



HAL
open science

Optimal Resurfacing Decisions for Road Maintenance: A POMDP Perspective

Mariem Zouch, Thomas G. Yeung, Bruno Castanier

► **To cite this version:**

Mariem Zouch, Thomas G. Yeung, Bruno Castanier. Optimal Resurfacing Decisions for Road Maintenance: A POMDP Perspective. the Annual Reliability and Maintainability Symposium, Jan 2011, United States. pp.978-1-4244-8855-1. hal-00632066

HAL Id: hal-00632066

<https://hal.science/hal-00632066>

Submitted on 13 Oct 2011

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

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

Optimal Resurfacing Decisions for Road Maintenance: A POMDP Perspective

Mariem Zouch, Thomas Yeung and Bruno Castanier
 Department of Automatic control & Industrial Engineering
 Ecole des Mines de Nantes / IRCCyN, F-44307 NANTES Cedex, France

Keywords: Road deterioration, maintenance optimization, Partially Observed Markov Decision process, grid-based approximation

Notations

a	maintenance action	i_θ	discrete value of θ
A_h	set of actions for a section thickness h	i_ρ	discrete value of ρ
\mathbf{b}	belief state	m	last maintenance type
\mathbf{b}^k	belief state immediately after action a_k	N	number of states
\mathbf{b}'	updated belief state after observation	N_h	number of available actions
\mathbf{b}^θ	belief state on θ	N_o	number of possible observations
B	set of state beliefs	N_ρ	number of possible values of ρ
c_0	inspection cost	N_θ	number of possible values of θ
c_f	total fixed cost	o	observation
c_n	fixed cost of milling	O	set of observations
c_p	fixed cost of resurfacing	r_a	DGR reduction factor
c_v	total variable cost	s	section state
c_{un}	unit cost of milling	s_a	section state after action a
c_{up}	unit cost of resurfacing	S	set of states
$C_q(\cdot)$	quality cost	V	Total expected discounted cost-to-go
h	section thickness	$\Omega(o \mathbf{b})$	probability of o at the decision epoch end when \mathbf{b} at its beginning
h_a	section thickness after action a	ρ	longitudinal cracking percentage
h_{max}	maximum road thickness	θ	deterioration growth rate
h_n	milling thickness	$\tilde{\theta}$	average deterioration speed
h_p	resurfacing thickness	τ	inter-inspection period
H_n	set of possible milling thicknesses	$\phi(\cdot)$	maintenance effect function on θ
H_p	set of possible resurfacing thicknesses		

1 Summary & Conclusions

We develop an optimal maintenance policy for a road section to minimize the total maintenance cost over the infinite horizon when some deterioration and decision parameters are not observable. Both perfect and imperfect maintenance actions are possible through the application of various thicknesses of resurfacing layers. We use a two-phase deterioration process based on two parameters: the longitudinal cracking percentage and the deterioration growth rate. Our deterioration model is a state-based model based on the state-dependent Gamma process for the longitudinal cracking percentage and the Bilateral Gamma process for the deterioration growth rate. Moreover the maintenance decision is constrained by a maximum road thickness that makes the maintenance decisions more complex as it becomes how much surface layer to add as well as to remove. Because only one of the two deterioration parameters is observable, we formulate the problem as a partially observed Markov decision process and solve it using a grid-based value iteration algorithm.

Numerical examples have shown that our model provides a preventive maintenance policy that slows down the initiation as well as the propagation of longitudinal cracks and that may ameliorate the road state to a better than as-good-as-new one by altering its composition through additive resurfacing layers.

2 Introduction & Literature review

During the last decades, the interest in maintenance optimization has increased considerably and several approaches have been developed for different fields of applications. For a survey of maintenance optimization models and application, refer to [4, 20]. Civil and transportation infrastructure networks is one of the fields where maintenance optimization is very important due to the importance to these networks to social and economical development of societies and to the large expenditure and construction time needed to construct new facilities [2].

Maintenance optimization approaches aim to determine optimal strategies, i.e. what maintenance action to perform and when in order to minimize variable costs as well as to ensure the proper functioning of infrastructure.

