ADAPTIVE OPTIMIZATION OF INFRASTRUCTURE MAINTENANCE AND INSPECTION DECISIONS UNDER PERFORMANCE MODEL UNCERTAINTY

Samer Madanat[1a], Pablo Durango-Cohen[b]
[a]Department of Civil and Environmental Engineering, University of California, Berkeley CA 94720, Tel: 1-510-643-1084
[b]Department of Civil and Environmental Engineering, Northwestern University, Evanston, IL 60208, Tel: 1-847-491-4008
e-mails : madanat@ce.berkeley.edu**,** pdc@northwestern.edu

Abstract

Decision-making models provide highway agencies with a plan for optimal decisions regarding maintenance and repair activities. The objective of these models is to minimize the total expected social cost of maintaining a system of facilities incurred by an agency and the users over a given planning horizon. Recent models take into account measurement error in the inspection process and optimize the inspection schedule. Other state-of-the-art models account for uncertainty in performance forecasting. Our research develops a model that jointly determines when to inspect and what maintenance activity to perform, while taking into account both uncertainty in the measurements and feedback in the estimation of the deterioration model parameters. A computational implementation is performed in order to study empirically the relative significance of uncertainties in the deterioration rate and the state of the system.

Keywords: Transportation infrastructure; Maintenance; Optimization; Performance model; Uncertainty; Adaptive Control
Topic Area: C5 Transport Systems Maintenance

1. Introduction

Infrastructure management is the process through which agencies collect and analyze data about the state of infrastructure systems and make decisions on maintenance, repair and reconstruction (MR&R) of the facilities comprised in their network over a given planning horizon. Bridge maintenance, road improvement, and highway rehabilitation are examples of MR&R activities.

Every time period, commonly every year, agencies face two types of decisions for each facility in the network: whether to inspect or not, and which MR&R action to perform, if any. In this role, they are supported by Infrastructure Management Systems (IMS) that provide them with tools to help them in this three-step process:

- data collection,
- interpretation,
- decision-making.

The agencies base their MR&R decision on performance models that forecast the behavior of the infrastructure facility under the effect of MR&R action and deterioration. The agencies' objective is to minimize the expected cost related to the facilities' use and maintenance over the planning horizon.

This problem has been extensively researched and is known in the literature as the Infrastructure Maintenance and Repair (IMR) problem.

---

[1] Corresponding author

The focus of this paper is on decision-making at a facility level. The methodology that we develop can be used for each facility an agency is responsible for, as long as it does not face any budget constraint. In reality, agencies operate under budget constraints and level-of-service constraints, in the context of a network of facilities.

The main motivation of this research is to examine the effect of relaxing the annual inspection constraint, in the case of infrastructure management under model uncertainty.

We develop a formulation that builds upon recent developments in the application of adaptive control (AC) schemes, i.e. decision-making algorithms that explicitly account for uncertainty in the performance models (Durango and Madanat, 2002). Furthermore, we relax the constraint of annual inspection that is typically imposed in MR&R models.

2. Literature review

The first row of Table 1 summarizes the main steps of the Event Chain in IMS. In the first phase, traffic, weather, and ageing contribute to facility deterioration. Then the agency observes the new state of the facility. The observation may or may not be error-free, depending on the inspection technology. Given the measured state of the system, the agency makes decisions concerning the actions to be taken at the end of the period. The decision rule in the state-of-the-art IMSs is to chose an action that will minimize the sum of expected future and present cost. The actions may include MR&R and inspection activities.

A review of existing research on which this research is based shows that, so far, no model has been developed to take into account uncertainty in all stages of this event chain. Table 1 summarizes the main features of those models. The following subsection describes these models in more detail.

Model	Facility deterioration	Observation	Decision	Action
MDP	Probabilistic	True state	MR&R	Annual Inspection
Joint decision	Probabilistic	True state	MR&R Inspection	Relaxation of the annual constraint
Latent MDP	Probabilistic	Measured state	MR&R Inspection	Relaxation of the annual constraint
AC	Probabilistic	True state	MR&R	Annual Inspection
Proposed model	Probabilistic	Measured state	MR&R Inspection	Relaxation of the annual constraint

Table 1. Features of IMR models

MDP- based models take into account the uncertainty in facility deterioration forecasting. The state of the facility is discretized and the deterioration process is represented by the transition probabilities defined below:

$$\pi_{ij}^{t}(a) = \Pr(x_{t+1} = j \mid x_t = i, a_t = a) \qquad (1)$$

where:

- x_{t+1} is the state of the facility at the beginning of period t+1,
- x_t is the state of the facility at the beginning of period t,
- a_t is the action taken in period t.

The Transition Probabilities Matrix (TPM) can be derived from empirical data. Several approaches to estimate the TPMs are present in the literature. One method that uses statistical estimation and time series data is found in Carnahan et al. (1987) and Olsonen (1988). Another is proposed by Madanat (1991) and is based on a performance model and the mathematical properties of Markov Chains. Madanat and Wan Ibrahim (1995) describe how Poisson regression and, more generally, negative binomial regression can be used to estimate the TPMs coefficient, which ensures more statistically sound estimation and recognizes the discrete nature of the condition state transitions. Finally, Mishalani and Madanat (2001) developed a stochastic duration-based methodology to estimate the transition probabilities that specifically takes into account the effect of causal variables, and recognizes the correlation between successive observations.

If the deterioration process that governs the TPMs is the same during each period, then the rate is homogeneous with respect to time. In this case, the transition probabilities are constant over time.

Then,
$$\pi_{ij}^{t}(a) = \pi_{ij}(a) \quad (2)$$

The MDP based formulation for the infrastructure maintenance and repair problem is solved by dynamic programming (DP) over a finite horizon. The objective function is the discounted cost-to-go until the end of the planning schedule. The cost incurred during each period includes the sum of the user costs and the cost of the MR&R activity chosen, which is incurred by the agency. The solution gives a set of actions to be performed at each time period for a given state of the facility.

