

PAVEMENT NETWORK MAINTENANCE OPTIMIZATION CONSIDERING MULTIDIMENSIONAL CONDITION DATA

Kenneth Kuhn; Department of Civil and Natural Resources Engineering, University of Canterbury, Christchurch, New Zealand; kenneth.kuhn@canterbury.ac.nz

ABSTRACT

A growing body of research seeks to optimize the selection and scheduling of maintenance, repair and rehabilitation activities for networks of sections of pavement. Such research typically relies on a composite condition index, a one-dimensional and often discrete measure of the overall structural health and/or serviceability of pavement. Pavement can suffer from a large number of related but distinct distresses. Difficulties associated with unobserved heterogeneity have hampered efforts to accurately model deterioration via composite condition indices. At the same time, optimization techniques used in pavement management have been shown both to be sensitive to deterioration model specification and to become computationally intractable as condition data increase.

This research describes how approximate dynamic programming can be used to manage a large network of related sections of pavement each one of which may be plagued by a number of different distresses. Approximate dynamic programming mitigates the curse of dimensionality that has haunted distinct Markov decision problem formulations of the infrastructure management problem and limited their complexity. A computational study illustrates how the proposed approach leads to more sophisticated maintenance decision rules, which can be used to ensure the suggestions of pavement management systems more closely match engineering best practices.

Keywords: pavement management, infrastructure management, deterioration modelling

INTRODUCTION

Pavement management systems inventory historical and current conditions of roadway networks, predict the future conditions of such networks and suggest schedules for maintenance activities. Data related to numerous potential functional deficiencies are collected for different sections of pavement. Examples of the described data include the percentage surface area of a roadway covered in cracks and the depth of ruts in vehicle wheel paths. The volume of data being collected is growing rapidly, in part due to the emergence of automated, often high-speed pavement distress survey techniques (see, for

instance, Tighe et al, 2008). The German authorities have been regularly using high-speed monitoring systems to collect multidimensional data on their federal road network for at least 15 years now (Burger et al., 1994). Data are often combined into a single composite index representing the overall condition of a section of pavement. The composite condition index is then used during deterioration modelling and suggested maintenance action selection within pavement management systems.

A recent study found that composite indices purporting to measure the same attributes of condition rated test sections of pavement substantially differently (Gharaibeh et al., 2010). This conclusion is troubling given the proliferation of different composite condition indices in use today. Difficulties arise even when a single condition index is used. When sections of pavement with similar condition ratings are subject to similar stresses, future condition ratings can “vary considerably” (Carnahan et al., 1987). Models that seek to capture stochasticity have been and are being developed. However, it has proved difficult to select appropriate forms and to accurately parameterize such models. This has led to a distinction between stochasticity and uncertainty, with the latter continuing to haunt pavement management systems (Kuhn and Madanat, 2005). One of the major causes of stochasticity and uncertainty in modelling composite condition index evolution is unobserved heterogeneity (Prozzi and Madanat, 2003). Sections of pavement that are identical in terms of data analyzed currently are, in fact, behaving quite differently.

This problem has been studied previously; see for instance, Madanat, 1993; Madanat and Ben-Akiva, 1994. The authors of the cited studies note that engineers make various observations when evaluating the health of any infrastructure facility. Each observation is “probabilistically related to the true condition of the facility”, which is not directly observable (Madanat, 1993). Madanat introduced a framework where numerous observations of pavement condition could be used to refine estimates of probabilities that the pavement was actually characterized as being in different *latent* condition states. The framework is motivated by the early work of Humplick (1992) modelling pavement condition measurement error. Techniques based on robust optimization have also been developed to optimize pavement management explicitly recognizing both deterioration stochasticity and model uncertainty (Kuhn and Madanat, 2005). The studies described above are based on dynamic programming, which becomes computationally intractable as the space of condition states grows.

In contrast to previous research, this research focuses on separately analyzing the severity and importance of different distresses within the context of network-level pavement maintenance, repair, and rehabilitation action optimization. In order to minimize uncertainty and problems associated with unobserved heterogeneity, it would seem logical to separately consider all available data during optimization. Such an approach is impractical within current pavement management systems given the widespread use of dynamic programming. This research seeks to select maintenance actions for a network of related facilities based upon the most detailed available data, avoiding the use of composite condition indices. Approximate dynamic programming is used to overcome challenges related to computational burden.