From the literature, different approaches have been developed and applied to many national networks such as in Netherlands [9, 19] and in the USA [6, 18]. Many approaches have been formulated as Markov decision processes (MDPs). This approach has the advantage to be dynamic and easy to solve. Nevertheless, MDPs assume that the deterioration parameters can be perfectly observed and measured. This assumption may not be realistic in some cases: even if advanced inspection technologies minimize measure errors, some deterioration parameters, such as the instantaneous speed of deterioration cannot be observed, but just estimated using available information on the system state. In order to take into account the non-observability of decision parameters, MDPs have been extended to partially observed Markov decision processes (POMDPs) [14, 8]. Instead of observed states, POMDPs work with belief states that are represented by probability distributions over all the states.

Several imperfect information problems have been treated in literature. Madanat *et al.* [11] present a dynamic programming method to determine optimal maintenance and inspec-

tion policy, in the presence of inspection error. In [7], perfect maintenance of an aircraft engine component subject to cracking is considered. The problem is formulated as a POMDP model since cracks are partially observable and the maintenance scheduling decision is based on other collected information. A similar problem is presented in [12] where a POMDP over the finite horizon is formulated for a maintenance problem using the failure rate to update the knowledge on the system state. In [13] an optimal stopping problem with partial information is formulated for a system with obvious failure and transformed to a problem with complete information. Ghasemi *et al.* [5] formulated a POMDP for a system subject to perfect condition-based maintenance. Their POMDP formulation is combined to a proportional hazards model for the system degradation.

Moreover, POMDPs offer an elegant way to take into account the uncertainty about the system state. However, their most important challenge is solution computation. Exact solutions are almost impossible except for problems of small sizes. In the literature, several approaches are combined with value iteration or policy iteration [17] algorithms to approximate POMDPs. Examples of such approaches are grid-based approximations ([10] and [22]) and point-based approximations ([15], [3], [21] and [16]). In this work, we extend the model presented in [23] by considering a non-observable deterioration growth rate. We present an optimization framework for a road section maintenance strategy using a POMDP formulation and show the interest of such formulation.

In a road maintenance context, our model has the advantage to take into account an underlying deterioration parameter that cannot be observed: the deterioration speed. This deterioration speed information is a key decision parameter in condition-based maintenance optimization. Indeed, the knowledge of the deterioration growth rate allows updating the maintenance decision based on the single deterioration level. Moreover, unlike most of developed approaches for partially observed problems, our model considers multiple imperfect maintenance actions comprised of two maintenance components.

The remainder of this paper is structured as follows. The model is formulated in Section 2. In Section 3, we discuss a numerical example. Finally a conclusion is presented in Section 4.

3 Model formulation

3.1 Problem statement

Consider a road section that is continuously and stochastically deteriorating. A road is a complex system with several deterioration modes. However, we focus in this work on the longitudinal fatigue cracking process because of its frequency and economical consequences. This deterioration mode is a two phase process: a non-observable and an observable phase. The deterioration begins at the bottom of the road layer and propagates until reaching the surface giving way to longitudinal cracks that propagate on the surface. The current metric available in the French IQRN database [1] to describe the deterioration level of the road section is the longitudinal cracking percentage (LCP) denoted ρ and represented in Figure 1 where it is given by $\rho = \sum_{i=1}^3 L_i/200$.

The LCP metric only reflects the observable deterioration level, especially for small val-




Figure 1: Construction of the longitudinal cracking percentage

ues of ρ . This indicator is not sufficient to predict the future behavior of the road cracking growth. Additional information has to be integrated in order to estimate the underlying deterioration. We introduce the deterioration growth rate (DGR), denoted θ , which is defined as the expected instantaneous deterioration speed and can be greater than zero when no cracks are observed, i.e., $\rho = 0$, to indicate cracking potential.

We consider, as in current practices, a constant maintenance decision interval equal to τ . At the beginning of each decision epoch, the road section is inspected at a cost c_0 and an observation of ρ is yielded. We assume that this observation is perfect in that the measure error is equal to zero. Unlike the LCP, θ can neither be observed nor measured. Only an average value can be estimated using successive observations of ρ . The complete deterioration is then partially observed.

