POMDP based Maintenance Optimization of Offshore Wind Substructures including Monitoring

Pablo G. Morato
PhD Candidate, ANAST, Naval & Offshore Engineering, ArGEnCo, University of Liege, Liege, Belgium

Jannie S. Nielsen
Assistant Professor, Dept. of Civil Engineering, Aalborg University, Aalborg, Denmark

Anh Q. Mai
PhD Candidate, ANAST, Naval & Offshore Engineering, ArGEnCo, University of Liege, Liege, Belgium

Philippe Rigo
Professor, ANAST, Naval & Offshore Engineering, ArGEnCo, University of Liege, Liege, Belgium

ABSTRACT: Sequential decision making under uncertainty is a complex task limited normally by computational requirements. A novel methodology is proposed in this paper to identify the optimal maintenance strategy of a structural component by using a point-based Partially Observable Markov Decision Process (POMDP). The framework integrates a dynamic bayesian network to track the deterioration over time with a POMDP model for the generation of a dynamic policy. The methodology is applied to an example quantifying whether a monitoring scheme is cost effective. This complex decision problem comprised of 200 damage states is solved accurately within 60 seconds of computational time.

1. INTRODUCTION

Offshore wind energy is a sustainable solution for energy generation. Further from shore, higher and steadier wind speeds can be harnessed and the visual impact is reduced as compared with onshore wind. However, offshore wind substructures are subjected to a harsh deterioration due to the combined action of fatigue and corrosion.

Besides, maintenance operations are complex and expensive. It is therefore of utmost importance to provide decision support to a decision maker (operator) who is taking the decisions under uncertainty. The maintenance strategy can be optimized by following a risk-based approach where an optimal balance is achieved between the maintenance efforts and the large consequences associated with a structural failure.

In addition to inspections, Structural Health Monitoring (SHM) can be employed to gather more information about the state of the structures. SHM techniques have improved considerable as more accurate and reliable sensors are available. Nevertheless, there is a cost associated with a SHM scheme due to the installation and operation of the system and a risk of increased costs, if too many inspections are initiated on the basis of false alarms. Then, the decision maker must face the decision whether to utilize and install a SHM scheme or not. This decision can be optimally chosen by quantifying the value of the information.

The concept of the Value of Information (VoI) was introduced by Raiffa and Schlaifer (1961), providing a theoretical framework to quantify the value of information within the Bayesian decision analysis. Based on this framework, a great number of research efforts have been devoted recently to quantify the value of monitoring for civil infrastructures, such as bridges or hydraulic structures. The reader is directed to Memarzadeh and Pozzi (2016) for a more exhaustive illustration on the VoI framework for sequential decision problems.

The main limitation of these methodologies strives on the assumptions and simplifications imposed due to the computational requirements involved in the solution of complex decision problems. For instance, the applications consider small state spaces or stationary decision rules such as "preset interference threshold" are imposed.

This work presents a methodology to quantify the value of monitoring by employing a point-based "Partially Observable Markov Decision Process" (POMDP). Since a POMDP point-based solver samples only a subset of the belief space, this methodology can be employed to generate dynamic maintenance policies, even when complex sequential decision problems are involved (Morato et al., 2018; Papakonstantinou and Shinozuka, 2014).

2. POMDP BASED METHODOLOGY

A novel methodology is presented hereby to quantify the value of monitoring. The expected maintenance costs are estimated separately for the case when only inspections are included, and for the case when a monitoring system is also included. Thereafter, the Value of Information (VoI) can be computed as the difference.

Concerning the inspection planning, the influence diagram (Fig.1) displays how this sequential decision problem is approached. The damage evolving over time is represented by the chance node D_t and it is possible to choose an inspection method (including no-inspection) by means of the decision node I_t.

The chance node Z_t indicates the quality of the inspection method. Additionally, the node E_t tracks the probability of being in the last damage state, or in other words, the failure probability. The utility nodes C_{F_t} and C_{I_t} assign a cost of failure and a cost of inspection, respectively.

Ultimately, the chance node R_t represents the decision of whether to perform a repair or not. If it is decided to make a repair, then the damage state will be transferred to a healthier state and it will have an associated cost of repair C_{R_t}.

