Optimal Inspection and Maintenance Policies for Infrastructure Systems Under Measurement and Prediction Uncertainty

KAREN SMILOWITZ
SAMER MADANAT
University of California, Berkeley

ABSTRACT

State-of-the-art infrastructure management systems utilize Markov Decision Processes as a methodology for maintenance and rehabilitation (M&R) decision-making. The underlying assumption in this methodology is that an inspection is performed at the beginning of every year, and that inspections reveal the true condition state of the facility, with no error. As a result, after an inspection, the decision maker can apply the activity prescribed by the optimal policy for that condition state of the facility.

Previous research has developed a methodology for M&R activity selection, which accounts for the presence of both forecasting and measurement uncertainty. This methodology is the Latent Markov Decision Process (LMDP), an extension of the traditional MDP that does not assume the measurement of facility condition to be necessarily error-free.

In this paper, we extend this methodology to include network level constraints. This can be achieved by extending the LMDP model to the network-level problem through the use of randomized policies. We present both finite horizon (transient) and infinite horizon (steady state) formulations of the network-level LMDP. A case study application demonstrates the expected savings in life-cycle costs that result from increasing the measurement accuracy used in facility inspections, and from scheduling inspection decisions optimally.

INTRODUCTION

State-of-the-art infrastructure management systems utilize Markov Decision Processes as a methodology for maintenance and rehabilitation (M&R) decision-making (Golabi et al. 1982, Carnahan et al. 1987, Carnahan 1988, Feighan et al. 1988, Harper et al. 1990, Gopal and Majidzadeh 1991). In this methodology, facility condition at any time is measured by a discrete state and the deterioration process is represented by discrete transition probabilities of the form:

$$p\left(x_{t+1}=j \middle| x_t=i, a_t\right); \quad 1 \le i, j \le n; \quad t=0,\ldots,T-1 \tag{1}$$

where:
x_t = condition state of the facility at the beginning of year t,
i, j = indices of elements in the set of discrete conditions,
a_t = M&R activity performed during year t,
n = number of possible states the facility can be in,
T = number of years in the planning horizon.

Given these transition probabilities, the optimal M&R policy can be solved for using Dynamic Programming. The underlying assumption in this methodology is that an inspection is performed at the beginning of every year, and that inspections reveal the true condition state of the facility, with no error. As a result, after an inspection, the decision maker can apply the activity prescribed by the optimal policy for that condition state of the facility.

There are two major limitations in this approach. First, it assumes that inspections are error-free. Second, it depends on a fixed, not necessarily optimal inspection schedule (an inspection has to be performed at the beginning of every time period).

The first assumption has been demonstrated to be incorrect in several empirical studies (for example, Humplick 1992): there is substantial measurement uncertainty in infrastructure inspection. This uncertainty affects M&R decisions because a measurement error will lead to the selection of a "wrong" activity if the prescribed M&R activity for the true condition and that prescribed for the measured condition are different. The second assumption reflects the absence of a systematic methodology for making inspection decisions in the field of infrastructure. Current methods of inspection planning are ad-hoc and subjective (for example, Shahin and Kohn 1981).

In previous papers, a methodology for M&R activity selection, which accounts for the presence of both forecasting and measurement uncertainty, was presented. This methodology is the Latent Markov Decision Process (Madanat 1993, Madanat and Ben-Akiva 1994). The Latent Markov Decision Process (LMDP) is an extension of the traditional MDP, but differs from it in one major aspect: it does not assume the measurement of facility condition to be necessarily error-free (Eckles 1968, Smallwood and Sondik 1973).

In this paper, we extend this methodology to include network level constraints. This can be achieved by extending the LMDP model to the network-level problem through the use of "randomized policies." In the following section, an overview of the LMDP is presented. Section 3 presents the extension of this methodology to incorporate network level constraints for the case of a finite horizon problem. Section 4 extends the network level problem for the infinite planning horizon problem. Section 5 presents a case study on the application of this methodology. Section 6 discusses some practical aspects of these models and concludes the paper.

