An Integrated Resource Allocation Model for Infrastructure Maintenance

Pekka Mild*
Systems Analysis Laboratory, Helsinki University of Technology, P.O. Box 1100, 02015 TKK, Finland
pekka.mild@tkk.fi

Pöyry Infra Ltd, P.O. Box 500, Jaakonkatu 3, 01621 Vantaa, Finland
pekka.mild@poyry.com

Ahti Salo
Systems Analysis Laboratory, Helsinki University of Technology, P.O. Box 1100, 02015 TKK, Finland
ahti.salo@tkk.fi

A key decision in infrastructure management is the allocation of resources to maintenance activities which consist of periodic rehabilitation actions and routine day-to-day operations. This decision involves investments into several products that improve the service quality of different infrastructure assets, contribute to different objectives, and differ from each other in other ways as well; yet all these products impact the same infrastructure system and compete for resources from the same budget. There are also important temporal tradeoffs: for example, although increased funding of routine operations may improve customer satisfaction in the shorter term, it may result in a lower funding of rehabilitation actions and erode the quality of assets in the longer term. In response to these challenges, we present a generic resource allocation model for which we developed for the Finnish Road Administration (Finnra) by building and interlinking (i) a preference model which yields the aggregate value of maintenance products by applying multi-attribute value functions to their quality distributions, (ii) a life-cycle model which captures the deterioration-improvement dynamics of rehabilitation actions, (iii) an optimization model which determines funding recommendations that maximize the aggregate long-term value of maintenance investments. The results were explored in facilitated workshops where 'on-the-fly' computations gave senior managers insights into how the recommended allocations depended on preferences and budget levels. The case study was awarded in Finnra's research program, and it was also recognized as a Finalist for the Decision Analysis Society Practice Award in 2007.

Keywords: Decision analysis, OR practice, multiple criteria, resource allocation

* Corresponding author

1. Introduction

The construction and maintenance of transportation infrastructures – such as road and rail networks – have long-term impacts on the traffic environment, the competitiveness of industries and social and regional development, among others. They also call for major investments, given that the book value of these infrastructures can be as much as 10 per cent of the gross domestic product in industrialized nations. Although the construction of new infrastructures tends to attract most public attention, it is often the maintenance and operation of existing infrastructures that consume more resources and have broader long-term impacts: indeed, in countries with fully developed transportation infrastructures, maintenance activities can consume as much as 60-70 per cent of budgets for these infrastructures (see, e.g., OECD, 2008).

In broad terms, planning decisions in the maintenance of transportation infrastructures are taken (i) at the programming-level, where the focus is on the life-cycle analysis of individual assets and, more specifically, the selection of actions to be funded within annual maintenance and repair programs with pre-established budgets; and (ii) at the network level, where the focus is on the aggregate quality distributions of assets and services, and the purpose is to guide the allocation of resources *among* different maintenance activities, infrastructure asset types, road districts, or other sub-networks. These network-level decisions are inherently more strategic, because they set the budgets for the optimization of programming-level decisions. From the viewpoint of decision support, network-level decisions are also more challenging, because they must address longer time horizons and a broader range of objectives. In effect, they are key to the sustainable development of the whole network.

In the context of road networks, maintenance activities can be organized as a hierarchy of products that belong to two categories, most notably (i) periodic rehabilitation and (ii) routine operations. Periodic rehabilitation actions restore or improve the condition of deteriorating physical assets (e.g., pavements, bridges, gravel roads and road equipment), while routine operations seek to secure the day-to-day trafficability of the road network (e.g., winter-time plowing and de-icing, minor repairs and maintenance of road surroundings, gravel roads surface treatments). The products in the two categories differ from each other in terms of their costs and the durability of their impacts; but they all contribute to the overall service level of same road system and compete for resources from the same budget. As a result, there is a need for systematic and transparent decision support tools that permit trade-offs among multiple objectives and integrate analyses across several assets and operations (Sinha,1989; NCHRP, 2005; Krugler et al., 2007).

Motivated by this need, we present an integrated decision model for the allocation of resources to different maintenance products. This model – which is quite generic even though it was initially developed for the Finnish Road Administration (Finnra) – consists of (i) a multi-criteria model for determining the value of quality distributions associated with maintenance products, (ii) a life-cycle model for describing the dynamics of quality distributions (i.e., deterioration of asset quality over time, improvement of asset quality due to maintenance), (iii) an optimization model for maximizing the long-term aggregate value of all products, subject to annual budgets, quality targets and other constraints. In the deployment of this model, we also conducted extensive sensitivity analyses, incorporated ordinal preference statements about the relative importance of criteria, and explored the implications of these statements for recommended resource allocations (cf. Kirkwood and Sarin, 1985; Hazen, 1986; Salo and Hämäläinen, 1992; 2001; Mustajoki, Hämäläinen and Salo, 2005; Salo and Punkka, 2005; Liesiö, Mild and Salo, 2007; 2008, among others).

We believe that our model, its enthusiastic uptake by senior managers, and subsequent development efforts at Finnra are one of the first successful attempts at integrated resource planning for road maintenance products. One plausible reason for this was that this model was built on existing data and formulated at a relatively high level of aggregation, wherefore it was transparent enough to support interactive analyses in facilitated workshops (instead of offering seemingly 'optimal' take-it-or-leave-it recommendations). These workshops allowed the managers to engage in exploratory, evidence-based discussions about the required funding levels in the long term (cf., e.g., Peerenboom, Buehring and Timothy, 1989; Gamvros, Nidel and Raghavan, 2006; Kleinmuntz, 2007; Phillips, 2007). The managers were particularly interested in exploring (i) how much funding should be given to different products when considering these products jointly without the constraints of prior budgeting practices, (ii) how the recommended funding levels would change in response to different preference statements, and (iii) how possible changes in the total level of funding would be reflected in funding priorities. The possibility to address 'softer' criteria explicitly – such as customer satisfaction – was helpful, too, because it highlighted what maintenance products contributed most to the attainment thereof, and how funding levels should be adjusted to account for them.