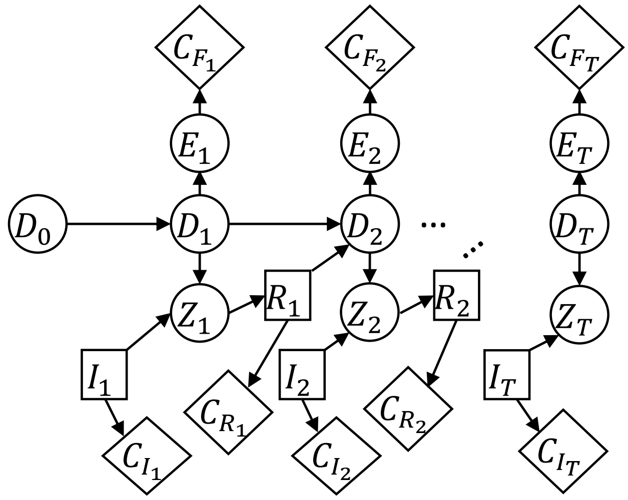

Figure 1: Influence diagram corresponding to the inspection planning decision problem

2.1. Building the model

A POMDP model is built in order to solve the maintenance decision problem. The outcome of the POMDP model is a policy which informs the optimal decision depending on the current belief state (probability distribution for the node D). This decision is comprised within the context of this framework by a combination of an action and an observation. Examples of actions are "do-nothing" or "repair" and it is possible to gather observations by "inspection", "monitoring", including also the case of "no-inspection". A decision could be for instance "do-nothing / inspection" or "do-nothing / no-inspection".

The input of a POMDP simulation includes therefore: (1) the transition probability $[T]$ from one damage state to another depending on the action chosen, (2) the observation likelihood [O] depending on the inspection type selected, (3) the rewards associated with the taken "action/observation" decision (including a discount factor), and (4) the initial damage state D_0.

2.1.1. Transition probabilities

The transition matrices are defined according to the associated action. The transition probabilities for the action "do-nothing" can be obtained as the conditional probabilities corresponding to the "Dynamic Bayesian Network" (DBN) shown in Fig. 1, where only the damage nodes are kept. Thus, T_{DN} is equivalent to the conditional probabilities of damage at time step "t" given the damage at the previous time step "$t-1$" (Eq. 1). If no time-invariant uncertainties are involved, the conditional probabilities can be obtained by Monte Carlo simulations.

$$T_{DN} = P(D_t|D_{t-1}) \tag{1}$$

For the case of the repair action, the transition matrix simply transfers the component to a healthier state, depending on the repair quality.

2.1.2. *Observation probabilities*

If an inspection or monitoring is performed, then an observation is gathered. The observation matrix conveys the quality of this information (likelihood). The observation matrix is constructed depending on the selected observation: 1)"No-inspection": No information is gained, thus, the belief state must remain unaltered; 2)"Inspection": The observation matrix is directly computed from a "Probability of Detection" curve; 3)"Monitoring": The observation matrix is obtained in a similar manner as for the inspection case.

2.1.3. *Decisions and associated rewards*

The rewards depend on the decision taken and this will be greatly influenced by the nature of the problem. For instance, the following approach can be taken: 1)"Do-nothing / No-inspection" ($DN-\bar{I}$): Only the failure cost is considered; 2)"Do-nothing / Inspection" ($DN-I$): The inspection cost is included along with the failure cost; 3)"Repair / No-inspection" ($R-\bar{I}$): Here the repair cost is considered; 4)"Do-nothing / Monitoring" ($DN-M$): As the value of information will be calculated a posteriori, monitoring costs are not included in the POMDP model.

If desired, more decisions can be added into the model, yet with an additional computational cost. Finally, a discount factor must be included $\gamma \in (0-1)$ to quantify the present value of money over time. This discount factor becomes necessary if an infinite horizon POMDP is employed.

2.2. *POMDP Simulation*

Once the POMDP input has been prepared by including: transitions, observations, rewards and the initial state; then, a point-based solver is employed to generate a "POMDP" maintenance policy. This framework allows the computation of large state POMDP spaces due to the fact that point-based solvers are able to efficiently compute large belief states within a reasonable computational time (Morato et al., 2018; Papakonstantinou and Shinozuka, 2014). In the application presented in this paper (Section 3), the solver "SARSOP" (Kurniawati et al., 2008) is selected; nevertheless, other POMDP solvers such as "PERSEUS" or "HSVI" can be used instead.