THE LATENT MARKOV DECISION PROCESS

When measurement uncertainty is introduced into the MDP, what the decision maker observes at the beginning of t becomes a measured state which is probabilistically related to the true state of the facility. This relation can be mathematically stated as:

$$q\left(\hat{x}_t = k | x_t = j\right) \quad 1 \le j,k \le n \quad t = 0,...,T-1 \tag{2}$$

where:
$\hat{x}_t$ = measured condition state of the facility at start of t,
x_t = true condition state of the facility at start of t,
j, k = indices of elements in the set of discrete condition states,
q = a known probability mass function.

These measurement probabilities can be obtained empirically, using the concept of measurement error models (see Humplick 1992, Ben-Akiva et al. 1993).

Due to the presence of measurement errors, the true state of the facility is no longer observed. Using state augmentation, a new state is defined to account for all the information available to the decision maker and relevant to future decisions (Bertsekas 1987). The information available to the decision maker at the beginning of t includes the entire history of measured states up to t and the decisions made up to $t-1$. Moreover, knowledge of the measured state is not sufficient for decision making. Thus, all the previous measured states and decisions can be relevant to future decisions, and have to be included in the augmented state. Denoting the new state by I_t, we have:

$$I_t = \{I_0, a_0, \hat{x}_1, ..., a_{t-1}, \hat{x}_t\} \quad t = 1,2,...T$$
$$I_0 = \{\hat{x}_{-\tau}, a_{-\tau}, ..., a_{-1}, \hat{x}_0\} \tag{3}$$

where:
τ = number of years between first inspection of facility and start of planning horizon. It follows that:

$$P(I_t | I_0, a_0, \hat{x}_1, ..., I_{t-1}, a_{t-1}) = P(I_t | I_{t-1}, a_{t-1}) \quad t = 1, ..., T \tag{4}$$

Assuming I_0 to be known, the transition probabilities $P(I_t|I_{t-1},a_{t-1})$ define the evolution of the state of information, and this evolution is Markovian.

We can thus write a Dynamic Programming formulation over the space of information states; for notational simplicity, this is done using the generic cost function $g(x,a)$. First, the cost function has to be rewritten in terms of the new variables. The cost per stage as a function of the new state, I, and of the activity a, is:

$$E_x\{g(x,a)|I\} = \sum_{t=1}^{n} p(x = i|I) * g(x,a) \tag{5}$$

To evaluate expression (5), we need to know the distribution of the true condition state at t conditional on I, that is:

$$p(x|I), \quad \forall x, \forall I \tag{6}$$

or, in vector form,

$$P|I, \quad \forall I \tag{7}$$

where:
$P|I$ = n-dimensional vector (the information vector).
If $P_0|I_0$ is known, then $P_t|I_t$ can be calculated recursively for all t, starting from $t = 1$ to $t = T$, using Bayes'' law, the known measurement probabilities and the known transition probabilities. Given $I_t = \{I_{t-1}, a_{t-1}, \hat{x}_t\}$, each element of $P_t|I_t$ can be calculated as follows:

$$p(x_t = j|I_t) = \frac{q(\hat{x}_t|x_t = j)\sum_i p(x_t = j|x_{t-1} = i, a_{t-1}) p(x_{t-1} = i|I_{t-1})}{\sum_k q(\hat{x}_t|x_t = k)\sum_i p(x_t = j|x_{t-1} = i, a_{t-1}) p(x_{t-1} = i|I_{t-1})} \tag{8}$$

It can be seen that the denominator of (8) yields the total probability of measuring $\hat{x}$ in time period t conditional on the information state in the previous time period (I_{t-1}) and the activity a performed in time period $t-1$. Since each combination of measured state $\hat{x}$ and activity a produces a new information state in time period t, the denominator of (8) is the transition probability from state I_{t-1} to state I_t conditional on activity a_{t-1}. We denote these transition probabilities by $P(I_t|I_{t-1}, a_t)$.

