following questions respectively:

3. Can we not accurately observe the exact condition (i.e. state) of the asset?

4. Is there an overhead measurement that provides a general impression of the asset's performance/condition?

The degradation process that is present by the second question should not be completely observable using traditional CBM and without any "investments". However, as explained, there should be an overhead measurement. Question 3 and 4 will cover these two requirements for POMDP.

GL6 points to the probabilistical relation between the belief state formed by means of the general impression, and the internal state. The overhead measurement should give an impression that can be related to the underlying state but not reveal the complete state. Therefore, the fifth question is defined as follows:

5. Does the overhead measurement give an impression that can be related to the internal state of the asset but not reveal the exact state?

One could ask if there can be, for example, be multiple causes for deviations in the observation of the current performance/condition leading to uncertainty about the state to clarify. Though, it should be noted that multiple causes are not mandatory for a POMDP as explained in Section 4.1.

Finally, the seventh and eighth guideline will lead to two additional questions:

6. Can a more in-depth inspection reveal the exact condition of the asset including the underlying failure mechanisms?

7. Is it beneficial to perform a more in-depth inspection at a cost over performing preventive maintenance?

These questions are important in the context of the RNLN but not formally required to model the situation as POMDP as explained in Section 4.1. They make it much more plausible for the RNLN that it will be economically interesting.

5 Case Study on Finding a Component for POMDP at the RNLN

In this chapter, the guidelines defined are tested on the diesel engine. This case study allowed checking if a component could be found for which the guidelines apply (RQ4) and, on top of that, verify the applicability of the guidelines. Section 5.1 gives a short description of the diesel engine, on which the case study is performed. Subsequently, the components of this system are studied with the help of the guidelines. Finally, the result of the case study is elaborated in Section 5.3 by giving a textual description of the model such that the fifth research question is answered (RQ5).

5.1 Diesel Engine

As explained earlier, the current research was scoped to the diesel engine of the vessel (see Section 1.5). Figure 9 shows an impression on how such a system typically looks. The diesel engine is part of the diesel electric transmission system on board which is one of the primary functions on board. Therefore, the diesel engine is a critical equipment of the vessel. It is a very complex reciprocating system with numerous components that work together to convert the potential energy from fuel into propulsion force. The cylinders within the diesel engine work according to a combustion cycle, see Armstrong Proctor (2019) for more details on the diesel combustion cycle.

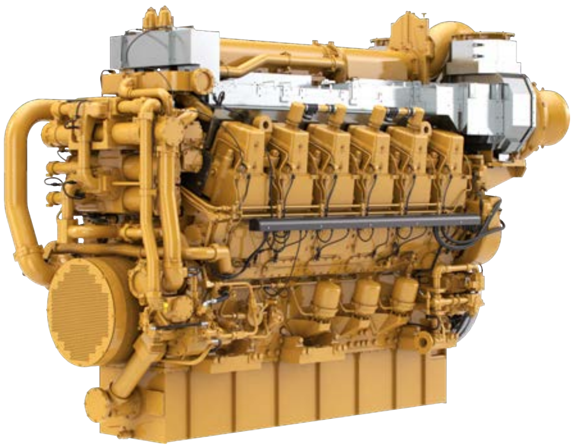

Figure 9: Illustration of a typical ship diesel engine (*Google Image Result*, 2020)

The Royal Netherlands Navy (RNLN) provided a failure mode, effects, and criticality analysis (FMECA) of the diesel engine as input for the case study (Tiddens, 2014). This FMECA was a result of an analysis conducted to "investigate" which components are of interest for a condition-based maintenance (CBM) policy within RNLN for the MaSeLMA project (Maintenance and Service Logistics for Maritime Assets). The output of this FMECA is a list of 95 components with each having a unique failure mechanism. The failure mechanisms, failure causes, downtime in hours and failure frequency is recorded for each of these components.

