Title
Model Uncertainty and the Management of a System of Infrastructure Facilities

Authors
Kuhn, Kenneth D.
Madanat, Samer M.

Publication Date
2005-09-01

Model Uncertainty and the Management of a System of Infrastructure Facilities

Kenneth D. Kuhn

Ph.D. Candidate, Univ. of California, Dept. of Civil and Environmental Engineering, 116 McLaughlin Hall, Berkeley, CA 94720

Samer M. Madanat

Professor, Univ. of California, Dept. of Civil and Environmental Engineering, 110 McLaughlin Hall, Berkeley, CA 94720

Abstract

The network-level infrastructure management problem involves selecting and scheduling Maintenance, Repair, and Rehabilitation (MR&R) activities on networks of infrastructure facilities so as to maintain the level of service provided by the network in a cost-effective manner. This problem is frequently formulated as a Markov Decision Problem (MDP) solved via Linear Programming (LP). The conditions of facilities are represented by elements of discrete condition rating sets, and transition probabilities are employed to describe deterioration processes. Epistemic and parametric uncertainties not considered within the standard MDP/LP framework are associated with the transition probabilities used in infrastructure management optimization routines. This paper contrasts the expected costs incurred when model uncertainty is ignored with those incurred when this uncertainty is explicitly considered using Robust Optimization. A case study involving a network-level pavement management MDP/LP problem demonstrates how explicitly considering uncertainty may limit worst case MR&R expenditures. The methods and results can also be used to identify the costs of uncertainty in transition probability matrices used in infrastructure management systems.

Key words: infrastructure management, epistemic uncertainty, robust optimization

Email addresses: kuhn@berkeley.edu (Kenneth D. Kuhn), madanat@ce.berkeley.edu (Samer M. Madanat).

Preprint submitted to Elsevier Science *27 September 2005*

1 AN OVERVIEW OF THE PROBLEM

Asset management is the process by which agencies monitor and maintain built systems of facilities, with the objective of providing the best possible service to the users, within the constraints of available resources. The primary decisions made by a public works agency are the selection and scheduling of Maintenance, Repair, and Rehabilitation (MR&R) actions to perform on facilities in the system during a specified planning horizon.

There are two forms of information on infrastructure condition: information on current condition, provided by facility inspection, and information on future condition, predicted by using a deterioration model. Deterioration models are mathematical relations having as a dependent variable the condition of a facility and as independent variables the facility's age, previous condition, level of utilization, environment, historical MR&R actions, etc. Expert judgment or empirical data are used to develop these deterioration models. The models, in turn, are used to produce inputs into the MR&R decision-making process. Network level infrastructure management problems are frequently formulated as Markov Decision Processes (MDPs) (Golabi et al (1982), Gopal and Majidzadeh (1991), Smilowitz and Madanat (2000)).

1.1 A Sample Network-Level Infrastructure Management Problem

A classic example of a network-level asset management optimization system is the Pavement Management System (PMS) used by the Arizona Department of Transportation (ADOT) (Golabi et al, 1982). This system minimizes the life cycle costs of the pavement sections managed by the ADOT by linear optimization techniques. Facilities are separated into groups based on their construction, environment, and traffic loading. ADOT's PMS recommends MR&R actions for fractions of facilities in the different groups. Pavement managers are given leeway in determining which specific facilities to choose to make up the fractions selected by the PMS.

An example of a mathematical programming formulation for an asset management problem, built using the ADOT PMS as a template, is presented here. In this example, it is assumed that there is a network of N facilities to be managed for a period of T years. Each facility, in any given year, is characterized by a condition state in a set I of possible states for this particular type of infrastructure facility. Every year, a single management action in a set A of possible management actions must be selected for every individual facility. (The set A may include an action representing the absence of taking one of the other actions, a "do nothing" option.) Each of the N facilities in this example is identical in the sense that the description of deterioration, as well as the costs, associated with different facilities are exactly the same.

Model Parameters

Let $p(j|i,a)$ represent the conditional probability of an infrastructure facility being in state j next year, given that it is in state i this year with action a being applied.

Let $ac(i,a)$ represent the agency costs for a facility in state i with action a being applied.

Let $tc(i,a)$ represent the total (user + agency) costs for a facility in state i with action a being applied.

Let $b(t)$ represent the agency budget in year t.

Let α be the discount amount factor that relates future costs to present dollars ($\alpha = \frac{1}{1+r}$).

Let $init_i$ represent the fraction of all facilities in condition state i at the beginning of the infrastructure management exercise.

Let X be a set of condition states that provide such poor service that these states are deemed unacceptable.

Decision Variables

Let $f_t(i,a)$ represent, in year t, the expected fraction of all facilities that is in condition state i with action a being applied.

Markov Decision Problem

$$\min_f \sum_{t=0}^{T} \alpha^t [\sum_{i\in I} \sum_{a\in A} tc(i,a) f_t(i,a) N]$$

$s.t.$

$$\sum_{a\in A} f_0(i,a) = init_i \qquad \forall i \in I \quad (1)$$

$$f_t(i,a) \geq 0 \qquad \forall i \in I, a \in A, t \in \{0,1,...,T\} \quad (2)$$

$$f_t(i,a) = 0 \qquad \forall i \in X, a \in A, t \in \{0,1,...,T\} \quad (3)$$

$$\sum_{i\in I} \sum_{a\in A} f_{t-1}(i,a) p(j|i,a) = \sum_{a\in A} f_t(j,a) \qquad \forall j \in I, t \in \{1,2,...,T\} \quad (4)$$

$$\sum_{i\in I} \sum_{a\in A} ac(i,a) f_t(i,a) N \leq b(t) \qquad \forall t \in \{0,1,...,T\} \quad (5)$$

Like the original ADOT PMS, a decision variable is used that produces fractions of facilities to which MR&R actions should be applied. This formulation does not offer explicit, facility by facility, guidance. However, defining the decision variable this way allows engineers with immediate knowledge of the facilities in question to choose a repair plan out of a set considered equivalent by this asset management system. Furthermore, the formulation shown here has the crucial advantage of allowing the scope of the asset management problem to be independent of the number of facilities in the network.