The rest of this paper is structured as follows. Section 2 discusses earlier models for infrastructure management and describes the context for our case study at the Finnish road administration. Section 3 presents the integrated resource allocation model, describes its key results and discusses how they were explored in interactive workshops. Section 4 discusses the broader relevance of our case study for decision analytic modeling and outlines extensions for future research. Section 5 concludes.

2. Background and objectives

2.1. Challenges in integrated resource allocation

Since the 1980's, transportation agencies in countries like United Stated, United Kingdom, Australia, Germany and Finland have developed network-level models to predict how the quality distribution of infrastructure assets will evolve in response to maintenance decisions (Golabi, Kulkarni and Way, 1982; Golabi and Shepard, 1997; among others). Most of these employ Markov Decision Processes (MDP) that describe the stochastic evolution of the assets' quality levels based on empirical data about deterioration rates and the impacts of preventive/corrective rehabilitation actions. These actions incur costs to the transportation agency while the lower quality of assets incurs greater costs to the users (e.g., greater risk of accidents, increased travel time, higher fuel consumption). The MDP models seek to minimize to the total socio-economic sum of all these resulting costs subject to quality constraints. Their solutions are derived from the steady state of the Markov chain, and they also show the most cost-effective path for reaching this state. These results support long-term budgeting by characterizing optimal funding levels for rehabilitation activities, and by conveying the corresponding distributions of asset quality.

Although technically and economically sound, cost minimization models have several deficiencies from the viewpoint of integrated resource planning:

- The MDP models usually focus on the pavements and bridges of the primary road network which has a high traffic volume and which is usually best understood and also most effectively managed. However, these models are often quite complex which makes it difficult to calibrate their parameters and communicate their assumptions: for example, their cost parameters and quality constraints may pertain to multiple objectives in complex and non-transparent ways, wherefore they are not ideal for explorative interactive analyses.
- Quality-dependent user-costs may not exist for all assets, particularly for routine operations or rehabilitation actions on the secondary road network with a lower traffic volume. Furthermore, routine operations are inherently different from rehabilitation actions in terms of work type and deterioration characteristics. Such differences among products undermine possibilities of using existing cost-based model structures for between-product comparisons.
- Increasing pressures towards the explicit recognition of multiple objectives in infrastructure management (e.g., increased safety, higher customer satisfaction, reduced environmental impacts) make it necessary to address impact dimensions beyond cost considerations (e.g., Li and Sinha, 2004; Kulkarni *et al*, 2004; NCHRP, 2005). Many of these objectives – such as customer satisfaction – are inherently subjective, wherefore there is a need to develop approaches that accommodate preferences explicitly.

In view of these considerations, cost-based MDP models do not really support comprehensive analyses where the managers must address the short and long term impacts of different products with regard to multiple objectives. This creates a need for integrated approaches that capture these multiple objectives, the decision makers' subjective preferences, and also the life-cycle dynamics of infrastructure assets, with the aim of supporting the allocation of resources to different products.

2.2. Finnish road administration (Finnra)

Finnra is responsible for the management of the Finnish public road network which has an asset value of 15 Billion Euros and comprises 78 000 kilometers (km) of roads and 14 000 bridges. Finnra's annual turnover is about 600 Million Euros of which some 400 Million

Euros is spent on maintenance. Due to geographical reasons, the Finnish road network has some particular characteristics, such as (i) an extensive active gravel road network of 23 000 km which must be maintained even at low traffic volumes, and (ii) snowy and icy conditions which call for routine winter-time operations during 3-6 months a year. These characteristics, however, do not limit the relevance of our resource allocation model for other countries, states or regions.

Finnra has a central administration and nine road districts. The central administration allocates annual budgets to the districts and sets performance targets to enforce national policies. The road districts manage infrastructure assets in their regions by developing annual maintenance plans and by implementing these plans by acquiring services from contractors through competitive tendering. Although the central administration sets strict budgets and performance targets, each district has some autonomy in adjusting national guidelines according to its needs. In the districts, different maintenance products have their product managers and designated experts with partly overlapping domains of expertise. The product managers work in close collaboration; but they also compete for resources.

Our resource allocation model was developed for the road district of South Eastern Finland in 2006. Its development was motivated by the question of what changes would be called for in the current allocation of resources to different products, subject to alternative assumptions about the relevance of the impact criteria. Earlier on, we had collaborated with Finnra and this road district when applying Robust Portfolio Modeling (RPM; Liesiö, Mild and Salo, 2007; 2008) to the development of bridge repair programs. The positive experiences from this collaboration contributed to the conjecture that decision analytic modeling could be helpful even in integrated resource planning.

The project team consisted of the authors of this paper, a senior manager from Finnra's central administration, a senior consultant from Pöyry Infra Ltd and three product managers from the South-East Finland road district. The development of the value model involved five further experts, and the workshops for the uptake of results were attended by senior managers from the road district, including its General Director and Financial Director. The project involved some 10 man-months of work and lasted about 8 months. Requisite data were obtained either from Finnra's databases (e.g., asset volumes) or estimated by experts in facilitated workshops. Standard software tools were employed to develop tools for preference elicitation (MS Excel®) and the computation and visualizations of results (Matlab®, Xpress-MP®).

3. Integrated multi-criteria model

3.1. Decision variables

The product structure in Figure 1 was derived from Finnra's standard structures, with the aim of balancing decision support needs with modeling issues (e.g., appropriate level of detail, data availability, required preference elicitation efforts). Out of the seven main products, four were periodic rehabilitation actions and three were routine operations. Where appropriate, the main products were divided into sub-products, for instance by regarding actions on pavements with high and low traffic volume as two separate products; this resulted in 13 sub-products. For each sub-product, five quality classes were defined to characterize different service levels.

INSERT FIGURE 1 ABOUT HERE

As in most network level analyses, the sub-products were analyzed by considering the distribution of total asset volume across quality classes. This volume was expressed in terms of basic asset units, either by counting individual items (bridges) or summing one-kilometer segments of roads (all other products). Specifically, for each sub-product, the quality distribution indicated what the volume of corresponding assets in the five quality classes was (e.g., 1400 bridges in quality class 1; or 1 500 kilometers of high traffic volume roads maintained by routine operations in quality class 5). Assets within each quality class were assumed homogeneous enough so that the same parameter values could be applied to all units contained in them.