5.2 Finding Components of the Diesel Engine for POMDP

To determine if the components of the diesel engine are suitable for a partially observable Markov decision process (POMDP), an interview with two diesel engine experts was conducted. For both experts it is assumed that they have sufficient technical knowledge on

Table 3: Components of the diesel engine satisfying GL1

	Q1 (GL1)	Q2 (GL2)	Q3 (GL4)	Q4 (GL5)	Q5 (GL6)	Q6 (GL7)	Q7 (GL8)	Q8 (GL3)
Air cooler	+	+	+	+	+	+	+	+
Turbo	+	+	+	+	+	+	+	+
Gear	+	±	+	−				
Foundation rubbers	+	+	±	±	+	+	−	
Main bearing	+	+	+	+	+	+	+	+
Connecting rod bearing	+	+	+	+	+	+	−	
Vibration damper	+	+	+	+	+	+	−	
Piston ring	+	+	+	+	−			
Piston pin bearing	+	+	+	−				
Piston lining	+	+	+	+	±	+	±	
Camshaft	+	+	±	−				

the diesel engine. Within the interview the flowchart was discussed step by step. However, the first question regarding the usefulness of CBM was already answered using the three-stage funnel-based selection method (see Section 4.2). This method is applied on the FMECA of the diesel engine by Tiddens (2014). The output of the selection method led to eleven components within the preventive maintenance quadrant thus satisfying the first guideline. These remaining eleven components were, in turn, questioned one by one to the experts to check for the other guidelines. Table 3 shows the results of the interview, it can be seen that three of the eleven components satisfy all guidelines and thus are suitable for a POMDP. This discovery means that we can answer the fourth research question positively as we found more than one component within the diesel engine for which all guidelines apply.

When looking in more detail in Table 3, it is noticeable that the majority of the components are not suitable because of GL5 or GL8. This means that for the gear, the piston pin bearing and the camshaft it is not possible to have an overhead measurement denoting that it is not possible to form a belief about the condition of these components. For the foundation rubbers, connecting rod bearing, vibration damper and piston lining it is not beneficial to perform an inspection compared to directly performing preventive maintenance. Therefore, it can be concluded that in terms of cost savings it would not be beneficial to model these components with a POMDP. Though, note that the liability for inspection is not a hard requirement for the formulation of a POMDP. GL8 eliminates these specific components but it could be argued that a POMDP is suitable. However, in the context of the RNLN, it is concluded that POMDP is not beneficial for them Therefore, they are not included in the list of components that is suitable nor recommended for POMDP.

5.3 Results of Case Study

The case study of applying the guidelines on the diesel engine led to three components for which it is found that POMDP is a beneficial method and thus recommended, namely the air cooler, the turbo and the main bearing. This section briefly discusses per component per guideline why these components are suitable showing how the model will look in short textual description. This description is established in cooperation with the experts on the diesel engine. After reading this section the fifth and final research question (RQ5) is answered.

Air Cooler

The first component that satisfies all guidelines is the air cooler. The air cooler is a heat exchanger that is used within the diesel combustion cycle to cool back the inlet air. In the turbo engine, the air is pressurized to allow air to enter the combustion chamber and inject more fuel. As a consequence, the diesel engine can deliver more power. The compression, subsequently, increases the temperature. This increase in temperature can be reduced by placing an air cooler between the turbocharger and the inlet manifold. The performance of the air cooler decreases over time as the air and waterside can be polluted. Also, the engineer states that the future state of the air cooler only depends on its present condition suggesting that the underlying degradation process follows a Markov process. This will lead to separable states for the condition of the air cooler. During common usage, it is not possible to accurately observe the exact condition of the air cooler because you cannot see the heat exchanging surface. The general impression only shows the efficiency based on the difference in temperature before and after the air cooler. This temperature is used to form a belief about the state of the air cooler. A lower temperature difference between in and outflow points to a decrease in cooling performance (i.e. lower efficiency). The belief state formed is related to the internal states of the air cooler. On the belief states the decision-maker will decide which action to choose and how to optimize the maintenance planning. However, it is possible to open the cooler and have a visual look by making use of an endoscope. This form of inspection will reveal the true state of the air cooler on which the engineer can confirm the state of his belief and then decide what action is optimal; do nothing or perform preventive maintenance. Performing an inspection is more cost-effective compared to directly performing preventive maintenance making the POMDP beneficial.