In the above formulation, constraint 1 requires that the current conditions of facilities are recognized. Constraint 2 ensures negative fractions of facilities are never considered. Constraint 3 deals with level of service considerations. There are many ways that level of service might be considered within an infrastructure management problem. In this example, we simply say that X is a set of condition states deemed unacceptable because of the poor service facilities in these states provide. The third constraint as laid out ensures that facilities are never allowed to reach condition states in X deemed unacceptable. The fourth constraint is the key to how deterioration is considered in decision-making. Decisions made in a given year must be based on the expected conditions of facilities, which are obtained from the decisions made the previous year and the transition probability matrix. Constraint 4 says that at a time t the expected number of facilities in a state j is equal to the sum across all states of the expected number of facilities in that state at time $t-1$ times the respective probabilities of transitioning into state j. (Note that this constraint, when used in conjunction with constraint 1, ensures that all facilities are accounted for.) Finally, the fifth constraint demands that the agency never runs over its budget in any given year. The objective function and constraints 3 and 5 could be rewritten any number of ways depending on how a planning agency wishes to consider level of service and budget constraints.

1.2 Epistemic and Parametric Uncertainty In Transition Probabilities

In general, infrastructure deterioration model forecasts are associated with a high degree of uncertainty due to the uncertainty in the following factors: (a) exogenous factors such as the environment and level of utilization, (b) endogenous factors such as facility design and materials, and (c) statistical factors such as the limited size and scope of data sets used to generate models, or the differences between data generated in a laboratory setting versus in the field. Although we can improve the quality of data by developing more advanced inspection methods and deterioration models, it is impossible to eliminate entirely the uncertainty associated with MR&R decision-making. In state-of-the-art asset management systems, the stochastic nature of a facility's deterioration process (intrinsic uncertainty) has been captured through the use of stochastic process models (such as Markov transition probabilities) as representations of facility deterioration. On the other hand, the determination of the parameters of these stochastic models is still subject to significant parametric uncertainty. Additional uncertainty known as epistemic uncertainty stems from a more fundamental inability to model deterioration precisely, even if parameters could be set properly.

For example, we consider a study by Prozzi and Madanat (2004) that used data from the American Association of State Highways Officials (AASHO)

Road Test to model pavement roughness. Statistical testing by Prozzi and Madanat (2004) revealed statistically significant "unobserved heterogeneity" in their model that "cannot be ignored" despite the fact that the model's prediction error was 50% smaller than that of earlier models developed with the same data. In other words, they were able to limit parametric uncertainty but still encountered significant errors in stochastic deterioration models due to epistemic uncertainty.

Recent research in the area of adaptive infrastructure management acknowledges the presence of parametric uncertainty in infrastructure decay modeling. Adaptive control methods use data as it becomes available to update the parameters of models that guide decision-making. For example, the popular bridge management system PONTIS updates transition probability matrices over time (Golabi and Shepard, 1997). Durango and Madanat (2002) have proposed a decision support system that uses Bayesian updating, while Madanat et al (2005) have made use of maximum likelihood estimation to update the parameters of transition probabilities. Durango-Cohen (2004) has used reinforcement learning techniques for the same purpose.

Adaptive control is able to reduce parametric uncertainty over time, but does not consider epistemic uncertainty. Additionally, open-loop feedback adaptive control approaches, like the one used in PONTIS, cannot guarantee that their decision support systems will converge to optimal policies. Not all actions are performed on assets in all states a sufficient number of times to ensure that the system converges to an accurate transition probability matrix (Madanat et al, 2005). The flip side of this argument is that decision support systems that can guarantee convergence to optimal policies must at times choose actions that are suboptimal in order to "probe" the system being learned (ibid). However, the most serious limitation to the effectiveness of the adaptive control approach is that while managing a network of infrastructure facilities, data on condition, deterioration, and the effectiveness of different MR&R actions accumulates slowly. Thus, adaptive control approaches require a long time to improve the precision of the transition matrices and, during this time, will incur high costs associated with transition matrix uncertainty.

2 THE PROPOSED SOLUTION: ROBUST OPTIMIZATION

Robust optimization is a modeling methodology to solve optimization problems in which the data are uncertain and only known to belong to some uncertainty set. Incomplete information and changing system dynamics often provide motivation for robust optimization (Averbakh, 2000). The objective is to seek optimal (or near optimal) solutions that are not overly sensitive to any realization of uncertainty. Recent reviews on this topic can be found in Mulvey et al (1995), Ben-Tal and Nemirovski (1999) and El Ghaoui (2003) among others. A robust feasible solution is one that tolerates changes in the problem data, up to a given bound known a priori, and a robust optimal solution is a robust feasible solution with the best possible value of the objective function. When the uncertainty set can be parameterized, the objective function may minimize expected costs across the uncertainty set. Assuming less available information regarding uncertainty sets, the objective may be to minimize 'worst-case' costs or 'regret,' a measure of opportunity cost (Averbakh, 2000). One interesting alternative proposed in Bertsimas and Sim (2004) minimizes costs while considering that a set number of parameters used in optimization are different from initially estimated values. By carefully constructing and efficiently solving robust problems, it is possible to obtain solutions that trade off performance vs. guaranteed robustness and reliability. One interesting successful example that deals specifically with a robust Markov Decision Problem in a transportation setting (air traffic control) is Nilim and El Ghaoui (2004).

