

Contents lists available at [ScienceDirect](http://www.sciencedirect.com)

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

O.R. Applications

Optimizing winter road maintenance operations under real-time information

Liping Fu, Mathieu Trudel, Valeri Kim *

Department of Civil and Environmental Engineering, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1

ARTICLE INFO

Article history:

Received 18 July 2007

Accepted 1 March 2008

Available online 8 March 2008

Keywords:

Linear programming

Integer programming

Transportation

Winter road maintenance

Real-time

ABSTRACT

This paper introduces a real-time optimization model that can be used by maintenance managers to develop and evaluate alternative resources allocation plans for winter road maintenance operations. The model takes into account a wide range of road and weather condition factors such as road network topology, road class, weather forecasts, and contractual service levels, and produces a vehicle dispatch schedule that is optimal with respect to operating costs and quality of service. The model is then used in an analysis on a realistic case to illustrate the potential impact of improved information on winter maintenance operations.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Winter maintenance is a significant challenge to most transportation agencies in North America. Transport Canada estimated that Canadian provincial and municipal governments spend more than \$1 billion dollars (CDN) every year to keep roads and highways clear of snow and ice and at safe travel conditions (Andrey et al., 1999). This includes dumping over 5 million tonnes of salt onto Canadian roads – causing a growing concern on its environmental impact. Under the mounting pressure of high demand for improved winter safety and mobility, tight budget constraint, and widespread concern on the environmental effect of salts, winter maintenance authorities are constantly seeking for technology-based solutions such as advanced road weather information systems (RWIS) for monitoring localized weather and road surface conditions, automated vehicle location (AVL) for tracking fleet operations and performance, and advanced anti-icing strategies for reducing salt usage. For example, a survey of winter maintenance operations undertaken in 2000 by the SHRP (2000) indicates that there were 134 reported RWIS towers operating in Canada, in eight of Canada's ten provinces and territories. As of the 2004 winter season, the number of RWIS stations in Ontario alone reached 112 (Ontario Ministry of Transportation, 2004).

However, beyond providing improved weather data and vehicle instrumentation for supervisors, most of these technologies have played a limited role in improving winter road operations at a

supervisory level, where great benefit can be expected. Decision-making at the supervisory level of winter maintenance operations are often complex and constrained by time and resources. For example, at the onset of a snow storm, maintenance managers and supervisors must decide when to start sanding and plowing, what operation routes and sequence to follow, and how much chemical agent to apply. The situation becomes even more challenging when an anti-icing strategy is deployed as the effectiveness of this strategy depends largely on the ability to make accurate prediction on weather and pavement conditions and to determine the precise timing and rate for applying treatment chemicals. Conventional approaches to these maintenance decision problems are highly empirical in nature. Decisions pertaining to when and where to deploy service vehicles are typically made by a human supervisor, based on mostly static weather forecasts, the first hand reports of deployed vehicles, and personal experience. As such, it is not only difficult for knowledge to be quantified and transferred between supervisors and between districts, but also to soundly compare alternative treatment plans, or to perform 'what-if' scenarios without actually carrying them out in real life.

There is an extensive literature of academic research on various issues related to the planning, design and management of winter road maintenance operations, as recently summarized by Perrier et al. (2006a,b, 2007a,b). Most of the existing research has however focused on the static problems of winter road maintenance such as depot location, vehicle routing and fleet sizing. Few have attempted to address the dynamic nature of winter maintenance operations in which fleet and other resources must be directed in

* Corresponding author. Tel.: +1 519 888 4567x33984; fax: +1 519 888 4349.

E-mail addresses: lfu@uwaterloo.ca (L. Fu), vkim@uwaterloo.ca (V. Kim).

real time to where they are most needed in response to unfolding and forecasting of snow storm events. Attempts have also been made in literature to develop decision support systems that can assist maintenance managers and operators to make more informed decisions. For example, the US FHWA (Federal Highway Administration) has developed a software tool called maintenance decision support system (MDSS) which is capable of presenting users with real-time information on the current and predicted weather and road conditions of a given road network as well as recommending treatment plans. However, the treatment plans generated by MDSS are created on the basis of the rules of the conventional maintenance manual without systematic optimization considering all inputs and resource constraints.