Turbo

The second component satisfying all guidelines is the turbo. The turbo compresses the inlet air of the diesel engine. It makes the compressed air available under high pressure. As with the air cooler, the engineer states that the efficiency of the turbo decreases over time, and that its future state only depends on the present condition of the turbo. Therefore, the second guideline is satisfied leading to ability to define states of which the intermediate transitions follow a Markov process. It is impossible to accurately observe the exact condition of the turbo as it must be completely dismantled. However, a general impression is present by measuring the pressure. A decrease in pressure will tell the engineer that the turbo is becoming less efficient and allows to form a belief of the current state. These belief states are related to the internal states because, for example, when the pressure becomes too low, the engineer decides that the condition is so poor that preventive maintenance should be executed. However, the pressure does not reveal the exact state of the turbo. There are multiple failure mechanisms possible such as pollution or wear and tear. Both failure mechanisms ask for a different way of performing preventive maintenance. Therefore, the engineer can choose to perform an inspection instead. The inspection possible is a visual for contamination and vibration measurement that show the unbalance and bearing defects. In terms of cost-effectiveness, the inspection will be beneficial with regard to directly performing preventive maintenance.

Main Bearing

The last and third component satisfying all guidelines is the main bearing. The main bearing supports the rotation of the main crankshaft of the diesel engine. It reduces friction before turning. The main bearings are cooled and lubricated by the engine oil. The condition of the main bearing will decrease over time as wear occurs and small particles of metal end up in the oil. This degradation of the condition can be modelled as separate states descending from healthy to failure. Consequently, it is assumed that the degradation process follows a Markov process as the expert states that the future condition does not depend on the past activities. Again, it is not possible to accurately observe the exact condition of the main bearing because the component is assembled within the diesel engine. Sensors installed measure the temperature of the main bearing. This temperature gives a general indication of the condition but cannot tell if the main bearing is failing or not. Only a belief can be formed based on the temperature behavior as presented by the sensors. The engineer can decide based on this belief what the best moment is to perform preventive maintenance. Though, an inspection will lead to a more optimal decision as it will reveal the exact state of the main bearing. The exact state of the oil can be observed with an oil analyses where the bearing can be visually inspected to determine its exact condition. In both cases, it is more expensive to perform directly preventive maintenance instead of inspecting first. These findings show that POMDP will be beneficial for the main bearing as well.

6 Conclusion and Future Research

In this chapter we conclude the research with a main conclusion that answers the main question formulated within this research. Next, Section 6.2 gives some recommendations and directions for future research.

6.1 Conclusion

Within this research, it was studied if a partially observable Markov decision process (POMDP) would be a beneficial method for some components to apply condition-based maintenance at the Royal Netherlands Navy (RNLN). After a literature review on previous applications of POMDPs in the context of maintenance, eight guidelines were defined. These guidelines have to be satisfied to model a system or component within the RNLN as POMDP. It was found that POMDP could be recommended for maintenance if one is not capable of having full information about the true state of the system but still needs to decide which of the following actions would be best to choose. POMDPs account for both state transition and observation uncertainty.

The guidelines were verified by experts and tested via a case study on the diesel engine that was used on the vessels of the RNLN. The case study showed that there were three components for which POMDP would be a recommended method to apply condition-based maintenance (CBM) at the RNLN, namely the air cooler, turbo and main bearing. It was found that for these components a degradation pattern is present and thus could be partially observed. An overhead measurement could give a general impression on the condition of the components leading to a belief on the state. With this belief, the engineer of the RNLN can decide what the optimal maintenance planning is. In addition, a more in-depth inspection can be performed to reveal the complete condition making the POMDP not only suitable but also beneficial. Thus, for the air cooler and main bearing, this inspection will show the progress of the degradation over time. A more in-depth inspection for turbo, on the other hand, will give an indication of which failure mode is occurring.

During the process of asking questions to the engineer, components get eliminated by one of the specific guidelines. The starting point of the case study included only the components that satisfy GL1; it has to be useful to perform CMB on the asset. This guideline was found to be true for only eleven of the 95 components. This large number of components being eliminated shows that the usefulness of CBM is an important factor for the application of POMDP within the RNLN. Additionally, GL7; a more in-depth inspection has to reveal the complete condition of the asset was not formally needed to model the situation as a POMDP. The same goes for GL8; it has to be beneficial (i.e. in terms of cost) to perform a more in-depth inspection compared to preventive maintenance. Both only ensure that a POMDP is beneficial for practitioners as the RNLN. It was found that the eighth guideline eliminates some specific components, however, it could be argued that a POMDP can be applied for them.

In conclusion, knowing that there are components found for which all guidelines satisfy, it is concluded that a POMDP would be a suitable and recommended method to apply CBM at the RNLN. With the guidelines, it is possible to find supplemental components of other systems at the RNLN for which POMDP could be useful and beneficial. Subsequently, one could model the specific components as a POMDP, and optimize its maintenance planning accordingly.