In infrastructure management, Kuhn and Madanat (2005) have applied Robust Optimization to single facility infrastructure management. The results that they obtained in their parametric studies demonstrated that robust optimization was able to limit excess costs associated with errors in infrastructure deterioration modeling. The robust single facility infrastructure management problem was solved in Kuhn and Madanat (2005) by considering potential actions in turn and, for each action, defining an appropriate robust cost. A similar approach was used in Nilim and El Ghaoui (2004) to solve more general robust dynamic programming problems. Neither of these approaches is well suited to dealing with multi-facility infrastructure management. When dealing with a large number of facilities, each of which may be maintained a large number of ways, the space of possible agency policies becomes quite large and dynamic programming approaches become computationally intractable. Linear programs, like the one outlined in section 1.1 of this paper, are more appropriate.

2.1 A Sample Problem Considering the 'Worst Case'

Robust optimization typically minimizes costs considering that nature will act as an opponent. Worst case conditions are considered, making it possible to make performance guarantees (bound worst case costs) across an entire uncertainty set. This method does not require decision-makers to estimate a probability distribution over the uncertainty set. Consideration of worst case conditions may be incorporated into the asset management problem formulated earlier in this paper.

New Model Parameters

Let δ represent the uncertainty level ($\delta \in [0, 1]$).

Let $q(j|i, a)$ represent the initially assumed conditional probability of an infrastructure facility being in state j next year, given that it is in state i this year with action a being applied.

New Decision Variables

Let $p(j|i, a)$ represent the conditional probability, as considered in the robust optimization, of an infrastructure facility being in state j next year, given that it is in state i this year with action a being applied.

Markov Decision Problem

$$\max_p[\min_f \sum_{t=0}^{T} \alpha^t (\sum_{i \in I} \sum_{a \in A} tc(i, a) f_t(i, a) N)]$$

$s.t.$

$$\sum_{a \in A} f_0(i, a) = init_i \qquad \forall i \in I \ (1)$$

$$f_t(i, a) \geq 0 \qquad \forall i \in I, a \in A, t \in \{0, 1, ..., T\} \ (2)$$

$$f_t(i, a) = 0 \qquad \forall i \in X, a \in A, t \in \{0, 1, ..., T\} \ (3)$$

$$\sum_{i \in I} \sum_{a \in A} f_{t-1}(i, a) p(j|i, a) = \sum_{a \in A} f_t(j, a) \qquad \forall j \in I, t \in \{1, 2, ..., T\} \ (4)$$

$$\sum_{i \in I} \sum_{a \in A} ac(i, a) f_t(i, a) N \leq b(t) \qquad \forall t \in \{0, 1, ..., T\} \ (5)$$

$$p(j|i, a) \geq 0 \qquad \forall i \in I, j \in I, a \in A \ (6)$$

$$\sum_{j \in I} p(j|i, a) = 1 \qquad \forall i \in I, a \in A \ (7)$$

$$|p(j|i, a) - q(j|i, a)| \leq \delta \qquad \forall i \in I, j \in I, a \in A \ (8)$$

Setting the uncertainty level δ to 0 implies no uncertainty, meaning a "likelihood region" is defined that includes only the transition probability matrix given by the initial model q. A likelihood region is a set of transition probability matrices, each of which may describe the deterioration of the facilities in the network. An uncertainty level of 1 would imply absolute uncertainty. In this case all transition probability matrices may define infrastructure deterioration and would be included in the likelihood region. According to constraint 8, a transition probability matrix is included in the likelihood region if and only if the difference between any element of the transition matrix and the corresponding element of the initially assumed matrix is less than or equal to the uncertainty level. Seen in this light, the uncertainty level represents how large an error in transition probabilities is considered possible.

This definition of uncertainty is quite simple and ignores the possibility of correlations among probabilities from different states and actions. Infrastructure deterioration models often rely heavily on expert judgment (Harper et al, 1990), and infrastructure management systems that consider uncertainty will have to characterize this uncertainty using expert judgment. Cases where uncertainty is likely to be an issue are exactly the cases where there is an absence of empirical data regarding deterioration. Parameterizing complex uncertainty models may be impossible or prone to uncertainty itself. In any event, the point of this work is to introduce the explicit consideration of uncertainty into the context of infrastructure management optimization.

One criticism of the mathematical program described above is that it is too conservative. This criticism is reinforced by the fact that certain transitions that might be considered impossible in real life may be given positive probabilities of occurring in this formulation. Fortunately, it is a relatively simple task to ensure that certain "impossible" transitions are never considered. For example, it is possible to force any transitions with zero probability in an initial model to have zero probability in any model included in the likelihood region. This constraint can be accomplished by fixing $p(j|i,a)$ to 0 whenever $q(j|i,a)$ is equal to 0. Note that the inclusion of the additional constraint actually reduces the computational burden of solving the mathematical program.

Fixing certain $p(j|i,a)$ to 0 restricts the likelihood region considered. However, the approach is still fairly conservative; nature is still seen as a malevolent opponent. While managing a large network of facilities, it may be too costly and unrealistic, to manage each facility under the assumption that nature is always malevolent. An alternate approach involves acting under the assumption that nature will work with decision-makers instead of against them. The most realistic point of view would be to recognize that nature will act neither as a perpetual adversary nor ally, but somewhere in between.

2.2 The Hurwicz Criterion

An attractive alternative to decision making focusing exclusively on worst case conditions involves applying the Hurwicz criterion (Hurwicz, 1951). The Hurwicz criterion allows a decision-maker to set his or her "optimism level." The optimism level must be a number between 0 and 1. The pessimism level is defined as 1 - the optimism level. Decisions are then made by selecting actions that minimize costs obtained by summing the optimism level times the costs under best case conditions and the pessimism level times the costs under worst case conditions. Setting the optimism level to 1 would therefore entail comparing alternatives based only on the lowest possible cost they might entail. Similarly, setting the optimism level to 0 would entail considering only worst case conditions, as was done in section 2.1 of this paper. Solving a Hurwicz-criterion-based asset management problem with an optimism level between 0 and 1 can yield MR&R policies that limit the sensitivity of maintenance costs to realized deterioration rates without being overly costly in normal or good conditions. To understand how the Hurwicz criterion might be applied to asset management, consider an infrastructure management problem with a short two-year management horizon.

