

UC Berkeley

Earlier Faculty Research

Title

Adaptive Optimization and Systematic Probing of Infrastructure System Maintenance Policies under Model Uncertainty

Permalink

<https://escholarship.org/uc/item/4fb7k5rc>

Authors

Madanat, S M
Park, Sejung
Kuhn, K D

Publication Date

2006-09-01

Adaptive Optimization and Systematic Probing of Infrastructure System Maintenance Policies under Model Uncertainty

Samer Madanat¹; Sejung Park²; and Kenneth Kuhn³

Abstract: We present an application of systematic probing for selecting optimal maintenance, repair, and reconstruction (MR&R) policies for systems of infrastructure facilities under model uncertainty. We use an open-loop feedback control approach, where the model parameters are updated sequentially after every inspection round. The use of systematic probing improves the convergence of the model parameters by ensuring that all permissible actions are applied to every condition state. The results of the parametric analyses demonstrate that the MR&R policies converge earlier when systematic probing is used. However, the savings in the expected total costs as a result of probing are minor, and are only realized when the optimal probing fractions are used. On the other hand, the additional costs incurred when the wrong probing fractions are used are significant. The major conclusion from this work is that state-of-the-art adaptive infrastructure management systems, that do not use probing, provide sufficiently close to optimal policies.

DOI: 10.1061/(ASCE)1076-0342(2006)12:3(192)

CE Database subject headings: Infrastructure; Maintenance; Optimization; Probabilistic models; Adaptive systems; Uncertainty principles.

Introduction

Infrastructure management systems (IMs) are decision-support tools that aid transportation and public works agencies in planning the maintenance activities of their facilities. IMs support the following tasks: Inspection of facilities, prediction of the deterioration of facilities through performance models, and selection of optimal maintenance, repair, and reconstruction (MR&R) policies over the planning horizon.

Facility deterioration is a probabilistic process, often represented by a stochastic model. In existing IMs, the parameters of these models were developed on the basis of expert judgment or empirical observations of select distresses, and assumed to be constant over the planning horizon. Expert judgment or empirical observations of select distresses, however, may not accurately represent the true facility deterioration process. This may lead to erroneous predictions and inappropriate MR&R policies, which result in increasing the total lifecycle costs for the agency and users. On the other hand, available condition data, collected

during the life of the facilities, can be used to improve the accuracy of the deterioration models.

The majority of state-of-the-art IMs use Markov decision processes as MR&R policy decision-making modules. In a MDP, deterioration is represented by transition probabilities. The transition probabilities can be determined based on expert judgment (Harper and Majidzadeh 1991) or empirical observations. In the latter case, statistical estimation with time series or panel data was used in Carnahan et al. (1987), Madanat et al. (1995), and Mishalani and Madanat (2002).

The objective of optimization models is to minimize the total costs associated with MR&R activities and user costs. Optimal maintenance policies are obtained through dynamic programming (DP) or linear programming (LP). DP has been used for single facility problems (Feighan et al. 1988; Madanat 1993; Madanat and Ben-Akiva 1994; Durango and Madanat 2002), and LP has been utilized for network-level problems (Golabi et al. 1982; Harper and Majidzadeh 1991; Smilowitz and Madanat 2000). This paper focuses on the network-level problem. There is uncertainty in the parameters of models used to represent facility deterioration. This is due to incomplete information regarding construction quality and material composition of the facilities, limited sample sizes of the data sets used for deterioration model calibration, and the use of laboratory testing as a substitute for field data when the latter are not available.

The developers of modern IMs have recognized the presence of uncertainty in the deterioration models used in practice. Therefore, they have included a model-updating step in these management systems, where data collected as part of condition surveys are used to update deterioration model parameters. For example, Harper and Majidzadeh (1991) used Bayesian methods to update the parameters of their deterioration models. Likewise, in the popular bridge management system Pontis, transition probability matrices are updated over time (Golabi and Shepherd 1997). Durango and Madanat (2002) proposed a decision-support

¹Professor, Dept. of Civil and Environmental Engineering, Univ. of California, Berkeley, 110 McLaughlin Hall, Berkeley, CA 94720 (corresponding author). E-mail: madanat@ce.berkeley.edu

²Dept. of Civil and Environmental Engineering, Univ. of California, Berkeley, 116 McLaughlin Hall, Berkeley, CA 94720. E-mail: sejpark@berkeley.edu

³Dept. of Civil and Environmental Engineering, Univ. of California, Berkeley, 116 McLaughlin Hall, Berkeley, CA 94720. E-mail: kuhn@berkeley.edu

Note. Discussion open until February 1, 2007. Separate discussions must be submitted for individual papers. To extend the closing date by one month, a written request must be filed with the ASCE Managing Editor. The manuscript for this paper was submitted for review and possible publication on May 6, 2004; approved on April 19, 2006. This paper is part of the *Journal of Infrastructure Systems*, Vol. 12, No. 3, September 1, 2006. ©ASCE, ISSN 1076-0342/2006/3-192-198/\$25.00.

system, where the uncertainty in the deterioration model is represented by a probability mass function of deterioration rates. Rather than updating the parameters of the deterioration model, it is this probability mass function of deterioration rates that is updated in light of inspection data in their system. Irrespective of these differences, the common element in these three systems is that they are adaptive; in that they combine model updating and optimization.