This operationalization of quality classes ensured that (i) each asset unit belonged to exactly one quality class at any given time, (ii) the extreme classes represented the worst and best plausible service levels, and (iii) differences between quality classes were significant enough in terms of their impacts and technical parameters (e.g., costs). Except for routine operations on road surroundings and gravel roads, we used Finnra's standardized quality classifications with pre-defined thresholds on technical condition parameters and service specifications. These classifications had been built earlier to establish a unified scheme where the same class descriptors represented comparable service levels across different products.

For routine operations on road surroundings and gravel roads, the service level attained through current funding level was taken as the mid-point quality. Then, funding levels at ±10% and ±20% were employed to define two inferior (-10%, -20%) and two superior (+10%, +20%) quality classes. These changes in funding levels complied with the three requirements (i)-(iii) above. For these two sub-products, Finnra's experts also specified the corresponding physical service levels which gave enough information for the assessment of their impacts. Overall, we thus had five well-defined quality classes for all sub-products so that the initial quality distributions represented the status quo of the road district, as obtained from Finnra's data records.

The decision variables consisted of (i) rehabilitation actions and (ii) levels of routine operations. Specifically, for rehabilitation products, they corresponded to decisions about how many asset units in quality class j should be raised to some higher quality class $j' > j$ in a given year t. For routine operations, they corresponded to decisions about the funding levels that these sub-products should be operated at.

3.2. Multi-criteria evaluation model

To support comparisons of different quality distributions within and between the sub-products in view of multiple criteria, we built an additive value function based on Multi-Attribute Value Theory (MAVT) (see, e.g,, Keeney and Raiffa, 1976). The four evaluation criteria, listed in Table 1, were derived from Finnra's mission. Apart from descriptive titles, the criteria were given detailed descriptions based on discussions with Finnra's experts so as to ensure that they would be interpreted consistently throughout model development and its use.

INSERT TABLE 1 ABOUT HERE

The scores of the evaluation model, denoted by $v_k^i(j)$, were elicited for every criterion k = 1,…,4, quality class j = 1,…,5 and sub-product i = 1,…,13. These scores were elicited through a two-phase procedure, consisting of (i) within-product evaluation which established the shape of the value function over the quality classes of a given sub-product, (ii) between-product evaluation in which the largest quality swings of sub-products were compared with each other.

In within-product score elicitation, the worst and best quality classes were assigned scores $\tilde{v}_k^i(1) = 0$ and $\tilde{v}_k^i(5) = 100$, respectively. The intermediate quality classes were evaluated with respect to these extreme classes and, for the purposes of validation, also to each other. The elicitation questions started by encouraging the respondents to state their ordinal preferences for quality differences (e.g., is the quality difference between class 3 and class 4 more or less significant than that between quality classes 2 and 3?). The responses were recorded into an Excel®-tool and they were validated and adjusted through visual inspection as well. This approach assigned a cardinal within-product score $\tilde{v}_k^i(j)$ in the of 0-100 range to each quality class (Figure 2).

INSERT FIGURE 2 ABOUT HERE

Under each criterion, the between-product weight elicitation focused on the comparison of quality improvements ('swings') that could be achieved by improving the quality of different products. Specifically, we applied SWING-weighting (von Winterfeldt and Edwards, 1986) by assigning a non-normalized product weight $\tilde{w}_k^i$ of 100 to the sub-product for which the quality improvement (per asset unit) from quality class 1 to class 5 was considered greater than that for any other sub-product. The product weights for the other sub-products were estimated by comparing the significance of their respective 1-to-5 quality swings in relation to this maximum swing, and by setting the corresponding product weight $\tilde{w}_k^i$ accordingly in the range from 0 to 100 (Figure 3). Throughout the elicitation of these product weights, the shapes of the value functions from within-product score elicitation were displayed to the respondents. The final scores from value aggregation were obtained by adjusting the within-product scores with the resulting product weights, i.e., $v_k^i(j) = \tilde{w}_k^i \tilde{v}_k^i(j)/100 \in [0,100]$.

INSERT FIGURE 3 ABOUT HERE

We explored uses of incomplete information through Preference Programming methods (Kirkwood and Sarin, 1985; Hazen, 1986; Salo and Hämäläinen, 1992; 2001; Mustajoki, Hämäläinen and Salo, 2005; Salo and Punkka, 2005; Liesiö, Mild and Salo, 2007; 2008, among others) so that – instead of focusing on a single vector of criterion weights – we considered a feasible weight set that that was consistent with ordinal preference statements about the relative importance of the criteria. This set was derived from an incomplete rank ordering provided by the Finnra's experts (cf., Salo and Punkka, 2005) who noted that R*oad safety* and *Asset value preservation* were the two most important criteria (without specifying which one would be the most important one), followed by *Customer satisfaction* as the third, and *Environmental aspects* as the fourth most important criterion. Technically, the corresponding criteria weights w_k could be interpreted in terms of preferences for the largest sub-product quality swings under the respective criteria.

The above elicitation procedure gave four criterion-specific scores for every sub-product and quality class. The quality distribution of the *i*-th product was characterized by $a_j^i(t)$ which denoted the volume of assets in quality class *j* at time *t*. The total value of the sub-products (or, more specifically, their quality distributions) was aggregated as follows:

- The multi-criteria value of *a single asset unit* of sub-product *i* in a quality *j* class was obtained as the weighted sum of its criterion-specific scores (i.e., $\sum_{k=1}^{4} w_k v_k^i(j)$).
- The value of *the quality distribution of subproduct i* at time *t*, denoted by $V^i(t)$, was obtained by multiplying the volume of assets in different quality classes by their respective multi-criteria values of asset units, i.e., $V^i(t) = \sum_{j=1}^{5} a_j^i(t) \sum_{k=1}^{4} w_k v_k^i(j)$.
- The aggregate value of *the whole system* at time *t* was computed by summing the values of the sub-products' quality distributions, i.e., $\sum_{i=1}^{13} V^i(t)$.

3.3. Life-cycle dynamics and optimization

The quality of assets for rehabilitation products deteriorates over time due to traffic and environmental loads. In terms of quality distributions, this implies that assets in the higher quality classes fall to the lower quality classes unless actions are taken to raise them back to the higher classes.