After inspection, the decision-maker can opt to do nothing (DN) or perform one of several maintenance actions that have different effects on the road system. A maintenance action consists in resurfacing the section to conceal cracks. However, the road is constrained by a maximum thickness h_{max} . Therefore, the maintenance decision for a road section of thickness h consists in how much thickness, h_n , to remove and how much thickness, h_p , to add so that the thickness constraint is respected. The future behavior of the cracking process is a complex function of the state of the road before maintenance, the new thickness of the road and the underlying composition of the road characterized by h_n and h_p .

We assume that only R thicknesses to add are available. Let $H_p = \{h_{p1}, \dots, h_{pR}\}$ be the set of available thicknesses to add. Therefore, given a current thickness h and a selected thickness to add h_p , the possible thicknesses to remove are given by $H_n(h, h_p) = [\min(0, h_p + h - h_{max}), h]$. We denote a maintenance action by $a = (h_p, h_n)$. Note that the DN action corresponds to $(h_p = 0, h_n = 0)$.

Each maintenance action incurs a fixed cost $c_f = c_n I_{\{h_n > 0\}} + c_p I_{\{h_p > 0\}}$ as well as a

variable cost $c_v(h_n, h_p) = c_{un} h_n + c_{up} h_p$ where I_X is the binary indicator function ($= 1$ if X and $= 0$ otherwise), c_n and c_p are the fixed cost of removing and adding asphalt layers, respectively, c_{un} and c_{up} are the unit costs of removing and adding, respectively. Moreover, a quality cost, denoted $C_q(\cdot)$ is incurred in order to prevent the decision maker from always select the DN action and to penalize reaching high deterioration levels.

We define the state of the road section by $s = (\rho, \theta, h, m)$, where ρ and θ are the current LCP and the DGR, respectively, h is the current section thickness and m is the last maintenance type performed on the section (m is equivalent to the thickness of the last added layer).

Let A_h be the set of the feasible actions when the current section thickness is h and let N_h denote its cardinality.

$$A_h = \{a = (h_p, h_n) ; h_p \in H_p, h_n = h_n(h_p) \in H_n(h, h_p)\}.$$

We assume that an action a is instantaneously selected and performed. Moreover, if the DN-action is selected, the system state just after decision remains in s . Otherwise, according to the maintenance assumptions, the LCP is reset to zero and the new value of the DGR is reduced to a function of both its state before maintenance and the type of the last maintenance, $\phi(s, a)$, and the new characteristic of the road defined by $h_a = h - h_n + h_p$. h_p determines the future deterioration behavior. Finally, the new state of the road is $s_a = (0, \phi(s, a), h_a, h_p)$.

In [23], the perfect information problem is formulated as an MDP where the instantaneous DGR is approximated by the average using two successive observations of ρ . The model is solved for the optimal policy that minimizes the discounted cost-to-go over the infinite horizon subject to the maximum thickness constraint. However, numerical examples showed that our policy is very sensitive to the DGR θ . In addition, we introduced the DGR as a deterioration parameter to complete the information that the LCP gives about the deterioration level of the section. This approximation limits the significance of θ . These are the main motivations to consider a partially observed problem in this work.

3.2 Decision process formulation

The state transitions during next decision epoch are given by the deterioration model defined by Zouch *et al.* [24]. This model is a two-phase deterioration model based on a Poisson process modeling the first crack arrival and a gamma process for the deterioration growth rate in the initiation phase. In the propagation phase, ρ is modeled a state-dependent gamma (SDG) process, i.e; the increments of the process depend on both the current ρ and θ , whereas the GDR is modeled as a bilateral gamma (BG) process. For details of the deterioration model refer to [24].

For numerical convenience, a state space discretization is considered. Denote $S = s_1, \dots, s_N$ the state set where $N = |S|$ and $s = (\rho_i, \theta_j, h, m)$ where $i = \{1, \dots, N_\rho\}$ and $j = \{1, \dots, N_\theta\}$. In practice, we can aggregate the possible levels of instantaneous speed of deterioration into qualitative states, e.g. *very slow, slow, fast, very fast*.