The approach followed by this methodology leads to the creation of an infinite horizon POMDP, where the obtained policy is applicable for any time step. If a finite horizon POMDP is preferred; then, time must be encoded within the transition matrices at the cost of significantly increasing the belief space and computational time (Morato et al., 2018; Papakonstantinou and Shinozuka, 2014).

2.3. *Post-processing*

After the simulation is conducted, a policy is obtained as a result of the POMDP model. Additionally, the POMDP solver provides for each computational time step: (1) expected costs (delimited by upper and lower bounds), (2) number of beliefs and α-vectors and (3) number of backups. The expected costs provides the main outcome for the decision problem, whereas the other parameters can be checked to understand more details about the generated policy and the complexity of the problem.

Furthermore, the obtained policy is comprised of a set of α-vectors (Γ), each of them associated to a decision. The optimal decision is the one which corresponds to the α-vector that maximizes the value function $V(b)$ as shown in (Eq. 2). Hence, the decision is chosen only based on the current belief state (b).

$$V(b) = \max_{\alpha \in \Gamma}(\alpha \cdot b) \tag{2}$$

Additionally, the influence diagram displayed in Fig. 1 can be used in combination with the generated policy to choose the optimal decision for a particular scenario. The inspection decision node I_t is then instantiated with the optimal decision by applying Eq. 2.

2.4. *Quantifying the value of monitoring*

If a monitoring system is installed, the uncertainties are reduced because an observation is continuously gathered (every time step). This will have an effect on the maintenance strategy as normally less inspections might be necessary. Therefore, the benefit of installing a maintenance scheme is quantified as the difference between: (1) the achieved reduction of expected costs (as additional information is provided by monitoring), and (2) the cost of the monitoring system. In other words, the objective is the quantification of the Value of Information (VoI) or in this case the value of monitoring.

The VoI is calculated as the difference between the expected costs if monitoring is not conducted ($\mathbb{E}(C_0)$) and the expected cost if monitoring is conducted ($\mathbb{E}(C_1)$). However, the cost of the monitoring system C_M is neglected for this calculation:

$$VoI = \mathbb{E}(C_0) - \mathbb{E}(C_1) \qquad (3)$$

Additionally, it is useful to introduce the concept of Net Value of Information (NVoI) which also includes the cost of monitoring:

$$NVoI = VoI - C_M = \mathbb{E}(C_0) - \mathbb{E}(C_1) - C_M \qquad (4)$$

The NVoI is very helpful for the decision maker because it is used to decide whether monitoring should be performed or not: if the NVoI is positive, then monitoring provides an added value; if the NVoI is negative, then the monitoring system is more expensive than the benefit gained by its installation.

3. APPLICATION

The value of monitoring is now quantified for a maintenance decision problem partially based on the Example presented by Nielsen and Sorensen (2015). However, here the decision maker must decide whether to install a monitoring system or not.

3.1. *Model*

The fatigue deterioration of a structural component is here modelled by a probabilistic fracture mechanic model based on the Paris' law (Eq. 5). Both the initial crack size and stress range are considered as random variables. The damage size or crack size is computed for each time step with the expression developed by Ditlevsen and Madsen (1996):

$$a_t = \left[\left(1 - \frac{m}{2}\right) C \Delta S^m \pi^{m/2} \Delta n + a_{t-1}^{1-m/2}\right]^{(1-m/2)^{-1}} \qquad (5)$$

Where a_t is the damage (crack size) at the time step "t", a_{t-1} is the damage at the previous time step "$t-1$", C and m are material parameters which condition the crack propagation, ΔS is the stress range and Δn is the number of cycles per time step.

Thus, given an initial crack size ($a_{t=0}$), the crack size distribution can be computed for the following time steps. In this example, the time step is considered to be one month. The values of the parameters are listed in Table 1. The component fails once the crack has reached the critical crack size a_c, which is considered here to be 9 mm.