To capture this deterioration-improvement dynamics, we built a simple life-cycle model based established Markovian principles (cf., Golabi, Kulkarni and Way, 1982; Golabi and Shepard, 1997, among others). For all rehabilitation products, a fixed percentage of the share of assets in a given quality class was assumed to fall to the quality class immediately below it in one year's time. Where appropriate, these sub-product and class-specific percentages were estimated using evidence about deterioration times and average life-cycles of assets and actions. The class-specific unit cost of raising a share of assets from a lower to higher class was estimated from past records and expert judgments. For routine operations, the decisions defined the quality distribution for each year, wherefore no deterioration models were needed. The unit cost of routine operations in different quality classes was derived from valid service contracts or estimated through expert judgments using quality class definitions.

In general, the sub-products exhibited very different cost characteristics and service lives: for instance, the unit costs of the actions varied from a few hundred euros to several hundreds of thousands of euros, while the improved service levels could be sustained in the range between 1 to 30 years. The time span of the model was set to 50 years, which was long enough to cover full dynamics of all sub-products and actions taken in the later years of the model run.

The multi-criteria value model and life-cycle dynamics were combined into an optimization model which yielded resource allocation recommendations for the funding of different sub-products. The objective function of this linear programming model was the discounted sum of the annual aggregate values of the sub-products' quality distributions. It was maximized by allocating funds to rehabilitation actions and routine operations. The life-cycle dynamics (e.g., deterioration of quality distributions in the absence of investments, opposing impacts of maintenance actions) were captured through linear constraints. Additional constraints were introduced to capture budget limitations and minimum threshold requirements on quality distributions. The principle of associating multi-criteria value with quality distributions is illustrated in Figure 4. The full model is described in the Appendix.

INSERT FIGURE 4 ABOUT HERE

3.4. Decision recommendations

The optimization yielded 'funding curves' which conveyed how much funding should be allocated to each sub-product on a year-to-year basis. Specifically, these recommended resource allocations were computed with and without sub-product-specific budget constraints. Without these constraints, the recommendations suggested some unrealistically large departures from the status quo, wherefore further analyses were carried out by setting upper and lower bounds on the funding levels that could be meaningfully be allocated to the sub-products. These bounds were derived from the supply capabilities of market contractors, and they prevented excessively large and sharp deviations from the initial funding levels. They also helped ensure all sub-products would fulfill their respective minimum quality requirements. Once these minimum requirements were fulfilled, the model could be used to explore possibilities for further improvements (cf., e.g., Kleinmuntz, 2007).

Incomplete preference information about the importance of the criteria was explored by using Monte Carlo sampling to generate criterion weights that complied with the ordinal preference statements. The base case solution – which is shown in Figure 5 – was obtained by generating 1 000 criterion weights and by averaging the corresponding optimal funding solutions on a year-by-year basis. In addition, the solutions for the extreme points of the feasible weight set were examined in detail. The recommended funding levels for the sub-products were aggregated to show the funding curves for the seven main products in Figure 1, partly because the emphasis was on the comparison of these products. For visualization purposes, the funding curves were smoothed by computing 3-year moving averages: this was motivated by the recognition that abrupt changes would be undesirable, because the contractors are unable to build up new service capacity very quickly, nor do they wish to experience sudden reductions in the demand for their services.

INSERT FIGURE 5 ABOUT HERE

3.5. Management insights and communication of results

The base case solution supported several important conclusions which also exemplify what kinds of results integrated resource allocation models can offer:

- *Maintenance backlog of bridges*. The funding curve for the rehabilitation of bridges indicated that there was a need for rapid corrective actions to improve the relatively poor *status quo* quality of bridges. This need had already been recognized in national policies and was not therefore particularly surprising; however, the model helped validate and thus provided additional justification for this shift in funding policies.
- *Dynamic funding patterns*. The detailed analysis showed that the funding curves for the different products were inherently *dynamic* so that the relative shares of funding allocated to different sub-products should indeed change over time to respond to the needs of the different products. This was an important finding, because it differed markedly from earlier funding policies where the different products had had relatively rigid and static shares of the overall budget from year to year.
- *Impacts on customer satisfaction*. For example, the model suggested that increased funding should be given to routine operations, because these tend to have a strong impact on customer satisfaction. The link had been implicitly recognized; but the model provided an analytical justification for how increased emphasis on customer satisfaction (interpreted in terms of a higher criterion weight) would be reflected in funding levels. Likewise, the model highlighted that improving the quality of road equipment (e.g., road markings, signs, guardrails) would indeed be a relatively inexpensive way to increase customer satisfaction.

The model results were communicated in an interactive facilitated workshop which was attended by the managers of the South-East Finland road district. In this workshop, the focus was on the funding curves and their corresponding quality distributions, whereby most attention was given to the dynamic behavior and relative magnitudes of funding curves. Specifically, the managers could pose questions about what different preferences or quality requirements would mean in terms of corresponding funding levels. Their questions were interpreted by changing the model parameters, by computing the results 'on the fly', and by presenting the corresponding results in a minute or less. This mode of truly interactive model-based analysis was a totally new way of working at Finnra: in particular, the ability to consider all products simultaneously provided fresh insights and catalyzed discussions about required policy changes.

3.6. Exploration of criterion weights

We also conducted extensive sensitivity analyses with regard to criterion weights, based on the ordinal preference statements about the importance of the four criteria. In particular, we examined (i) minimum and maximum bounds on funding levels, determined for each year as the minimum and maximum level of funding proposed by the funding curves that corresponded to the 1000 sampled weight vectors, and (ii) the extreme funding curves which corresponded to the extreme points of the feasible weight set.

The lower and upper bounds from sampled weights showed which products were sensitive to criterion weights. Products with wide bounds tended to perform well only on one or two criteria and were hence more sensitive to the weights. Products with more robust funding curves (in terms of relatively tight bounds around the base case solution) performed relatively well on most criteria and/or exhibited high benefit-cost ratios (i.e., although the value of some products was sensitive to criterion weights, they had low enough costs so that they offered superior benefit-cost ratios in comparison with other products and merited relatively high levels of funding).