Using the elements $p(x_t|I_t)$ calculated above, the Dynamic Programming recursion over the space of information states can be written as:

$$J_t(I_t) = \min_a \left\{ \sum_i p(x_t = i|I_t) g(x_t, a_t) + \alpha \sum_i p(x_t = i|I_t) \left[\sum_j p(x_{t+1} = j|x_t = i, a_t) \right. \right.$$

$$\left. \left. \sum_k q(\hat{x}_{t+1} = k|x_{t+1} = j) * J_{t+1}(I_{t+1}) \right) \right] \right\} \qquad \forall I_t, t = 0, \ldots, T-1 \tag{9}$$

where:
α = discount amount factor, and
$I_{t+1} = \{I_t, a_t, k\}$

The model defined by this formulation is referred to as the Latent Markov Decision Process with annual inspections, because it assumes that the state of the facility is latent, and because it assumes that a measurement of facility condition is available at the start of every period t.

For computational reasons, it is of interest to replace I_t with a quantity of smaller dimension, but which has the same information content. Such a quantity is referred to as a "sufficient statistic" for I_t. Since, as observed earlier, I_t affects decisions only through $p(x_t|I_t)$, an ideal sufficient statistic is the information vector $P_t|I_t$.

The advantage of using the vector $P_t|I_t$ is that it allows for direct comparison among states at a given t. Whereas it is impossible to directly compare the states I_t, it is possible to compare the information vectors $P_t|I_t$, by pair-wise comparison of corresponding elements. When two states are found to have equal, or almost equal, values of $P_t|I_t$, they can be combined into a single state, which reduces the number of times equation (9) has to be applied. The subject of combination of different states is covered in (Madanat 1991) in more detail.

We assume that $P_0|I_0$ is given, in order to be able to calculate $P_t|I_t$, for all t, using (8). This initial vector, the prior distribution of the true states, represents the prior belief of the decision-maker regarding the state of the facility.

Solving program (9) for a specified planning horizon T will yield the minimum expected cost $J_0(P_0|I_0)$ and the optimal sequence of policies $\pi^* = \{\mu_0^*(P_0|I_0), \ldots, \mu_{T-1}^*(P_{T-1}|I_{T-1})\}$, where $\mu_t^*(Pt|It)$, the optimal policy for period t, specifies the optimal M&R activity for each possible information vector at time t, $P_t|I_t$. A "possible" information vector is one which can be reached with non-zero probability at time t, given $P_0|I_0$, and given the optimal policies up to t.

The components of the cost function $g(x,a)$ are:

- the cost of performing an activity, which depends on the type of activity and the condition state of the facility; this cost is denoted by $c(a, x)$;

- the user cost, which depends on the condition state of the facility; this cost is denoted by $u(x)$.

EXTENSION TO THE NETWORK LEVEL PROBLEM: THE FINITE HORIZON CASE

The type of policies produced by the LMDP (and by the ordinary MDP) are non-randomized policies. This is because, given a state of the system, the model specifies a single M&R activity. A randomized policy does not specify a single optimal activity for each state of the system. Instead, it specifies optimal probabilities for different activities for each state of the system.

At the facility level, the concept of the probabilities of different activities is ambiguous. On the other hand, if we are dealing with a number of facilities in the same information state, these probabilities can be interpreted as fractions. In other words, the optimal policies would specify the optimal fractions of activities to be applied to facilities in each state. The choice of the specific activity to be applied to each facility in this state is left to the engineer. This procedure recognizes that there exist other considerations in the choice of M&R activities which are not captured by the model; for example, materials and labor availability, traffic disruption, etc. Thus, it is necessary to allow for some flexibility in the model recommendations, which can be exploited by the engineer in charge of these facilities.

By using randomized policies, it is possible to include in the LMDP formulation budget constraints, a feature which would make it possible to use the new model for the solution of the network-level problem. This is achieved by formulating the problem as a Linear Program, where the decision variables are the fractions of activities to be applied to facilities in each state of the information. The concept of randomized policies has been used extensively in the context of classical MDP models; see Golabi et al (1982), Harper et al (1990) and Gopal and Majidzadeh (1991).