Broadly speaking, there are two types of adaptive optimization routines. Open-loop feedback control (OLFC) methods alternate between updating model parameters and optimizing decision making with respect to the most recent estimates of parameters. Closed-loop control improves on this methodology by explicitly considering the future updating of deterioration model parameters within present time MR&R optimization. Unfortunately, consideration of the many ways a network of facilities may deteriorate and how this will lead to different updated deterioration model parameters is not possible within a LP framework. For this reason, adaptive infrastructure management systems that deal with a network of related facilities are based on OLFC approaches.

In the present paper, model uncertainty is modeled by treating the transition probabilities as continuous random parameters. The successive updating of deterioration model parameters improves the representation of the actual deterioration process, only if the transition probabilities converge to their true values. For this to happen, a large number of state transitions have to be observed for every state and MR&R action combination. This means that all MR&R activities must be applied to every condition state a sufficient number of times. This may not happen in adaptive IMSs because the optimization process will tend to select only a subset of MR&R activities to apply to each condition state. As a result, the transition probabilities for state-action pairs that are not selected a sufficient number of times may converge to incorrect values. This is a limitation of all OLFC-based adaptive optimization models that exist in the infrastructure management literature, but one that has not been addressed systematically before.

One way of resolving this issue is to randomly assign all activities to a small number of facilities, thus guaranteeing that all activities are applied to every condition state. This approach is known as systematic probing. In this paper, an approach for the optimization of MR&R policies under model uncertainty by using systematic probing is presented.

This paper is organized as follows. The following section describes the MR&R optimization formulation in more detail. The subsequent section, entitled “Updating Transition Matrices and OLFC,” explains the method used for updating the deterioration models with field data and the OLFC approach. The next section, entitled “Systematic Probing” describes the application of systematic probing. Finally, the last section—entitled “Parametric Study”—shows the result of a computational study and discusses the usefulness of this approach.

MDP-Based MR&R Optimization Model

The key assumption of MDP-based MR&R optimization is that facility deterioration is a Markovian process, which means that the state of the facility in a period only depends on the state of the facility and the MR&R action taken in the preceding period. The deterioration process is represented by transition probabilities. It is assumed that inspections are performed periodically, and that the true state of a facility is revealed through inspections. The following model represents the transition probabilities:

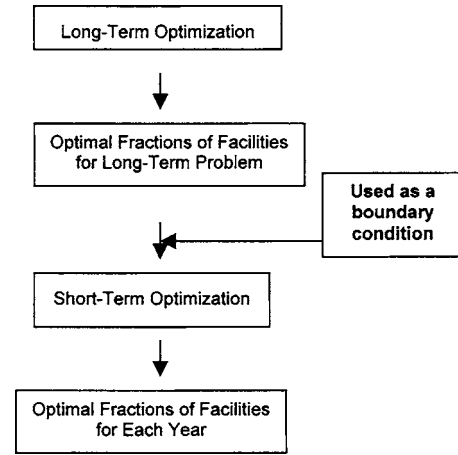


Fig. 1. Relation between long-term and short-term optimization

$$\pi_{ij}(a) = P(s^{t+1} = j | s^t = i, c^t = a) \quad (1)$$

$$1 \leq i, j \leq K, t = 0, 1, \dots, T - 1, a \in A$$

where $\pi_{ij}(a)$ =transition probability of the facility changing state i to j under maintenance activity a ; s^t =condition state of a facility in year t ; i and j =indices of the state of a facility; k =number of discrete states of a facility; c^t =maintenance activity performed in year t ; a =index of a maintenance activity; A =set of maintenance; and T =number of years in the planning horizon.

The transition probabilities are arranged in transition matrices. A transition matrix is shown next

$$\underline{\Pi}(a) = \begin{bmatrix} \pi_{11}(a) & \pi_{12}(a) & \dots & \pi_{1K}(a) \\ \pi_{21}(a) & \pi_{22}(a) & \dots & \pi_{2K}(a) \\ \dots & \dots & \dots & \dots \\ \pi_{K1}(a) & \pi_{K2}(a) & \dots & \pi_{KK}(a) \end{bmatrix} \quad (2)$$

where $\underline{\Pi}(a)$ =transition matrix given that maintenance activity a is performed.

The optimal MR&R policy for a network of similar infrastructure facilities can be obtained by using LP. The objective of this network-level optimization is to minimize the expected cost associated with performing MR&R activities, and the associated user costs subject to budget and level of service constraints. The decision variables are the fractions of the facilities in the network that are in various states, and to which different MR&R actions should be applied. The LP provides randomized MR&R policies rather than a single deterministic policy for each state. This approach was first proposed by Golabi et al. (1982).