The extreme funding curves showed how the products would be prioritized as a function of different criterion weights. To some extent, this prioritization might have deduced from the evaluation scores; but the funding curves were backed up by more evidence because they also accounted for the initial quality distributions, deterioration-improvement dynamics and costs of maintenance activities. The funding curves for weight vectors where some criteria had no weight at all tended to be volatile and infeasible for practical implementation; nevertheless, they were useful for validation purposes and for catalyzing discussions.

Overall, the exploration of different criterion weights was very instructive. Indeed, our experiences from the workshop suggest that it may be beneficial to conduct such analyses in a truly explorative manner, instead of first computing the 'optimal' solution, followed by limited sensitivity analyses on selected model parameters, although the numerical results from the two approaches may not differ considerably. In our case study, the exploration of incomplete information was highly appreciated by the managers, because it helped them understand *why* the results looked the way they did. This, we believe, made it easier for them to accept the model and its results.

3.7. Changing budget levels

The model was also used to analyze how the prioritization of products would change in response to changes in the overall budget. Specifically, because the budgets are expected to be on the decline, we explored how cuts in the budget, introduced at consecutive steps of 5% reductions, would be reflected in the funding curves of the individual products.

These budget cuts had different implications for different products, for the funding of some products was cut dramatically while that of others remained close to the base case solution. There were also changes in the temporal patterns of funding levels, because some products were subjected to cuts during the first few years, but then received increased funding so that the long term average remained much the same. More generally, our analyses helped establish product funding priorities that showed when and how the products' funding level should be changed in response to decreases (or also to increases) in the overall budget. In this context, sensitivity analyses also showed how the resulting priorities depended on the criterion weights.

Overall, insights into product prioritization as a function of changing budget levels and criteria weights were among the most significant results. Indeed, although much of the annual funding is fixed based on national guidelines, long-term programs and service contracts, road districts may have to adjust their budgets by ±1-10% due to needs that arise at a short notice. Towards this end, a prior consensus on product priorities, achieved through interactive analyses, can be very useful in implementing these adjustments.

3.8. Subsequent developments

Overall, our pilot model was very well accepted, even though Finnra had not had much prior experience on the deployment of similar decision analytic approaches. First, this model was the first one to help Finnra consider funding priorities among rehabilitation actions and routine maintenance in an integrated manner. Second, although the model was not very detailed (due to the fact that it integrated very different types of products with varying levels of data and modeling detail), it provided a transparent framework that offered valuable management insights and enabled a variety of interactive analyses. In fact, our experiences suggest that the lack of excessive technical detail may make it easier to communicate results and to develop exploratory guidelines into how funding patterns should be altered in response to shifting preferences or budget changes.

At Finnra, our case study has sparked interest in the further development of decision analytic approaches for resource allocation, as evidenced by a larger follow-up project whose objectives illustrate possibilities of extending the model:

- The model will be extended to the whole nation to cover all nine road districts. This makes it possible to examine how resources should be allocated among the districts which have different amounts of road infrastructure and different quality distributions as well. Provisional plans have been made to consider investments into the building of new infrastructures along with the maintenance of existing infrastructures.
- The model, and also the processes of parameter elicitation and communication of its results, will be revisited to establish a well-defined and repeatable decision support process. Most probably, some parts of the model may undergo major revision; but its main structure is expected to remain the same.
- The communication of results will be enhanced by developing 'priority maps' which help convey graphically which products/districts merit increased/decreased funding when the overall budget is raised/cut by some percentage points. For example, once the average funding needs for the first three and ten years have been computed, such a map can be presented as a table where the products and evaluation criteria correspond to rows and columns, respectively, and the cells contain colored symbols that how into what direction and by how much the product's funding level should be altered when increased attention is given to some criterion. Such displays sum up many key results from weight-based prioritization of products.

4. Discussion

The maintenance of many other infrastructure systems – such as railroads, waterways and buildings – involve similar questions about how several maintenance products, multiple criteria, subjective preferences and deterioration dynamics can best modeled, often in complex planning environments where complete preference or technical information may be difficult to acquire. It is widely accepted that the exhaustive modeling of these infrastructure systems may not be possible; but nor is this necessary for the purpose of building models that serve as validated frameworks in support of structured and repeatable decision making processes (cf., Peerenboom, Buehring and Joseph, 1989; Clemen and Kwit, 2001; Belton and Stewart, 2002; Gamvros, Nidel and Raghavan, 2006; Kleinmuntz, 2007; Phillips, 2007; among others).

It is against this background that we believe that our model is generic enough and hence applicable to other infrastructure systems as well. There are also several avenues for extending this model:

- The sensitivity analyses – which pertained mostly to questions about how changes in the overall budget or different criterion weightings would impact the results – can be extended to analyze how long-term gradual developments or even abrupt changes may influence the required total level and allocation of resources for maintenance activities. For instance, one could build alternative scenarios (cf. Bunn and Salo, 1993) to examine what implications alternative assumptions about traffic growth in different sub-networks, the deterioration of assets due to climate change, or the completion of complementary infrastructures such as railroads, will have for maintenance activities as a whole.
- As indicated by the plans to extend our model to the whole nation, integrated resource allocation models can support comparative analyses across several geographical regions (or other relevant groupings according to which the products are administered). For instance, these analyses may convey information about how sensitive the overall value of maintenance is to budget levels in the different regions; or what regional funding levels would be required to achieve goals for improved service levels.
- In our model, the correspondence between maintenance products and infrastructure assets was one-to-one so that each maintenance product had impacts on one asset only. But there are other contexts – such as environmental systems – where maintenance actions have consequences on several dimensions: for instance, the reduced use of nutrients in agricultural systems impact not only crop yields but also environmental loads. Even though the modeling of these causal relations may call for more complex impact models, and also for the consideration of the synergistic or antagonistic impacts that the products may have on different criteria, the overall model framework will still remain relevant.

5. Conclusions

We have described a case study for the Finnish Road Administration where we developed an integrated decision model for the allocation of resources to maintenance products. This model – which is generic and hence relevant to the maintenance of other infrastructure assets and decision contexts – features modules for the evaluation of quality class distributions with regard to multiple criteria, the characterization of deterioration-improvement dynamics of product quality, and the determination of recommended product funding levels over a long time horizon. Although some parts of this model were inspired by earlier cost-based models, it is nevertheless very different as its main purpose is to support the *allocation* of scarce resources among different products (rather than to guide the socio-economically optimal level of *spending* of resources on any individual maintenance product). By doing so, it provides strategic decision support to administrators who must continually re-prioritize resource allocation levels.