For a finite horizon of T periods, a Linear Program can be used to solve for the optimal fraction of facilities in each state of the information receiving each of the maintenance and rehabilitation activities. As in the facility-level LMDP, information states are represented by the vector I_t, or by the sufficient statistic $P_t|I_t$. The decision variables are the fractions of facilities in each information state I receiving activity a in time t, and are denoted by W_{aIt}. Included in the set of activities is the decision to inspect; each activity can be performed with inspection or without. Since not inspecting essentially yields the same information as inspecting with a technology of infinite variance, the decision to inspect is modeled as a choice between performing an inspection and incurring the additional cost of inspection or inspecting with a technology of infinite variance and incurring no inspection cost (Madanat and Ben-Akiva 1994).

The formulation of the Linear Program is as follows:

$$\min \sum_t \sum_I \sum_a \left[\sum_k \left(c(a,k) + u(k) \right) * p\left(x_t = k | I_t \right) \right] * W_{alt} \qquad (10)$$

Subject to:

$$W_{alt} \geq 0 \qquad \forall a, I, t \qquad (11)$$

$$\sum_a \sum_I W_{alt} = 1 \qquad \forall t \qquad (12)$$

$$\sum_a W_{aJ(t+1)} = \sum_a \sum_I W_{alt} * P\left(J_{t+1} | I_t, a_t\right) \qquad \forall J, t \qquad (13)$$

$$\sum_a \sum_I W_{alt} * p\left(x_t = k | I_t\right) \geq PMIN_k \qquad \forall k, t \qquad (14)$$

$$\sum_a \sum_I W_{alt} * p\left(x_t = k | I_t\right) \leq PMAX_k \qquad \forall k, t \qquad (15)$$

$$\sum_a \sum_I \left[\sum_k c(a,k) * p\left(x_t = k | I_t\right) \right] * W_{alt} \geq BMIN_t \qquad \forall t \qquad (16)$$

$$\sum_a \sum_I \left[\sum_k c(a,k) * p\left(x_t = k | I_t\right) \right] * W_{alt} \leq BMAX_t \qquad \forall t \qquad (17)$$

where:

W_{alt} = the fraction of facilities in information state I receiving activity a in time t,
$PMIN_k$ = minimum fraction of facilities allowed in condition state k,
$PMAX_k$ = maximum fraction of facilities allowed in condition state k,
$BMAX_t$ = budget maximum for time period t,
$BMIN_t$ = budget minimum for time period t.

The objective function (10) minimizes the total expected cost to both the agency and the users by selecting the optimal values of the decision variables, W_{alt}, over the planning horizon T.

The first constraint (11) limits all decision variables to non-negative values. The unity constraint (12) defines each W_{alt} as a fraction of the total number of facilities in the network. The Chapman-Kolmogorov equations (13) provide for conservation of facilities over time, where facilities move from one information state I_t to another state J_{t+1} with transition probabilities $P(J_{t+1}|I_t, a_t)$.

Additional constraints (14–17) impose network-level restrictions on the optimal solution. Performance level constraints (14,15) set acceptable ranges for the fractions of facilities in the best and the worst condition states. Agency budget constraints (16,17) require yearly expenditures to fall within a specified budget.

EXTENSION TO THE NETWORK LEVEL PROBLEM: THE INFINITE HORIZON CASE

It is desirable, from the highway agency's perspective, to achieve a constant distribution of network conditions and M&R expenditures. To satisfy this objective, a steady-state

formulation of the network level LMDP is used. In this formulation, it is assumed that after an initial transient period, an optimal steady-state distribution of network conditions and activity mix will be reached. A steady-state distribution of network conditions can be reached if the optimal policies used produce a probabilistic evolution of states that corresponds to an ergodic Markov chain (Gallager 1995). Therefore, an infinite planning horizon formulation is used, in which the objective is to minimize the expected cost per time period (Kulkarni 1984). In this infinite horizon case, the fractions of facilities in each information state that receives a given activity do not vary from year to year. The resulting Linear Program can be written without time indices. The formulation of the LP reduces to:

$$\min \sum_a \sum_I \left[\sum_k (c(a,k) + u(k)) * p(x = k|I) \right] * W_{aI} \tag{18}$$