The primary objective of this research is to develop a decision support model for winter road maintenance activities that can be used both to derive optimal service plans in real time for a road network, and to perform 'what-if' analysis on the effect of different service plans under specific road weather conditions. This paper is divided into two sections. The first section introduces the winter road maintenance scheduling model, explains its structure, and shows how data flows through the model. The second section presents a case study using the model to assess the effects of weather forecast quality on realized treatment levels.

2. WRMSP – Problem formulation and solution approach

2.1. Problem definition and formulation

The winter road maintenance scheduling problem (WRMSP) can be formally defined as follows. Consider a road network consisting of a set of road segments. The network has been covered by a set of pre-defined maintenance routes with each representing an itinerary that a service vehicle may take through the road network, i.e. the vehicle routing problem is not considered as part of this research. It is assumed that all routes are represented as an ordered list of road segments, and are cyclic with a single patrol yard, that is, they begin and end at the same location (their home patrol yard). Note also that routes are not necessary mutually exclusive, that is, there could be overlapped portions with some road segments belonging to multiple routes. A fleet of vehicles is available to provide winter maintenance services such as plowing and salting during a snowstorm. The problem is to develop an operations plan for the available service vehicles that specify route assignment, service type, and the corresponding start time. The scheduling solution should take into account the following requirements:

1. The total network wide level of service is maximized while a minimal level of service is guaranteed for each class of roads.
2. The total operating costs are minimized.
3. Total negative environmental effects (salt usage) are minimized (not considered in this research).

In order to model the WRMSP mathematically, we must first address the issue of representing road conditions and the impacts of weather and maintenance treatment events on road conditions in quantitative terms. Winter road conditions usually feature different attributes, requiring multiple measures for complete description, such as snow depth, snow coverage, and pavement friction. Among all these measures, friction has been suggested in literature as the most appropriate as it reflects the road driving conditions most closely from a user (driver)'s point of view. Currently, however, there are no reliable performance models available that relate friction to various quantifiable factors such as snow depth, temperature and service events. Without loss of generality, in this research

we assume that service conditions on a road can be represented by the average snow depth on the road. This metric of road condition could be changed in the future by replacing it with a more comprehensive metric such as road friction coefficient.

As we consider snow depth to be the metric of interest in this report, weather and treatment events are assumed to affect road conditions by altering snow depth. The impact of weather events can be expressed as changes in snow depth on the basis of snow fall or melt. Modeling the impacts of treatments could however become very complex, depending on snow and ice control techniques such as plowing, salting and anti-icing, and weather conditions. For example, the Ontario Ministry of Transportation's DART database (Perchanok, 2002), seeks to do exactly that, and is a large-scale ongoing research project. To simplify the mathematical formulation, we assume a single type of snow removal service with a known relationship. Specifically, snow removal is assumed to take place by plowing alone, and a single plowing operation removes all snow on the roadway, up to a given maximal depth (as will be introduced shortly). We also assume that the traffic volume does not affect the road condition. The possible traffic delays due to the plowing operations are also ignored.

To capture the dynamic nature of the problem, we consider a planning horizon divided into a set of uniform time intervals. These intervals represent the shortest intervals which can be scheduled, or during which any activity can take place. The following notations are introduced:

(1) Index sets:

- i = index of road segments, $i = 0, 1, \dots, n - 1$, where n is the number of road segments.
- k = index for maintenance routes, $k = 0, 1, \dots, m - 1$, where m is the number of routes.
- t = index of time period, $t = 0, 1, \dots, T - 1$, where T is the number of time intervals. Note that all time intervals are set of have an equal length, which is also used to define some of the other system parameters such as link length and route length.