The primary assumptions of this model are:

- the true state of the facility is observed during inspections,
- the evolution of facility condition depends only on the previous state and the last action,
- an inspection is performed every year.

MDP formulations fail to :

- include the possibility of flexible inspection schedule,
- account for uncertainty in the inspection process,
- incorporate information about the actual deterioration process[2].

The incorporation of joint decisions that includes inspection and MR&R activity is relatively straightforward. The DP objective function depends on three variables: time, state, length since last inspection was performed. This issue has been investigated by several researchers, such as Klein (1962) and Mine and Kawai (1982). Nevertheless these models and their variations fail to account for uncertainty in the measurement process.

Research by Humplick (1992) has shown that there are significant measurement errors in existing infrastructure inspection technologies. This can lead to taking the wrong MR&R activity if the prescribed activities for the true condition and the measured condition are different. A second limitation of traditional MDP formulations is the lack of systematic methodology for inspection decision-making. Current methods for inspection scheduling are ad hoc and subjective (for example Shahin and Kohn, 1981).

The purpose of the LMDP is to account for uncertainty in the inspection process and to allow MR&R decision making at each stage –not only when an inspection is performed. This entails the violation of a basic assumption of the MDP methodology: perfect information about the state. Therefore the technique of state augmentation for DP, whose theory is described in Bertsekas (1987), was used by Madanat and Ben Akiva (1994). This technique takes advantage of the fact that at any point in time, the decision-maker knows the history of the past actions taken and observed states. In the LMDP, the measured state is related to the true state of the facility by measurement probabilities.

[2] In practice, agencies update their deterioration model by using the observed transitions. A failure to account for this updating in the DP formulation results in suboptimal policies.

State augmentation

Under the state augmentation technique, the state of the system at stage t takes into account all the information available to the decision maker, since the beginning of the planning horizon, and that is relevant for decision-making. This is summarized into the information set I_t. This set includes at any point in time the information about previous actions and observations i.e.

$$I_t = \{I_0, a_0, \hat{x}_0, a_1, ..., \hat{x}_{t-1}, a_{t-1}, \hat{x}_t\}, \qquad (3)$$

where I_0 is the initial information available at the start of the planning horizon and $\hat{x}_t$ is the observed state during period t.

By recursion, we have $I_t = \{I_{t-1}, a_{t-1}, \hat{x}_t\}$. This shows that the evolution of the information set follows a Markovian Process:

$$\Pr(I_t \mid I_0, a_0, \hat{x}_1, I_1, a_1, ..., \hat{x}_{t-1}, a_{t-1}, I_{t-1}) = \Pr(I_t \mid I_{t-1}, a_{t-1}) \qquad (4)$$

Therefore, we can use a DP formulation over the space of the information sets. The objective function is generalized into the minimum discounted <u>expected</u> cost-to-go given the information available. The cost for each stage is generalized from $g(x_t, a_t)$ to $\tilde{g}(I_t, a_t) = \underset{x_t}{E}[g(x_t, a_t) \mid I_t]$, which includes the following conditional probability distribution: $\Pr(x_t \mid I_t)$ i.e. in vector notation $\underline{P_t \mid I_t}$.

$\underline{P_t \mid I_t}$ is known in the literature as the information vector.

Under these assumptions, the observed state is now only probabilistically related to its true state. The distribution of the measurement relative to the true state is known and depends on the technology used. Although only the observed state is available, forecasting models are still based on the true state.

The model assumes the availability of the following measurement probabilities:

$$\varepsilon_{jk}^r = \Pr(\hat{x}_t = k \mid x_t = j, R_t = r) \qquad (5)$$

where:

- $\hat{x}_t$ and x_t are respectively the observed and true condition state of the facility,
- j, k are discrete condition states,
- R_t is the technology used for the measurement.

Measurement probabilities can be derived empirically using measurement error models. Such models relate pavement performance to specified indicators. Commonly used indicators in the field of pavement management are the Pavement Condition Index (PCI) as detailed in Shahin and Kohn (1981) and the Present Serviceability Index (HRB 1962). In the remainder of this paper we will only consider the PCI scale. Madanat (1991) derives a relationship between the true value of an indicator and its measured value using a given technology. It assumes normality in the error distribution and ignores bias in measurement error. Humplick (1989) shows that biases can be statistically estimated, so they are neglected in this model, since they can be corrected for.

The DP formulation is implemented in the following manner: at the beginning of the planning horizon, the agency determines an initial information vector according to its beliefs about the state of the facility. The state vector is then updated forward in time, according to Bayes' rules, taking into account:

- the previous information vector,
- the forecasting model,

- the observation made,
- uncertainty in the inspection process.

At each stage of the DP problem, the decision variables are now whether to inspect or not and what MR&R activity to perform.

The DP solution is in the form of a policy μ^* i.e. $\{\mu_0^*(I_0),...,\mu_{T-1}^*(I_{T-1})\}$, where T is the end of the planning horizon and $\forall t \in [0,...,T]$, $\mu_t^*(I_t) = (a_t^*, R_t^*)$.

LMDP models address the issue of uncertainty in the inspection process and include flexible joint inspection and MR&R decision making, but:

- they do not allow for uncertainty in the choice of performance models,
- they fail to take into account possible feedback from observations to improve the performance models.

The models described in the preceding paragraphs all refer to a single set of TPM for performance forecasting. Yet Carnahan (1988) acknowledges that those models are subject to significant uncertainty particularly in the exogenous factors and are often updated in the course of operations by the agency. Hence the need to include intertemporal feedback and performance model updating in the original planning process.