Composite Condition Indices

Pavement can suffer from a large number of related but distinct distresses. Depressions may form in vehicle wheel paths, a process known as rutting. Fractures may develop on the pavement surface, which is known as cracking, in a variety of different ways. There may be variations on the surface of the roadway large enough to impact ride quality or raise vehicle operating costs, indicating roughness. There may be smaller variations associated with the macrotexture of the roadway surface, which can impact the skid resistance of the roadway in wet conditions. There may be even smaller variations associated with the microtexture of the roadway surface, which can impact skid resistance in dry conditions. Raveling, flushing, shoving, and spalling are additional potential distresses. When investigating the deterioration of a specific section of roadway and selecting an appropriate maintenance or rehabilitation activity, engineers will often consider separate pieces of data related to separate deficiencies. There have been and continue to be significant research efforts to build mechanistic and mechanistic-empirical models describing how individual distresses evolve over time in a variety of contexts. For example, Henning (2008) models rutting and cracking on New Zealand state highways. Madanat et al. (1995) presents a methodology for modelling the initiation and progression of distresses.

At public agencies worldwide, data regarding different distresses are combined into composite condition indices. Some commonly used examples include the pavement condition index, the present serviceability index, the pavement quality index, the pavement overall index, and the distress manifestation index (Golroo and Tighe, 2009). These indices allow concise representation of the overall condition of pavement. For instance, a particular section of road can be rated 4 out of 5 without having to mention the depth of ruts in vehicle wheel paths, the extent of cracking on the roadway surface, etc.. The composite indices are used to quickly evaluate the condition of a roadway network and to make large-scale maintenance funding decisions. The indices are especially useful for communication with political decision-makers and the general public.

Here, we take the distress manifestation index as an example of a composite condition index. The distress manifestation index is used in Ontario and is based on the results of visual condition surveys describing the progression of several distresses. Measurements of the severity and density of individual distresses are weighted, to reflect their relative importance, and summed. Table 1, from Tighe et al., 2008, shows the weighting factors associated with various distresses, for various types of pavement.

Table 1: Surface distress and weighting factors for distress manifestation index calculation.

Asphalt Concrete Pavement	W	Portland Cement Concrete Pavement	W	Composite Pavement	W	Bituminous Surface Treatment Pavement	W
Raveling	3	Raveling	0.5	Raveling	3	Raveling	3
Flushing	1.5	Polishing	1.5	Flushing	1.5	Flushing	2
Rippling or shoving	1	Scaling	1.5	Spalling	2	Streaking	1
Rutting	3	Potholing	1	Tenting or cupping	2.5	Potholing	1
Distortion	3	Joint cracking or spalling	2	Rutting	3	Rippling and shoving	2
Multiple cracking	1.5	Faulting	2.5	Joint failure	3	Rutting	3
Alligator cracking	3	Distortion	1	Distortion and settlement	1	Distortion	3
Meander mid-lane cracking	1	Joint failure	3	Longitudinal meander	2	Longitudinal cracking	1
Transverse alligator	3	Longitudinal meander	2	cracking		Pavement edge breaking	2
Centerline alligator	2	cracking		Transverse cracking	1	Alligator cracking	3
Pavement edge single/multiple	0.5	Transverse cracking	2	multiple			
Pavement edge alligator	1.5	Sealant loss	0.5	Transverse joint	2		
		Diagonal corner/edge	2.5	reflective cracking			
		cracking		Centerline cracking single	0.5		
				Centerline cracking multiple	1.5		

Equation 1, also from Tighe et al., 2008, shows exactly how the weighting factors are used to calculate the distress manifestation index.

$$DMI = 10 \left(\frac{G - \sum_{i=1}^n W_i (s_i + d_i)}{G} \right) \quad (1)$$

where

i = an index of distress type,

n = total number of distress types considered, for the given pavement type,

W_i = weighting factor of distress type i (Table 1),

s_i = observed severity of distress type i ,

d_i = observed density of distress type i occurrence, and

G = maximum value of weighted summation of distress severity and density observations, for the given pavement type.

Most composite condition indices, like the distress manifestation index, are based on a weighted summation of distress-specific evaluations. As was mentioned earlier, different indices will rate the relative health of different sections of pavement differently (Gharaibeh et al., 2010). According to the cited study, the disparity “can be attributed to differences in the distress types considered, weighting factors, and the mathematical forms of the indices” (Ibid.). Unfortunately, there are no clear correct mathematical forms, weighting factors, or even sets of distress types to survey when constructing a composite condition index.