Table 1: Parameters for the fracture mechanics model

Parameter	Distribution	Mean	StDev
a_0	Exponential	0.2	-
a_c	Deterministic	9	-
$ln(C)$	Deterministic	-33.5	-
m	Deterministic	3.5	-
ΔS	Normal	60	10
Δn	Deterministic	10^6	-

The limit state is then formulated in Eq. (6) where the failure probability is computed as the probability of the limit state being negative.

$$g_{FM}(t) = a_c - a(t) \qquad (6)$$

In principle, the failure probability can be estimated by using a crude Monte Carlo simulation. Nevertheless, a "Dynamic Bayesian Network" (DBN) is here proposed because it will be the basis to both define the transition probabilities for the actions where no maintenance actions are involved and for evaluating the obtained policy.

3.1.1. Building the DBN model

A discretization scheme is used to convert the deterioration model from a continuous space to a discrete space so as to facilitate the inference of the DBN (Eq. 7). The DBN is derived from the influence diagram presented in Fig. 1 to track the damage. The crack size is represented by the nodes a_t. If an inspection is performed, a node Z_t is incorporated into the network as shown in in Fig. 2. Additionally, the node E_t collects the failure probability which it is equivalent to the probability of being in the last damage state.

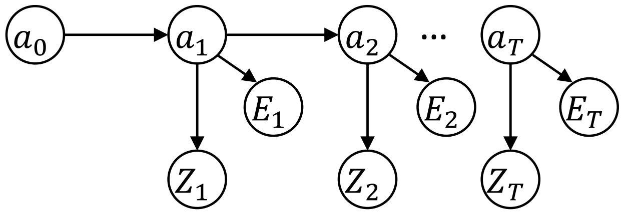

Figure 2: DBN Model

Due to the fact that point-based POMDP solvers are able to solve large belief spaces in a reasonable computational time, the number of states for the damage size is not limited here. It is chosen in such a way that the computed failure probability by the DBN model is similar as the result obtained by a crude "Monte Carlo Simulation" (MCS).

$$a \in (0, exp[ln(10^{-5}) : \frac{ln(9) - ln(10^{-5})}{states - 2} : ln(9)], \infty) \quad (7)$$

As it can be seen in Fig. 3, a discretization with 200 states provides enough accuracy for the DBN as the failure probability is in good agreement with the result from the crude MCS.

3.1.2. POMDP model including inspections

A POMDP model is built by defining the transition probabilities, observation or emission probabilities, rewards and the initial state. For this case, three possible decisions are included:

- Do-nothing / No-inspection $(DN - \bar{I})$: comprised of the transition matrix "do-nothing" T_{DN} and the observation matrix "no-inspection" $O_{\bar{I}}$.

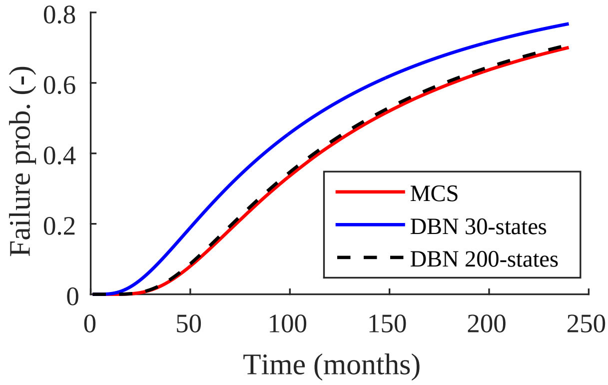

Figure 3: DBN Model - discretization accuracy

- Do-nothing / Inspection $(DN - I)$: composed of the transition matrix "do-nothing" T_{DN} and the observation matrix "Inspection" O_I.
- Repair / No-inspection $(R - \bar{I})$: comprised of the transition matrix "repair" T_R and the observation matrix "No-inspection" $O_{\bar{I}}$.