New Model Parameters
Let β represent the optimism level ($\beta \in [0,1]$).

New Decision Variables
Let $p_b(j|i,a)$ represent the best case conditional probability of an infrastructure facility being in state j next year, given that it is in state i this year with action a being applied.
Let $p_w(j|i,a)$ represent the worst case conditional probability of an infrastructure facility being in state j next year, given that it is in state i this year with action a being applied.
Let $f_0(i,a)$ represent, in year 0, the expected fraction of all facilities that is in condition state i with action a being applied.
Let $f_{1b}(i,a)$ represent, in year 1, best case conditions, the expected fraction of all facilities that is in condition state i with action a being applied.
Let $f_{1w}(i,a)$ represent, in year 1, worst case conditions, the expected fraction of all facilities that is in condition state i with action a being applied.

Markov Decision Problem

$$\max_{p_w} \min_{p_b}$$

$$\min_f \sum_{i\in I}\sum_{a\in A} tc(i,a)f_0(i,a)N + \alpha[\beta\ tc(i,a)f_{1b}(i,a)N + (1-\beta)tc(i,a)f_{1w}(i,a)N]$$

$s.t.$

$$\sum_{a\in A} f_0(i,a) = init_i \qquad \forall i \in I \ (1)$$

$$f_t(i,a) \geq 0 \qquad \forall i \in I, a \in A, t \in \{0,1b,1w\} \ (2)$$

$$f_t(i,a) = 0 \qquad \forall i \in X, a \in A, t \in \{0,1b,1w\} \ (3)$$

$$\sum_{i\in I}\sum_{a\in A} f_0(i,a)p_b(j|i,a) = \sum_{a\in A} f_{1b}(j,a) \qquad \forall j \in I \ (4.1)$$

$$\sum_{i\in I}\sum_{a\in A} f_0(i,a)p_w(j|i,a) = \sum_{a\in A} f_{1w}(j,a) \qquad \forall j \in I \ (4.2)$$

$$\sum_{i\in I}\sum_{a\in A} ac(i,a)f_t(i,a)N \leq b(t) \qquad \forall t \in \{0,1b,1w\} \ (5)$$

$$p_b(j|i,a) \geq 0 \qquad \forall i \in I, j \in I, a \in A \ (6.1)$$

$$p_w(j|i,a) \geq 0 \qquad \forall i \in I, j \in I, a \in A \ (6.2)$$

$$\sum_{j\in I} p_b(j|i,a) = 1 \qquad \forall i \in I, a \in A \ (7.1)$$

$$\sum_{j\in I} p_w(j|i,a) = 1 \qquad \forall i \in I, a \in A \ (7.2)$$

$$|p_b(j|i,a) - q(j|i,a)| \leq \delta \qquad \forall i \in I, j \in I, a \in A \ (8.1)$$

$$|p_w(j|i,a) - q(j|i,a)| \leq \delta \qquad \forall i \in I, j \in I, a \in A \ (8.2)$$

Note that now two sets of transition probabilities have to be considered, one for best case conditions and one for worst case conditions. Furthermore, expected conditions and costs have to be calculated in best and worst case conditions. The objective function is now quite complex, involving a cost minimization problem nested within both maximization and minimization problems.

The new mathematical program appears quite complex and it only involved a two-year planning horizon. For a longer planning horizon, applying the Hurwicz criterion to facility management involves considering a combinatorial number of cases. The MR&R decision in the first year depends upon best and worst case costs, so there need to be variables considering best and worst condition states and actions taken in the second year. In order to consider a management strategy that consistently applies the Hurwicz criterion, the actions to take in the best and worst cases of the second year will each individually be based on separate best and worst case comparisons for the third year, and so on.

Decisions in all years must be made simultaneously. Facility conditions and maintenance activities to be performed in later years will be determined by which maintenance actions were performed in earlier years. Decisions about maintenance actions to perform in later years will depend on expected future costs, which are functions of future conditions and maintenance actions. This problem's complexity grows exponentially with the time frame to be considered. However, the Hurwicz criterion remains an interesting approach to consider ranges of infrastructure decay rather than focusing on only one deterioration model.

2.3 Setting Up and Solving Robust Problems

The robust mathematical programs presented here require the specification of an initial transition probability matrix, an uncertainty level, and (in the case of Hurwicz optimization) an optimism level. While it may be difficult to specify an initial transition probability matrix, it is worth noting that the robust model will yield results that are less sensitive to errors in this initial model than more traditional non-robust optimization. Parametric uncertainty in deterioration models can be described statistically in certain situations. For instance, when dealing with parameters of deterioration models calibrated with empirical data, it is possible to study the standard errors associated with these parameters. In other situations, epistemic uncertainty cannot be described using statistical techniques. In such cases, domain-specific expert opinion may be used to identify realistic bounds on the values of transition matrices. Additionally, it is possible to solve a particular robust asset management problem numerous times with various uncertainty (and optimism) levels to see how performance and reliability guarantees can be traded off, leading to a more informed discussion of uncertainty and its costs.

The key to solving the bi-level mathematical programming problems presented in this paper is recognizing that in infrastructure management it is obvious that slower deterioration rates yield lower costs. In such circumstances, the problems of finding worst and best case transition probability matrices can be separated from the problems of selecting policies optimal for the given transition probability matrices. Probability must be 'shifted' from better condition states to worse condition states to come up with a worst case transition probability matrix, and vice-versa to find a best case transition probability matrix. The simple model used here to describe uncertainty allows the probabilities originating from different initial state-action pairs to be considered separately.

Once best and worst case transition probability matrices have been found, it is necessary to select the optimal management policy f. This problem is essentially the same as the original multi-facility infrastructure management problem; there is a given transition probability matrix/matrices and linear programming can be used to reveal the optimal policy in response to this description of deterioration. However, the transition probability matrix used to generate the management policy f will not describe the actual expected realization of deterioration processes. So every year, new information on the condition of infrastructure facilities will have to be collected, and a new management policy will have to be found from that time through the end of the planning horizon, incorporating the conditions of the infrastructure facilities according to their most recent inspection.

