Microcomputer Optimization of Light-Pavement Rehabilitation

PAUL D. THOMPSON, RAIMO O. TAPIO, AND JUHA ÄIJÖ

A Markovian decision model on a personal computer is being developed for Finland's Roads Administration to optimize the allocation of funding among rehabilitation actions, geographic regions, and traffic volume classes on roads with oil gravel and other light pavements. In the long term, the model recommends a funding allocation and condition distribution that minimizes the sum of agency and user costs; in the short term, it recommends a strategy to maximize annual progress toward the long-term goal. As traffic volume and network length grow, the system simulates their effect on rehabilitation policy.

The Finnish National Roads Administration, with support from the consulting firms Viasys Oy of Finland and Cambridge Systematics, Inc., of the United States, is developing a national economic model to optimize light-pavement rehabilitation policy and the allocation of funding. The model is implemented as an extension to the administration's Highway Investment Programming System (HIPS) (*1*), which has been operational for all of the country's asphalt-surfaced roads for 2 years. It is able to analyze general rehabilitation procedures ranging from general patching to total reconstruction for all of the nation's 51,601 km of light pavements, evaluate the need for converting oil gravel pavements to asphalt concrete, and analyze the distribution of funding between the two types of pavement. Because it is a network-level model, the system analyzes policies at an aggregate level, considering only classes of roads, or subnetworks. A separate project analysis system (PAS) (*2*), now used in the district offices, relates the policy and budget recommendations of the network optimization to specific data bases of road segments.

Central to the optimization model is a Markov dynamic program, which has been formulated as a linear programming problem for solution by off-the-shelf software. The dynamic program categorizes pavements into 108 condition states and five actions, and represents deterioration as the probability of making transactions among all possible pairs of states over 1 year. An agency cost model estimates the cost of each possible action, and a user cost model evaluates the results in terms of travel time, fuel consumption, and vehicle depreciation. In selecting optimal actions for each possible condition state, the model tries to find a level of rehabilitation that balances the higher user costs of poor maintenance against the higher agency costs of good maintenance.

Separate models are available for each permutation of two climatic regions and three traffic volume classes; the Markov model optimizes budget allocations within each of these six subnetworks, and an incremental benefit-cost procedure based on parametric analysis results optimizes the allocation of funding among them and between light pavements and asphalt concrete.

A complete data set for the models has been prepared, and the necessary modifications to HIPS are nearing completion. As the section on performance and implementation issues will show, data collection efforts have been comparable in scope with any other pavement management methodology. Software performance on the microcomputers has been excellent for the asphalt models and is expected to be even better on the light pavement models. The system shows great promise for improving the economic basis for capital budgeting decisions on Finland's low-volume roads.

OVERALL MODEL STRUCTURE

The Finnish National Road Administration and the 13 district offices are highly independent in their day-to-day activities, with the central administration playing an administrative role and providing consulting services in new technologies and road and traffic research. Each year, the administration prepares rehabilitation and maintenance budgets and negotiates objectives with the districts and the Ministry of Transportation.

Annual and long-term road management objectives are set for each district, and these are strictly tied to the coming budget. The districts execute all maintenance actions, with the central office taking little interest in the specific actions chosen as long as overall objectives are met. Objectives are set for 1- and 5-year periods by mutual negotiation with the director general of the administration and the districts' chief engineers; they are then approved by the Ministry of Transportation. The results are monitored over the following year. As an example of such maintenance objectives, the 1990 light-pavement objectives have been as follows:

- The bearing capacity of oil gravel roads should not get any worse.
- The road condition (defects) index should be at least 3.5 on roads where traffic volume is above 800, and 2.5 for roads under 800.

Using the HIPS models, the optimal objectives, road policy, and strategy are defined for the long and short term, to bring the road network to desired condition levels. The districts then prepare the capital program, including the definition of

P. D. Thompson, Cambridge Systematics, Inc., 222 Third Street, Cambridge, Mass. 02142. R. O. Tapio, Finnish National Road Administration, Opastinsilta 12A, 00520 Helsinki, Finland. J. Äijö, Viatek Oy/Viasys, Ahventie 4A, SF-02170 Espoo, Finland.

projects, actions, locations, and costs. Accomplishment of the objectives is judged by annual measurements taken in the fall.

The actual implementation of maintenance measures is done either by the districts' own forces or by contractors. The central administration sets maintenance and design standards for this work and provides funding consistent with these standards.

To address this division of responsibilities, the pavement management system consists of two separate software packages:

- Network level, embodied in HIPS, which addresses abstract categories of roads and the allocation of funding among them. HIPS helps the central administration determine the funding level for each district and develops policy guidelines on how the money is to be used.
- Project level, embodied in PAS, which guides the district engineer in matching the network-level policies with detailed data bases of roads that already exist in each district.