When the last performed maintenance is of type m , the transition probability from state $s = (\rho, \theta, h, m)$ at the beginning of the decision epoch t to $s' = (\rho', \theta', h, m)$ at its end $t + \tau$

is given by

$$\begin{aligned} p_{ss'}^m &= Pr\{s_{t+\tau} = s' | s_t = s\} \\ &= Pr\{\rho_{t+\tau} = \rho' | \rho_t, \theta_t, h, m\} Pr\{\theta_{t+\tau} = \theta' | \rho'_t, \rho_t, h, m\} \\ &= Pr\{\rho' | \rho, \theta, h, m\} Pr\{\theta' | \rho', \rho, h, m\}, \end{aligned}$$

where $Pr\{\rho' | \rho, \theta, h, m\}$ is given by the density of the SDG process and $Pr\{\theta' | \rho', \rho, h, m\}$ by the density of the BG process [24].

As the section state is partially observed, the knowledge on the section state is represented by a probability distribution over the real states, referred to as the belief state and denoted $\mathbf{b} = (b_1, \dots, b_N)$, where $b_j, j \in \{1, \dots, N\}$ is the probability that the real state (deterioration level) is s_j . More specifically, if $s = (\rho, \theta)$ and the observed LCP level is different from ρ , then $b_j = 0$, otherwise b_j is the probability that true DGR level is θ given the LCP observation ρ . In this last case, \mathbf{b} is equivalent to the N_θ -vector \mathbf{b}^θ , the belief on the DGR θ . Finally, the belief state verifies $b_j \geq 0$ for $j \in \{1, \dots, N\}$ and $\sum_{j=1}^N b_j = 1$. Let B be the set of all possible beliefs, i.e., the set of probability distributions over the section states.

Based on this knowledge, the decision-maker selects an action $a = (h_p, h_n)$. As the DN action does not alter the current state, it will also not alter the belief. Any other action changes the belief into a new belief \mathbf{b}^a given by a deterministic function of the performed action and the belief just before action. At the end of the decision epoch, given the last maintenance type m the true state s of the road section evolves to a new state s' with probability $p_{ss'}^m$ defined in Equation (1) and the section is inspected to yield an observation of ρ . This LCP observation is used to estimate an average deterioration speed $\tilde{\theta}$ on the last section. Let o be the state observation $(\rho, \tilde{\theta})$. Let O be the set of observation and assume that there is a finite number N_o of possible observations of (ρ, θ) .

For $s \in S$ and $o \in O$, let q_{so} be the probability observing o when the real section state is s . Let $Q = [q_{so}]$ be the $N \times N_o$ observation matrix.

Given the last maintenance type m , let $\Omega(o|\mathbf{b})$ be the probability of observing o at the end of the current decision epoch when the belief is \mathbf{b} at its beginning.

$$\Omega(o|\mathbf{b}) = \sum_{s \in S} \sum_{s' \in S} b_s p_{ss'}^m q_{s'o},$$

$\Omega(o|\mathbf{b})$ is the sum of all possible transitions from s to s' that may yield the observation o weighted by the belief distribution. Based on the observation o , the belief state is updated to \mathbf{b}' using the Bayes formula,

$$\mathbf{b}'_{s'} = \frac{\sum_{s \in S} b_s p_{ss'}^m q_{s'o}}{\sum_{s \in S} \sum_{s' \in S} b_s p_{ss'}^m q_{s'o}}, \quad \forall s'.$$

3.3 Cost decision criterion

The objective is to find the optimal policy that minimizes the discounted cost-to-go over the infinite horizon.

$$V^*(\mathbf{b}, h, m) = c_0 + \min \{DN(\mathbf{b}, h, m), MX(\mathbf{b}, h, m)\} \quad (1)$$

where

$$DN(\mathbf{b}, h, m) = \sum_{o \in \mathcal{O}} \Omega(o|\mathbf{b}) \{C_q(\mathbf{b}, \mathbf{b}') + \lambda V^*(\mathbf{b}', h, m)\} \quad (2)$$

$$MX(\mathbf{b}, h, m) = \min_a \{MX_{a_1}(\mathbf{b}, h, m), \dots, MX_{a_2}(\mathbf{b}, h, m)\} \quad (3)$$

and for $k \in 1, \dots, N_h$ and $a_k = (h_p, h_n)$,

$$MX_{a_k}(\mathbf{b}, h, m) = c_v(h_n, h_p) + DN(\mathbf{b}^k, h - h_n + h_p, k).$$

where λ is the discount factor and \mathbf{b}^k is the belief state just after a maintenance action is performed at a current belief \mathbf{b} .