The "do-nothing" transition matrix T_{DN} is easily defined by utilizing the conditional probabilities used for the development of the DBN, as shown in Eq. 8. For the transition corresponding to the repair action T_{RP}, the damage is transferred to a healthier state (initial damage size), independently of the current damage state.

$$T_{DN} = \begin{bmatrix} p(a^1_{t+1}|a^1_t) & 0 & \dots & 0 \\ p(a^2_{t+1}|a^1_t) & p(a^2_{t+1}|a^2_t) & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ p(a^n_{t+1}|a^1_t) & p(a^n_{t+1}|a^2_t) & \dots & p(a^n_{t+1}|a^n_t) \end{bmatrix} \quad (8)$$

Inspection outcomes are normally defined by "Probability of Detection" curves (PoDs) within the context of traditional risk-based inspection methods. PoDs determine the measurement uncertainty by assigning the probability of detection given the damage size. This is can be translated to the DBN because PoDs are equivalent to the conditional probabilities corresponding to the inspection node Z_t given the damage a_t. Furthermore, in this example, an inspection can lead to six different outcomes, each of them depending on the damage size. Table 2 states the probability of obtaining each outcome as lognormal distributions, defined similarly as (Nielsen and Sorensen, 2015).

Table 2: Inspection: measurement uncertainty

State	Description	Mean	COV
1	no detection	-	-
2	mild damage	2.0	1.0
3	some damage	4.0	0.8
4	significant damage	6.0	0.6
5	severe damage	8.0	0.4
6	failure	9.0	0.0

The measurement uncertainties are therefore employed to define the observation matrix for the case where an inspection is performed O_I:

$$O_I = \begin{bmatrix} p(ins_1|a^1) & p(ins_2|a^1) & \dots & p(ins_m|a^1) \\ p(ins_1|a^2) & p(ins_2|a^2) & \dots & p(ins_m|a^2) \\ \vdots & \vdots & \vdots & \vdots \\ p(ins_1|a^n) & p(ins_2|a^n) & \dots & p(ins_m|a^n) \end{bmatrix} \tag{9}$$

If the component is not observed (inspection is not planned $O_{\bar{I}}$); then, the observation matrix is defined as shown in Eq. 10. By using this observation matrix, the belief state prevails invariable. Since the belief state remains unaltered, it is equivalent to the case where no information is obtained.

$$O_{\bar{I}} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & 0 \\ 1 & 0 & \dots & 0 \end{bmatrix} \tag{10}$$

Once, the transition and observation probabilities are stated, the next step is to define the rewards. In this example, an inspection is associated with a cost of 1 money unit, a repair costs 50 money units and if the failed state is reached, a penalization of 500 money units must be paid as listed in Table 3. A exact definition of the costs is not crucial for a risk-based analysis whereas the relative difference between the cost associated to each decision (action/observation) is very important because it conveys the preference of the decision maker. Finally, the discount factor is defined as $\gamma = 0.95$.

3.1.3. *POMDP model including inspections and monitoring*

A POMDP is now built for the case when both inspections and monitoring are incorporated. The

Table 3: Rewards

	State 1	...	Failed state
Do-nothing	0	0	500
Repair	50	50	50
Inspection	1	1	1

transition probabilities, rewards and initial state are defined in the same manner as for the case when only inspections were included. However, the decision "Do-nothing / No-inspection" ($DN-\bar{I}$) is here replaced by the decision "Do-nothing / Monitoring" ($DN-M$) as the structure is monitored continuously. Thus, three decisions are now possible: (1) "Do-nothing / Monitoring" ($DN-M$), (2) "Do-nothing / Inspection" ($DN-I$) and (3) "Repair / No-inspection" ($R-\bar{I}$).

Table 4: Monitoring: measurement uncertainty

State	Description	Mean	COV
1	no alarm	-	-
2	low alarm	2.0	1.0
3	high alarm	5.0	1.0
4	failure	9.0	0.0

The observation probabilities for the case when monitoring is performed (O_M) correspond to the conditional probabilities of obtaining each monitoring outcome given the damage size ($mon|a$). Hence, the observation matrix is defined according to Eq. 11. The probability of obtaining each outcome is modelled by a lognormal distribution and it is presented in Table 4.

$$O_M = \begin{bmatrix} p(mon_1|a^1) & p(mon_2|a^1) & \dots & p(mon_m|a^1) \\ p(mon_1|a^2) & p(mon_2|a^2) & \dots & p(mon_m|a^2) \\ \vdots & \vdots & \vdots & \vdots \\ p(mon_1|a^n) & p(mon_2|a^2) & \dots & p(mon_m|a^n) \end{bmatrix} \tag{11}$$

3.2. *Results*

Both POMDP models (only inspections / monitoring and inspections) are simulated with the point-based solver "SARSOP". Firstly, the expected costs resulting from each POMDP model are presented; secondly, the value of information is computed, and finally, the application of the POMDP policy for a particular case is conducted.