Although both the districts and the central administration are potential users of the entire system, the central office is the primary user of the relatively abstract modeling and allocation procedures in HIPS, whereas the districts are the primary users of the more detailed PAS. Several levels of analysis are provided in HIPS to address the capital programming policy questions of interest to the highway administration. These are as follows:

- Pavement type level distinguishes asphalt concrete pavements from light pavements.
- Region level distinguishes the cold and relatively dry interior and northern parts of the country from the more temperate and moist southern coastal area. Climatic differences affect the behavior of pavement deterioration and the choice of maintenance policy.
- Volume class level affects the rate of deterioration of pavements as well as the level of user costs associated with pavement condition. Low-volume, oil gravel roads are those with average daily traffic below 350, high-volume roads are those above 800, and medium-volume roads are those in between. As traffic volume on a road increases above 800 toward 1,500, it receives consideration for paving with asphalt concrete. Much of the modeling work occurs on the level of the 12 permutations of pavement, region, and volume class, termed the "P/R/VC" level.
- National level accumulates the results from the pavement types, regions, and volume classes to address national funding levels and funding allocations.
- District level reflects the 13 maintenance districts that are the recipients of capital funding and that carry out all of the programmed work. Each district is contained in one region and has roads representing all three volume classes.

Table 1 summarizes the light-pavement network at the region/volume class level.

As in any far-sighted capital programming process, the roads administration is concerned with the long-range goals that should be established for the highway network, and also with the steps needed to proceed from the current situation toward the long-range goals. Because the projects addressed by the pavement management system are largely ones of in-kind facility replacement and major maintenance, the roads administration would like decision making in this area to concentrate on direct economic benefits and costs to road users. This expectation leads to the next important division within HIPS:

- The long-term model analyzes possible long-term goals and tries to find a future policy that minimizes social costs (the sum of user and agency costs) and is sustainable indefinitely into the future. The long-term model is not tied to the current condition of the network and imposes no requirements on which specific year it should be achieved.
- The short-term model's first priority is to find the quickest practical means of achieving the level of network condition that would make the long-term policy possible, and its second priority is to minimize the social cost incurred in the short-term period between now and the time when the long-term goals are achieved.

As shown in Figures 1 and 2, the flow of activities in using HIPS starts at the most abstract level and ends at the most concrete level. The long-term model analyzes the general behavior and cost structure of roads in each of the six subnetworks of region and volume class, and then uses an incremental benefit-cost model to determine the best allocation of funding among them. It defines goals broadly and at some undetermined time in the future. This then proceeds to the short-term model (Figure 2), which is more concrete because it is explicitly tied to the current observed condition of the road network (from the Road Data Bank). Following this, the analysis becomes even more concrete in the short-term budget allocation steps, an activity of immediate interest to the managers within the administration and the districts. Finally, the least abstract activity is the definition of actual projects on specific roads, in PAS. This flow of abstraction follows the general flow of the administration's planning process, and provides a way in which the economic merits and costs of rehabilitation policy can be conveniently merged with the noneconomic and political considerations that also determine the ultimate budget allocations and capital program.

TABLE 1 LIGHT-PAVEMENT CHARACTERISTICS IN FINLAND

Average Daily Traffic (vpd)	Road Length (km)		Pavement Width (m)		Average Speed (kph)	
	North	South	North	South	North	South
<350	15,363	7,424	5.7	6.0	45	45
350 to 800	12,309	7,828	6.3	6.3	62	45
>800	5,258	3,419	6.6	6.6	67	55

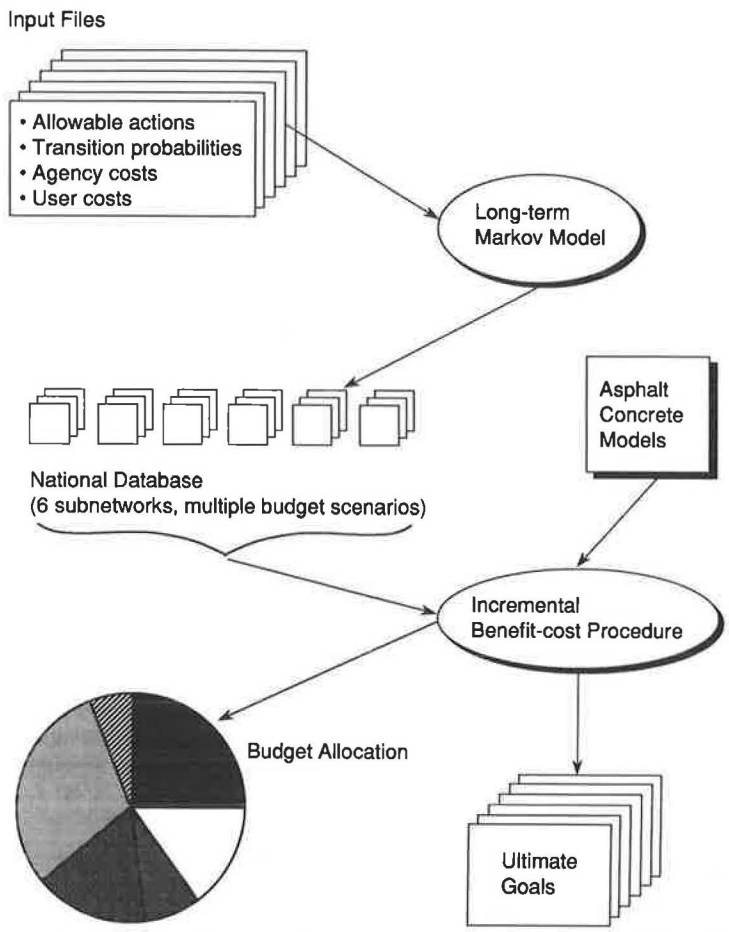

FIGURE 1 Long-term goal setting.