Equation (2) states that following the DN-action when the current belief is \mathbf{b} incurs the expected quality cost plus the cost-to-go to the updated belief \mathbf{b}' by the end of the decision epoch. The action a_k in Equation (3) incurs a fixed and a variable maintenance costs plus the DN-cost of a system with a new belief after maintenance.

4 Solution procedure & Numerical example

4.1 Solution procedure

POMDPs offer the possibility to take into account the non observability of some decision variables. However, they are very hard to solve and computing the optimal policy is challenging because of the curse of dimensionality [15]. The main problem with solving exactly POMDPs is that it is impossible to update all the belief values. Approximating a POMDP solution consists of applying value updates in only some specific beliefs. Different approaches to select these specific beliefs to update have been developed such as grid-based approximation ([10] and [22]) and point-based approximation ([15, 3, 21] and [16]). Once the beliefs are selected, their value updating routine is standard [15]. In this work, we use the value iteration algorithm with the regular grid-based approximation proposed by Lovejoy (1991) [10].

4.2 Numerical example

In this section we discuss our model using a numerical example. For the transition probabilities, we consider the same example deterioration model as proposed in [23] and [24]. We assume that, given the real DGR, the observed one $\tilde{\theta}$ is normally distributed with mean equal to θ and a given variance σ^2 .

$$Pr\{\tilde{\theta} | \theta\} = N(\theta, \sigma^2).$$

The maintenance effect function on the DGR θ is given for $s = (\rho, \theta, h, m)$ by

$$\phi(s, a) = r_a\left(\theta + \frac{\rho}{\tau} + \frac{1}{h}\right)$$

where r_a is a reduction factor depending on the performed action. Cost parameters are given by the following values (in Euros)

$$c_n = 200 ; c_p = 600 , c_{un} = 30 ; c_{up} = 70.$$

The quality cost function is given by

$$C_q(\tau; \mathbf{b}) = \tau(800 \mathbf{b} + \frac{1000}{h})$$

The discount factor is $\lambda = 0.95$. We also consider an action set presented by the following Table (h_p and h_n are given in centimeters)

Table 1: Definition of the action set by the associated removed and added layers thicknesses

a	1	2	3	4	5	6	7	8	9	10
h_p	0	5	5	5	5	10	10	10	10	15
h_n	0	0	2.67	5.33	8	3	4.67	6.33	8	8

Finally, we obtain the policy presented in Figure 2. The optimal policy is given here for a (6x5) state set and evaluated with the Value Iteration algorithm based on a grid based approximation with a grid resolution parameter $M = 3$. The decision matrix gives the optimal maintenance action for each observed ρ and each belief on θ when the current thickness of the road section is $h = 8$ and the last added layer is $h_p = 5$.

For example (frames in red full lines in Figure 2), for an observed $\rho = 0.4$ and a belief on θ , $\mathbf{b}^\theta = (0, 0, 1, 0, 0)$, the optimal maintenance action is ($h_p = 10, h_n = 8$), i.e., it is optimal to remove all the road and rebuild a thicker one. When $\mathbf{b}^\theta = (0.67, 0, 0.33, 0, 0)$, the optimal action is to remove nothing and add a layer of thickness 5.

Note that, due to the deterioration growth rate parameter, the provided policy is a preventive maintenance policy: even when the road is free of cracks, it is optimal to maintain it for high values of θ . The GDR is then a risk measure. The presence of action 10 in the optimal policy ensures that the road will be totally renewed at least once.

The five first rows of the decision matrix corresponds to the perfect information case, i.e, when θ is observable. Note that considering a probability distribution over θ may make some savings. For example (frames in green dashed lines in Figure 2), the optimal action when $\rho = 0.2$ and $\theta = 0.6$ with probability 1 is action 6. Whereas, for the same level of cracking but with θ equal to 0.6 with probability 0.67 and equal to 0.4 with probability 0.33, the action 2 is recommended.