Two types of LP optimizations are solved: A long-term and a short-term optimization. The long-term optimization model is based on an infinite planning horizon. It seeks optimal policies that minimize the average cost per period for a steady-state distribution of the facilities and maintenance activities. The short-term optimization minimizes the total cost over a predetermined and finite planning horizon. In the short-term optimization model, the steady-state distribution obtained in the long-term optimization is used as a boundary condition for the distribution of facilities and MR&R activities at the end of the planning horizon. This is illustrated in Fig. 1. The two models are described next.

Long-Term Optimization

$$\text{Min.} \quad \sum_{i=1}^K \sum_{a \in A} w(i,a) \cdot [g(i,a) + u(i) + r] \quad (3)$$

$$s.t. \quad w(i,a) \geq 0 \quad \forall i,a \quad (4)$$

$$\sum_{i=1}^K \sum_{a \in A} w(i,a) = 1 \quad (5)$$

$$\sum_{a \in A} w(j,a) = \sum_{i=1}^K \sum_{a \in A} w(i,a) \cdot \pi_{ij}(a) \quad \forall j \quad (6)$$

$$B_{\min} \leq \sum_{i=1}^K \sum_{a \in A} w(i,a) \cdot g(i,a) \leq B_{\max} \quad (7)$$

$$Q_{\min,i} \leq \sum_{a \in A} w(i,a) \leq Q_{\max,i} \quad \forall i \quad (8)$$

where $w(i,a)$ =fraction of facilities in the infrastructure network that is in state i and receives maintenance activity a ; $g(i,a)$ =cost associated with performing MR&R activity a when a facility is in state i ; $u(i)$ =user cost when a facility is in state i ; r =inspection cost; B_{\min} =lower limit of budget; B_{\max} =upper limit of budget; $Q_{\min,i}$ =lower limit of the fraction of facilities allowed in state i ; and $Q_{\max,i}$ =upper limit of the fraction of facilities allowed in state i .

The objective function given by Eq. (3) minimizes the expected cost per year. Listed in Eqs. (4)–(8) are the constraints necessary for this minimization problem. Constraints (4) and (5) specify that each decision variable (i.e., each fraction of facilities) should be non-negative, and that the sum of all the fractions should be equal to 1.0. Constraint (6) shows the Chapman-Kolmogorov equation. This equation, in this long-term problem, forces the distribution of facilities to remain constant over time. The budget constraints are given by Eq. (7), which allow both minimum and maximum values. The level of service constraint (8) forces the condition of the system to fall in an acceptable range. This constraint means that a minimum fraction of facilities, defined in the lower limit, $Q_{\min,i}$, should be in good states; and that the fraction in poor states should not exceed the upper limit, $Q_{\max,i}$. Inspection cost is a constant in this formulation, but could be modified to be a function of condition state or management policy. The inclusion of inspection cost is meant to draw attention to its importance in the asset management process.

It should be noted that user costs are typically not included in these formulations, because they are difficult to quantify. This is the reason that level of service constraint (8) is used, so as to incorporate a policy-makers' desire for having pavements in acceptable condition without introducing the complications of user costs.

Short-Term Optimization

$$\text{Min.} \quad \sum_{t=1}^T \sum_{i=1}^K \sum_{a \in A} w^t(i,a) \cdot [g(i,a) + u(i) + r] \quad (9)$$

$$s.t. \quad w^t(i,a) \geq 0 \quad \forall i,a,t=1,2,\dots,T \quad (10)$$

$$\sum_{i=1}^K \sum_{a \in A} w^t(i,a) = 1 \quad \forall t=1,2,\dots,T \quad (11)$$

$$\sum_{a \in A} w^1(i,a) = q^0(i) \quad \forall i \quad (12)$$

$$\sum_{a \in A} w^{t+1}(j,a) = \sum_{i=1}^K \sum_{a \in A} w^t(i,a) \cdot \pi_{ij}(a) \quad \forall j,t=1,2,\dots,T \quad (13)$$

$$\sum_{i=1}^K \sum_{a \in A} w^T(i,a) \cdot \pi_{ij}(a) = \sum_{a \in A} w(j,a) \quad \forall i \quad (14)$$

$$B'_{\min} \leq \sum_{i=1}^K \sum_{a \in A} w^t(i,a) \cdot g(i,a) \leq B'_{\max} \quad \forall t=1,2,\dots,T \quad (15)$$

$$Q'_{\min,i} \leq \sum_{a \in A} w^t(i,a) \leq Q'_{\max,i} \quad \forall i,t=1,2,\dots,T \quad (16)$$

where $w^t(i,a)$ =fraction of facilities that is in state i and receives maintenance activity a in year t ; $q^0(i)$ =initial fraction of facilities in state i ; $w(i,a)$ =steady-state fraction of facilities in the infrastructure network that is in state i and receives maintenance activity a , obtained from the long-term optimization; B'_{\min} =lower limit of budget for year t ; B'_{\max} =upper limit of budget for year t ; $Q'_{\min,i}$ =lower limit of the fraction of facilities allowed in state i for year t ; $Q'_{\max,i}$ =upper limit of the fraction of facilities allowed in state i for year t .