3 A CASE STUDY INVOLVING PAVEMENT MANAGEMENT

In order to illustrate the application of robust algorithms to multi-facility infrastructure management problems, an example is presented here. A highway pavement network is managed according to policies obtained from both traditional linear programming and robust methodologies. Previous research (Golabi et al (1982), Madanat (1993), Durango and Madanat (2002)) provides a ready source of data for how pavement deterioration can be modeled via static transition probabilities. However, given the uncertainty in these transition probabilities, potential cost savings can be achieved by applying robust optimization to this problem. It is worth noting that uncertainty is of more concern for infrastructure assets that have less refined deterioration models than pavement sections. Thus robust optimization may actually be better suited to the management of infrastructure assets like underground pipelines and drainage systems.

3.1 Problem Specification

A network of 10,000 square yards of highway pavement will be managed for five years in this example. The decisions to be made include when and how to maintain, overlay, or reconstruct the pavement. In the example presented here, it is assumed that the choices of actions to take in any given year are those presented by Durango and Madanat (2002). These actions include: (1) do nothing, (2) routine maintenance, (3) 1-in overlay, (4) 2-in overlay, (5) 4-in overlay, (6) 6-in overlay, and (7) reconstruction. A section of pavement is said to be in state 1 if it is unusable and in state 8 if it is brand new, with the intermediate states representing intermediate condition ratings, following the convention of Durango and Madanat (2002). The costs of the actions presented here are derived from empirical work by Carnahan (1987) and are included in Table 1. User costs vary according to the condition state of the pavement, while agency costs vary by action chosen and condition state.

In describing the effects of the 7 maintenance actions outlined above, Durango and Madanat present three sets of transition probability matrices. The matrix that describes a section of pavement deteriorating at a "medium" rate is meant to reflect the current best estimate of how a given section of pavement will deteriorate. The inclusion of alternative "fast" and "slow" rates of deterioration draws attention to the fact that this estimate may under or over estimate decay in meaningful ways. For the purposes of the present example, Durango and Madanat's medium deterioration rate transition probabilities are used to initialize the robust programming application used in this paper.

Various uncertainty levels between 0 and 1 are considered. Costs are then calculated for the cases where the actual transition probabilities are as predicted by the initial model, as well as the worst and best possible (i.e. fastest and slowest deterioration). It is assumed that all pavement sections begin the computational study in like new condition and deteriorate according to the same transition probabilities. Both robust and non-robust linear programming algorithms are applied to guide the management of the infrastructure facilities, each being re-run each year in the five year planning horizon as new condition state information becomes available.

The optimization problems solved to guide decision-making are based on those outlined in this paper: an original linear program, a modified mathematical program that considers worst case conditions, and a Hurwicz criterion based mathematical program. The second and third of these programs fix at 0 those transition probabilities equal to 0 in the original medium decay rate. The worst condition state for the pavement, state 1, is deemed unacceptable. A maximum budget of 25,000 dollars each year is set. Given costs as outlined above, this is sufficient to maintain the network if wise investments in maintenance are made, even in the worst case. However, the budget is small enough to make it impossible to perform the actions that yield lowest possible user + agency costs on all facilities at all times.

The focus of this paper is on showing why robust optimization may be worthwhile in the case of infrastructure management, rather than developing any computational tricks for robust optimization. However, it is worthwhile to present broadly how the optimization problems used in this computational study were solved. It is apparent that, in the pavement management example here, better condition states are associated with lower management costs. This makes it relatively easy to describe transition probability matrices associated with best and worst case conditions, as was discussed. A program was written in C++ to perform this task which has an order of complexity equal to the number of condition states possible in this problem. The yearly selection of management policies required the optimization of a linear program equivalent to that shown in section 1.1 of this paper. This was done using the AMPL programming language. A few different workstation computers were used, but in all cases, solutions were found in between 5 and 3,000 seconds.

3.2 Results of the Computational Study

The policies prescribed by the robust linear programming algorithm that considers worst case conditions are able to achieve significant (user + agency) cost savings when compared to traditional non-robust optimization. Figure 1 shows the accumulated five year infrastructure management costs accrued to both the users of the system and the planning agency in charge.

If the uncertainty level exceeds 0.8 then the worst-case cost of non-robust optimization is undefined. Given the available budget, non-robust optimization is unable to meet the level of service requirement that the worst condition state be avoided. At the same uncertainty levels, robust optimization is able to meet budget and service requirements. The reason is that the non-robust management scheme either chooses not to spend or inefficiently spends its available budget in the first years of asset management. In the problem studied here, budget not spent one year is lost rather than carried over to the following year. Thus, the belief that decay occurs slower than it actually does leads to a crisis in later years when the planning agency suddenly realizes in one year that given that year's budget alone it cannot maintain the system above the minimum service requirement.

Figure 1 shows the worst-case costs of asset management, both using an approach that seeks to minimize those costs and the current standard approach, as a function of uncertainty. Seen in this light, the difference between cost figures can be seen as the maximum cost of uncertainty in the facility performance models. Naturally, non-robust optimization produces lower costs if decay proceeds as anticipated by the initial model. It is interesting to look at the additional cost of optimization that considers worst-case conditions in this scenario (shown in Figure 2), and in particular to compare this potential extra cost of robust management with the potential savings observed in worst-case conditions (Figure 1).

Discontinuities in costs in Figure 2 are associated with the fact that at certain threshold levels of uncertainty, robust asset management schemes alter their maintenance schedule recommendations. The costs of asset management in the expected case were plotted on the same scale as the costs in the worst case shown before. Note that the potential savings achieved by robust optimization in worst case conditions are of a larger magnitude than the extra costs in expected case conditions. This was by no means guaranteed; robust optimization does not even consider costs for expected case conditions. For further demonstration of this phenomenon, consider the total cost ranges possible under traditional and robust asset management, as shown in Figure 3.