Current research on the application of AC schemes to the IMR problem has focused on the characterization of performance forecasting by deterioration rates (Durango and Madanat, 2002). As the information about the true deterioration rate is uncertain, it can be described by a discrete probability distribution $\underline{Q_t} = \{Q_t^1, Q_t^2, ..., Q_t^D\}$ where D is the number of possible deterioration rates. More specifically, at the beginning of each stage, the decision maker has the following beliefs about the deterioration rate:

$$\forall d \in [1,..,D], \quad Q_t^d = \Pr_t(\tilde{d} = d \mid I_t) \tag{6}$$

where $\tilde{d}$ is the true deterioration rate.

To account for uncertainty in the deterioration rate in a DP formulation, the method of state-augmentation was used by Durango and Madanat (2002). The augmented state at the beginning of stage t is now defined by the true state of the system, which is revealed at the end of the previous stage, and the beliefs about the rate: $\{x_t, \underline{Q}_t\}$. The objective function is the discounted, expected (over the distribution of the deterioration rates) cost-to-go. The goal is to provide a set of policy functions that give the optimal MR&R activity at each stage, for any state of the system. The DP solution will give $\forall t \in [1,..,T-1], \forall \{x_t, \underline{Q}_t\}$ a_t^*, the action that minimizes the discounted expected cost-to-go.

Two AC schemes have been studied: the Closed Loop (CL) scheme that includes Bayesian updating in the argument of the objective function, and a suboptimal Open-Loop Feedback (OLF) scheme that takes into account the agency's beliefs about the rate only in the expected cost. The major difference between CL and OLF is that the latter assumes that the current updating will be the only updating performed.

The CL and OLF schemes implemented in former studies did not allow flexible inspection and did not account for uncertainty in the measurement process. In spite of the efficiency of DP, the computer implementation raises the issue of discretization and coarseness of the discrete grid.

3. Formulation

In the remainder of this paper, the following additional notation will be used:

N number of possible states

α discount factor

The definition of the information set is similar to Equation (3) but includes information about past inspections i.e.

$$I_t = \{I_0, a_0, \hat{x}_0, R_0, a_1, ..., \hat{x}_{t-1}, a_{t-1}, \hat{x}_t, R_t\} \quad (7)$$

or in the recursive version, assuming I_0 is known,

$$I_t = \{a_{t-1}, \hat{x}_t, R_t, I_{t-1}\} \quad (8)$$

The elements of the set of beliefs about the deterioration rate that we introduced in equation (6) have to be modified to include this new information set.
To simplify the notation, we replace $Q_t^d\left(R_t, \hat{x}_t, \underline{P_{t-1} \mid I_{t-1}}, \underline{Q_{t-1}}, a_{t-1}\right)$ by Q_t^d.

Figure 1 summarizes the main decision-making steps that are accounted for in the DP formulation that we present later. The first event at the beginning of a time period t is the usage of the facility, which implies a user cost that is a function of the true state of the facility. Then an inspection can be made according to the decision made at the previous stage, which increments the information vector available to the agency by the observed state. Once the decision maker has this information, she can update her beliefs about the deterioration rate, and the state of the facility. The decision making involves the choice of action to perform during time period t, as well as whether to inspect or not at the beginning of the next period –note that stage t+1 inspection cost is incurred during this time period t. Before the start of the next stage, the true condition state of the facility experiences deterioration under the effect of weather, traffic and ageing.

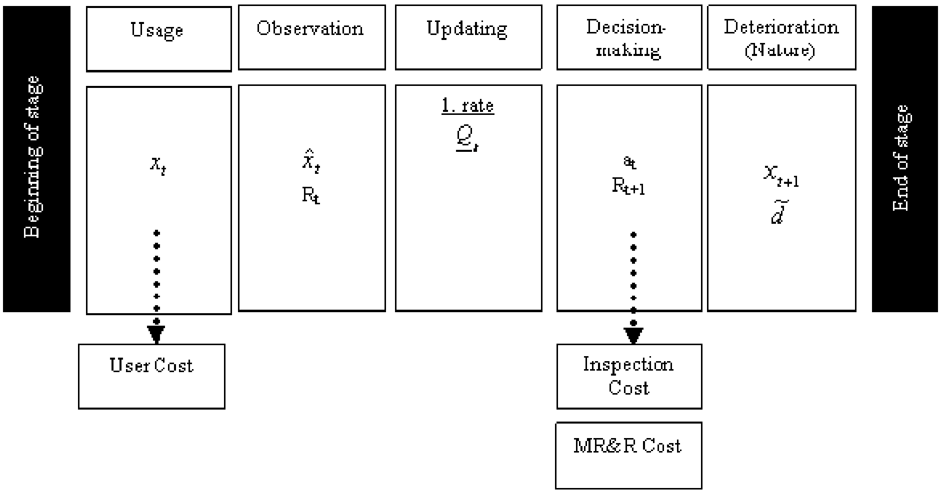

Figure 1. The decision-making framework

Before presenting the recursive formula that is at the core of the DP algorithm, we will present the costs, performance models and measurement error formulations that constitute the model specification.

Transition Probability Matrices

The transition probability matrices that are used to forecast the effect of each action on the true state of the facility, depending on the true deterioration rate are the different performance forecasting models that we consider. Now the TPMs depend on the pair initial true state/MR&R activity and on the deterioration rate, and they are therefore denoted by:

$$\pi_{ij}^{d}(a_t) = \Pr\left(x_{t+1} = j \mid \tilde{d} = d, x_t = i, a_t\right) \qquad (9)$$

Measurement error

The notation for modeling measurement error is the same as the one presented for the LMDP model, but we need to be more specific about the choice of technology. The decision of whether to inspect or not can be modeled by a choice between two technologies: one that has the measurement precision of the actual inspection technology, as described above, and the other with a measurement error of infinite variance. The model therefore inherently allows for using a set of different technologies. However in our model we reduce this to a binary choice denoted here by R_t=1 for inspecting using the available technology or R_t=0 for not inspecting.