Subject to:

$$W_{aI} \geq 0 \qquad \forall a, I \tag{19}$$

$$\sum_a \sum_I W_{aI} = 1 \tag{20}$$

$$\sum_a W_{aI} = \sum_a \sum_I W_{aI} * P(J|I,a) \qquad \forall J \tag{21}$$

$$\sum_a \sum_I W_{aI} * p(x = k|I) \geq PMIN_k \qquad \forall k \tag{22}$$

$$\sum_a \sum_I W_{aI} * p(x = k|I) \leq PMAX_k \qquad \forall k \tag{23}$$

$$\sum_a \sum_I \left[\sum_k c(a,k) * p(x = k|I) \right] * W_{aI} \geq BMIN \tag{24}$$

$$\sum_a \sum_I \left[\sum_k c(a,k) * p(x = k|I) \right] * W_{aI} \leq BMAX \tag{25}$$

The objective function (18) in the infinite horizon formulation minimizes the total expected cost for one year, since expenditures will be equal in each year in the steady-state. Again, decision variables are non-negative and must sum to one. With the absence of time indices, the Chapman-Kolmogorov equations (21) guarantee that the distribution of facilities will remain constant, although individual facilities may change states over time. Network level constraints (22–25) set the range of acceptable annual performance and budget levels.

CASE STUDY

Combining the finite and infinite horizon formulations, the network-level LMDP offers a complete methodology for optimizing M&R and inspection activities for a network of

facilities in the presence of measurement and forecasting uncertainty. The methodology was used to evaluate the expected life cycle cost reduction that results from a reduction in the uncertainty in the measurement output used in facility inspection.

Specification

In this case study, both transient and steady-state optimal M&R and inspection policies were determined for a network of highway pavements. Facility condition is represented by a set of three discrete states: 1, 2, and 3 (3 being the best condition). The M&R activities consisted of no action, routine maintenance and resurfacing. Associated with each M&R activity was a set of costs which varied by condition state and a set of condition state transition probabilities. These costs and transition probabilities are presented in Tables 1 to 4. At each time period, the M&R activity for the current period and the inspection decision for the next period are jointly selected. Table 4 presents the cost of inspection as well; this cost was assumed to be independent of measurement accuracy.

For computational simplicity, the information state space was discretized using an interval of 0.25. The allowable values of the conditional probabilities $p(x = k|I)$ were therefore limited to the following: 1.0, 0.75, 0.5, 0.25, and 0. As a result, there were only

Table 1: Deterioration Transition Matrix ("Do Nothing" Alternative)

$x(t+1)$		3	2	1
$x(t)$	3	0.5	0.5	
	2		0.5	0.5
	1			1

Table 2: Routine Maintenance Transition Matrix

$x(t+1)$		3	2	1
$x(t)$	3	0.8	0.2	
	2		0.8	0.2
	1			1

Table 3: Resurfacing Transition Matrix

$x(t+1)$		3	2	1
$x(t)$	3	1		
2		0.8	0.2	
1		0.6	0.4	

Table 4: Costs of M&R Alternatives and Inspection (in Dollars per Square Yard)

State	Routine Maintenance	Resurfacing	Inspection
3	5	20	0.1
2	8	20	0.1
1	15	20	0.1

15 points in the grid of the state space, P|I, in each time period. Each information vector computed by using equation (8) was mapped to the numerically closest of the fifteen predetermined information states. As a result, several of the actual information states are aggregated into a single point in the grid. This aggregation leads to a substantial reduction in computational costs, albeit at the cost of some loss of computational accuracy.

For the finite time horizon case, a planning horizon (T) of 15 years was used. The initial distribution of facilities reflects a fairly new network of pavements: 90% of facilities were in state 3, 9.75% in state 2 and 2.25% were in state 1. The maximum performance level constraint specified that, over all possible actions and information states, the probability of facilities being in state 3 must be at least 45%. The minimum performance constraint limited the probability of facilities being state 1 over all activities and information states to, at most, 5%. The range of acceptable budget levels and user costs was predetermined. In the finite planning horizon, these constraints were relaxed to allow for greater flexibility in the short term solution.