Maintenance Policy Optimization

The condition indices described are used a few different ways during the selection and scheduling of suggested pavement maintenance, repair and rehabilitation activities. The indices help define user costs and are used in deterioration modelling. An example of a typical approach to maintenance optimization is presented below.

In the example, a network of roadway pavements is to be managed over a finite planning horizon. The condition of each section of pavement in the network is rated each year according to a discrete composite index. The overall condition of the network is defined in terms of the conditions of each of the sections of pavement, and is thus itself a state in a discrete state space. Each year, one action plan from of a set of feasible maintenance action plans is performed on the network. The action plan describes maintenance actions to take on each separate section of roadway in the network. In most cases, budget and resource constraints define the set of feasible action plans and ensure that our decision problem is not separable into section-specific sub-problems.

The social costs of the network being in any condition state with any maintenance action plan applied are either known or can be quickly estimated. These costs account for agency costs

performing relevant maintenance and user costs the public incur using the facility. User costs can incorporate diverse considerations, for example delay to motorists caused by speed limit restrictions reflecting various roadway surface imperfections, accident risk (related to roadway surface macro- and micro-texture) and vehicle ride quality (related to roughness). Future costs are discounted. The goal will be to minimize the sum of discounted costs over time, i.e., to find the social optimal maintenance policy.

Transition probabilities capture the probabilities that the network will be in different condition states one year after being in a given state with a given action applied. These values serve as the model of deterioration. Based on the complex (and somewhat arbitrary) nature of their construction, composite condition indices essentially prohibit the use of mechanistic deterioration models. Established models of distress evolution are typically abandoned and empirical models, often based on quite limited data sets, predominate.

Given these definitions, Bellman's equation is as in equation 2.

$$V_t(i) = \min_{a \in A} \left[c(i, a) + \alpha \sum_{j \in I} p(j | i, a) V_{t+1}(j) \right] \quad \forall t \in \{0, 1, 2, \dots, T-1\}, i \in I \quad (2)$$

where

i, j = indices of (network) condition state,

I = the set of all (network) condition states,

a = an index of maintenance action plan,

A = the set of all feasible maintenance action plans,

t = an index of year number,

T = the length of the planning horizon, in years,

$c(i, a)$ = the (single-year) social cost of being in condition state i and applying action plan a ,

$p(j | i, a)$ = the probability of being in state j one year after being in state i with a applied,

α = the discount factor, and

$V_t(i)$ = the expected discounted future costs when in condition state i in year t .

Assume salvage values, terms representing the costs of ending the planning period with the network in a given condition state, are known for all condition states i in I . Set $V_T(i)$ terms equal to the salvage values. Then use equation (2) to find costs from year $T-1$ onward, for all possible network condition states in I . Repeat the procedure for year $T-2$, and so on. For year 0, solve equation (2) given the current (assumed known) condition of the facility to provide optimal social costs through the planning horizon. The actions that have been chosen along the way to minimize costs provide the optimal maintenance policy, describing what maintenance action to take as a function of year and network condition.

Note that the computational burden rises extraordinarily rapidly as the set I grows. For each combination of year, action plan and network condition state, equation (2) requires summing over the set of network condition states. Note also that the set I grows extraordinarily rapidly as the description of each pavement section's condition becomes more detailed. The network condition state describes each pavement section's condition and thus a network of n

sections which may each be described m ways yields m^n network condition states. In 1982 the state of Arizona was reported to be responsible for 7,400 sections of pavement related by a common maintenance budget (Golabi et al., 1982). The approach to maintenance optimization outlined above is an example of dynamic programming. Concerns regarding computational burden reflect the associated curse of dimensionality.

Past research efforts have yielded techniques for mitigating the curse of dimensionality. Golabi et al. (1982) uses decision variables that describe the fraction of all facilities that are in different condition states. The size of the associated optimization problem is made independent of the number of facilities being managed (though still dependent upon the number of condition states considered). In a precursor to this study (Kuhn, 2010), this author introduces the use of approximate dynamic programming in maintenance optimization.

AN APPROXIMATE DYNAMIC PROGRAMMING FRAMEWORK

Given infinite computing resources, we could use the optimization framework described above involving equation (2) and generalize the notion of a condition state. The condition of our pavement network could contain any and all data relevant to the prediction of future deterioration or estimation of future costs, rather than simplified composite condition index data. This leads us quite naturally to the framework of approximate dynamic programming (as described in Powell, 2007).