The LP formulation for the short-term optimization model is similar to the long-term optimization model. The difference is that the decision variables and the budget and level of service constraints are time dependent in the short-term model. This time dependency is denoted by the superscript t in the model. The objective function (9) minimizes the total cost over the short-term horizon. The required constraints for this minimization problem are listed in Eqs. (10)–(16). Constraint (10) guarantees the non-negativity of fractions. The sum of the fractions should sum to 1.0, per constraint (11). Constraint (12) guarantees that the fraction of facilities in state i in year 1 is equal to $q^0(i)$, which is known. The Chapman-Kolmogorov equation (13) guides the deterioration processes of the different facilities after year 1. The optimal solution obtained in the long-term optimization, $w(i,a)$, acts as a boundary condition on the distribution of states at the end of the planning horizon in constraint (14). The budget constraint and level of service constraint are stated in Eqs. (15) and (16). Again, user costs are usually not included in the formulation if a level-of-service constraint is used.

The formulations above assume that all facilities can be represented by the same set of transition probabilities, i.e., that we have a homogeneous network. This simplification was made for clarity of exposition.

In reality, an agency solves a number of MR&R optimizations, one for each homogeneous group of facilities, and each with its own budget constraint. To find the optimal allocation of the total budget across groups, an agency uses the following procedure: Each group's optimization is solved for a range of budget constraints. For each group, the change in the value of the objective function at optimality for a unit change of the budget is the value

of the shadow price of the budget constraint at that point. The optimal allocation of the total budget across groups is attained at the points where the shadow prices are equal across groups, and the groups' budget constraints sum up to the total agency budget.

Updating Transition Matrices and OLFC

This section describes how the transition probabilities can be updated with new data provided by inspections. It is assumed that an inspection of all facilities is performed at the beginning of each year, revealing the true condition of the facilities.

Notation

- $x_{ij}^t(a)$: Number of facilities that are in state i at the beginning of year t to which maintenance activity a is applied, and are in state j at the beginning of year $t+1$;
- $X_i^t(a) = \sum_{j=1}^k x_{ij}^t(a)$; $\forall i, a, t$;
- $n_{ij}^t(a) = \sum_{\tau=0}^t x_{ij}^{\tau}(a)$; $\forall i, j, a, t$;
- $N_i^t(a) = \sum_{\tau=0}^t X_i^{\tau}(a)$; $\forall i, a, t$;
- $\pi_{ij}^t(a)$: Estimated transition probability from state i to j under maintenance activity a in year t ; and
- $\underline{\Pi}^t(a)$: Estimated transition matrix for maintenance activity a in year t .

Through inspections at the beginning of year $t+1$, we observe the number of facilities whose conditions change from state i in year t to state j in year $t+1$ under maintenance activity a . This is denoted by $x_{ij}^t(a)$ for all i, j, t , and a . Then, $n_{ij}^t(a)$ and $N_i^t(a)$ are calculated. With this information, the transition probabilities are updated at the beginning of year $t+1$ by maximum likelihood estimation (MLE). The MLE for transition probabilities, at the beginning of year $t+1$, is

$$\pi_{ij}^t(a) = \frac{\sum_{k=0}^t x_{ij}^k(a)}{\sum_{k=0}^t X_i^k(a)} = \frac{n_{ij}^t(a)}{N_i^t(a)} \quad \forall i, j, a, t \quad (17)$$

The Bayesian updating of the transition probabilities provides the same result as in Eq. (17) (DeGroot 1970).

Procedure

For year t ,

- Inspect all facilities;
- Observe $x_{ij}^{t-1}(a)$; $\forall i, j, a$;
- Update the transition probabilities

$$\pi_{ij}^{t-1}(a) = \frac{n_{ij}^{t-1}(a)}{N_i^{t-1}(a)} \quad \forall i, j, a \quad (18)$$

- Obtain the updated transition matrices $\underline{\Pi}^{t-1}(a)$; $\forall a$.

This process is summarized in Fig. 2.