The model has allowed Finnra managers to consider multiple criteria, to address trade-offs among these, and to perform extensive sensitivity analyses in facilitated workshops. Indeed, these interactive workshops with end-user managers have been a crucial part of model deployment, because they have allowed the managers to understand, explore and debate the implications of different modeling assumptions. In effect, these workshops have been 'hands-on' learning sessions which have generated important management insights and sparked broader interest in decision analytic modeling activities at Finnra. To quote the General Director of the South-East Finland road district, the model and the interactive analyses provided "an innovative and practical platform for thinking and communication to support network level strategy work".

Acknowledgements

This research has been supported by the Finnish road administration and the Academy of Finland. The computations were carried out with Xpress-MP© software provided by Dash Optimization Inc. through their Academic Partner Program.

References

Belton, V., Stewart, T.J. 2002. *Multiple Criteria Decision Analysis: An Integrated Approach*, Kluwer Academic Publishers, Dordrecht.

Bunn, D.W., Salo, A. 1993. Forecasting with Scenarios, *European Journal of Operational Research* **68**(3) 291-303.

Clemen, R.T., Kwit, R.C. 2001. The Value of Decision Analysis at Eastman Kodak Company, 1990-1999, *Interfaces* **31**(5) 74-92.

Gamvros, I., Nidel, R., Raghavan, S. 2006. Investment Analysis and Budget Allocation at Catholic Refief Services, *Interfaces* **36**(5) 400-406.

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

Golabi, K., Shepard, R. 1997. Pontis: A System for Maintenance Optimization and Improvement of US Bridge Networks, *Interfaces* **27**(1) 71-88.

Hazen, G. 1986. Partial Information, Dominance, and Potential Optimality in Multiattribute Utility Theory, *Operations Research*, **34**(2), 296-310.

Keeney, R., Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*, John Wiley & Sons, New York.

Kirkwood, C.W., Sarin, R.K. 1985. Ranking with Partial Information: A method and an Application, *Operations Research*, **33**(1), 38-48.

Kleinmuntz, D. 2007. Resource Allocation Decisions, in Edwards, W., Miles Jr., R.F., von Winterfeldt, D., eds. *Advances in Decision Analysis – From Foundations to Applications*, Cambridge University Press, New York.

Krugler, P.E., Chang-Albitres, C.M., Pickett, K.W., Smith, R.E., Hicks, I.V., Feldman, R.M., Butenko, S., Kang, D.H., Guikema, S.D. 2007. *Asset Management Literature Review and Potential Applications of Simulation, Optimization and Decision Analysis Techniques for Right-of-way and Transportation Planning and Programming*. Report FHWA/TX-07/0-5534-1, Texas Transportation Institute, Austin, Texas.

Kulkarni, R.B., Miller, D., Ingram, R.M., Wong, C-W, Lorenz, J. 2004. Need-based Project Prioritization: Alternative to Cost-benefit Analysis, *Journal of Transportation Engineering* **130**(2) 150-158.

Liesiö, J., Mild, P., Salo, A. 2007. Preference Programming for Robust Portfolio Modeling and Project Selection, *European Journal of Operational Research*, **181**(3) 1488-1505.

Liesiö, J., Mild, P., Salo, A. 2008. Robust Portfolio Modeling with Incomplete Cost Information and Project Interdependencies, *European Journal of Operational Research*, **190**(3) 679-695.

Mustajoki, J., Hämäläinen, R.P., Salo, A. 2005. Decision Support by Interval SMART/SWING - Incorporating Imprecision in the SMART and SWING Methods, *Decision Sciences* **36**(2) 317-339.

NRCHRP. 2005. *Analytical Tools for Asset Management*, NCHRP report 545, Transportation Research Board, Washington D.C.

OECD (2008). *Transport Infrastructure Investment: Options for Efficiency*, OECD Publishing, Paris.

Peerenboom, J.P., Buehring, W.A., Timothy, W.J. 1989. Selecting a Portfolio of Environmental Programs for a Synthetic Fuels Facility, *Operations Research* **37**(5) 689-699.

Phillips, L.D. 2007. Decision Conferencing, in Edwards, W., Miles Jr., R.F., von Winterfeldt, D., eds. *Advances in Decision Analysis – From Foundations to Applications*, Cambridge University Press, New York.

Salo, A., Hämäläinen, R.P. 1992. Preference assessment by imprecise ratio statements, *Operations Research* **40**(6) 1053-1061.

Salo, A., Hämäläinen, R.P. 2001. Preference Ratios in Multiattribute Evaluation (PRIME) - Elicitation and Decision Procedures under Incomplete Information, *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* **31**(6) 533-545.

Salo, A., Punkka, A. 2005. Rank Inclusion in Criteria Hierarchies, *European Journal of Operational Research* **163**(2) 338-356.

Von Winterfeldt, D., Edwards, W. 1986. *Decision Analysis and Behavioral Research*, Cambridge University Press, Cambridge, UK.

Appendix: The resource allocation model

Index sets

REH	Index set of the rehabilitation sub-products
ROU	Index set of the routine operations sub-products

Decision variables

$a_j^i(t)$	Volume of sub-product i asset units that are operated at the level of quality class j at time t (for routine operations $i \in ROU$)
$x_{j,j'}^i(t)$	Volume of sub-product i asset units that are moved from quality class j to quality class $j' > j$) at time t (for rehabilitation sub-products $i \in REH$)

Constants

d_j^i	Percentage of the total volume of assets units $a_j^i(t)$ whose quality declines to the next lower quality class in one period (for rehabilitation sub-products $i \in REH$)
$c_{j,j'}^i$	Cost of raising one asset unit from quality class j to $j' > j$ (for rehabilitation sub-products $i \in REH$)
c_j^i	Cost of operating one asset unit at the level of quality class j for one year (for routine operations $i \in ROU$)
$B(t)$	Budget in year t
δ	Upper bound on the share of assets units whose quality level changes from one year to the next
$v_k^i(j)$	Score of sub-product i for quality class j with regard to criterion k
w_k	Weight of criterion k
r	Annual discount rate