The "no inspection" decision (R_t=0) refers to a technology where for any true state the measurement probability is uniformly distributed:

$$\forall j, k \in [1, \ldots, N], \quad \varepsilon_{jk}^{0} = \Pr\left(\hat{x}_t = k \mid x_t = j, R_t = 0\right) = \frac{1}{N} \qquad (10)$$

This case, where every condition state is equally likely to be observed regardless of the true state, is shown to be equivalent to not inspecting in Madanat and Ben Akiva (1994). The cost associated with such an inspection technology will obviously be zero.

Cost

In the remainder of this section we will use the notation $g(x_t, a_t, R_{t+1})$ for the generic stage cost incurred during period t for performing activity a_t on a facility in state x_t and choosing to use inspection technology R_{t+1} at the beginning of next period. The breakdown of this stage cost is:

User cost: a function of the condition state of the facility; it cannot be easily quantified, it is often replaced by a minimum allowable standard; it is denoted by $uc(x_t)$.

Inspection cost: a detailed study of how inspection cost is related to measurement error standard deviation can be found in Madanat (1991); we will denote it by m.

MR&R cost: depends on both the activity performed and the true condition state of the facility; it is denoted by $mrrc(x_t, a_t)$.

Salvage value: at the end of the planning horizon, an infinite cost is assigned if the facility does not reach a minimum standard level. We will denote by $s(x_t)$ the salvage value corresponding to condition state x_t

The above list can be summarized by the generic definition of the stage cost:

$$g(x_t, a_t, R_{t+1}) = uc(x_t) + m \cdot R_{t+1} + mrrc(x_t, a_t) \qquad (11)$$

The DP formulation of the facility infrastructure maintenance and repair problem under measurement error and uncertainty about the rate includes, like any DP, a recursive formula and boundary conditions.

The arguments of the objective function can be found when asking "the consultant question", as in Dreyfus and Law (1977). The procedure is the following: if we were to be hired as consultants and asked to minimize the expected discounted cost-to-go, what information would we need? The answer is: $\underline{P_t \mid I_t}$ and $\underline{Q_t}$. Madanat and ben Akiva (1994) emphasize that $\underline{P_t \mid I_t}$ is a sufficient statistic for I_t.

Recursive formula

At each stage, the minimum expected cost-to-go $f_t\left(\underline{P_t \mid I_t}, \underline{Q_t}\right)$ can be expressed in terms of the stage cost and the expected cost-to-go in the next stage.

$$f_t\left(\underline{P_t \mid I_t}, \underline{Q_t}\right) =$$
$$\underset{a_t, R_{t+1}}{Min}\left[\sum_{i=1}^{N} \Pr_t(x_t = i \mid I_t) \cdot \left(g(x_t, a_t, R_{t+1}) + \alpha \cdot \sum_{d=1}^{D} Q_t^d \cdot \sum_{j=1}^{N} \pi_{ij}^d(a_t) \cdot \sum_{k=1}^{N} \varepsilon_{jk}^{R_{t+1}} \cdot f_{t+1}\left(\underline{P_{t+1} \mid I_{t+1}}, \underline{Q_{t+1}}\right)\right)\right] \quad (12)$$

Note that this formula presents a clear explanation of why the inspection decision about whether to inspect or not during period t+1 is made in period t: the inspection in t+1 directly influences the information vector $\underline{P_{t+1} \mid I_{t+1}}$, which is used for the recursive computation of the objective function in period t and the measurement probabilities $\varepsilon_{jk}^{R_{t+1}}$.

The recursive formulation is obtained via successive probability conditioning of:

$$f_t\left(\underline{P_t \mid I_t}, \underline{Q_t}\right) = \underset{a_t, R_{t+1}}{Min}\; E_{x_t \mid I_t, \hat{x}_{t+1} \mid I_t}\left[g(x_t, a_t, R_{t+1}) + \alpha \cdot f_{t+1}\left(\underline{P_{t+1} \mid I_{t+1}}, \underline{Q_{t+1}}\right)\right] \quad (13)$$

Boundary condition

The boundary condition is:

$$f_T\left(\underline{P_T \mid I_T}, \underline{Q_T}\right) = E_{x_T \mid I_T}\left[s(x_T)\right] = \sum_{i=1}^{N} \Pr_T(x_T = i \mid I_T) \cdot s(i) \quad (14)$$

Bayesian updating

Our model specifies that the beliefs about the rate $\underline{Q_t}$ are updated before the beliefs about the facility condition state $\underline{P_t \mid I_t}$.

<u>Updating the beliefs about the facility deterioration rate</u>

Updating the beliefs about the rate takes into account:

- the current inspection decision R_t,
- the observed state $\hat{x}_t = k$,
- the past beliefs about the rate and the state, i.e. $\underline{P_{t-1} \mid I_{t-1}}$ and $\underline{Q_{t-1}}$,
- the last action taken a_{t-1}.