6.2 Future Research

The next step in the current research would be the actual implementation of POMDPs to be able to optimize the maintenance planning of the RNLN accordingly. However, to model the situation of the components found as a POMDP the underlying degradation model should be deducted from the data. This deduction includes a complex statistical analysis that the RNLN should consider closely. The major question is how we find the underlying failure behavior of each of the suitable components on which CBM can be applied. It would be interesting to investigate these underlying failure behaviors, and how they are related to general observations that can be made.

Currently, a case study was performed on the main engine of the vessel that consists of 95 components. From these components, only eleven functioned as the starting point for applying the guidelines because those satisfy the first guideline (i.e. CBM should be useful). Though, as stated in Section 1.5, each ship of the RNLN consists of multiple thousands of components meaning that only a small fraction of the components present by the RNLN was evaluated. Therefore, the RNLN can consider applying the guidelines to other systems to find more suitable components for POMDP.

Besides, it was found that many components are eliminated because CBM seems not useful to apply. The RNLN might take a closer look at why so many components are not useful for CBM. Currently, the guidelines are specifically defined for application at the RNLN. Thus, it would be interesting to see how they perform when applied to other practitioners. It might occur that different guidelines will lead to the majority of the elimination of components for these practitioners. Consequently, the guidelines can be formulated in general instead of focusing on a specific context.

References

Arts, J. (2017). *Maintenance Modeling and Optimization* (Tech. Rep.). Eindhoven University of Technology. doi: 10.1007/978-1-4615-4329-9

Cassandra, A. R. (1997). A Survey of POMDP Applications. *Uncertainty in Artificial Intelligence*, 472–480.

Cassandra, A. R. (2020). *POMDP FAQ.* Retrieved from http://pomdp.org/faq.html

Celen, M., & Djurdjanovic, D. (2020). Integrated maintenance and operations decision making with imperfect degradation state observations. *Journal of Manufacturing Systems*, *55*(May), 302–316. doi: 10.1016/j.jmsy.2020.03.010

Derman, C. (1963, 4). Optimal Replacement and Maintenance Under Markovian Deterioration with Probability Bounds on Failure. *Management Science*, *9*(3), 478–481. doi: 10.1287/mnsc.9.3.478

Giorgio, M., Guida, M., & Pulcini, G. (2011, 9). An age- and state-dependent Markov model for degradation processes. *IIE Transactions (Institute of Industrial Engineers)*, *43*(9), 621–632. doi: 10.1080/0740817X.2010.532855

Google Image Result. (2020). Retrieved from https://www.google.com/imgres?imgurl=https://www.pngkey.com/png/full/892-8924657_our-maritime-business-magazines-caterpillar-c280-12.png&imgrefurl=https://xixingseng.cn/aeecu/engine-derate-caterpillar.html&tbnid=FPlPTtJ-uTgx8M&vet=1&docid=hIv-TchPrhwdlM&w

Ibe, O. C. (2013). *Markov Processes for Stochastic Modeling* (Second ed.). Elsevier Inc. doi: 10.1016/C2012-0-06106-6

Karabağ, O., Eruguz, A. S., & Basten, R. (2020). Integrated optimization of maintenance interventions and spare part selection for a partially observable multi-component system. *Reliability Engineering and System Safety*, *200*(March), 106955. doi: 10.1016/j.ress.2020.106955

Karremann, J. (2020). *Marineschepen van de Koninklijke Marine.* Retrieved from https://marineschepen.nl/schepen/marineschepen-nederland.html

Kim, M. J., & Makis, V. (2013). Joint optimization of sampling and control of partially observable failing systems. *Operations Research*, *61*(3), 777–790. doi: 10.1287/opre.2013.1171

Krishnamurthy, V. (2016). *Partially Observed Markov Decision Processes: Structural Results.* doi: 10.1017/cbo9781316471104.012

Labib, A. W. (2004). A decision analysis model for maintenance policy selection using a CMMS. *Journal of Quality in Maintenance Engineering*, *10*(3), 191–202. doi: 10.1108/13552510410553244

LaRiviere, J., McAfee, P., Rao, J., Narayanan, V. K., & Sun, W. (2016). *Where Predictive Analytics Is Having the Biggest Impact.* Retrieved from https://hbr.org/2016/05/where-predictive-analytics-is-having-the-biggest-impact