Constraints

The total volume of infrastructure assets does not change.

$$\sum_j a_j^i(t) = \sum_j a_j^i(t_0) \;\; \forall i,t$$

Moves are non-positive; the sum of moves from class *j* cannot exceed the volume of assets that are in this class.

$$x_{j,j'}^i(t) \ge 0 \,, \;\; \sum_{\{j' | j' > j\}} x_{j,j'}^i(t) \le a_j^i(t) \;\; \forall i,j,j',t$$

Deterioration-improvement dynamics

$$a_j^i(t+1) = (1-d_j^i)a_j^i(t) + d_{j+1}^i a_{j+1}^i(t)$$
$$- \sum_{j'>j} x_{j,j'}^i(t) + \sum_{j'<j} x_{j',j}^i(t) \;\; \forall i,j,t$$

Budget constraint(s); may be specified separately for different sub-products and time periods

$$\sum_{i \in REH} \sum_{j,j'} c_{j,j'}^i x_{j,j'}^i(t) + \sum_{i \in ROU} \sum_j c_j^i a_j^i(t) \le B(t) \;\; \forall t$$

Bounds on the relative change of asset volumes in one year (optional)

$$(1-\delta)a_j^i(t) \le a_j^i(t+1) \le (1+\delta)a_j^i(t) \;\; \forall i,j,t$$

Feasible criterion weights

$$S_w = \left\{ w = (w_1,...,w_4)^T \mid w_1, w_2 \ge w_3 \ge w_4 \,, \sum_k w_k \right.$$

Objective function

Sum of the discounted aggregate multi-criteria value achieved through the sub-products' quality distributions

$$\max_{\substack{a_j^i(t) \\ x_{j,j'}^i(t)}} \sum_t \frac{\sum_{i,j} a_j^i(t) \sum_k w_k v_k^i(j)}{(1+r)^t}$$

Figures and tables

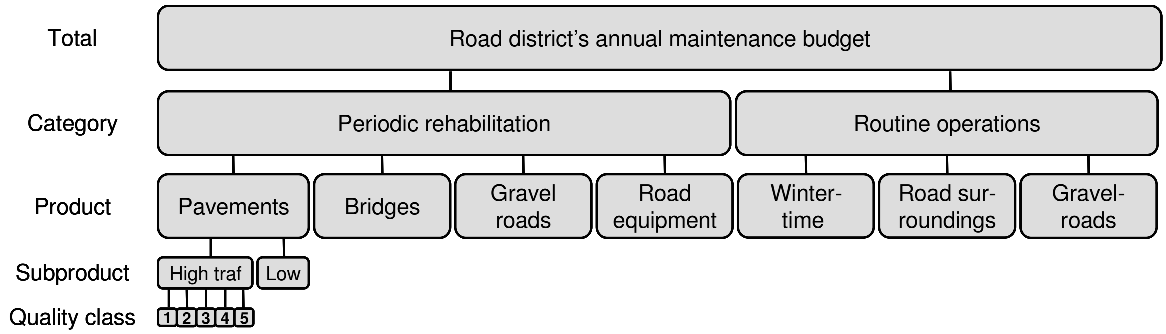

Figure 1. Product structure.

Table 1: The four evaluation criteria used.

Criterion	*Brief description*
Road safety	Frequency and severity of road accidents due to poor quality of assets and routine operations.
Asset value preservation	Monetary decline in asset value (cf. increase in maintenance and repair costs) due to deterioration of asset quality.
Customer satisfaction	Service level as perceived by road users in terms of driving comfort, sustained driving speed, safety confidence; customer feedback.
Environmental aspects	Adverse effects of road chemicals, traffic noise and dusting; aspects of nature conservation, tidiness and appearance.

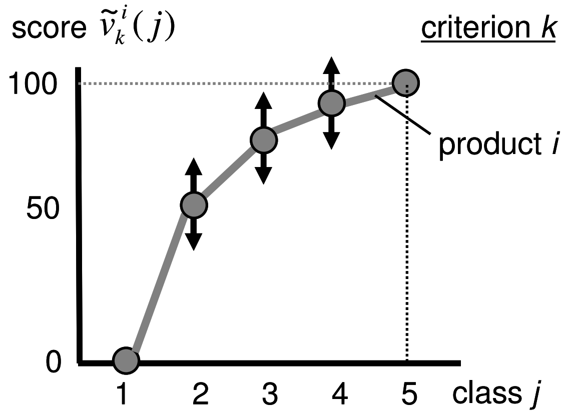

Figure 2. Within-product value function for sub-product *i* with regard to criterion *k* – determination of intermediate raw scores.

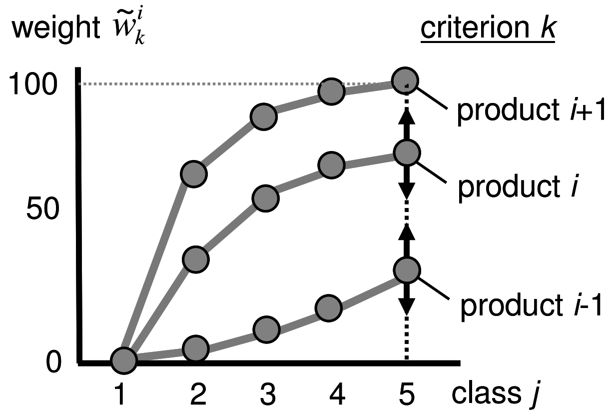

Figure 3: Between-product evaluation of sub-products with regard to criterion k – determination of between-product weights from the consideration of maximum swings.