After the observation phase in period t, the beliefs about the rate are modified so that:

$$Q_t^d = \Pr\left(\tilde{d} = d \mid I_t\right) = \Pr\left(\tilde{d} = d \mid \hat{x}_t, a_{t-1}, R_t, I_{t-1}\right)$$
$$= \frac{\Pr\left(\hat{x}_t \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right) \cdot \Pr\left(\tilde{d} = d \mid a_{t-1}, R_t, I_{t-1}\right)}{\sum_{d'=1}^{D} \Pr\left(\hat{x}_t \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right) \cdot \Pr\left(\tilde{d} = d' \mid a_{t-1}, R_t, I_{t-1}\right)}$$

As $\Pr\left(\tilde{d} = d \mid a_{t-1}, R_t, I_{t-1}\right) = \Pr\left(\tilde{d} = d \mid I_{t-1}\right) = Q_{t-1}^d$, we can write:

$$Q_t^d = \frac{\Pr\left(\hat{x}_t \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right)\cdot Q_{t-1}^d}{\sum_{d'=1}^{D}\Pr\left(\hat{x}_t \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right)\cdot Q_{t-1}^{d'}}$$

If we observe that:

$$\begin{aligned}\Pr\left(\hat{x}_t = k \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right) &= \sum_{j=1}^{N}\Pr\left(\hat{x}_t = k \mid x_t, \tilde{d}, a_{t-1}, R_t, I_{t-1}\right)\cdot \Pr\left(x_t = j \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right) \\ &= \sum_{j=1}^{N}\varepsilon_{jk}^{R_t}\cdot\sum_{i=1}^{N}\Pr\left(x_t = j \mid \tilde{d}, x_{t-1}=i, a_{t-1}\right)\cdot\Pr\left(x_{t-1}=i \mid I_{t-1}\right)\end{aligned}$$

It follows that:

$$\Pr\left(\hat{x}_t = k \mid \tilde{d}, a_{t-1}, R_t, I_{t-1}\right) = \sum_{i,j}\varepsilon_{jk}^{R_t}\pi_{ij}^{d}\left(a_{t-1}\right)\cdot \Pr_{t-1}\left(x_{t-1}=i \mid I_{t-1}\right) \quad (15)$$

As $\Pr_{t-1}\left(x_{t-1} = j \mid I_{t-1}\right)$ is a component of the state vector from the previous stage, we have a recursive expression for Q_t^d.

$$Q_t^d = \frac{\sum_{i,j}\varepsilon_{jk}^{R_t}\pi_{ij}^{d}\left(a_{t-1}\right)\cdot \Pr_{t-1}\left(x_{t-1}=i \mid I_{t-1}\right)\cdot Q_{t-1}^d}{\sum_{d'=1}^{D}\sum_{i,j}\varepsilon_{jk}^{R_t}\pi_{ij}^{d'}\left(a_{t-1}\right)\cdot \Pr_{t-1}\left(x_{t-1}=i \mid I_{t-1}\right)\cdot Q_{t-1}^{d'}} \quad (16)$$

where $k = \hat{x}_t$.

Note that when the "no inspection" decision has been made, the beliefs about the rate are not updated. Indeed, no additional information is available to the decision-maker.

Updating of the beliefs about the facility condition state

Given the updated beliefs about the deterioration rate, the decision-maker then revises her beliefs about the new state of the facility condition, taking into account:

- the current inspection decision R_t and observed state $\hat{x}_t = k$,
- the new beliefs about the deterioration rate $\underline{Q_t}$,
- the past beliefs about the facility condition state $\underline{P_{t-1} \mid I_{t-1}}$,
- the last action taken a_{t-1}.

The beliefs about the facility state are defined by the state vector $\underline{P_t \mid I_t}$:

$$\begin{aligned}\Pr_t\left(x_t = j \mid I_t\right) &= \sum_{d=1}^{D}\Pr\left(x_t = j \mid \tilde{D}=d, I_t\right)\cdot\Pr\left(\tilde{D}=d \mid I_t\right) \\ &= \sum_{d} Q_t^d \frac{\Pr\left(x_t = j \mid \tilde{D}=d, a_{t-1}, R_t, I_{t-1}\right)\cdot\varepsilon_{jk}^{R_t}}{\sum_{j'}\Pr\left(x_t = j' \mid \tilde{D}=d, a_{t-1}, R_t, I_{t-1}\right)\cdot\varepsilon_{j'k}^{R_t}}\end{aligned}$$

In the same fashion as earlier, we write:

$$\begin{aligned}\Pr\left(x_t = j \mid \tilde{D}=d, a_{t-1}, R_t, I_{t-1}\right) &= \sum_{i=1}^{N}\Pr\left(x_t = j \mid x_{t-1}=i, \tilde{D}=d, a_{t-1}, R_t, I_{t-1}\right)\cdot\Pr\left(x_{t-1}=i \mid I_{t-1}\right) \\ &= \sum_{i=1}^{N}\pi_{ij}^{d}\left(a_{t-1}\right)\Pr_{t-1}\left(x_{t-1}=i \mid I_{t-1}\right) \quad (17)\end{aligned}$$

Again, we notice that $\Pr_{t-1}(x_{t-1} = j \mid I_{t-1})$ has already been determined, so:

$$\Pr_t(x_t = j \mid I_t) = \sum_d Q_t^d \frac{\sum_i \pi_{ij}^d(a_{t-1})\Pr_{t-1}(x_{t-1} = i \mid I_{t-1}) \cdot \varepsilon_{jk}^{R_t}}{\sum_{j',i} \pi_{ij'}^d(a_{t-1})\Pr_{t-1}(x_{t-1} = i \mid I_{t-1}) \cdot \varepsilon_{j'k}^{R_t}} \qquad (18)$$

One should notice that if the agency's beliefs about the deterioration rate are such that an observation is considered impossible, the beliefs are not updated; i.e. if the observation is inconsistent with the available performance forecasting models, the agency does not take it into account for the next decision.

In the case where no inspection has been performed at the beginning of the time period, the state vector is updated such that the new probabilities are the weighted transition probabilities. Although no new information is available, the decision-maker updates her beliefs using the performance forecasting models.

Expected cost

In the computational study and for simulation purposes, we will be interested not in the minimum total expected cost as defined in the objective function, since it does not describe properly the costs actually incurred by the agency. Our focus will be on the expected cost given that we perform the optimal policy and that the true rate is actually $\hat{d}$.