Robust optimization dramatically shrinks the range of potential costs associated with asset management in cases of decay rate uncertainty. Cost uncertainty is undesirable to planning agencies, particularly given the political environment in which they operate. It is worth noting that the extra costs of robust management as compared to traditional management in best case conditions are actually slightly greater than those seen earlier in expected case conditions. Robust management is more conservative than traditional asset management schemes and incurs higher costs more or less in proportion to how benevolent conditions are. Thus, consideration of the whole range of costs that are possible may suggest the use of Hurwicz style robust optimization, which considers both minimum and maximum potential costs. The cost ranges obtained by using the Hurwicz criterion are presented in Figure 4, for the case of an optimism level of 0.5. An optimism level of 0.5 was chosen because this weights best and worst case costs equally and thus results in the narrowest possible total range of costs. Using this level of optimism would be an entirely logical option for the responsible planning agency.

Note that Hurwicz robust asset management is able to limit worst case costs well below those of the non-robust scheme. At the same time, the best-case costs of Hurwicz robust management are in line with those of non-robust management. Thus, Hurwicz robust optimization may be seen to have the benefits of robust optimization, while incurring costs commensurate with non-robust optimization under more benevolent conditions.

4 DISCUSSION

Statistician Leo Breiman once noted that he was "deeply troubled by the current and past use of data models in applications, where quantitative conclusions are drawn and perhaps policy decisions are made" since "conclusions are about the model's mechanism, and not about nature's mechanism" (Breiman, 2001). Robust optimization offers a way to mitigate against the effects of uncertainty in the deterioration models that underlie asset management systems. Applying robust optimization to the management of multiple infrastructure facilities can achieve significant cost savings. These savings can also be thought of as the costs associated with uncertainty in the transition probability matrices used in modern asset management systems.

Monte-carlo simulation of the computation of pavement deterioration transition probabilities, based on the statistical uncertainty surrounding parameters used in this process (as found in Mishalani and Madanat (2002)) reveals standard errors in the range of 0.22 to 0.35. Consideration of these standard errors supports uncertainty levels of 0.6 or greater. The computational study undertaken in this paper, also dealing with pavement deterioration, reveals the potential for significant benefits associated with considering uncertainty of this magnitude.

It is worth noting that this paper focuses upon uncertainty in infrastructure deterioration modeling, but that there may also be uncertainty in the costs of performing different maintenance actions on different infrastructure facilities. The techniques employed to deal with uncertainty in deterioration modeling in this paper could equally be used to address concerns of cost uncertainty. Similarly, although the models presented here are benchmarked in a specific pavement management problem, they are generalizable to a wide range of problems regarding the management of different infrastructure assets.

ACKNOWLEDGEMENTS

The University of California Transportation Center (UCTC) provided funding for this project to the first author through a Dissertation Research Grant and to the second author through a Faculty Research Grant. Discussions with Professors Andrew Lim and Raja Sengupta of UC Berkeley helped guide portions of the research.

References

Averbakh, I. (2000) Minmax Regret Solutions for Minimax Optimization Problems with Uncertainty. *Operations Research Letters*, 27(2), 57-65.

Ben-Tal, A. and Nemirovski, A. (1999) Robust Solutions of Uncertain Linear Programs. *Operations Research Letters*, 25, 1-13.

Bertsimas, D. and Sim, M. (2004) The Price of Robustness. *Operations Research*, 52(1), 35-53.

Breiman, L. (2001) Statistical Modeling: The Two Cultures. *Statistical Science*, 16(3), 199-231.

Durango-Cohen, P. (2004) Maintenance and Repair Decision Making for Infrastructure Facilities without a Deterioration Model. *Journal of Infrastructure Systems*, 10(1), 1-8.

Carnahan, J.V., Davis, W.J., Shahin, M.Y., Kean, P.L. and Wu, M.I. (1987), Optimal Maintenance Decisions for Pavement Management, *Journal of Transportation Engineering*, **113**(5), 554-572.

Durango, P.L. and Madanat, S. (2002) Optimal Maintenance and Repair Policies in Asset Management Under Uncertain Facility Deterioration Rates: An Adaptive Control Approach. *Transportation Research*, 36A, 763-778.

El Ghaoui, L. (2003) *A Tutorial on Robust Optimization.* presentation at the Institute of Mathematics and Its Applications, March 11, 2003. Available at http://robotics.eecs.berkeley.edu/ elghaoui.

Golabi, K., Kulkarni, R., and Way, G. (1982) A Statewide Pavement Management System. *Interfaces*, 12(6), 5-21.

Golabi, K. and Shepard, R. (1997) Pontis: A System for Maintenance Optimization and Improvement of US Bridge Network. *Interfaces*, 27 (1), 71-78.

Gopal, S. and Majidzadeh, K. (1991) Application of Markov Decision Process to Level-of-Service-Based Systems. *Transportation Research Record*, 1304, 12-17.

Harper, W., Majidzadeh, K. and Hurst (1990) Application of Expert Opinion in Customized Pavement Management Systems. *Highways and Data Processing*, 1304, 549-557.

Hurwicz, L. (1951) Optimality Criteria for Decision Making Under Ignorance. *Discussion Paper: Statistics*, 370, Cowles Commission.

Kuhn, K. and Madanat, S. (2005) Robust Maintenance Policies in Asset Management. accepted for publication in *Computer-Aided Civil and Infrastructure Engineering.*

Madanat, S. (1993) Optimal Infrastructure Management Decisions Under Uncertainty. *Transportation Research Part C*, 1, 77-88.