OLFC

An OLFC approach was used to incorporate the updated transition probability matrices into the MR&R optimization models. In the OLFC, short-term and long-term optimizations are rerun after transition matrices are updated. At the start of each year, the set of optimal MR&R policies over the planning horizon is obtained based on the updated transition matrices, but only the optimal policy for the current year is performed. It may seem odd to be constantly rerunning "long-term optimizations," but this ensures that actions are always taken in accordance with current best

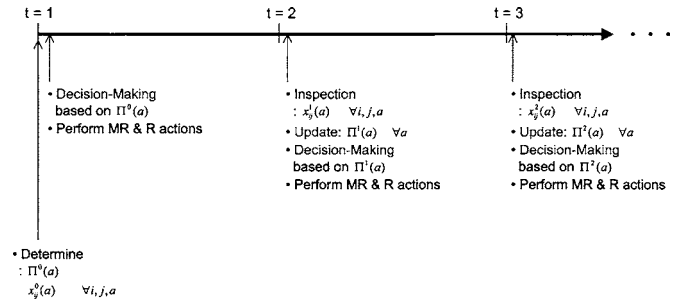


Fig. 2. Performance model updating procedure

estimates of transition probabilities. As transition probabilities converge, so too will optimal long-term policies. The OLFC algorithm is described next.

Algorithm

For year t ,

- Inspect all facilities at the beginning of year t ;
- Update transition matrices;
- Use the updated transition matrices in the LP to determine the fractions of the facilities in the network in different states on which maintenance activity a is performed for all years in the planning horizon

$$\Rightarrow \{w^t(i, a), \dots, w^T(i, a)\}; \quad \forall i, a; \text{ and}$$

- Apply $w^t(i, a) \quad \forall i, a$ for the network in year t only.

The linear program that is applied in the OLFC is the same as the linear program shown earlier. The only difference is that the transition matrix is updated every year. Fig. 3 illustrates the OLFC optimization procedure.

Systematic Probing

The successive updating of deterioration models will improve the representation of the real deterioration process only if the transition probabilities converge to the true values. As shown in Kumar and Varaiya (1986) and Bertsekas (2000), it is possible that, in certain cases, the parameters converge to the wrong values.

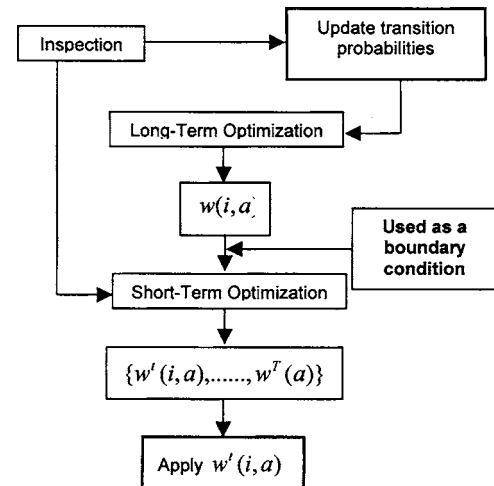


Fig. 3. Open-loop feedback MR&R optimization

Maximum likelihood produces consistent estimates. Therefore, the transition probabilities for a policy that is selected periodically by the optimization will converge to the true values. In other words, if each MR&R action is selected a large number of times for every condition state, then convergence to the true values is guaranteed. On the other hand, there is no similar guarantee for the transition probabilities corresponding to policies that are not selected periodically by the optimization.

The problem with OLFC methods used in adaptive infrastructure systems is that they always select actions that are “optimal” given the current deterioration model, and never select actions that may appear suboptimal, but that produce data that may reveal more information about systems dynamics (i.e., the effectiveness of various MR&R actions). In the language of Sutton and Barto (1998), at the core of adaptive control, there is a dilemma between exploitation and exploration. OLFC methods overlook the latter in favor of the former. If the initial performance model does not select the true optimal management policies, the optimization will tend to favor alternate policies. Data will be stockpiled about these alternate policies. By contrast, the true optimal policies will be avoided and little knowledge will be gained about their effectiveness. In circumstances like this, the infrastructure management system will take a relatively long time to identify the true optimal management policies, if it ever does so.

The solution is to ensure that all MR&R actions are applied to every condition state so that all transition probabilities are updated a sufficient number of times. This is especially important in the early years of infrastructure management; the faster information is gathered, the faster true optimal policies can be identified. One way to balance exploitation and exploration is to use systematic probing (Sutton and Barto 1998).

In this paper, the early period of the planning horizon is assigned for probing. During the probing period, every MR&R action is applied to every state, in order to produce a sufficiently large number of observations to update all transition probabilities. In year t , $(1-\epsilon^t) \times 100\%$ of the facilities receive the MR&R actions that are selected by the optimization while the remaining $\epsilon^t \times 100\%$ receive random MR&R actions. The probing fractions ϵ^t can either be a decreasing function of t or a constant. As a result of probing, all transition probabilities are updated frequently at the beginning of the planning horizon because all MR&R actions have been applied to all states a relatively large number of times.

From a practical point of view, it is difficult to imagine an agency applying randomly selected MR&R policies to even a small subset of its infrastructure facilities. But these random policies need not be applied to in-service facilities. Many state highway agencies in the U.S., and many in other industrialized countries, have developed full-scale experimental test facilities for research purposes. Examples include the MnRoad test site in Minnesota, which comprises experimental test sections and in-service pavement sections; the WesTrack test site in Nevada, a full-scale experiment that consists of a specially built track; and test sections subjected to heavy vehicle simulators in California, Indiana, Texas, and other states. As part of these experiments, researchers can, and in some cases already do, apply randomly selected MR&R activities to such pavement sections.