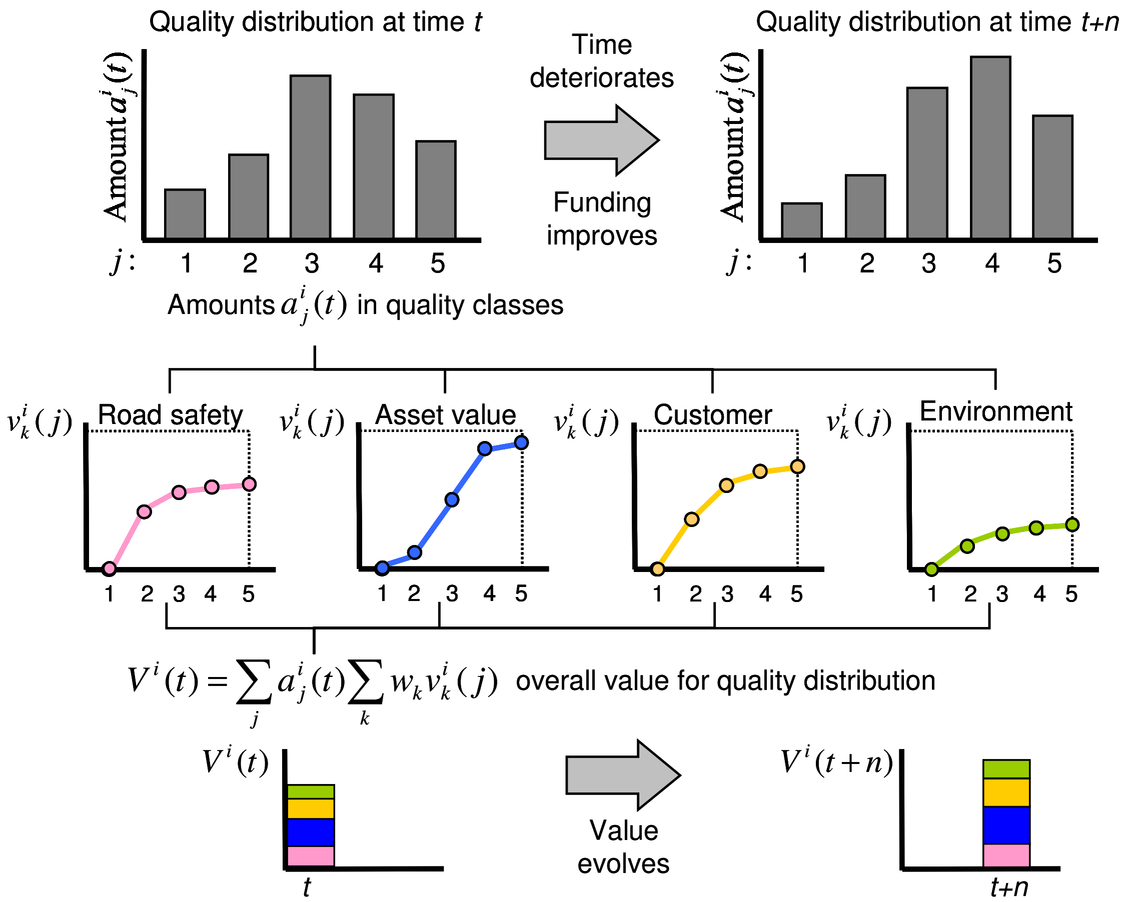

Figure 4: Schematic summary of the model components.

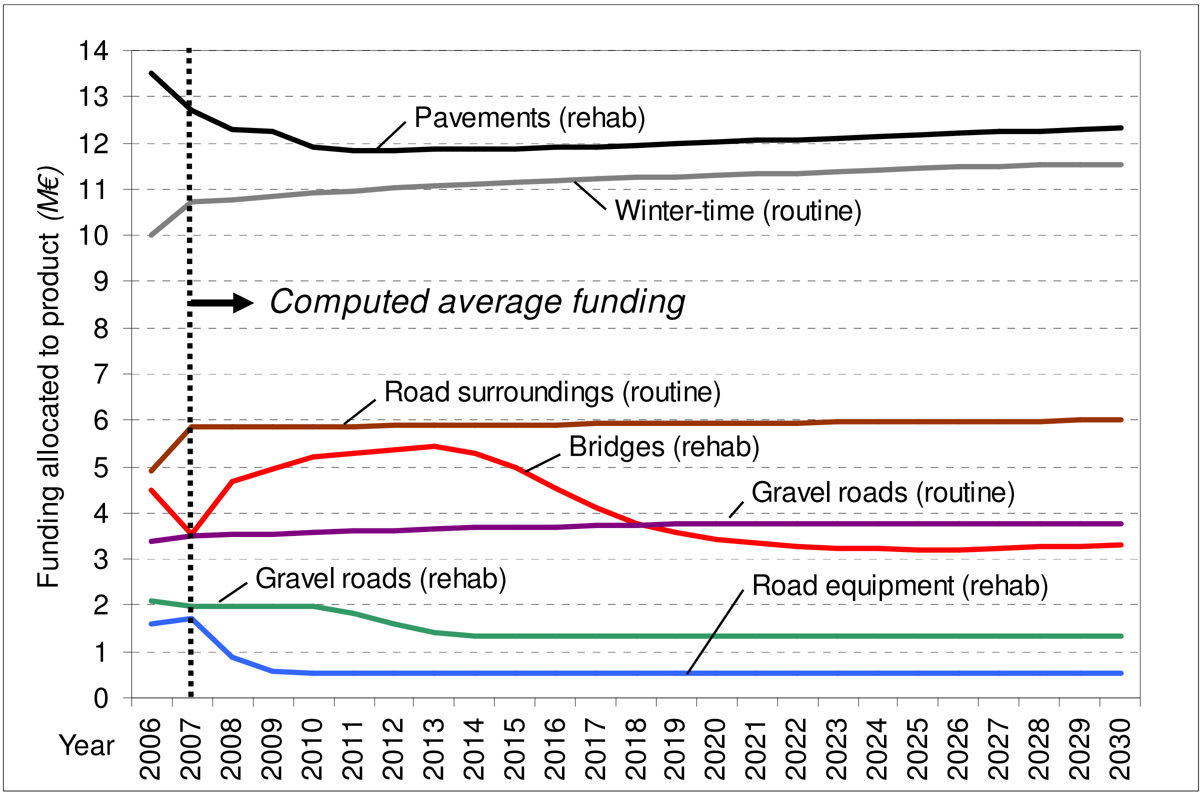

Figure 5: The base-case funding curves for a total annual funding of 40 Million Euros.