Madanat, S., Park, S. and Kuhn, K. (2004) Adaptive Optimization and Systematic Optimization of Infrastructure System Maintenance Policies under Model Uncertainty. accepted for publication in *Journal of Infrastructure Systems.*

Mishalani, R. and Madanat, S. (2002) Computation of Infrastructure Transition Probabilities Using Stochastic Duration Models. *Journal of Infrastruc-*

ture Systems, 8(4), 139-148.
Mulvey, J.M., Vanderbei, R.J., and Zenios, S.A. (1995) Robust Optimization of Large-Scale Systems. *Operations Research*, 43(2), 264-281.
Nilim, A., and El Ghaoui, L. (2004) Robust Solutions to Markov Decision Problems with Uncertain Transition Matrices. *EECS ERL Memo* MO 04/26.
Prozzi, J.A. and Madanat, S. (2004) Development of Pavement Performance Models by Combining Experimental and Field Data. *Journal of Infrastructure Systems*, 10(1), 9-22.
Smilowitz, K. and Madanat, S. (2000) Optimal Inspection and Maintenance Policies for Infrastructure Networks. *Computer-Aided Civil and Infrastructure Engineering*, 15(1), 5-13.

Table 1: Costs (dollars per lane-yard), both to users and agency, of performing different M&R actions.

Agency Costs	Condition									User Costs
Action to take		1	2	3	4	5	6	7		
	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

Fig. .1. The worst case costs of worst case robust and non-robust pavement management.

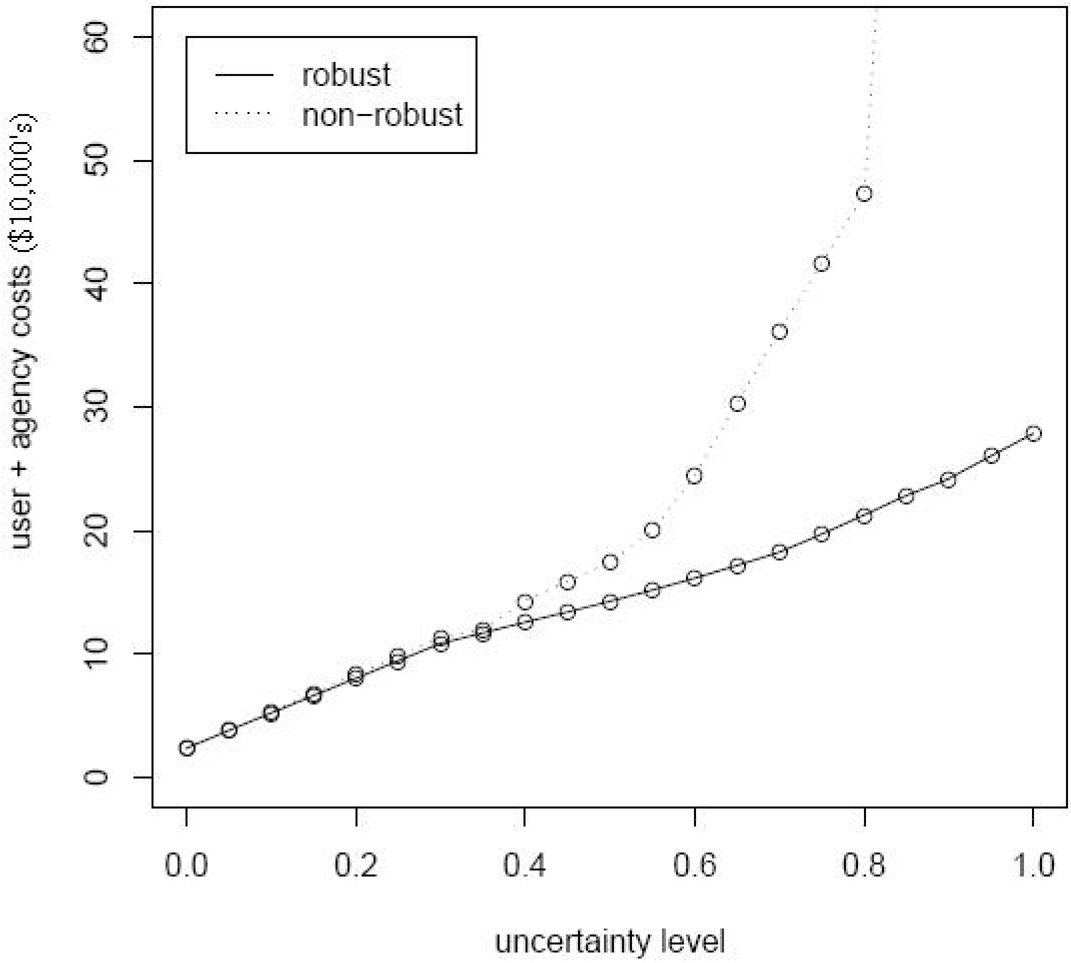

Fig. .2. The expected case costs of worst case robust and non-robust pavement management.