Parametric Study

A parametric study was conducted to compare the OLFC optimization with and without systematic probing. The savings achieved

Table 1. Cost (\$/Lane-Yard)

| Condition of pavement | MR&R activities | | | | | | | User costs |
|-----------------------|-----------------|------|-------|-------|-------|-------|-------|------------|
| | 1 | 2 | 3 | 4 | 5 | 6 | 7 | |
| 1 | 0.00 | 6.90 | 19.90 | 21.81 | 25.61 | 29.42 | 25.97 | 100.00 |
| 2 | 0.00 | 2.00 | 10.40 | 12.31 | 16.11 | 19.92 | 25.97 | 25.00 |
| 3 | 0.00 | 1.40 | 8.78 | 10.69 | 14.49 | 18.30 | 25.97 | 22.00 |
| 4 | 0.00 | 0.83 | 7.15 | 9.06 | 12.86 | 16.67 | 25.97 | 14.00 |
| 5 | 0.00 | 0.65 | 4.73 | 6.64 | 10.43 | 14.25 | 25.97 | 8.00 |
| 6 | 0.00 | 0.31 | 2.20 | 4.11 | 7.91 | 11.72 | 25.97 | 4.00 |
| 7 | 0.00 | 0.15 | 2.00 | 3.91 | 7.71 | 11.52 | 25.97 | 2.00 |
| 8 | 0.00 | 0.04 | 1.90 | 3.81 | 7.61 | 11.42 | 25.97 | 0.00 |

as a result of using systematic probing are quantified.

We used a case study in the field of pavement management. The data, such as transition probability matrices, costs of performing MR&R activities, and user costs, were taken from Durango and Madanat (2002). The user costs represent the vehicle operating costs associated with the condition of pavement. The condition of pavements was expressed by discrete states from one to eight, as in Carnahan et al. (1987). A higher number means a better condition. Seven MR&R activities exist: Do nothing (1), routine maintenance (2), one-in overlay (3), two-in overlay (4), four-in overlay (5), six-in overlay (6), and reconstruction (7). The costs associated with performing MR&R activities and the user costs are shown in Table 1, which were taken from Carnahan et al. (1987). The user costs represent vehicle operation costs that depend on the condition of pavement. The inspection cost was assumed to be 10% of the average MR&R costs. For the short-term optimization, the planning horizon (T) was assumed to be 25 years.

The following initial distribution of facilities was assumed: 60% of facilities are in State 8, 20% in State 7, 10% in State 6, 5% in State 5, 3% in State 4, 2% in State 3, and 0% in States 2 and 1. The minimum level of service constraint restricts, such that no facilities should be in State 1. The budget was limited to \$100/lane-yard-year. Like inspection costs, the inclusion of a budget constraint aims to draw attention to its importance in network-level adaptive infrastructure management.

For computational reasons, we included some tolerance limits in constraint (14) so that the steady-state distribution of facilities and maintenance activities is attained with the specified tolerance. We used a tolerance ϕ of 0.01. Constraint (14) was modified to

$$\sum_{a \in A} w^T(i, a) \cdot \pi_{ij}(a) \geq \sum_{a \in A} w(i, a) \cdot (1 - \phi) \quad \forall i \quad (19)$$

$$\sum_{a \in A} w^T(i, a) \cdot \pi_{ij}(a) \leq \sum_{a \in A} w(i, a) \cdot (1 + \phi) \quad \forall i \quad (20)$$

OLFC Optimization without Systematic Probing

The basic logic of the computational study is that the initial transition matrices do not represent the true deterioration process, and the infrastructure facilities perform according to their true deterioration characteristics. The initial and true matrices were selected from among three categories of deterioration of Durango and Madanat (2002): Slow, medium, and fast. For example, in one trial, the slow deterioration rate was selected as the initial process estimate; while the real process was described by the

Table 2. Total Costs with Constant Probing Fraction

| Deterioration process | | Probing fraction ($\epsilon \times 100\%$) | | | | | |
|-----------------------|---------|--|--------|-------|--------|--------|--------|
| Real | Initial | 0% | 2% | 4% | 6% | 8% | 10% |
| Fast | Medium | 101.54 | 102.75 | 99.31 | 104.29 | 105.47 | 107.89 |
| | Slow | 84.34 | 86.48 | 81.59 | 86.22 | 89.47 | 92.94 |
| Medium | Fast | 55.78 | 59.04 | 55.77 | 60.89 | 65.71 | 70.49 |
| | Slow | 41.89 | 47.83 | 49.56 | 52.03 | 54.21 | 57.68 |
| Slow | Fast | 47.14 | 46.63 | 50.25 | 51.69 | 55.72 | 58.94 |
| | Medium | 35.49 | 35.44 | 38.05 | 42.8 | 45.11 | 47.39 |