In both the long-term and short-term analyses, the central feature is an economic model of pavement behavior, rehabilitation policy, and their combined social cost implications. The main components of the economic model are as follows:

- Agency cost model, giving the average construction costs for five general categories of rehabilitation action, from do nothing to total reconstruction.
- User cost model, which quantifies in economic terms the increase in travel time, fuel consumption, and vehicle wear-and-tear associated with deteriorated road condition. Even though the absolute level of user costs on light pavements is quite low, because of low traffic volumes, it is expected that the high dependence of travel speeds on rehabilitation policy will still have a significant effect on the selection and timing of actions.
- Deterioration model, describing the process by which a road deteriorates and thereby causes higher user costs to be incurred. Similarly, it also describes the improvements that can be expected after each of the general rehabilitation actions is applied.

As the expenditure of agency costs increases, the resulting level of user costs is expected to decline, as long as the available money is always used in the most cost-effective manner; also, as agency costs decrease, user costs go up. The economic optimization framework assumes that there is an intermediate point at which social cost, the sum of user and agency costs, is minimized. Policy questions that are addressed in the framework are as follows:

- What is the optimal (minimum social cost) level of expenditure on oil gravel pavement rehabilitation and on major maintenance on the nationwide road network, and what are the optimal expenditures within selected subnetworks?
- At funding levels that do not minimize social costs, what is the optimal allocation of funding among subnetworks, and what is the most cost-effective means of spending the available money? What is the best overall allocation among action types, and what specific actions should be applied to the various kinds of roads?
- To what extent do budget constraints increase the level of costs borne by road users, and what does this imply about the importance to society of user costs relative to agency costs?
- How much better is the long-term optimal solution than the current situation, and how long must it take to achieve the long-term goals?

It should be emphasized that the issue of accessibility, which is the major justification given for the construction of oil gravel roads, is not a factor in the choice of actions on existing

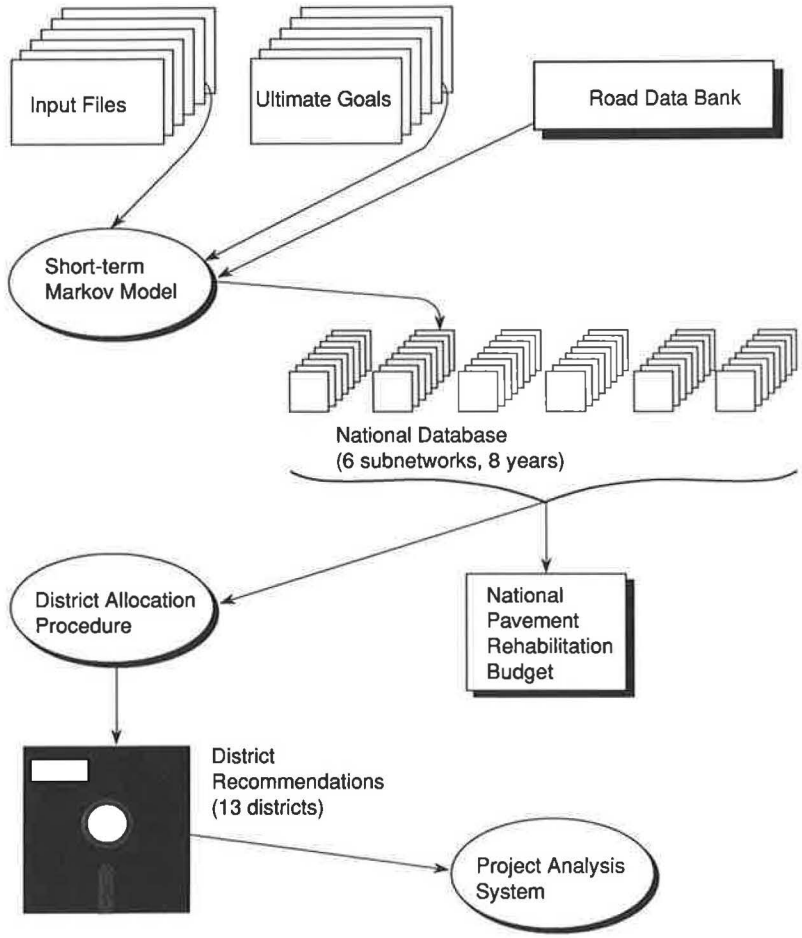

FIGURE 2 Short-term budgeting.

light surfaces. It is assumed in the models that existing oil gravel roads will be kept at least passable to all traffic.