This confirms the importance of the growth rate as a deterioration parameter and the sensitivity of our model to it. From numerical examples, we can notice that obtained policies present a structural properties with respect to the stochastic ordering of each belief parameter as well as the belief, i.e., the action type increases as the LCP and the belief on θ increase, as well as when the belief state increases. Moreover, the structure of optimal policies depends especially on the maintenance effect function. This function reflects the preference between a

	(0,0,0.33,0.33,0.33)	6	6	9	9	9	9
	(0,0.33,0,0.33,0.33)	2	2	2	9	9	9
	(0,0.33,0.33,0,0.33)	2	2	2	2	9	9
	(0,0.33,0.33,0.33,0)	2	2	2	2	9	9
	(0.33,0,0,0.33,0.33)	2	2	2	2	9	9
	(0.33,0,0.33,0,0.33)	2	2	2	2	2	2
	(0.33,0,0.33,0.33,0)	2	2	2	2	2	2
	(0.33,0.33,0,0,0.33)	2	2	2	2	2	2
	(0.33,0.33,0,0.33,0)	2	2	2	2	2	2
	(0.33,0.33,0.33,0,0)	2	2	2	2	2	2
	(0,0,0,0.67,0.33)	6	6	6	9	9	9
	(0,0,0.67,0,0.33)	2	2	2	9	9	9
	(0,0.67,0,0,0.33)	2	2	2	2	2	2
	(0.67,0,0,0,0.33)	2	2	2	2	2	2
	(0,0,0,0.33,0.67)	9	9	9	9	9	9
	(0,0,0.67,0.33,0)	2	2	2	9	9	9
	(0,0.67,0,0.33,0)	2	2	2	2	2	2
bθ	(0.67,0,0,0.33,0)	2	2	2	2	2	2
	(0,0,0.33,0,0.67)	9	9	9	9	9	9
	(0,0,0.33,0.67,0)	6	6	2	9	9	9
	(0,0.67,0.33,0,0)	2	2	2	2	2	2
	(0.67,0,0.33,0,0)	2	2	2	2	2	2
	(0,0.33,0,0,0.67)	9	9	9	9	9	9
	(0,0.33,0,0.67,0)	2	2	2	2	2	2
	(0,0.33,0.67,0,0)	2	2	2	2	2	2
	(0.67,0.33,0,0,0)	2	2	2	2	2	2
	(0.33,0,0,0,0.67)	9	9	9	9	9	9
	(0.33,0,0,0.67,0)	2	2	2	2	2	8
	(0.33,0,0.67,0,0)	2	2	2	2	2	2
	(0.33,0.67,0,0,0)	1	2	2	2	2	2
	(0,0,0,0,1)	9	9	9	9	9	10
	(0,0,0,1,0)	6	6	6	9	9	9
	(0,0,1,0,0)	2	2	2	8	9	9
	(0,1,0,0,0)	1	1	1	2	2	2
	(1,0,0,0,0)	1	1	1	1	1	1
		0	0.03	0.2	0.4	0.6	0.8
				p			

Figure 2: Optimal policy for a (6x5) problem where $h_{max} = 15$, $h = 8$ and $h_p = 5$.

newer and a thicker road.

A direct comparison with the current French practices in road maintenance is not provided here. The current road maintenance strategies are defined in a complex process based on the evaluation of various scenarios taking into account many external parameters. It may be noticed that the decision is based, among other things, on the observation of the deterioration level, i.e. ρ . This policy would be equivalent to a policy that only one action would be identified in each column of the decision matrix. The proposed optimized policy differentiates maintenance actions according to \mathbf{b}^θ highlighting the non-optimality of the current practices.

5 Conclusions

We propose a condition-based maintenance optimization approach for road section to minimize the total maintenance cost over the infinite horizon with perfect and different imperfect maintenance actions are available. Our approach is based on a two-phase deterioration process with two deterioration parameters: the longitudinal cracking percentage and the deterioration growth rate. Our deterioration model has the advantage to be state-dependent. Moreover, the introduction of the speed of deterioration as a decision parameter allows the model to take into account the evolution of the underlying deterioration and so to the prediction of the deterioration evolution to be more precise.