3.2.1. Expected costs and policies

The total expected costs are presented in Fig. 4. As expected, the total costs are higher for the case where only inspections are included. It is interesting to notice that the POMDP solution provides an upper and lower boundary for the total expected costs.

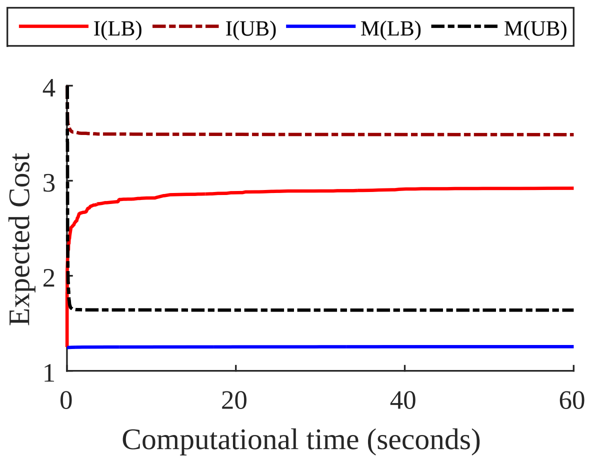

Figure 4: POMDP results: Expected costs

Due to the nature of the algorithm, new belief states are sampled and evaluated over time as shown if Fig. 5. The precision can be improved if the simulation is run for a longer time (new beliefs and α vectors will be generated); however, the accuracy is considered acceptable for this example within 60 seconds of CPU time, with an Intel Core I9 7900X @3.0 GHz and RAM 64GB.

3.2.2. Quantification of the value of information

It is possible at this point to provide decision support under uncertainty by quantifying the value of monitoring. Fig. 6 represents the upper and lower boundaries of the Net Value of Information (NVoI) for a given monitoring cost. The monitoring system is considered as economically feasible if the NVoI is positive and it is infeasible if the NVoI is negative.

The result suggests that it is cost-effective to install the monitoring system if its cost is lower than approximately 1.3 money units. Since the expected

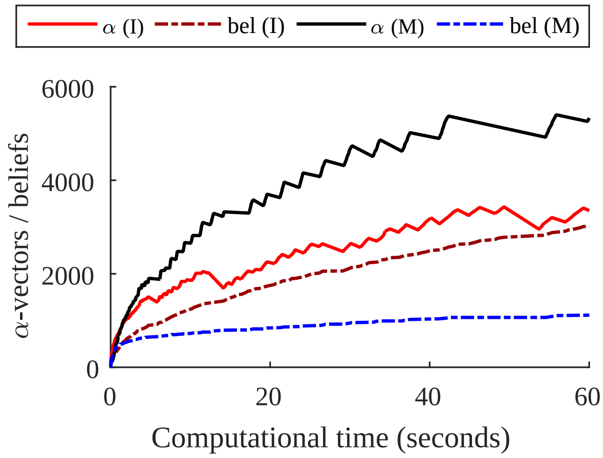

Figure 5: POMDP results: α vectors and beliefs

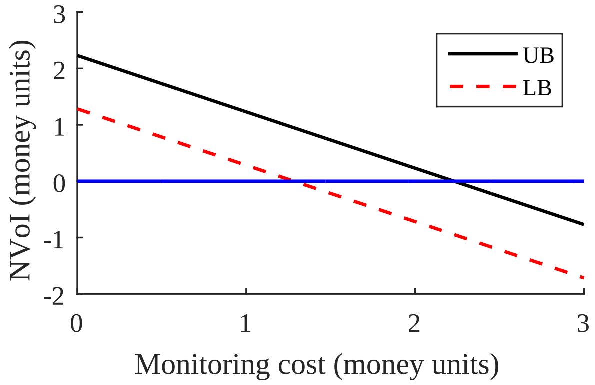

Figure 6: Quantification of the value of information

costs are delimited by upper and lower bounds, the NVoI is also delimited by bounds. A better accuracy can be achieved by increasing the simulation time, however, the precision is considered acceptable for this example.