medium deterioration rate. In this case, we predicted the performance of infrastructure facilities, and selected MR&R activities with the matrices corresponding to the slow deterioration in the first year. The optimal policy for the first year was applied, but the deterioration of infrastructure facilities after implementing the optimal policy was simulated based on the medium transition matrices. The inspection results at the beginning of the second year were generated using the medium transition matrices. We updated the initial transition matrices with the information from the inspections. Then, the optimal MR&R policies for the second year were selected by using the updated transition matrices. The selected activities were performed and the results of inspection at the beginning of the following year were generated based on the medium transition matrices. This procedure was repeated similarly over the planning horizon. The total cost was the sum of costs actually incurred over the planning horizon.

Open-Loop Feedback Control Optimization with Systematic Probing

The basic logic is the same as the OLFC optimization without systematic probing. However, as explained earlier, during the probing period, $\epsilon \times 100\%$ of the facilities receive random actions. The facilities that receive random actions are evenly chosen from all states. Two types of systematic probing were used in the computational study: A constant probing fraction strategy and diminishing probing fraction strategy. With the constant probing fraction strategy, ϵ is constant. Five cases were implemented: ϵ ranged from 0.02 to 0.10 in increments of 0.02. In the diminishing probing fraction strategy, ϵ was reduced every 2 years from 0.10 to 0.0 in increments of 0.02.

Table 3. Year to Convergence in Policies with Constant Probing Fraction

| Deterioration process | | Probing fraction ($\epsilon \times 100\%$) | | | | | |
|-----------------------|---------|--|-----|----|----|----|-----|
| Real | Initial | 0% | 2% | 4% | 6% | 8% | 10% |
| Fast | Medium | n/a | n/a | 24 | 23 | 21 | 19 |
| | Slow | n/a | n/a | 23 | 23 | 21 | 18 |
| Medium | Fast | 14 | 13 | 11 | 10 | 8 | 8 |
| | Slow | 18 | 18 | 17 | 17 | 13 | 9 |
| Slow | Fast | 15 | 12 | 11 | 9 | 9 | 8 |
| | Medium | 4 | 3 | 3 | 3 | 3 | 3 |

Table 4. Total Costs with Diminishing Probing Fraction

| Deterioration process | | Probing strategy | |
|-----------------------|---------|------------------|----------|
| Real | Initial | Diminishing | Constant |
| Fast | Medium | 96.02 | 99.31 |
| | Slow | 81.73 | 81.59 |
| Medium | Fast | 56.07 | 55.77 |
| | Slow | 40.56 | 41.89 |
| Slow | Fast | 49.08 | 46.63 |
| | Medium | 43.09 | 35.44 |

Total Cost Comparison

The total costs in each case with constant probing fractions are shown in Table 2. In all cases except one (when the real deterioration rate is medium and the initial deterioration rate is slow), lower costs were obtained by using systematic probing. This result is related to the time needed for convergence in policies. These times are shown in Table 3.

Systematic probing allows every state to be visited and every action to be applied. Thus, all transition probabilities are updated. Intuitively, when the fraction is larger, convergence in the transition probabilities is faster. This is why the convergence in policies occurred earlier with the larger probing fractions. The cases where convergence was not achieved within the planning horizon are denoted as n/a in the Table 3.

There is a tradeoff in total costs. With larger probing fractions, the policies converge faster and the MR&R costs after convergence are reduced. On the other hand, the costs of systematic probing are higher. This is because a larger fraction of suboptimal MR&R actions are applied during the probing period.

In the second part of the parametric study, a diminishing probing fraction strategy was used. The total costs are shown in Table 4. As can be seen, in certain cases, the total costs were reduced further relative to the constant probing fraction strategy. This was true in the case where the real deterioration rate is fast, the initial rate is medium, and—in the case—where the real deterioration rate is medium and the initial deterioration rate is slow. In these scenarios, it is important to quickly account for the fact that deterioration is proceeding faster than anticipated, but further model refinement is less critical. The number of years to reach convergence in policies is shown in Table 5.

The cost savings achieved by using systematic probing are shown in Table 6. The cost savings represent the differences between the cost incurred without systematic probing and the cost incurred by the most efficient probing strategy. The parametric study showed that using systematic probing, in the cases examined, reduces total costs if the optimal probing fractions are used.