Let $Y_t(i_1, i_2, \dots, a)$ represent the value of observing our network in year t and having characterized the status of various distresses using the data sets i_1, i_2, \dots and having selected, but not performed, the maintenance action plan a . Y is similar to V in equation (2) above, and is known as the value function.

Estimating the Value Function

It is common in approximate dynamic programming to model the value function as a linear combination of basis functions (Powell, 2007) as in equation 3.

$$\hat{Y}_t(i_1, i_2, \dots, a) = \sum_{s=1}^S w_s \phi_s(i_1, i_2, \dots, a, t) \quad (3)$$

where

i_1, i_2, \dots are a set of values for (multi-dimensional) network condition data,

s is an index of basis function number,

S is the number of basis functions considered,

$\phi_s(i_1, i_2, \dots, a, t)$ is basis function number s , evaluated in year t when the network is characterized by condition data i_1, i_2, \dots and action a has been selected,

w_s is the weight attached to basis function number s , and

$\hat{Y}_t(i_1, i_2, \dots, a)$ is a function that estimates the expected discounted future costs at year t given network condition data i_1, i_2, \dots and maintenance action plan a .

The basis functions are fixed and capture important attributes of the conditions of the pavement sections being managed, the actions selected and the year of analysis. Weights associated with basis functions are to be learned, generally using simulation. Note that in order to define the estimate of the value function we are required only to set S parameters. This is considerably simpler than setting the value function for each combination of year and observed network condition data as required by equation (2). Actually the number of functions considered and weights to be set is up to us; the computational complexity of the task can be controlled. Naturally, the choice of appropriate ϕ_s functions is very important.

Simulation can be used to set the w_s terms when high quality models are available describing how condition data evolve. Unlike in distinct approaches to maintenance optimization, very few restrictions on the form of deterioration models are necessary. Considering complex models describing different distresses appearing and propagating, according to distinct but inter-related mechanisms is not a problem. Data regarding the condition of a section of pavement need not be limited to a discrete, or even one-dimensional, variable.

Say that in year t of simulation run z , we take action plan a^1 when our pavement network is characterized by condition data i_1^1, i_2^1, \dots . The model would then approximate the future costs of managing the pavement network to be the current estimate of $\hat{Y}_t(i_1^1, i_2^1, \dots, a^1)$. The following (simulated) year we find the network in condition state i_1^2, i_2^2, \dots and apply action plan a^2 . We would like to refine w_s terms so that the estimate of the value function in the initial state moves closer to a new estimate based on what was observed during the simulation. To do so, an error term is defined via equation 4.

$$e = c(i_1^1, i_2^1, \dots, a^1) + \alpha \hat{Y}_{t+1}(i_1^2, i_2^2, \dots, a^2) - \hat{Y}_t(i_1^1, i_2^1, \dots, a^1) \quad (4)$$

where

e is the difference between our original estimate of future costs and the best available approximation one simulated year later, and

$c(i_1^1, i_2^1, \dots, a^1)$ is the (single-year) social cost of having the network in state i_1^1, i_2^1, \dots and applying action a^1 .

It is possible to use a gradient-based approach to correct the estimate of the value function. Note that calculating the gradient is straight-forward given our definition of the value function estimate, as in equation 5.

$$\nabla_w \hat{Y}_t(i_1^1, i_2^1, \dots, a^1) = [\phi_1(i_1^1, i_2^1, \dots, a^1, t) \phi_2(i_1^1, i_2^1, \dots, a^1, t) \dots \phi_S(i_1^1, i_2^1, \dots, a^1, t)]^T \quad (5)$$

The distance to move when refining w_s terms is determined by a step size. Assume the step size varies by simulation run. Then we update each w_s according to equation 6.

$$w_s \leftarrow w_s + \gamma_z e \phi_s(i_1^1, i_2^1, \dots, a^1, t) \quad \forall s \in \{1, 2, \dots, S\} \quad (6)$$

where

z is an index of simulation run number, and

γ_z is the step size during simulation run z .