(2) Input parameters:

- l_i : Represents the length of road segment i , expressed as the number of time intervals required to traverse it. We assume it stays the same whether the vehicle is plowing or just traversing the segment.
- L_k : Represents the length of a route k , expressed as the number of time intervals required to traverse it. This parameter is computed implicitly from the values of l_i given above, by summing the values for all constituent road segments in a route.
- λ_i : Represents a weight factor ascribing 'importance' values to road segment i . This allows the model to consider various classes of roads, and to differentiate between them for the purposes of establishing service priorities.
- ϕ_i : Represents a level of service threshold for road segment i which must be maintained at all times. In this research, minimum level of service is represented by the maximum allowable snow accumulation. While the formulation in this paper uses a common threshold value for all time periods, the model does not require this. Thus, the threshold could easily be made to vary by time of day as well as by road segment by including an index for time interval.
- $w_{i,t}$: Represents the change in road condition due to weather for road segment i at time interval t . In this research it represents the additional snow accumulation on pavement caused by snowfall at time period t . These values are expected to be obtained from a weather model which black boxes these predictions.

- $r_{k,i}$: Describes the incidental relationship between routes and road segments. It takes on the ordinal value of the road segment i 's position in route k . Road segments that do not appear in a route are denoted by the value 0. For example, if road segment i is the 4th segment in route k , then the value of $r_{k,i}$ will be 4. Let $R = \{r_{k,i}\}$.
 - d : An integer value representing the size of the fleet available for dispatch in this patrol region. No more than this many vehicles are available for use at any given time.
- (3) Decision-related variables:

- $x_{k,t}$: A binary decision variable representing a dispatch decision where entries are set to 1 if and only if a service vehicle is to depart on route k at the beginning of time interval t , and 0 otherwise. Let $X = \{x_{k,t}\}$.
- $y_{i,t}$: An intermediate variable representing how many service vehicles are on road segment i at time interval t . Let $Y = \{y_{i,t}\}$. These variables are used to reflect the fleet size restriction. The value of these variables can be determined as a function of scheduling decisions (X), route-link incidental relationship (R) and link service time (l). That is

$$Y = f(X, R, l)$$

Note that the inner workings of the function f are described in Appendix A, as they are somewhat complicated mathematically and are of peripheral relevance to this paper. It suffices to note that f is linear in its input.

- $y'_{i,t}$: An intermediate variable representing how many service vehicles are actually servicing road segment i at time interval t . Let $Y' = \{y'_{i,t}\}$. Due to the specifics of the treatment modeling in this research, there is a distinctive difference between $y_{i,t}$ and $y'_{i,t}$. That is, the treatment of road segment i by a particular service vehicle takes place only at a single point in time, namely when a service vehicle enters road segment i . Similarly to $y_{i,t}$, the values of these variables can be determined as a function of scheduling decisions (X), route-link incidental relationship (R) and service speed or time (l). That is

$$Y' = f'(X, R, l).$$

f is linear in its input (See Appendix A for the inner workings of the function f).

- $c_{i,t}$: An intermediate variable representing the estimated pavement condition of road segment i at the end of time interval t , after all treatment and weather events have been taken into account. The value of these variables is determined recursively as a function of the condition level in the previous time interval, plus any weather related changes, minus any changes due to service events. That is

$$c_{i,t} = \max(0, c_{i,t-1} + w_{i,t} - \Delta \cdot y'_{i,t}),$$

where Δ in the above equation represents the maximum depth of snow that can be cleared in a single pass of a service vehicle, and $c_{i,0}$ is the initial road surface condition at the beginning of time interval 0, which is assumed to be known.