This expected cost can be defined recursively as the expected cost-to-go after the observation phase of stage t, given that the actual deterioration rate is $\hat{d}$ and we perform the optimal policy given by the DP formulation: $E_t\left(\underline{P_t \mid I_t}, \underline{Q_t}, \hat{d}\right)$.

$$E_t\left(\underline{P_t \mid I_t}, \underline{Q_t}, \hat{d}\right)$$
$$= \sum_{i=1}^{N} \Pr_t(x_t = i \mid I_t) \cdot \left[g\left(x_t, a_t^*, R_t^*\right) + \alpha \cdot \sum_{j=1}^{N} \pi_{ij}^{\hat{d}}\left(a_t^*\right) \cdot \sum_{k=1}^{N} \varepsilon_{jk}^{R_{t+1}^*} \cdot E_{t+1}\left(\underline{P_{t+1} \mid I_{t+1}}, \underline{Q_{t+1}}, \hat{d}\right) \right] \qquad (19)$$

where:

- $\left(a_t^*, R_t^*\right) = \mu^*\left(\underline{P_t \mid I_t}, \underline{Q_t}\right)$ is the optimal policy given by the DP algorithm,
- $\left(\underline{P_{t+1} \mid I_{t+1}}, \underline{Q_{t+1}}\right)$ are updated according to the Bayesian rules given above.

The actual cost incurred at each stage is not available to the decision maker during the planning phase and not used in the optimization process; therefore we did not include it in our computational studies. Nevertheless, this information becomes available in the implementation phase, and thus should be accounted for during that process. However, we did not address this issue in our research.

Computational complexity

If we use the notation GRS, and GRD, for the number of points used to discretize the beliefs about the state, and the rate respectively; the computational complexity associated with the calculation of the optimal policy is in the order of $O\left(T \cdot GRS^{N-1} \cdot GRD^{D-1} \cdot A \cdot \left(N^3 \cdot D + N^2 \cdot D\right)\right)$. As this is exponential in both GRS and GRD, we have subsequently limited our computational implementation to relatively small values for those parameters.

4. Computational study

We considered a finite horizon IMR problem, where the planning horizon is 15 years, the discount rate being $\alpha = \dfrac{1}{1-d}$ where d is 5%.

The condition state of the facility is discretized into 8 states, each representing 12 PCI points on a scale of 100, as described by Carnahan (1987). State 8 indicates that the facility is in excellent condition. The agency has 7 possible MR&R activities ranging from no action (1), routine maintenance (2) to overlays of different depths (3-4-5-6) and reconstruction (7).

Three possible deterioration rates are considered, slow medium and fast. With each deterioration rate is associated a set of 7 of Transition Probability Matrices, each corresponding to a possible activity. The TPMs are taken from Durango and Madanat (2002). The deterioration process is assumed to exhibit a normal distribution, and common-sense rules the overall process in the following way:

- Each MR&R action has a mean effect on the transition from a given state,
- As the deterioration rate increases, these mean effects are negatively impacted,
- Faster deterioration rates have higher variance in forecasting.

In our computational studies, the measurement error is assumed to be zero i.e. $\varepsilon^{1}_{jk} = 1$ if k=j, 0 otherwise. If an inspection is performed, the agency has perfect information. This assumption was made to reduce the number of parameters and simplify the interpretation of the results.

As mentioned earlier, the stage cost depends on the cost of inspection, the user cost and the cost of the MR&R activity. The values used for these costs were taken from Durango and Madanat (2002). As in Madanat and Ben Akiva (1994), the inspection cost is assumed to be 0.065$/sq yard. Although this does not reflect the actual cost of the assumed error-free process, it is only used for comparison purposes.

Figures 2 and 3 show a comparison of the expected cost when the true rate is either Slow or Fast. For these figures, the initial beliefs about the state are $\underline{P_0 \mid I_0}$ = (0, 0.1, 0.1, 0.2, 0.4, 0.2, 0, 0); "slow" indicates that the initial set of beliefs about the rate is $\underline{Q_0}$ =(0.8, 0.1, 0.1), "fast" indicates that it is $\underline{Q_0}$ = (0.1, 0.1, 0.8), whereas "no" stands for non-informative initial beliefs i.e. $\underline{Q_0}$ = (0.3, 0.4, 0.3).

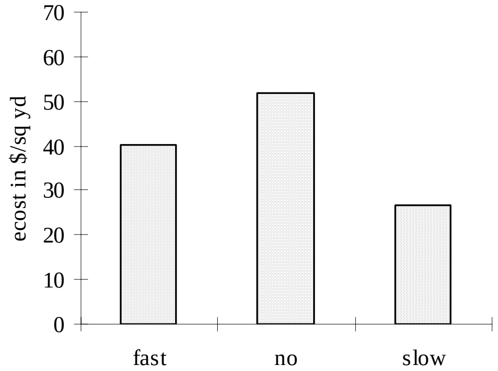

Figure 2. Expected cost: True Rate = Slow

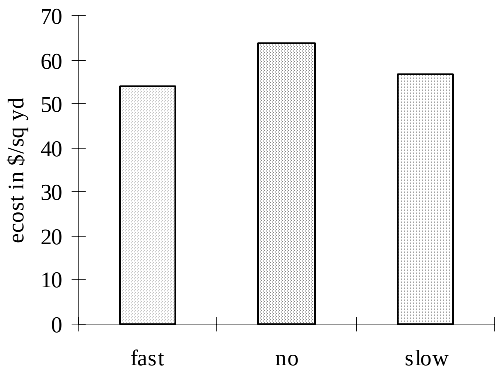

Figure 3. Expected cost: True Rate = Fast

As expected, in both instances –whether the true rate is Slow or Fast, when the initial beliefs match the actual rate, the expected cost is the lowest. Furthermore when the true rate is actually fast, all the expected costs are higher. What is surprising on the other hand is that the non-informative case is the worst case in both instances: it seems counter intuitive to have a lower expected cost when the initial beliefs are wrong. Similar qualitative observations can be found in Durango and Madanat (2002).