Many different modeling methodologies can be applied to these questions. At the beginning of the earlier development effort for asphalt concrete, a thorough review was conducted of the experiences of U.S. and international organizations using many different techniques. The methodology finally selected was an adaptation of Markov dynamic programming, a technique that had been used in optimization applications as diverse as fleet replacement, catalog mailing list selection, timing of bond calls, and purchase of satellites (*3*). At the time, the only full-scale implementation of the technique in pavement management was in Arizona (*4*). Attributes of the Markov model that made it attractive were as follows:

- It is intuitively appealing to top managers in the administration, as it describes the behavior of pavements in a manner that is simple and fits well the organization's structure and decision-making processes. Also, it is an innovative approach that can be useful to other countries if it works out well.
- It explicitly recognizes the stochastic nature of pavement behavior, and therefore expresses its conclusions in a form that, in comparison to deterministic models, can more readily be defended against anecdotal evidence. This feature is especially important when introducing the technique to skeptical district engineers.
- Creative and relatively simple techniques can be used to aggregate or allocate the conclusions of the models to higher and lower levels of abstraction for different purposes.
- It can be implemented on a microcomputer.

The point concerning microcomputers was not immediately obvious for a road network of the size in Arizona when the model must be solved by formulating it as a linear program using specialized computationally intensive software. However, the key to a successful microcomputer solution is the separation of the abstract, network-level model from the detailed data base of potential actual projects, so that the computational requirements of the two do not compound each other. The network-level model optimizes a small number of general categories of roads, defined by their geographic region, volume class, and condition classes. The definition of categories must be designed carefully to give the needed policy sensitivity, but they operate at a level of aggregation with which the central administration is most comfortable. After the optimization is completed, the resulting recommendations are transmitted on a floppy disk to the district offices, where they are applied to the large detailed data base of actual road segments using less computationally intensive techniques in PAS.

Both HIPS and PAS are menu oriented and user friendly. All through the system, analytical and summary reports and screen displays are available to show how the models are developing. A convenient filing system manages the input and output files on the hard disk and allows the user to keep track of multiple versions of an analysis archived on floppy disks or hard disk subdirectories. All activities involved in using HIPS are available from a menu hierarchy; at each leaf of the menu tree, a screen representing the module or report explains the purpose of the module and allows the user to fill in various options that are available for that module.

INPUT DATA AND ORGANIZATION OF THE MARKOV MODELS

The formulation and organization of the Markov models is loosely based on the formulation used in Arizona, but conditions are categorized in a much different way, and user costs are included explicitly in the objective function. Also, the desire to implement the system on microcomputers has introduced both design constraints and design opportunities that, in the end, have led to a system that is unique. Markov dynamic programming can be distinguished from other optimization approaches by several features:

- Problems are structured into multiple stages, which are solved one stage at a time. All the stages are structured identically to each other, and all have the same possible outcomes. The stages are evenly separated from each other, usually in time, in a uniform progression.
- The range of possible outcomes of each stage is expressed as a set of discrete states. It must be possible to write out a reasonably short list containing every possible result.
- The outcome of any particular stage depends stochastically on the outcome of the stage before it, and not on the outcomes of any of the other stages. Because of this, Markov dynamic programs are said to have "one-step memories."

By applying a Markov model recursively over a series of stages, it is possible to predict probabilistically the outcome of any future stage. Such a series of Markov predictions is called a Markov chain.

For the purposes of Finland's pavement management system, each stage is a description of the condition of the road network in a given year in terms of the distribution of roads among the set of possible states, combined with the choice of action taken in that year. Figure 3 shows how, in a system of five states, a Markov chain of deterioration plays out for a road starting in the highest state. As expected, the road ends up in the lowest possible state eventually, but the path it uses to get there may vary. Although such a stochastic prediction may be of limited use in designing the treatment for any particular road, it is useful for characterizing a whole road network.

Pavement deterioration and resulting capital needs are driven by the effects of climate and traffic, which cause both visible damage to the pavement and hidden loss of resistance to further deterioration. The visible damage—roughness and rutting—increases user costs and makes it necessary to invest in maintenance. Measurements of all of these factors must exist in any complete pavement management system. HIPS makes an important separation between factors that are exogenous and factors that are affected by maintenance policy. The endogenous factors measured in Finland are

- Longitudinal and transverse roughness, which most directly affect user costs;
- Distress index, a weighted average of several distress measures, which is visible evidence of pavement weakening and damage that may lead to future roughness; and