The optimal M&R policies were obtained for four levels of measurement accuracy. Measurement errors were assumed to be normally distributed with a mean of zero and four possible standard deviations: 0, 0.375, 0.75, and 1.125. Standard deviations are in units of condition states. Table 5 presents the four levels of accuracy evaluated in this case study.

Results

The Linear Program was solved with AMPL, a mathematical programming language, for both the finite and infinite horizon cases and the results are shown in Figures 1 and 2, respectively. Figure 1 shows the effect of measurement uncertainty on the minimum total expected cost for the transient problem, using a planning horizon of 15 years. Figure 2 shows the effect of measurement uncertainty on the minimum expected cost per year for the steady-state problem with an infinite planning horizon. The results show that, in both cases, the expected costs increase as measurement uncertainty increases. The slope of the expected cost curve yields a marginal benefit from decreasing measurement uncertainty. From these figures, one can quantify the value of better measurement uncertainty for both the short and long term planning horizons.

CONCLUSIONS

In this paper, we have presented a network-level extension of the Latent Markov Decision Process (Madanat and Ben-Akiva 1994). This extension is based on a Linear

Table 5: Precision of Measurement

Standard deviation	$q(\hat{x} = j-1 \mid x = j)$	$q(x = j \mid x = j)$	$q(x = j+1 \mid x = j)$
0	0	1	0
0.375	0.0925	0.815	0.0925
0.750	0.24	0.52	0.24
1.125	0.33	0.34	0.33

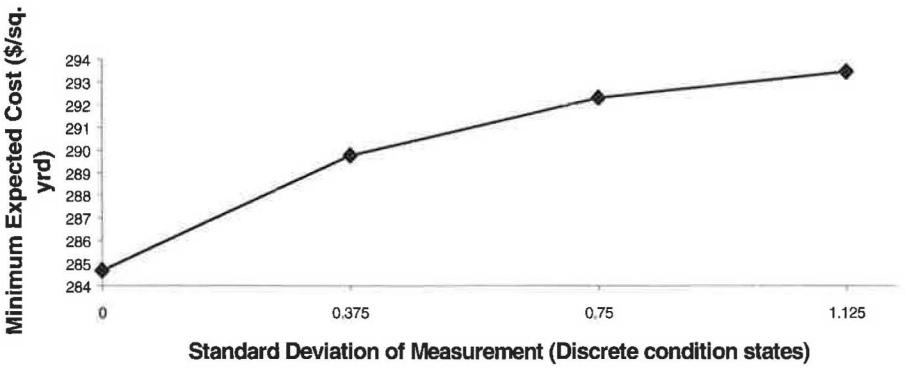

Figure 1: Rise in minimum expected cost with increase in measurement uncertainty (Finite Planning Horizon).

Programming formulation that uses randomized policies in conjunction with network level constraints, such as budget constraints and condition standards.

Two complementary versions of the joint optimal inspection and M&R policies problem were presented: a transient, finite-horizon formulation and a steady-state, infinite horizon formulation. An agency would use this methodology in the following manner:

- first, the infinite horizon problem would be solved to obtain the optimal steady-state distribution of facilities in each information state receiving each of the activities;
- then, using this optimal steady-state distribution as a target to be reached in T periods, the finite horizon problem would be solved to obtain the optimal transient distributions of facilities in each information state receiving each of the activities for each period from $t = 0$ to T.

A case study was performed to evaluate the effect of increasing the uncertainty in facility condition measurement on the minimum expected cost of both the finite horizon and infinite horizon problems. The results of the case study concur with earlier results in