An explanation can be found in Figure 4., which presents the result of a simulation performed according to the optimal policy given by our model. As above, the initial beliefs about the state are $\underline{P_0 \mid I_0}$ = (0, 0.1, 0.1, 0.2, 0.4, 0.2, 0, 0) and the actual initial state is assumed to be 5. The actual rate is assumed to be fast. The beliefs about the rate are averaged over the 1,000 simulations. We plot the trajectory of the average Q_t^3=Pr(Actual Rate = Fast | information), where average stands for mean across instances of the experiment.

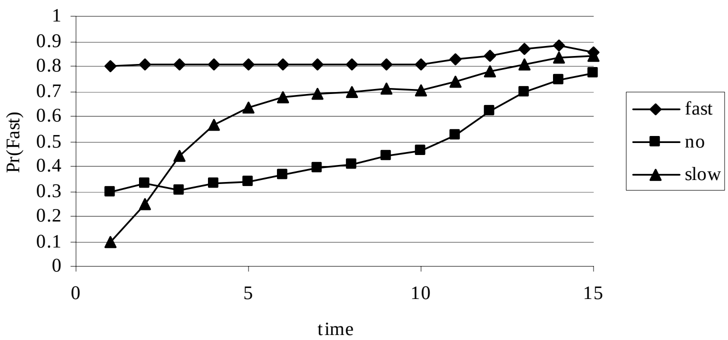

Figure 4. Trajectories of Q_t^3 - True Rate = Fast

As Q_t^3 converges much faster when the initial beliefs are “slow”, i.e. wrong, than when they are non-informative, the actions taken in the non-informative case cannot be as close to optimality as those taken when the initial beliefs are wrong. Hence the higher expected cost when the initial beliefs about the rate are more spread out.

The faster convergence of the beliefs in the wrong case compared to the non-informative case can be explained qualitatively by the contrast between the observations and the expectations. This contrast is augmented by the action taken in both cases: when the initial beliefs are wrong, i.e. “slow” the MR&R actions are expected to be mild compared to the “no” case. Therefore, low condition states are more likely to be observed than when a severe action is taken. Such unexpected outcomes provide feedback that leads to drastic revision of the beliefs in the “slow” case. Quantitatively, this contrast is expressed by the weight of the Slow deterioration rate in the denominator of the Bayesian updating formula being equal zero.

Açıklama [VG1]: give a formula #??

A primary objective of this research is to understand the relative role of uncertainties in deterioration rate and condition state of the facility. In order to investigate the effect of variance in the initial set of beliefs $\underline{Q_0}$ and $P_0|I_0$, we conducted a case study where $\underline{P_0 \mid I_0}$ can have a low or high variance around an initial state 5, i.e. $\underline{P_0 \mid I_0}$ = (0, 0, 0, 0, 1, 0, 0, 0) and $\underline{P_0 \mid I_0}$ = (0, 0.1, 0.1, 0.2, 0.4, 0.2, 0, 0). The initial uncertainty about the rate can be high, in other words $\underline{Q_0}$ =(0.3, 0.4, 0.3), or low. We only considered the effect of decreasing uncertainty about the deterioration rate in the correct direction, i.e. $\underline{Q_0}$ =(0.8, 0.1, 0.1) when the actual deterioration rate is Slow, and $\underline{Q_0}$ =(0.1, 0.1, 0.8) when it is actually Fast.

Figures 5 and 6 present the expected costs in the Slow and Fast deterioration rates respectively. The Δs represent the expected benefit in moving in the direction of the arrows. The double-lined arrows indicate the largest expected benefit in each case, if we start from the upper right corner of the figure, i.e. when both types of uncertainty are initially high.

<u>Expected cost in $/sq yd</u>

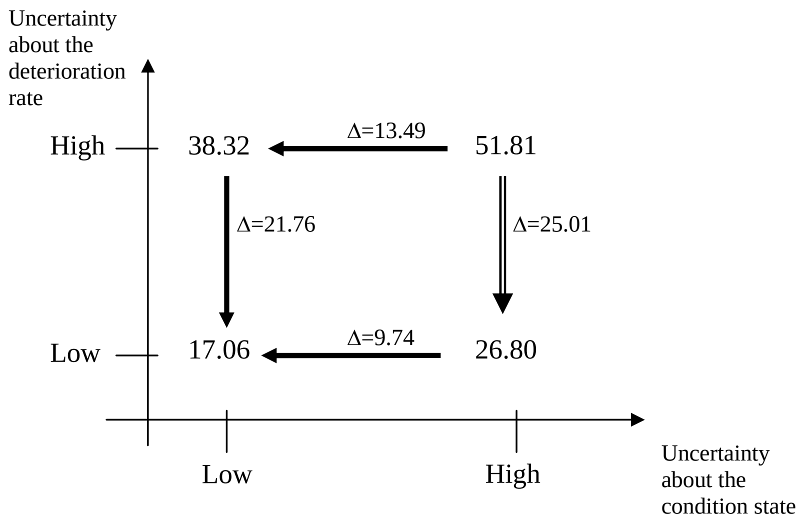

Figure 5. Relationship between expected cost and sources of initial uncertainty:
True Rate = Slow