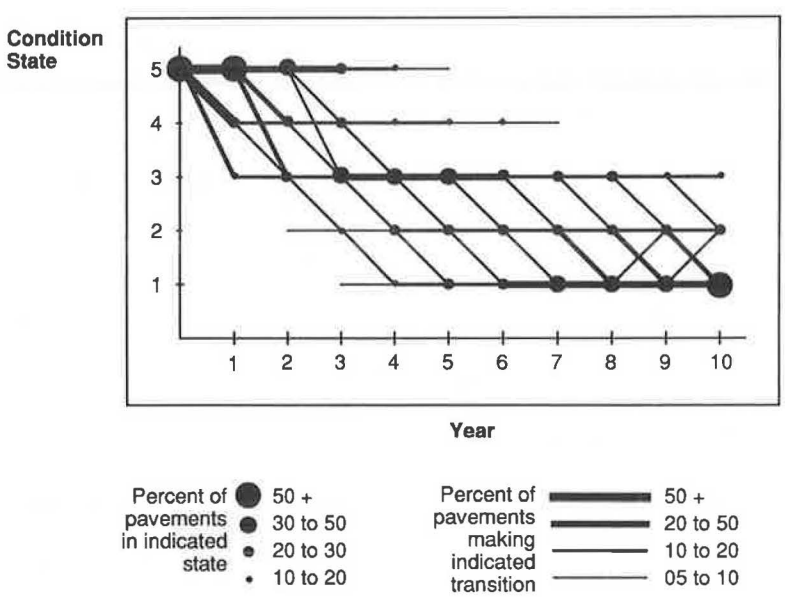

FIGURE 3 Markovian deterioration.

- Bearing capacity, which at lower values presages future distress and roughness.

Roughness is measured by a bump integrator (longitudinally) and a rod and ruler (transversely); distress is estimated visually. Finland is unusual in the large extent to which it measures bearing capacity, driven by the need to monitor the damage of annual freeze-thaw cycles. In countries that do not experience such damage, monitoring of bearing capacity is far less important.

To set up the optimization data base, Finland's oil gravel roads are surveyed for these four condition variables, and then placed in discrete categories: three categories each for longitudinal roughness, transverse roughness, and distress; and four categories for bearing capacity. The combination of these categories yields the 108 condition states.

Climate and traffic, the exogenous factors, are incorporated in the models by classifying all roads into two regions and three volume classes. Each of the six combinations of these variables is modeled separately; a set of optimal investment policies, with resulting user costs, is generated at several alternative budget levels for each one. An incremental benefit-cost procedure uses the alternative budget levels to choose the most equitable national allocation of funding. Traffic growth is modeled by the projected change in kilometers in each traffic volume class.

One strength of the Markovian methodology is that traffic growth and network size have no impact on the size or complexity of the optimization models; they work just as well on large networks or small ones, and policy recommendations within each traffic volume class are not affected by growth. Only after the optimal policies are determined is it necessary to consider the impact on individual roads, and this separation between optimization and data base manipulation is what makes the process simple and implementable on microcomputers.

Data have been developed by the administration to estimate the deterioration and cost models. For the do nothing action category (which includes routine maintenance), transition probabilities are directly estimated from the Road Data Bank; for other actions, probabilities are developed from a Delphi process of interviewing expert engineers. Because of the memory-less property of Markov models, only 2 years of data on a sample of roads is required for model estimation, and the only data items required in Finland's models are traffic volume, roughness, defects, and bearing capacity.

Model estimation is a simple matter of observing the roads in each condition state in the first year, and observing the fraction of them that transition to each other state in the second year. Monitoring, validation, and updating of the models is equally easy and is a necessary part of routine pavement management. Agency cost factors are developed from the administration's own historical records, whereas user costs are estimated by the administration's research center from experimental evidence and economic data resulting from more than 10 years of research (*5*).

LONG-TERM MARKOV MODEL

Mathematically, the long-term Markov model assumes a steady state distribution of pavements among the 108 condition states. This does not mean that each road is always in the same condition, but it does mean that, in every year, the same overall fraction of roads may be found in the same state. It also means that the same fraction of roads undergoes the same general action each year. This is all part of the requirement that the long-term program be not only optimal, but also sustainable, indefinitely.

Of course, the condition distribution among the 108 condition states measured today is generally not equal to the long-term optimal distribution. In fact, the long-term model is not in any way tied to the current condition distribution or current rehabilitation policy, but represents instead a goal that might be attained at some time in the future. What makes the goal desirable is that it minimizes social costs. For the purposes of translating the model to a linear program, the social cost of the long-term program is calculated in the following manner as the objective function to be minimized:

$$\text{Social Cost} = \sum_a \sum_i W_{ai}(C_{ai} + \phi U_{ai}) \tag{1}$$

where

W_{ai} = decision variable, the fraction of all pavements that are in State i and have Action a applied to them (there are 108 states and 8 actions);
U_{ai} = user cost factor in marks per kilometer;
C_{ai} = agency cost factor in marks per kilometer; and
ϕ = degree of user cost contribution to the objective function, usually 1. (It can be varied in an automated parametric analysis provided in HIPS.)

To prevent leakage from the system, and to give scale to the W_{ai} decision variables, the first constraint on the linear program is a definitional unity constraint:

$$\sum_a \sum_i W_{ai} = 1 \tag{2}$$

The most important element of the formulation is a constraint that combines the Markov model with the requirement of a steady state:

$$\sum_a \sum_i W_{ai}P_{aij} = \sum_a W_{aj} \qquad \text{for all } j \tag{3}$$

where P_{aij} is the transition probability of going to State j in Year $t + 1$ (given State i in Year t, when Action a is applied in Year t), which does not depend on t.