However, the deterioration speed cannot be observable. We model the problem as a partially observed Markov decision process and solve it using a grid-based value iteration algorithm. This algorithm is highly time-consuming when the size of the state space is big because it generates a large belief grid. In our case, the size of the grid is reduced because only one of two decision parameters is not observable.

In practice, excepted the difficulties to estimate the parameters of the two-phase deterioration model, our approach can easily be used to provide optimal policies where the maintenance decision is not based on the pre-definition of deterioration level thresholds, which is not possible in the initiation phase. We propose optimal preventive policies that allow to maintain the road when it is free of cracks.

References

- [1] *Elaboration d'une politique routière de maintenance par niveaux de service - Guide méthodologique, Documentation des techniques routières françaises*, 1994.
- [2] S. D. Boyles, Z. Zhang, and S. T. Waller. Optimal maintenance and repair policies under nonlinear preferences. *Journal of Infrastructure Systems*, 16(1):11–20, 2010.
- [3] H. Cheng. *Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, University of British Columbia. School of Commerce, 1988.
- [4] D. M. Frangopol, M. J. Kallen, and J. M. van Noortwijk. Probabilistic models for life-cycle performance of deteriorating structures: review and future directions. *Structural Engineering and Materials*, 6:197–212, 2004.
- [5] A. Ghasemi, S. Yacout, and M. S. Ouali. Optimal condition based maintenance with imperfect information and the proportional hazards model. *International journal of production research*, 45(4):989–1012, 2007.
- [6] K. Golbai, R. B. Kulkarni, and R. G. Way. A statewide pavement management system. *Interfaces*, 12:5–21, 1982.
- [7] W. J. Hopp and Y. L. Kuo. An optimal structured policy for maintenance of partially observable aircraft engine components. *Naval Research Logistics*, 45(4):335–352, 1998.
- [8] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.
- [9] M. J. Kallen. *Markov processes for maintenance optimization of civil infrastructure in the Netherlands*. PhD thesis, Delft University of Technology, Delft, 2007.
- [10] W. S. Lovejoy. Computationally feasible bounds for partially observed markov decision processes. *Operations Research*, 39(1):162–175, 1991.
- [11] S. Madanat and M. Ben-Akiva. Optimal inspection and repair policies for infrastructure facilities. *Transportation Science*, 28(1):55–62, 1994.
- [12] L. M. Maillart. Maintenance policies for systems with condition monitoring and obvious failures. *IIE Transactions*, 38:463–475, 2006.
- [13] V. Makis and X. Jiang. Optimal replacement under partial observations. *Mathematics of Operations Research*, 28(2):382–394, 2003.
- [14] G. E. Monahan. A survey of partially observable markov decision processes: Theory, models, and algorithms. *Management Science*, 28(1):1–16, 1982.
- [15] J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1025–1032, Acapulco, Mexico, 2003.

-
- [16] K. M. Poon. A fast heuristic algorithm for decision-theoretic planning. Master's thesis, Hong Kong University of Science and Technology, 2001.
- [17] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley and Sons Inc., New York, 1994.
- [18] K. Smilowitz and S. Madanat. Optimal inspection and maintenance policies for infrastructure networks. *Computer-Aided Civil and Infrastructure Engineering*, 15(1):5–13, 2002.
- [19] J. M. van Noortwijk and D. M. Frangopol. Two probabilistic life-cycle maintenance models for deteriorating civil infrastructures. *Probabilistic Engineering Mechanics*, 19(4):345 – 359, 2004.
- [20] H. Wang. A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3):469 – 489, 2002.
- [21] N. L. Zhang and W. Zhang. Speeding up the convergence of value iteration in partially observable markov decision processes. *Journal of Artificial Intelligence Research*, 14:29–51, 2001.
- [22] R. Zhou and E. Hansen. An improved grid-based approximation algorithm for POMDPs, 2001.
- [23] M. Zouch, T. G. Yeung, and B. Castanier. Optimal condition-based resurfacing decisions for roads. In *Proceedings of the Annual Conference of European Safety and Reliability Association*, pages 1379–1384, Rhodes, Greece, 2010. ESREL.
- [24] M. Zouch, T. G. Yeung, and B. Castanier. Two-phase state-dependent deterioration model for maintenance optimization. *Submitted to Naval Research Logistics*, 2011.