3.2.3. Application of the POMDP policy for a particular case

The generated policy by the POMDP where only inspections are included is utilized now to select the optimal decisions for three particular cases. The result is depicted in Fig. 7. For the case 1, it is assumed that the inspection outcome is always "No-detected"; for the case 2, the inspection outcome is "No-detected" up to the year 14, after, the outcome is "Mild damage"; and for the case 3, the inspection outcome is "Severe damage" after year 14.

Although the decision "do-nothing" dominates for the case 1, "inspections" are also scheduled. Re-

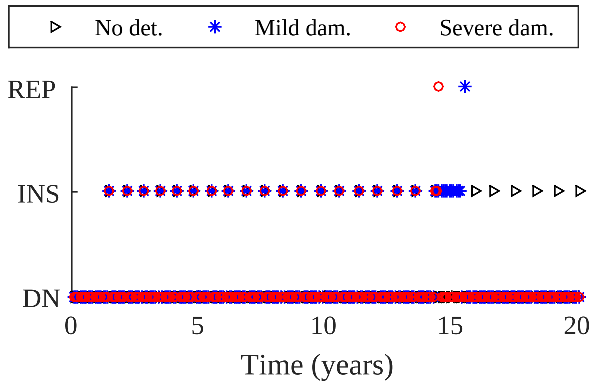

Figure 7: Policy evaluation for a particular case

garding the case 2, a repair is performed between the years 14 and 15. This repair is undertaken after successive inspections where the assumed outcome is "mild damage". However, the repair is selected only after one inspection for the case 3.

This result is reasonable because a repair is planned if a severe damage is found, whereas several inspections are necessary before the repair, if the outcome is mild damage. It is demonstrated by this example how the generated policy can be used in a dynamic fashion, providing support for complex sequential decision making, where different inspection outcomes can be expected.

4. CONCLUSIONS

Structural Health Monitoring (SHM) provides information about the state of the structural components leading to a reduction of uncertainty. As the uncertainties are reduced, better decisions can be taken. Within the context of maintenance planning, if information is gathered by monitoring, less inspections might be necessary, becoming especially important for the case of offshore wind structures, where inspections are complex and expensive.

However, there is a cost associated to the installation of a monitoring scheme and it must be decided whether it is cost-effective or not. A methodology is proposed in this paper to quantify the benefit or Value of Information (VoI) achieved by monitoring.

The methodology is applied to identify the optimal maintenance policy for a structural element subjected to fatigue deterioration. The policy is generated within a reasonable computational time for this complex case, where the damage size is discretized into 200 states and different outcomes of inspections and monitoring are possible.

In the future, efforts should be made to enable the use of time-invariant uncertainties within the deterioration modelling. Besides, the maintenance will be more optimal if it is performed at the system level incorporating the correlations or dependencies amongst the involved random variables. The development of hierarchical models is encouraged.

5. ACKNOWLEDGEMENTS

This research is funded by the National Fund for Scientific Research in Belgium F.R.I.A. - F.N.R.S.

6. REFERENCES

Ditlevsen, O. and Madsen, H. O. (1996). *Structural reliability methods*, Vol. 178. Wiley New York.

Kurniawati, H., Hsu, D., and Lee, W. S. (2008). “Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces..” *Robotics: Science and systems*, Vol. 2008, Zurich, Switzerland.

Memarzadeh, M. and Pozzi, M. (2016). “Value of information in sequential decision making: Component inspection, permanent monitoring and system-level scheduling.” *Reliability Engineering & System Safety*, 154, 137–151.

Morato, P. G., Nielsen, J. S., Mai, A. Q., and Rigo, P. (2018). “Point-based POMDP risk based inspection of offshore wind substructures.” *IALCCE 2018-The Sixth International Symposium on Life-Cycle Engineering*.

Nielsen, J. S. and Sorensen, J. D. (2015). “Risk-based decision making for deterioration processes using POMDP.” *ICASP12 International Conference on Applications of Statistics and Probability in Civil Engineering*, Civil Engineering Risk and Reliability Association.

Papakonstantinou, K. G. and Shinozuka, M. (2014). “Planning structural inspection and maintenance policies via dynamic programming and markov processes. part i: Theory.” *Reliability Engineering & System Safety*, 130, 202–213.

Raiffa, H. and Schlaifer, R. (1961). *Applied Statistical Decision Theory*. Cambridge University Press, Cambridge, Mass.