Based on the objectives and conditions of the scheduling operations, the WRMSP can be formulated as the following Integer Program:

$$\begin{aligned} \text{Min} \quad & \sum_{\forall i} \sum_{\forall t} \lambda_i c_{i,t} + \beta \sum_{\forall k} \sum_{\forall t} L_k x_{k,t} & (1) \\ \text{Subject to} \quad & c_{i,t} = \max\{0, c_{i,t-1} + w_{i,t} - \Delta \cdot y'_{i,t}\} & (2) \\ & \text{for all } i, \text{ for all } t, & (2) \\ & c_{i,t} \leq \phi_i \quad \text{for all } i, \text{ for all } t, & (3) \\ & \sum_{\forall i} y_{i,t} \leq d \quad \text{for all } t, & (4) \\ & Y = f(X, R, l), & (5) \\ & Y' = f'(X, R, l). & (6) \end{aligned}$$

The objective function (1) seeks to minimize a weighted total of both the overall negative impact of service, represented by total road condition index or snow depth adjusted by road class weighting factor, and the cost of providing service, represented as a function of service kilometers, where β is a factor assigning relative importance to these two goals as well as converting the two cost elements into a common unit. Eq. (2) populates the condition transition based on a recursive formula outlined previously. Eqs. (3) and (4) enforce service level and fleet size restrictions, respectively. Eqs. (5) and (6) create the location indices Y and Y' from the values of X , based on internal tables derived from the data in l and R as discussed previously.

2.2. Treatment effectiveness model

Any treatment effectiveness model must quantify effectiveness with respect to a specific metric. For example, when considering plowing operations, the average depth of snow on a section of road would be a suitable candidate metric, as it reflects the main purpose and effect of plowing. The selection of which metric to use to evaluate a given maintenance operation on the network is critical; selecting a metric that does not capture the actual effect of the operation introduces both a smaller measurement range for the effect of the treatment, but also has the potential of allowing secondary causes into the effect measurement which may not be captured by the model. For this reason, it is imperative to select a performance metric that is directly linked to the treatment being performed. For example, Keyser (1981) noted the existence of time varying measures of road friction as a result of various chemical applications, but did not control for temperature or other critical variables.

Plowing effectiveness models seek to relate the operation of plowing with both the immediate and ongoing changes to road condition as a result of the plowing operation. The immediate effects on road condition are obvious – as a plow passes an area of a road, snow is removed down to a certain depth determined by road and plow geometry, and the composition of the snow and ice overlying the roadway. The ongoing effects are a direct result of this change in snow depth; if pavement (particularly darker asphalt) is uncovered as a result of plowing, then solar heating and environmental exposure of this surface will affect pavement temperature, and subsequently the trend and rate of snow depth on the roadway. A number of studies have been undertaken to improve the physical composition and geometry of plow blades in an effort to improve their effectiveness and reduce their maintenance costs. One of the largest recent studies in this area was undertaken as part of the US Strategic Highway Research Project (Pell, 1994), which evaluated potential improvements in the classic displacement plow blade design. Unfortunately, this study, while exhaustive in the areas it was focused on, does not provide quantitative evidence of the amount of snow left behind on a roadway for a given set of parameters. Regarding the ongoing effects of plowing on roadway conditions, we turn to a first-principle energy balance model of temperatures of snow-covered surface, as described in Jordan (1991). The model described in this work, termed