This constraint starts with the distribution of condition states and actions in Year i, applies the transition probabilities to get the distribution in the next year, Year j, and then says that this resulting distribution must be the same as in the previous year, so that all years have the same condition distribution. Loosely interpreted, this means that, for every kilometer of road leaving any given state, another kilometer must arrive in that state to replace it. Next comes an optional budget constraint, which can force the agency cost total to a level either higher or lower than the social-cost minimization level.

$$\text{BMIN} \le \sum_a \sum_i W_{ai}C_{ai} \le \text{BMAX} \tag{4}$$

where BMAX and BMIN are budget constraints (marks/km) that can be varied in an automated parametric analysis.

Finally, the long-term model has optional condition constraints. To make these most relevant and usable, condition constraints are applied to condition classes, rather than states. There are 4 bearing capacity classes, 3 defects classes, 3 rutting classes, and 3 roughness classes, for a total of 13. Each state belongs to four different classes, one in each condition variable. The constraints are

$$\mathrm{CMIN}_c \leq \sum_{i \in c} \sum_{a} W_{ai} \leq \mathrm{CMAX}_c \qquad \text{for all } c \tag{5}$$

where CMAX_c and CMIN_c are fractions that represent the limits on the total fraction of pavements allowed to be in each class c.

The parametric analyses on user cost contribution and budget constraints are important. Both capabilities produce a series of different scenarios in which different levels of agency costs are incurred and different levels of user costs result. These scenarios are used in the allocation of resources among the subnetworks, as described in the next section.

ALLOCATION AMONG REGIONS AND VOLUME CLASSES

By setting up and running a long-term Markov model for each permutation of pavement, region, and volume class, the analyst can determine the social cost–minimizing pavement rehabilitation policy for each subnetwork. But the agency cost required to maintain this policy may not be within the range of funding that the administration can expect. To make the model more realistic, budget constraints are required. The question is, what constraint should be applied to each subnetwork? or, put another way, how should funding be allocated among them?

A way of looking at this issue is to assume that the national funding level for pavement rehabilitation represents the relative value that society places on user costs as opposed to agency costs. In HIPS, this value can be analyzed by varying the factor ϕ in the long-term objective function. One way to conduct budget allocation is to use the same ϕ in all 6 models (12 models when asphalt concrete roads are included) and to experiment with the value of ϕ until a set of solutions is found whose agency cost equals the expected national funding level.

HIPS conducts an analysis similar to this in a systematic manner by using a variation of the incremental benefit-cost technique. As Figure 4 shows, a long-term parametric analysis (either on the user cost contribution ϕ or on the budget constraints) produces a user cost versus agency cost curve that is always concave upward. Between any two parametric steps, the quotient of change in user cost divided by change in agency cost is called the "shadow price." The shadow price is always negative and always becomes less negative as agency cost increases.

A criterion for an economically efficient allocation of resources among the 6 or 12 subnetworks, at funding levels that do not minimize social costs, is that the shadow price is equalized among the subnetworks. In an allocation where this is true, it is not possible to save additional user costs by shifting money from any one P/R/VC to any other. If the user executes a parametric analysis for all 6 or 12 models, then the incremental benefit-cost procedure can use the algorithm shown in Figure 5 to create sets of allocations at varying budget levels that roughly equalize the shadow prices. In practice, the user first sets up a broad parametric analysis for each P/R/VC, including the highest and lowest conceivable budget levels, to give a rough allocation of funding over a wide range of budget scenarios. Then the parametric analyses are conducted again for the smaller budget ranges suggested by the first report, and the allocation report is reprinted to give more precise recommendations.

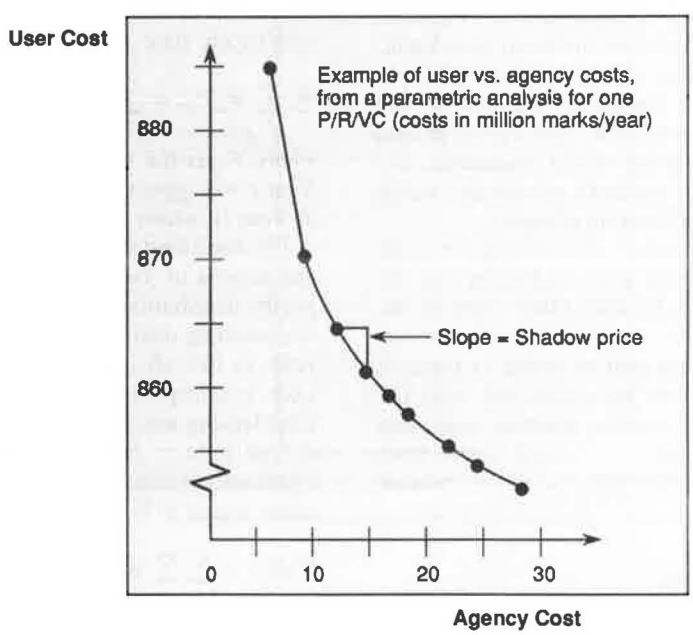

FIGURE 4 Long-term parametric analysis.