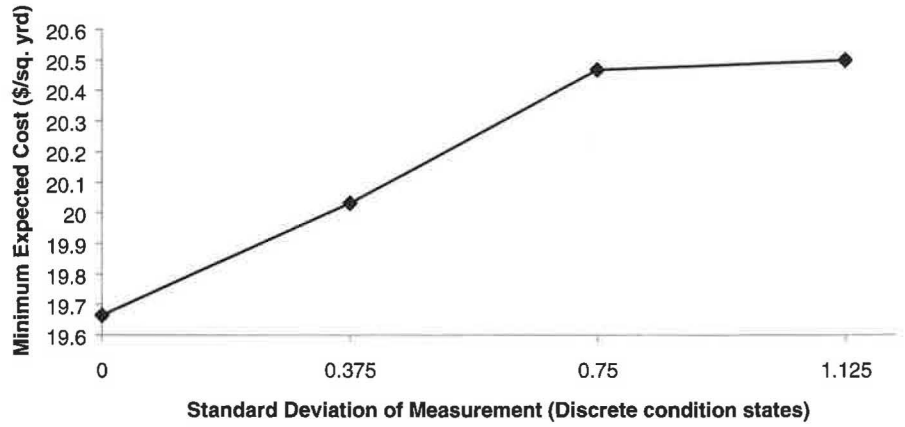

Figure 2: Rise in minimum expected cost with increase in measurement uncertainty (Infinite Planning Horizon).

the literature (Madanat 1993, Madanat and Ben-Akiva 1994) and with a priori expectations; even limited reductions in measurement error are expected to yield significant savings in the life cycle costs of a network of highway pavements. These results underscore the importance of improving the accuracy of infrastructure facility inspections.

REFERENCES

M. Ben-Akiva, F. Humplick, S. Madanat, and R. Ramaswamy, "Latent performance approach to infrastructure management," *Transportation Research Record 1311*, 188–195, 1991.

M. Ben-Akiva, F. Humplick, S. Madanat, and R. Ramaswamy, "Decision-making under uncertainty in infrastructure management: the latent performance approach," *Journal of Transportation Engineering*, Vol. 119, No. 1, ASCE, 1993.

D.P. Bertsekas, *Dynamic Programming, Deterministic and Stochastic Models*, Prentice-Hall Inc., Englewood Cliffs, NJ, 1987.

J.V. Carnahan, W.J. Davis, M.Y. Shahin, P.L. Keane, and M.I. Wu, "Optimal maintenance decisions for pavement management," *Journal of Transportation Engineering*, Vol. 113, 554–572, ASCE, 1987.

J.V. Carnahan, "Analytical framework for optimizing pavement maintenance," *Journal of Transportation Engineering*, Vol. 114, 307–322, ASCE, 1988.

J.E. Eckles, "Optimum maintenance with incomplete information," *Operations Research,* Vol. 16, 1058–1067, 1968.

K.J. Feighan, M.Y. Shahin, K.C. Sinha, and T.D. White, "Application of dynamic programming and other mathematical techniques to pavement management systems," *Transportation Research Record 1200,* 1988.

R. Gallager, *Discrete Stochastic Processes,* Kluwer Academic Press, Norwell, MA, 1995.

K. Golabi, R.B. Kulkarni, and G.B. Way, "A statewide pavement management system," *Interfaces*, Vol. 12, 5–21, 1982.

S. Gopal and K. Majidzadeh, "Application of Markov decision process to level-of-service-based maintenance systems," *Transportation Research Record 1304*, 1991.

W. Harper, J. Lam, A. Al-Salloum, S. Al-Sayyari, S. Al-Theneyan, G. Ilves, and K. Majidzadeh, "Stochastic optimization subsystem of a network-level bridge management system," *Transportation Research Record 1268*, 1990.

F. Humplick, "Highway pavement distress evaluation: modeling measurement error," *Transportation Research,* Vol. 26B, 135–154, 1992.

S. Madanat, "Optimizing Sequential Decisions under Measurement and Forecasting Uncertainty: Application to Infrastructure Inspection, Maintenance and Rehabilitation," Ph.D. Dissertation, Department of Civil Engineering, MIT, 1991.

S. Madanat, "Optimal infrastructure management decisions under uncertainty," *Transportation Research*, Vol. 1C, No. 1, 1993.

S. Madanat and M. Ben-Akiva, "Optimal inspection and repair policies for transportation facilities," *Transportation Science,* Vol. 28, No. 1, 1994.

R. Smallwood and E. Sondik, "The optimal control of partially observable Markov Processes over a finite horizon," *Operations Research*, 21, 1071–1088, 1973.