Table 5. Year to Convergence in Policies with Diminishing Probing Fraction

| Deterioration process | | Year to convergence |
|-----------------------|---------|---------------------|
| Real | Initial | |
| Fast | Medium | 22 |
| | Slow | 20 |
| Medium | Fast | 10 |
| | Slow | 15 |
| Slow | Fast | 11 |
| | Medium | 3 |

Table 6. Cost Savings

| Deterioration process | | Cost savings | |
|-----------------------|---------|--------------|------|
| Real | Initial | \$/lane-yard | % |
| Fast | Medium | 5.52 | 5.44 |
| | Slow | 2.75 | 3.26 |
| Medium | Fast | 0.01 | 0.02 |
| | Slow | 1.33 | 3.17 |
| Slow | Fast | 0.51 | 1.08 |
| | Medium | 0.05 | 0.14 |

Conclusion

An adaptive infrastructure MR&R optimization methodology that incorporates updating of transition matrices was presented. The methodology is based on an OLFC approach. Parametric studies showed that using systematic probing can lead to cost savings over the OLFC without systematic probing, if the optimal probing fractions are selected. The parametric studies also showed that the cost savings are related to the time needed for convergence in policies. The MR&R policies converged faster when probing was used.

As argued earlier in this paper, systematic probing can be achieved by performing a full set of MR&R activities on test facilities as part of experimental studies. Such experiments are underway in several states, as well as in other countries. The results of the parametric studies performed in this paper indicate that there are significant savings that can be achieved by using the results of such experiments jointly with the results of inspections of in-service facilities to update the parameters of infrastructure performance models. These results dovetail with those obtained in empirical studies (such as, Prozzi and Madanat 2004), which indicated that models developed by combining experimental and field data had higher accuracy than those developed with only one source of data, because experimental data included a broader range of action-condition combinations, while field data more closely represented the process of deterioration in the field.

On the other hand, the results of the previous section also showed that probing with the wrong probing fractions could significantly increase the expected cost. Given that highway agencies are unlikely to know what the optimal probing fractions are for their systems, it may be safer to avoid systematic probing until further research is conducted. In the meantime, the major conclusion from this work is that state-of-the-art adaptive IMSs, that do not use probing, provide sufficiently close to optimal policies.

Acknowledgments

Partial funding for this research was provided by a University of California Transportation Center (UCTC) research grant to the

first writer. The writers benefitted from discussions with Carlos Daganzo. The comments of two anonymous referees were instrumental in improving the quality of this paper.

References

- Bertsekas, D. P. (2000). *Dynamic programming and optimal control*, Athena Scientific, Belmont, Mass., Vols. 1 and 2.
- Carnahan, J. V., Davis, W. J., Shahin, M. Y., Keane, P. L., and Wu, M. I. (1987). "Optimal maintenance decisions for pavement management." *J. Transp. Eng.* 113(5), 554–572.
- DeGroot, M. H. (1970). *Optimal statistical decisions*, McGraw-Hill, New York.
- Durango, P., and Madanat, S. (2002). "Optimal maintenance and repair policies in infrastructure management under uncertain facility deterioration rates: An adaptive control approach." *Transp. Res., Part A: Policy Pract.*, 36A(6), 768–769.
- Feighan, K. J., Shahin, M. Y., Sinha, K. C., and White, T. D. (1988). "Application of dynamic programming and other mathematical techniques to pavement management systems." *Transportation Research Record. 1200*, Transportation Research Board, Washington, D.C., pp. 90–98.
- Golabi, K., Kulkarni, R., and Way, G. A. (1982). "Statewide pavement management system." *Interfaces*, 12(6), 5–21.
- Golabi, K., and Shepard, R. (1997). "Pontis: A system for maintenance optimization and improvement of U.S. bridge networks." *Interfaces*, 27(1), 71–88.
- Harper, W. V., and Majidzadeh, K. (1991). "Use of expert opinion in two pavement management systems." *Transportation Research Record. 1311*, Transportation Research Board, Washington, D.C., pp. 242–247.
- Kumar, P., and Varaiya, P. (1986). *Stochastic systems: Estimation, identification, and adaptive control*, Prentice-Hall, Englewood Cliffs, N.J.
- Madanat, S. (1993). "Incorporating inspection decisions in pavement management." *Transp. Res., Part B: Methodol.*, 27(6), 425–438.
- Madanat, S., and Ben-Akiva, M. (1994). "Optimal inspections and repair policies for infrastructure facilities." *Transp. Sci.*, 28(1), 55–62.
- Madanat, S., Mishalani, R., and Wan Ibrahim, W. H. (1995). "Estimation of infrastructure transition probabilities from condition rating data." *J. Infrastruct. Syst.* 1(2), 120–125.
- Mishalani, R., and Madanat, S. (2002). "Computation of infrastructure transition probabilities using stochastic duration models." *J. Infrastruct. Syst.*, 8(4), 139–148.
- Prozzi, J., and Madanat, S. (2004). "Development of pavement performance models by combining experimental and field data." *J. Infrastruct. Syst.* 10(1), 9–22.
- Smilowitz, K., and Madanat, S. (2000). "Optimal inspection and maintenance policies for infrastructure networks." *Comput. Aided Civ. Infrastruct. Eng.*, 15(1), 5–13.
- Sutton, R., and Barto, A. (1998). *Reinforcement learning*, MIT Press, Cambridge, Mass.