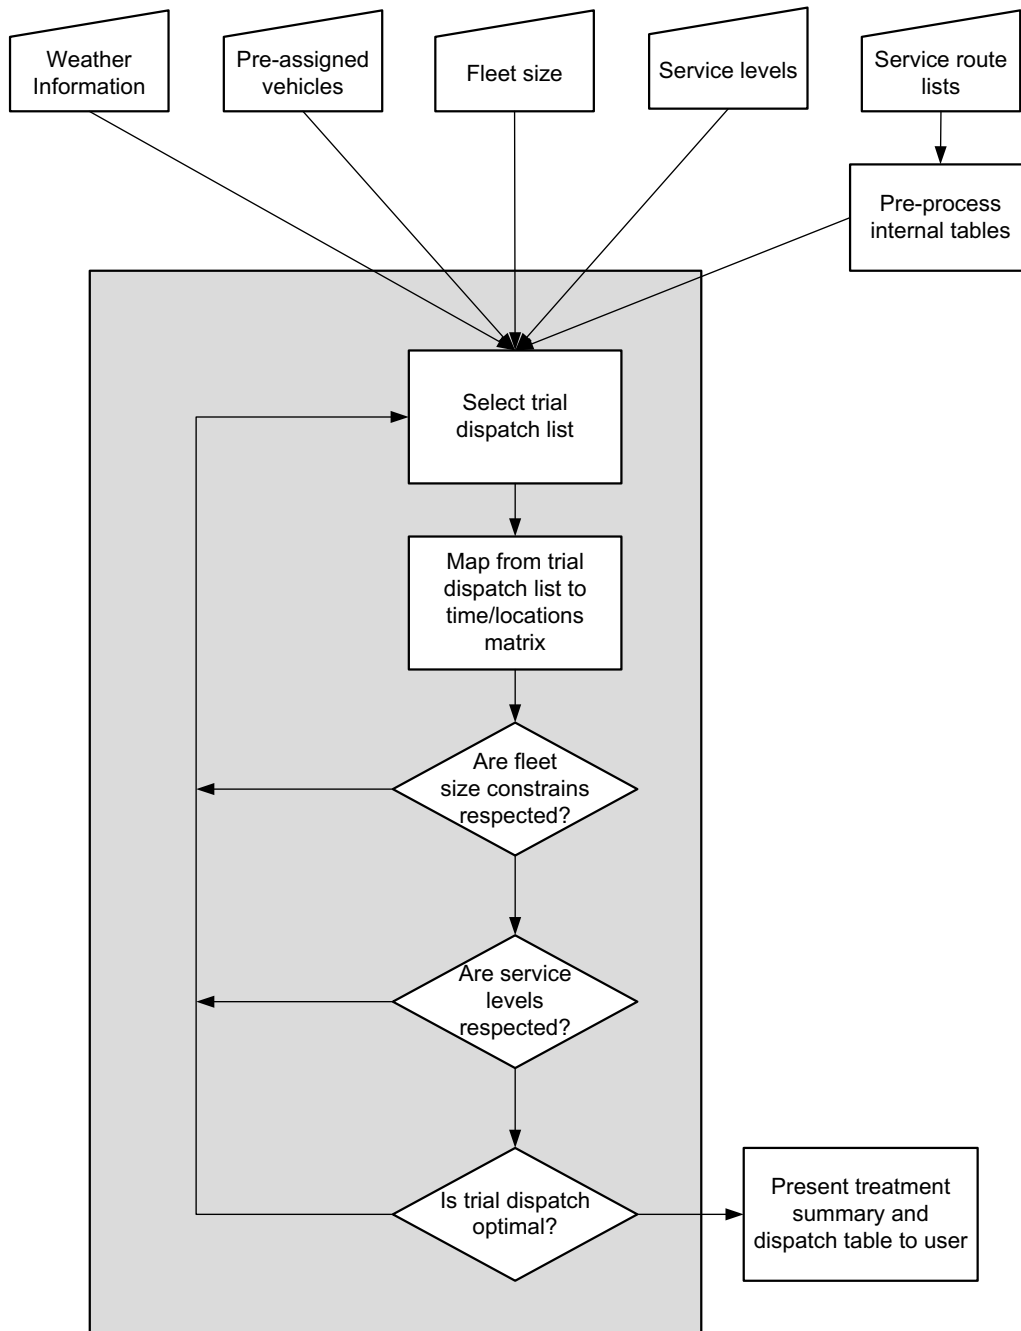
SNTHERM, is an energy balance model that predicts temperatures at arbitrary depths within a gradient of snow, pavement, and soil. This model is produced from a physicist’s point of view, and as such its internal details are far beyond the scope of this review. The algorithm itself is implemented in FORTRAN, and is suitable for use as a ‘black-box’ predictor of pavement temperature given a suitable set of conditional parameters.