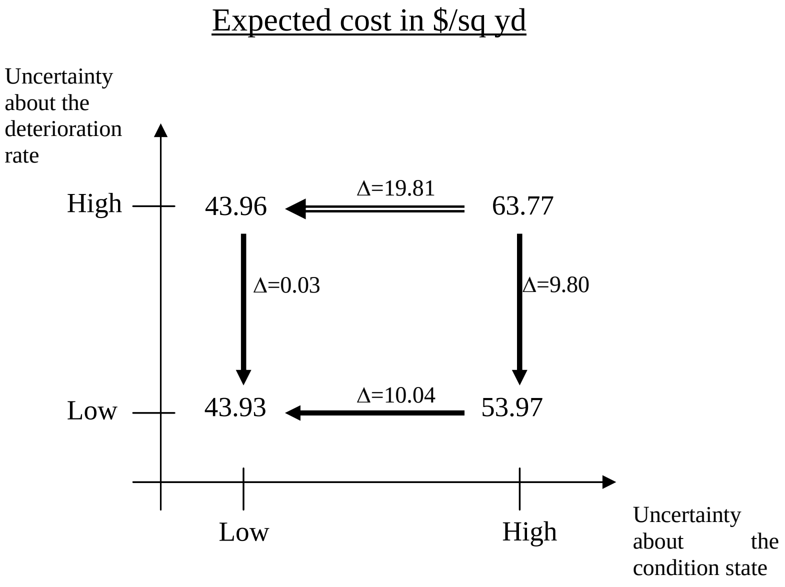

Figure 6. Relationship between expected cost and sources of initial uncertainty: True Rate = Fast

We can observe that a reduction in initial uncertainty always results in a decrease in the expected cost. If the rate is actually fast, there is more value in first decreasing the uncertainty about the state , as can be seen in Figure 6. This is because the beliefs about the rate converge faster when the actual rate is fast.

Açıklama [VG2]: see Durango

Assuming that the agency is initially in a situation where it has high uncertainty about both the rate and the state, Figure 5 shows that when the actual rate is Slow, reducing first uncertainty about the rate brings more value, whereas Figure 6 recommends a reduction in state uncertainty as the first step, since it achieves higher benefits. If we consider that the true deterioration rate is a variable revealed only at the end of the planning process, a cautious (Maximin) strategy vis-à-vis the value of information consists of reducing the uncertainty about the state first.

In practice, agencies have fairly advanced measurement technologies whereas they rarely have accurate and precise sets of performance forecasting models in their IMR planning decisions. It can be concluded from the above figures that the incorporation of improved performance forecasting models, that reduces the uncertainty about deterioration, in the planning process always provides value.

5. Conclusions

This paper has presented an IMS model that explicitly allows a flexible inspection schedule, takes into account measurement error, and includes feedback from measurements to improve the characterization of the deterioration rate.

The results show that the least expected cost is observed when the initial beliefs about the deterioration rate are correct, which is an intuitive result. However, the case of non-informative initial belief leads to the highest expected costs. The computational study clarified the role of the convergence of the beliefs about the deterioration rate in the expected cost. We therefore recommend that an agency should not initialize the implementation phase with the probability of all the performance models being equal. Even a wrong initialization would lead to smaller expected cost.

Finally, as the proposed model accounts for both state and performance model uncertainties, it was possible to determine the relative value of decreasing each source of uncertainty. Results showed that, if an agency is assumed to have high variance in both initial beliefs about the rate and the state, a cautious recommendation is to decrease uncertainty about the state first.

The scope of this research was intentionally limited to the facility-level of the IMR problem. An immediate extension is to see how the problem translates to the network-level, with budget constraints. The issue of implementation and "real-time" control can also be investigated in this context.

References

Bertsekas, D., 1987. Dynamic Programming: Deterministic and Stochastic Models, Prentice-Hall Inc., New York.

Carnahan, J., Davis, W., Shahin, M., Keane, P., Wu, M. 1987. Optimal Maintenance Decisions for Pavement Management, Journal of Transportation Engineering, 113 (5) 554-572.

Carnahan, J. 1988. Analytical Framework for Optimizing with Incomplete Information, Operation Research, 16.

Dreyfus, S., and Law, A., 1977. The Art and Theory of Dynamic Programming, Academic Press

Durango, P., Madanat, S., 2002. Adaptive Control Models in Infrastructure Management, forthcoming in Transportation Research Part A.

HRB (Highway Research Board), 1962. The AASHO Road Test, Report 5: Pavement Research, Special Report No. 61E.

Humplick, F., 1989. Theory and Methods of Analyzing Infrastructure Inspection Output: Application to Highway Pavement Surface Distress Evaluation, Ph.D. dissertation, Department of Civil Engineering, M.I.T.

Humplick, F.,1992. Highway Pavement Distress Evaluation: Modeling Measurement Error, Transportation Research, 26 B, 135-154.

Klein, M. 1962. Inspection-Maintenance-Replacement Schedules under Markovian Deterioration, Management Science, 9, 25-32.

Madanat, S., 1991. Optimizing Sequential Decisions under Measurement and Forecasting Uncertainty: Application to Infrastructure Inspection, Maintenance and Rehabilitation, Ph.D. dissertation, Department of Civil Engineering, M.I.T.

Madanat, S., Ben-Akiva, M., Optimal Inspection and Repair Policies for Infrastructure Facilities, Transportation Science, 28 (1) 55-62.

Madanat, S., and Wan Ibrahim, W.H., 1995. Poisson Regression Models of Infrastructure Transition Probabilities, Journal of Transportation Engineering, 121 (3).

Mine, H., Kawai, H., 1982. An optimal Inspection and Maintenance Policy of a Deteriorating System, Journal of the Operations Research Society of Japan, 25 (1).

Mishalani, R., and Madanat, S., 2001. Estimating Infrastructure Transition Probabilities using Stochastic Duration Models, submitted to the ASCE Journal of Infrastructure Systems.

Olsonen, R., 1988. Finland's Pavement Management System: Data and Models for Condition of Roads, Roads and Waterways Administration, Finland.

Shahin, M., Kohn, S. 1981. Pavement Maintenance Management for Roads and Parking Lots, Technical report M-29, Construction Engineering Research Lab, US Army Corps of Engineers.