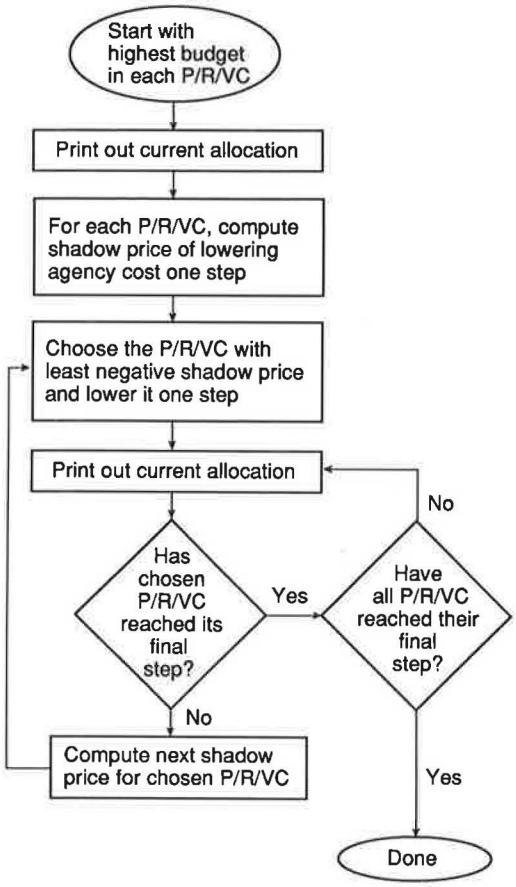

FIGURE 5 Incremental benefit-cost allocation.

SHORT-TERM MARKOV MODEL

All of the analysis described so far focuses on an optimal scenario that may take place at an undetermined time in the future. What is needed next is a capability to model the steps required to bring the existing road network to the condition level needed by the chosen long-term solution. The purpose of the short-term model is to find the shortest practical means of achieving the long-term condition distribution, minimizing as much as possible the social costs incurred along the way.

This is done by modeling each year separately in the short term, with the objective function trying to minimize the deviation between the condition distribution at the end of the year and the long-term optimal distribution. In other words, the short-term model tries to maximize the amount of progress made each year. Mathematically, the objective function is

$$\text{Min} \sum_c (K_c X_c + K_c Y_c) \tag{6}$$

where

X_c = amount by which the fraction of pavements in Class c at the end of the year is above the long-term optimal fraction (0 if it is below the long-term fraction);

Y_c = amount by which the fraction of pavements in Class c at the end of the year is below the long-term optimal fraction (0 of it is above the long-term fraction); and

K_c = average social cost (marks/km) of Class c. K_c is computed from the long-term steady state condition distribution and the unit agency and user costs by the expression

$$K_c = \left[\sum_{i \in c} \sum_a W_{Tai}(C_{ai} + U_{ai})\right] \Big/ \left[\sum_{i \in c} \sum_a W_{Tai}\right] \tag{7}$$

where

W_{Tai} = long-term fraction of pavements in State i with Action a,

U_{ai} = user cost factor (marks/km), and

C_{ai} = agency cost factor (marks/km).

The weights K_c placed on the objective function variables ensure that the classes having the highest social costs receive the highest priority. Using the X_c and Y_c variables allows the Markov constraints to be expressed in terms of the long-term condition distribution, as follows:

$$\sum_{j \in c} \sum_a \sum_i W_{ai} P_{aij} - X_c + Y_c = \varepsilon_c \qquad \text{for all } c \tag{8}$$

where

P_{aij} = transition probability of going to State j in Year $t + 1$, given State i in Year t, when Action a is applied in Year t; and

ε_c = short-term ultimate distribution among condition classes, usually computed as a summation of the long-term state distribution in the same manner as the long-term condition constraints.

Because this formulation is run separately for each year, a constraint is needed for the initial condition distribution at the beginning of the year. In the first year, this constraint is the known distribution that has been measured in the field and is maintained in the Road Data Bank. In the subsequent years, the constraint is the distribution output from the preceding year's solution.

$$\sum_a W_{ai} = q_i \qquad \text{for all } i \tag{9}$$

where q_i is the condition distribution at the beginning of the year.

The short-term model also has optional budget constraints, which are identical in structure to those used in the long-term model. A report is provided in HIPS to allocate a fixed short-term budget level, above and beyond the amount needed to maintain the long-term distribution, according to the amount of improvement needed to reach the long-term solution; the amount of improvement needed is the Year 0 value of the short-term objective function. Normally, the analyst's judgment and noneconomic considerations also play a major role in setting the short-term budget levels.

A file of ultimate goals for the short term is prepared from the long-term solution, expressed as the fraction of pavements in each of the 13 condition classes. The user is free to edit these goals and the weights associated with each condition class. To run the short-term model, the user gives a series of maximum and minimum budget levels for each year, and then

starts the optimization process. HIPS automatically executes a sequence of eight optimizations, one for each year, and places all of the results in a short-term results file for later reporting. Over the short-term period, the condition distribution approaches the long-term optimal distribution asymptotically, as long as the short-term budget level remains above the long-term agency cost.