Chemical agent treatment refers to methods of winter road treatment that involve the application of a free-point suppressant or other chemical to the roadway. Removal of snow and ice is thus affected through a chemical process. Chemical agent effectiveness models seek to model the effects of such treatment on an ongoing basis. Such a model is critical in evaluating the effectiveness of various

de-icing and anti-icing policies. Unfortunately, existing efforts have not yet resulted in a satisfactory model that can be integrated into the proposed optimization system. The MDSS system developed by FHWA includes a sophisticated Chemical Concentration Model component (MDSS National Laboratory Consortium, 2004), which is still under validation for its applicability for different environmental and geographical conditions.

2.3. Model implementation

The WRMSP formulated in this work is an Integer Linear Program (ILP) which takes as input a set of routes, weather forecasts, previously dispatched vehicles, and other ancillary information.



Branch and Bound Integer Linear Program

Fig. 1. Implementation of the WRMSP.

The model then produces as output a summary of treatment effectiveness, and a dispatch list describing when and where to deploy vehicles. The model is constrained by fleet size and service level restrictions to only produce results that are feasible with respect to these conditions.

The structure of the ILP that implements the WRMSP is described as a flowchart in Fig. 1. The ILP is implemented in the AMPL Linear Programming language (Fourer et al., 2003), and is solved using a branch-and-bound method with bounds computed using the simplex algorithm from the GNU Linear Programming Kit (GLPK) Library (GNU Software Project, 2006). In selecting a solver library to use, both computational speed and the level of ease with which the library could be integrated within a more complete data management environment were taken into account. The majority of the shell logic around the ILP core of the program is implemented in Java, with a JNI native interface to the GLPK library, which is implemented in readily portable C. This approach provides a great deal of functionality such as SQL connectivity, Inter-

net based services, and easy portability, without losing the raw speed of having a C based ILP solver doing the processor intensive work (refer to Trudel (2005) for implementation details).

ILP remains a hard problem. As such, the best we can typically hope for is an approximation solution. The WRMSP is not immune from this limitation; for all but the smallest problems, the WRMSP will not find a perfectly optimal dispatch schedule, but rather an approximation thereto. In practice, the solutions found by the WRMSP are an approximation obtained after allowing the solver to run for a fixed period of time. By solving the ILP as a relaxed Linear Program however, we can obtain an upper bound on the best possible solution to a given ILP, and thus establish an approximation bound for a given solution. For the remainder of this report, whenever we speak of ‘optimal solutions’, we implicitly mean the best solution the solver obtained in a fixed execution time (1000 seconds).

Before actual solving of the ILP takes place, the routing information which is contained in the matrix $R = \{r_{k,i}\}$ and the list l is used