It is worth noting that learning is being done online and it is only ever necessary to store in memory one set of S weights, two sets of condition data and two actions (the condition data and actions relating to the current and immediately previous year of the current simulation run). Memory requirements pose a significant challenge for dynamic programming approaches to infrastructure management and limit the scope of condition data considered. Approximate dynamic programming approaches to infrastructure management do not have the same difficulties associated with computer memory requirements. It is also worth noting that here weights are being adjusted to reflect maintenance actions being taken. Thus, the estimate of the value function being developed reflects decision-making during the simulation. The procedure outlined here is a heuristic, one way to come up with a reasonable estimate of the value function Y . Further research to identify alternate approaches to value function estimation would be worthwhile.

Selecting Maintenance Actions

Decision making both during simulation runs, when the value function is being estimated, and in practice, when suggested maintenance action plans are being selected, proceeds the same way. The goal is to minimize social costs recognizing the limited resources available to public agencies. Here, this translates into selecting the action plan, in the set of feasible plans, that minimizes the value function. Mathematically, in a year t observing condition data i_1, i_2, \dots we solve the optimization problem posed by equation 7.

$$\min_{a \in A} \hat{Y}_t(i_1, i_2, \dots, a) = \sum_{s=1}^S w_s \phi_s(i_1, i_2, \dots, a, t) \quad (7)$$

The set A is typically restricted by budget and resource constraints. For instance, it may only be feasible to reconstruct a certain length of roadway in one year. Kuhn (2010) notes that under such conditions, it is possible to choose basis functions that are both well-suited to estimating the value function and ensure the optimization problem faced during maintenance action plan selection can be solved efficiently. That approach is used here, with minor modifications.

The idea is to use basis functions that count the numbers of sections of pavement that can be characterized by different combinations of (section-specific) condition data and maintenance action. A condition state can refer to a specific combination of measurements of different distresses (i.e., between 4 and 5 % of the roadway surface cracked and 6-7 mm of rutting in the near-side vehicle wheel path). The maintenance action selection problem becomes as follows.

$$\min_x \sum_{f \in F} \sum_{b \in B} w_{j(f),b} x_{f,b} \quad (8)$$

s.t.

$$\sum_{b \in B} x_{f,b} = 1 \quad \forall f \in F \quad (9)$$

$$\sum_{f \in F} \sum_{b \in B} r_{j(f),b}^k x_{f,b} \leq s^k \quad \forall k \in K \quad (10)$$

where

f is an index of roadway section/facility,

F is the set of all sections/facilities being managed,

b is an index of (section-specific) maintenance action,

B is the set of all (section-specific) maintenance actions,

k is an index of resource type,

K is the set of all resource types considered,

$j(f)$ is the condition of section/facility f ,

s^k is the available budget for resource type k ,

$r_{j,b}^k$ is the amount of resource type k consumed when performing action b on a section of pavement in condition j ,

$w_{j,b}$ is the weight attached to the count of facilities in condition j with action b having been selected, and

$x_{f,b}$ is a binary variable to be set to 1 if the pavement section f will be assigned action b and 0 otherwise.

This formulation is actually a Multichoice Multidimensional Knapsack Problem (MMKP) which has been studied previously by experts in the field of Operations Research, and for which highly efficient exact and approximate solution search strategies exist.

COMPUTATIONAL STUDY

A computational study was run to demonstrate the framework introduced in this work. In this study, simulations of management of 1,000 sections of pavement were run. Ruts developed in the vehicle wheel paths and cracks formed and spread on the surfaces of the pavement. Only two potential distresses were considered for the sake of simplicity.

Cracking was modelled using the work of Madanat et al. (1995), which is based on data collected by the World Bank over seven years of study of highways in Brazil. A joint discrete-continuous model was developed; cracking on a particular section of pavement at a particular time has either been initiated or not and, if initiated, progresses stochastically at a given rate. Variables found to be significant in modelling cracking include the thickness of pavement surfacing, the structural number of the pavement, and the traffic using the pavement. Given these terms, the probability of crack initiation was found to be given by equation 11.

$$\Pr(\text{cracking initiation}) = \Phi \left[1.925 - 0.041h + 0.0000906h^2 - 2.198\hat{s}nc + 1.871(\text{traf} / \hat{s}nc) \right] \quad (11)$$

where

h is the thickness of pavement surfacing,

$\hat{s}nc$ is an estimate of the structural number of the pavement,

traf is the traffic in number of million 80-kN equivalent single axle loads, and

Φ is the standard normal cumulative distribution function.