An alternative short-term approach, in principle, is to add to the long-term formulation a state variable representing time. For the administration's 8-year planning period, then, there would be 108 condition states multiplied by 5 actions, multiplied by 8 years, for a total of 4,320 states. The model could plot a series of yearly policies in a Markov chain, culminating in a constraint at the end of the eighth year that would force the model to reach the optimal long-term distribution. A short-term model of this sort, for a period of 5 years, has been developed and solved in Arizona's system on a mainframe computer. Although the approach chosen for HIPS does not have the technical guarantee of optimality that the alternative would, it has shown in 2 years of use with asphalt pavements that it produces practical recommendations that are consistent with the long-term optimum, and it has the distinct advantage of being operable on a microcomputer.

IMPLEMENTATION AND PERFORMANCE ISSUES

A concern frequently voiced with Markov models for pavement management is the perception that the models require too much data and computing resources. The HIPS development team shared this concern at the outset of its asphalt concrete system in 1985, but has since proven that the requirements do not have to be so onerous.

The nature of the five data items collected for Finland's roads to support HIPS (longitudinal and transverse roughness, defects, bearing capacity, and traffic volume) is by no means unusual (*6*). Measurement of roughness, defects, and traffic are standard for most pavement management systems; bearing capacity is uncommon, but can be replaced by other more easily obtained information such as Arizona's index to first crack and change in cracking (*4*) if bearing capacity is unavailable or (for warmer countries) irrelevant. Estimation of deterioration models requires a smaller time series (2 years) than any other model form (*7*). User cost factors are the most difficult to estimate, but the administration's research (*5*) has given reliable data to supplement earlier work in Kenya (*8*), Brazil (*9*), the United States (*10*), and the original Winfrey tables (*11*).

Performance of the models on personal computers has also been better than expected. On a 33-MHz IBM-compatible computer, solution times for the 1989 long-term asphalt concrete models was 10 to 20 min; the short-term models took about half as long. With only 108 states (as opposed to 135) and 5 possible actions (as opposed to 8), the light-pavement models are expected to have solution times one-third to one-half as long as the asphalt models. Because the costs of developing and modifying software on personal computers are far less than on other platforms (*12*), there has been no need to consider using a more powerful machine.

CONCLUSIONS

The asphalt concrete models in HIPS have been in use in the roads administration since the summer of 1988, and the oil gravel models are scheduled to be completed in the fall of 1990. The asphalt models have proven to be stable and to give reasonable and useful recommendations, which have been heavily relied on in the administration's budgeting process. Managerial support of the system and confidence in its recommendations continues to be high.

The Finnish experience with Markov models in pavement management shows that they can be implemented on a microcomputer, in a form that is fast, user friendly, and understandable. Their most attractive feature is the stochastic nature of their performance models and recommendations. Perceived drawbacks to the approach, such as the computational intensity of solution procedures, can be overcome. A continuing challenge now facing the administration is to continue to improve the models as new data are collected, and to monitor the implementation of the recommended rehabilitation programs to see if the long-term condition objectives are achieved.

REFERENCES

1. P. Thompson, Olsonen, A. Talvitie, and R. Tapio. A Micro-Computer Markov Model for Optimal Pavement Rehabilitation Policy. In *Proc., 5th World Conference on Transport Research*, Yokohama, Japan, 1989.
2. P. Thompson, R. Tapio, and J. Äijö. A User-Friendly Capital Programming Tool for Highway Maintenance Districts. In *Proc., International Conference on Microcomputers in Transportation*, San Francisco, Calif., 1989.
3. D. J. White. Real Applications of Markov Decision Processes. *Interfaces*, Vol. 15, No. 6, Nov.–Dec. 1985.
4. K. Golabi, R. B. Kulkarni, and G. B. Way. A Statewide Pavement Management System. *Interfaces*, Vol. 12, No. 6, Dec. 1982.
5. *User Costs in Road Traffic*. Vol. 12, Roads and Waterways Administration, Finland, 1989.
6. R. L. Lytton, M. I. Darter, M. Y. Shahin, and J. B. Rauhut. State-Level Pavement Monitoring Program. In *Transportation Research Record 846*, TRB, National Research Council, Washington, D.C., 1982.
7. J. J. Hajek, W. A. Phang, A. Prakash, and G. A. Wrong. Performance Prediction for Pavement Management, In *Proc., North American Pavement Management Conference*, Toronto, Canada, 1985.
8. H. Hide. *Special Report 160: Investigation of Vehicle Operating Costs in Kenya*. TRB, National Research Council, Washington, D.C., 1975.
9. T. Watanatada, A. M. Dhareshwar, and P. R. S. R. Lima. *Models for Predicting Vehicle Speeds and Operating Costs Based on Mechanistic Principles: Theory and Quantification*. Highway Design and Maintenance Standards Study, Vol. II, The World Bank, Washington, D.C., 1985.
10. G. Zaniewski, B. Butler, G. Cunningham, G. Elkins, M. Paggi, and B. Machemehl. *Vehicle Operating Costs, Fuel Consumption, and Pavement Type and Condition Factors*. FHWA/PL/82/001. FHWA, U.S. Department of Transportation, 1982.
11. R. Winfrey. *Economic Analysis for Highways*. International Textbook Co., 1969.
12. B. Boehm. *Software Engineering Economics*. Prentice-Hall, Englewood Cliffs, N.J., 1981.