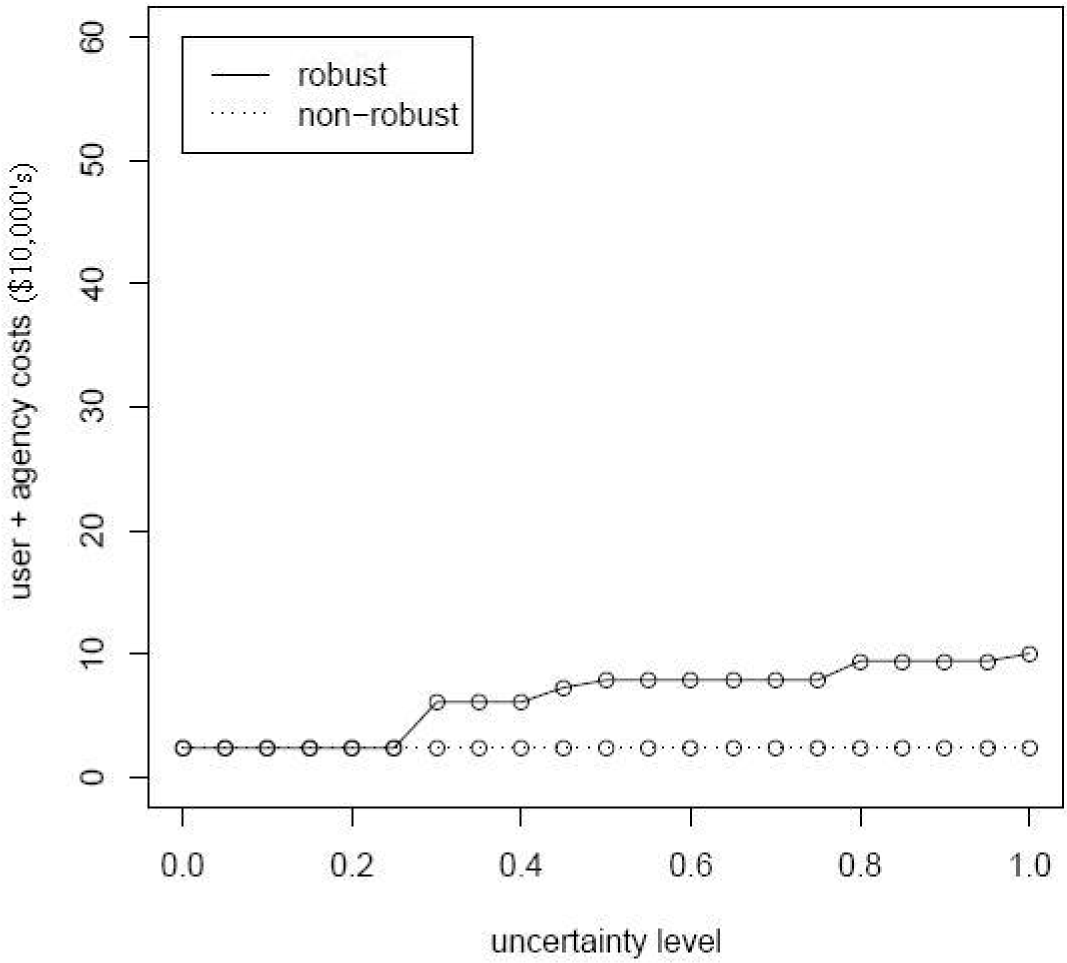

Fig. .3. The possible cost ranges of worst case robust and non-robust pavement management.

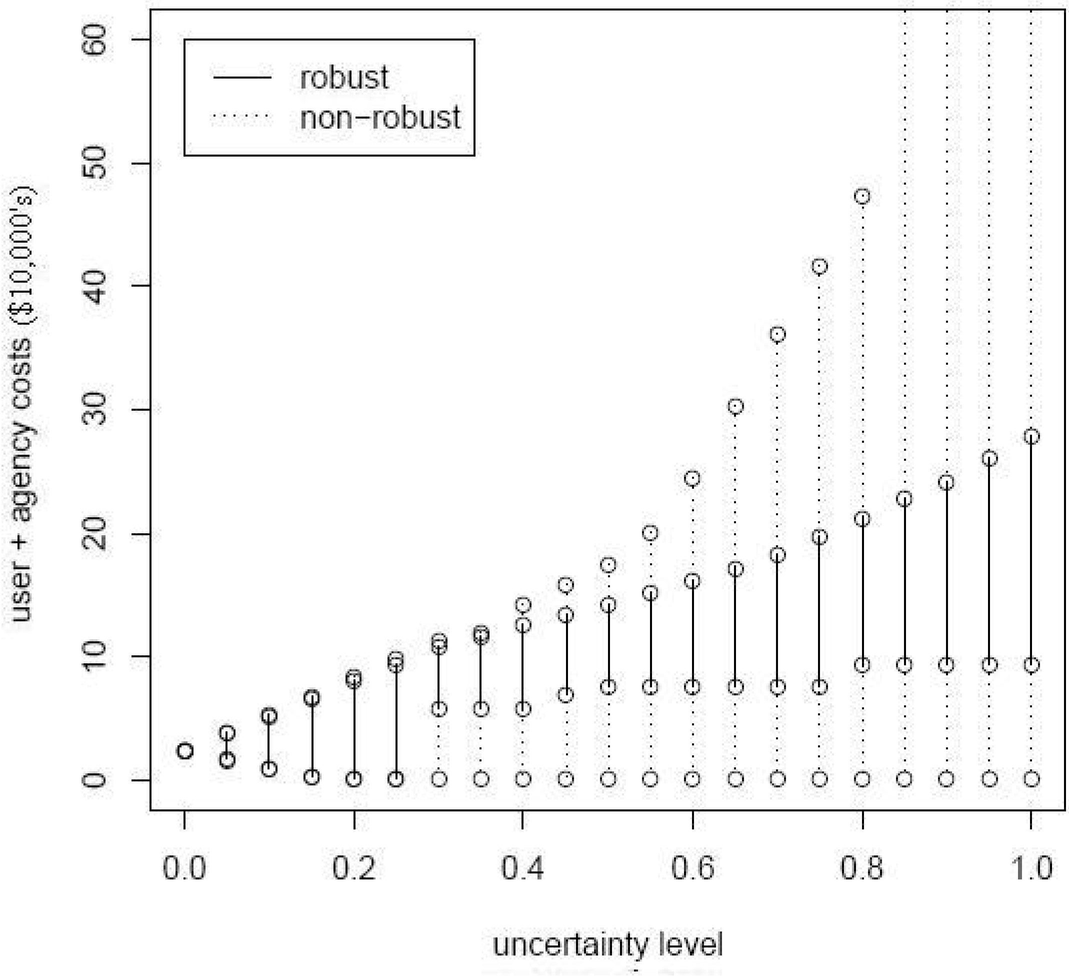

Fig. .4. The possible cost ranges of hurwicz robust (for $\beta = 0.5$) and non-robust pavement management.