Given that cracking has been initiated, the expected percent of total area covered by cracks increases each year according to equation 12.

$$E[\text{cracking (\% area per year)} | \text{cracking initiated}] = 6.41 - 0.787\hat{s}nc + 1.645(\text{traf} / \hat{s}nc) - 0.781\lambda \quad (12)$$

where

λ is an error correction term to be set to the standard normal probability density function divided by the standard normal cumulative distribution function, when both are evaluated at the value used in equation 11.

Rutting was modelled using the work of Henning (2008), based on data collected from specially designated long term pavement performance modelling sections of the New Zealand state highway system. In terms of the variables defined above, the model predicts that pavements with a thickness of more than 150mm will develop ruts that progress each year, in terms of mm of added rut depth, according to equation 13.

$$\text{rut progression (mm per year)} = (9.94 - 1.38\hat{s}nc)(\text{traf}) \quad (13)$$

In the computational study described here, all sections of pavement were assumed to be 200mm thick and to have a structural number of 3. Yearly traffic volumes were described as random variables uniformly distributed between 10,000 and 20,000 80-kN equivalent single axle loads. There were only two options available to the responsible public agency regarding a particular section of pavement during a particular year. The pavement could be reconstructed, or left alone. All assumptions were made again primarily for the sake of simplicity. The goal was not to develop a highly realistic simulation of a pavement management system, but only to show some of the pitfalls of the common approach to infrastructure management.

No models could be found describing the costs users incur driving over pavement, as a function of the severity and density of distinct distresses. In the study run here, it was assumed that ruts of a depth greater than 15mm and cracks that covered more than 20% of the surface area of a section of pavement indicated substandard conditions. This assumption was based on discussions with traffic engineers in New Zealand. The likely results of having substandard pavement include decreased ride quality, decreased vehicle speeds, and increased delay. There may also be an increased accident risk associated with substandard pavement, depending upon the severity of cracking and rutting as well as the

strength of the statistical links between cracking, rutting and other roadway surface imperfections. A cost ten times as great as the cost of roadway reconstruction was incurred for every year a section of pavement was characterized as being substandard.

Two trials were run using the approximate dynamic programming framework for pavement management described previously. Both trials involved describing the condition of a section of pavement in a year as being in one of 10,000 possible states. In trial 1, cracking and rutting were each evaluated and assigned one of 100 possible ratings. In trial 2, data regarding cracking and rutting were combined into a hypothetical composite index. Weighting factors of 7 for cracking and 3 for rutting were used (based loosely on Table 1), although it will be shown later that the weights are largely irrelevant to the points the results bring up. Over the course of 100,000 simulation runs, value functions were estimated and optimal policies developed. Value functions represented the expected discounted future costs incurred when a section of pavement is in a condition state defined by cracking and rutting progression ratings. The entire study took less than four minutes when run in C on a 2.53GHz Intel Core 2 Duo Mac mini personal computer with 4GB memory. Figure 1 shows a graphical representation of the optimal policies found.




Figure 1: Optimal policies with and without using a composite condition index.

The graph on the left-hand side of Figure 1 represents the maintenance policy suggested after trial 1, and the graph on the right-hand side that suggested after trial 2. In both graphs, the extent of rutting increases as you go to the right while the extent of cracking increases as you go down the graph. Sections of the graphs coloured black indicate pavement conditions that would trigger the recommendation of roadway reconstruction. Sections of the graph coloured white indicate conditions that would not trigger such a recommendation. The small irregularity in the policy recommended when not using a composite condition index highlights the fact the value function was estimated using simulation and can include small irregularities. It also indirectly shows the level of detail with which it is possible to

characterize the condition of a section of pavement when using an approximate dynamic approach to infrastructure management.

It is clear from Figure 1 that, in this computational study, the use of a composite condition index resulted in road reconstruction being recommended even when there was no danger of breaching the desired service standards. The condition index is unable to tell the difference between sections of pavement with worrying amounts of cracking but very little rutting and sections of pavement with less severe cracking but some rutting. Reconstruction is recommended for all sections of pavement where the condition index *could have been caused by* excessive cracking. There is a discrepancy between what engineers have defined as substandard and how the optimization framework has been set up. Note that the one-dimensional composite condition index yields a one-dimensional decision rule. A straight line separates conditions where reconstruction is recommended from conditions where no maintenance is recommended. Altering the weights attached to rutting and/or cracking would alter the slope of the line, but not the form of the decision rule. The result is a recommended maintenance regime that does not match the prerogatives of engineers. There is the potential for public agencies to overspend on maintenance or for engineers to lose faith in the pavement management system.