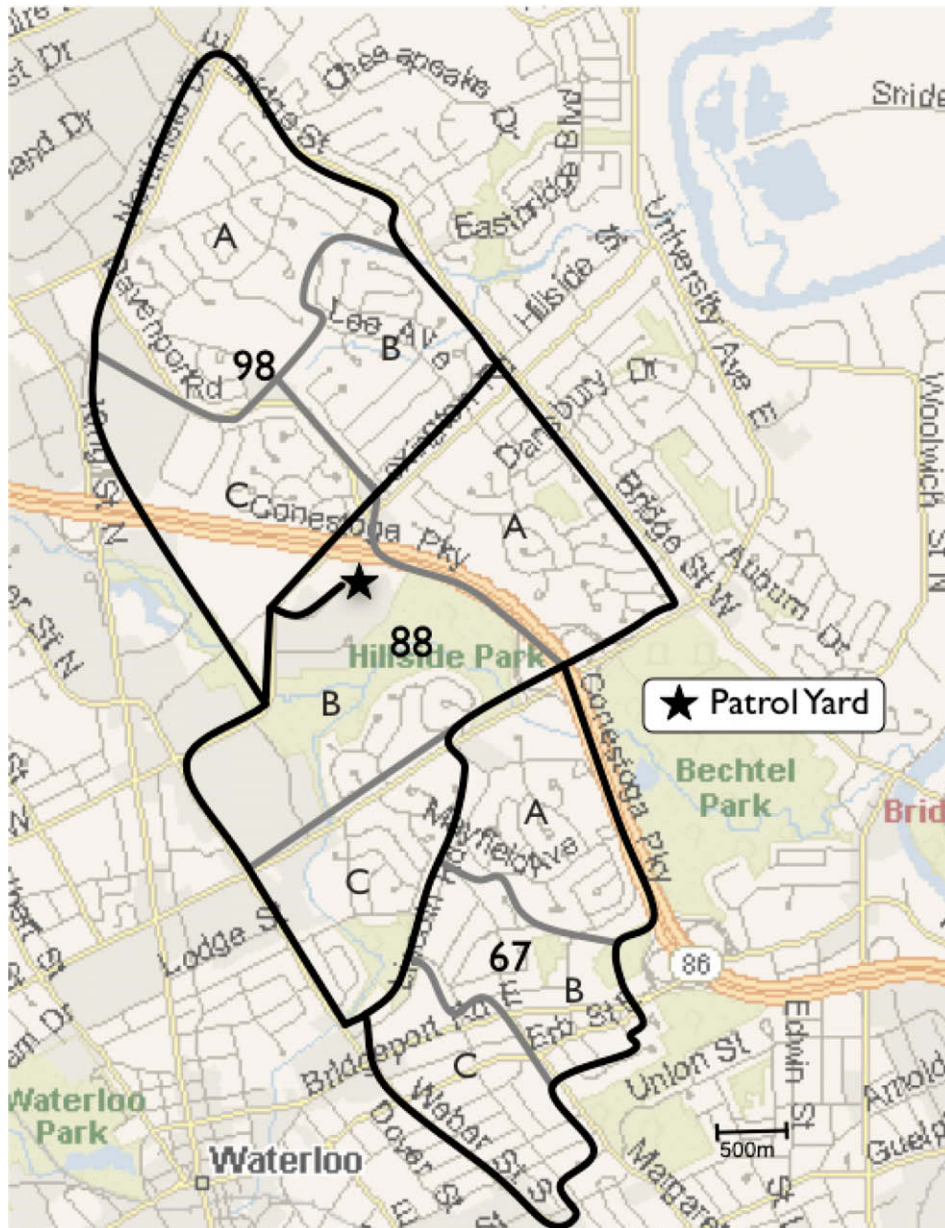


Fig. 2. Map of case study coverage area.

to pre-calculate several internal tables used by the function f in Eq. (5). As described previously, f is an ancillary function that takes as input a prospective dispatch table X that relates route departures to time, and produces a table Y that defines the number of vehicles on each road segment during each time interval that would be produced by this prospective dispatch table. This calculation is expressed as a computed parameter in AMPL and as such is completely automatic.

When running the WRMSp on a data set, an execution time limit is typically set. The shell program transforms user-entered data into a format expected by the ILP solver, and sets the solver running. Once an optimal solution is found or (as is more typical) the specified time limit expires, the best solution found to date is output, in the form of a dump of the X and C matrices. This data can then be used to dispatch vehicles (in the case of a field application of WRMSp), or analyzed further (as is the case in the following case study).

3. Value of real-time information – A case study

3.1. Case study overview

The case study seeks to quantify the operational benefit that could be obtained by using weather forecasts of different levels of accuracy as a way of illustrating the value of real-time information for winter road maintenance operations. We formulate four different sets of forecast information of the same storm event, with each set representing a particular forecast quality. The four weather information sets each represent a data set that reflects a varying quality of weather data. We consider cases ranging from perfect knowledge (where we know the exact condition of every segment of road at every point in time) to the most minimal amount of knowledge on which we can base a decision (where we have a single average condition value covering all road segments and all points in time). The WRMSp is run using each of these forecast sets, and the effectiveness of the indicated schedules from each forecast set are compared. In this way, we can quantify the benefit to be obtained from an improvement in weather forecast quality.