Lee, J., Liao, L., Lapira, E., Ni, J., & Li, L. (2009). Informatics platform for designing and deploying e-manufacturing systems. In *Collaborative design and planning for digital manufacturing* (pp. 1–35). Springer London. doi: 10.1007/978-1-84882-287-0{_}1

Lin, X., Basten, R. J., Kranenburg, A. A., & van Houtum, G. J. (2017). Condition based spare parts supply. *Reliability Engineering and System Safety*, *168*(May), 240–248. doi: 10.1016/j.ress.2017.05.035

Littman, M. (2001, 1). Markov Decision Processes. In *International encyclopedia of the social & behavioral sciences* (pp. 9240–9242). Elsevier. doi: 10.1016/b0-08-043076-7/00614-8

Maltby, H., Pakornrat, W., Jackson, J., Hernandez, A., Williams, C., Lin, C., & Khim, J. (2020). *Markov Chains — Brilliant Math & Science Wiki.* Retrieved from `https://brilliant.org/wiki/markov-chains/`

Margés, J. (2019). Voorkomen efficiënter dan genezen. *Alle Hens 07, Defensie Magazine*. Retrieved from `https://magazines.defensie.nl/allehens/2019/07/06_data_gestuurd-onderhoud`

Ministerie van Defensie. (2020). *Ministerie van Defensie — Defensie.nl.* Retrieved from `https://www.defensie.nl/`

Mobley, R. K. (2002). *An Introduction to Predictive Maintenance* (Second ed.) (No. 3). Butterworth-Heinemann. doi: 10.1016/b978-0-7506-7531-4.x5000-3

Morato, P. G., Nielsen, J. S., Mai, A. Q., & Rigo, P. (2019). POMDP based maintenance optimization of offshore wind substructures including monitoring. *13th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP 2019*(May).

Ohnishi, M., Mine, H., & Kawai, H. (1986). An Optimal Inspection and Replacement Policy under Incomplete State Information. *European Journal of Operational Research*, *R-24*(27), 117–128. doi: 10.1109/TR.1975.5214918

Papakonstantinou, K. G., & Shinozuka, M. (2014). Optimum inspection and maintenance policies for corroded structures using partially observable Markov decision processes and stochastic, physically based models. *Probabilistic Engineering Mechanics*, *37*, 93–108. doi: 10.1016/j.probengmech.2014.06.002

Powell, W. B. (2014, 9). Clearing the Jungle of Stochastic Optimization. In *Bridging data and decisions* (pp. 109–137). INFORMS. Retrieved from `https://pubsonline.informs.org/doi/abs/10.1287/educ.2014.0128` doi: 10.1287/educ.2014.0128

Tiddens, W. (2014). *FMECA worksheet Pon CAT.* Maintenance and Service Logistics for Maritime Assets.

Tiddens, W. (2018). *Setting sail towards predictive maintenance predictive maintenance* (Doctoral dissertation, University of Twente). Retrieved from `https://ris.utwente.nl/ws/portalfiles/portal/51158164/Tiddens_WW_PhD_Thesis_2018.pdf`

Tijms, H. C. (1986). *Stochastic Modelling and Analysis: A Computational Approach.* John Wiley & Sons, Inc. Retrieved from `https://dl.acm.org/doi/abs/10.5555/7374`

Tinga, T., Tiddens, W. W., Amoiralis, F., & Politis, M. (2017, 6). Predictive maintenance of maritime systems: Models and challenges. In *Safety and reliability - theory and applications - proceedings of the 27th european safety and reliability conference, esrel 2017* (pp. 421–429). CRC Press/Balkema. doi: 10.1201/9781315210469-56

Ton, B., Basten, R., Bolte, J., Braaksma, J., Di Bucchianico, A., Van De Calseyde, P., … Stoelinga, M. (2020, 11). PrimaVera: Synergising Predictive Maintenance. *Applied Sciences*, *10*(23), 8348. doi: 10.3390/app10238348

Van Oosterom, C., Maillart, L. M., & Kharoufeh, J. P. (2017). Optimal maintenance policies for a safety-critical system and its deteriorating sensor. *Naval Research Logistics*, *64*(5), 399–417. doi: 10.1002/nav.21763

Van Staden, H. E., & Boute, R. N. (2020). The effect of multi-sensor data on condition-based maintenance policies. *European Journal of Operational Research*. doi: 10.1016/j.ejor.2020.08.035