The relative importance of the points brought up here depends upon a few factors. The first of these factors is the extent to which the evolution patterns of distinct distresses are related. If cracking, rutting, flushing, shovelling, spalling, etc. are all symptoms of one underlying deterioration process then the use of a composite condition index makes some sense. Further research is needed to examine how inter-related distresses are, from a statistical point of view. A (related) factor worth considering is the extent to which engineering decision-making at the facility-level, as well as roadway user costs, can be described using one dimensional measures of pavement condition. Network-level pavement management systems for managers model decision-making at the facility-level done by engineers, and the fidelity of the model depends upon how well the nuances of engineering judgement are captured. The techniques introduced here could be used immediately to revamp decision making within network-level pavement management systems. The magnitude of the benefits of doing this are difficult to accurately estimate at this point.

CONCLUSION

Much of the literature on optimizing the selection of pavement maintenance activities revolves around one-dimensional and discrete measures of overall pavement condition. Deterioration is modelled using probabilities describing how composite condition indices evolve over time, but it has proven difficult to accurately establish such probabilities. In truth, pavement deteriorates as related but distinct distresses accumulate and propagate. Distress-specific, rather than overall condition, data is often collected. This work introduces a framework for separately considering the evolution of different distresses in a network-level pavement management system. Approximate dynamic programming is used to ease the computation burden, which would be too great if conventional approaches to infrastructure management were simply extended.

There appears to be a widely-held belief that existing mechanistic and mechanistic-empirical models of how facilities deteriorate must be simplified during maintenance, repair and rehabilitation strategy optimization. Composite condition indices are likely useful in communications with politicians and the general public, but they are not required for maintenance optimization. All data available should be considered, until such a point as some is found to be redundant, when selecting maintenance action plans. This is what engineers do in practice when examining individual sections of pavement. Further work is needed to model costs associated with various distresses, to forecast the evolution of individual pavement distresses and to establish basis function for use in approximate dynamic programming based infrastructure management.

REFERENCES

- Burger, W., Canisius, P.P. and P. Sulten (1994). Toward a new pavement management system in Germany: Organization, data collection, experiences, and innovations.
- Carnahan, J., Davis, W., Shahin, M., Kean, P. and M. Wu (1987). Optimal maintenance decisions for pavement management. *J. of Trans. Eng.* 113, 554-572.
- Gharaibeh, N. G., Zou, Y. and S. Saliminejad (2010). Assessing the agreement among pavement condition indices. to appear in *J. of Trans. Eng.*
- Golabi, K., Kulkarni, R. and G. Way (1982). A statewide pavement management system. *Interfaces* 12, 5-21.
- Golroo, A. and S. Tighe (2009). Fuzzy set approach to condition assessments of novel sustainable pavements in the Canadian climate. *Can. J. of Civ. Eng.* 35, 754-764.
- Henning, T. (2008). The development of pavement deterioration models on the state highway network of New Zealand. PhD thesis, University of Auckland.
- Humplick, F. (1992). Highway pavement distress evaluation: modelling measurement error. *Trans. Res. B* 26, 135-154.
- Kuhn, K. (2010). Network-level infrastructure management using approximate dynamic programming. to appear in *J. of Inf. Sys.*
- Kuhn, K. and S. Madanat (2005). Robust maintenance and rehabilitation policies for a system of infrastructure facilities. *Trans. Res. C* 13, 391-404.
- Madanat, S. (1993). Optimal infrastructure management decisions under uncertainty. *Trans. Res. C* 1, 77-88.
- Madanat, S. and M. Ben-Akiva (1994). Optimal inspection and repair policies for infrastructure facilities. *Trans. Sci.* 28, 55-62.
- Madanat, S., Bulusu, S. and A. Mahmoud (1995). Estimation of infrastructure distress initiation and progression models. *J. of Inf. Sys.* 1, 146-150.
- Powell, W. (2007). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. John Wiley & Sons, New York.
- Prozzi, J. and S. Madanat (2003). Incremental nonlinear model for predicting pavement serviceability. *J. of Trans. Eng.* 129, 635-641.
- Tighe, S., Ningyuan, L. and T. Kazmierowski (2008). Evaluation of semiautomated and automated pavement distress collection for network-level pavement management. *Trans. Res. Rec.* 2084.