The case study is based on the winter maintenance plan of the City of Waterloo, Ontario, Canada, which includes a total of 17 service zones with each divided into three subzones. Current practice within the city is to dispatch service vehicles at the subzone level only. Once dispatched to a subzone, the vehicle operator is allotted a rough time interval in which to service the subzone, and decides on their own in what order to treat the residential streets that comprise the subzone. To control the scale of our analysis, we extracted three zones covering approximately 8 km², as shown in Fig. 2. The three zones are broken into nine subzones, each of which roughly encompasses a neighborhood of the city: Zone 1 = {98A, 98B, 98C}; Zone 2 = {88A, 88B, 88C}; and Zone 3 = {67A, 67B, 67C}.

To cast the case to our mathematical model, we must first identify the possible maintenance routes and the associated road segments. Because there are no dedicated routes defined in current practice, we made a simplification by consolidating the individual streets in a subzone as a single loop with a fixed entry and exit node. The length of each loop is equal to the allotted time to treat the corresponding subzone. On their way to and from a subzone (loop), operators may be directed to treat certain major arterial routes along the way. As such, major arterials that link the subzones are considered separately in this model, and are represented by individual edges in the network. Fig. 3 shows the network representation of the region. With this network representation, alternative treatment routes from and to the service depot can be

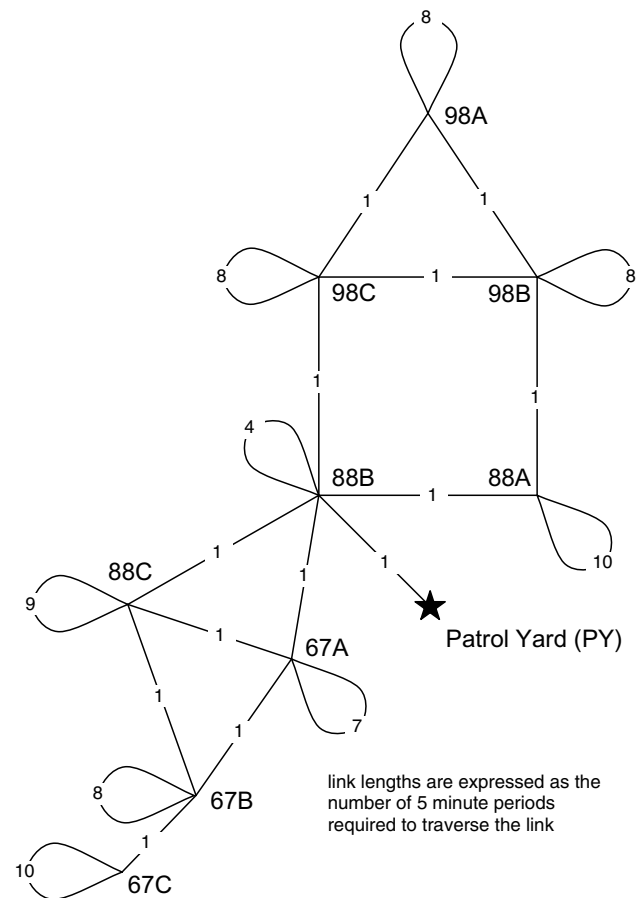


Fig. 3. Case study network diagram.

created. In this research, a total of 11 routes were generated and used as input to our scheduling model.

Table 1 summarizes the attributes of these routes.

We simulate the effects of a winter snow storm, and consider dispatch decisions made in the period before and after the peak of the storm. Snowfall is assumed to be centered on the three subzones, and totals 30cm in Zone 1 (98A, 98B, 98C), and all roads to and from these regions, and 20 cm in the intervening subzones and arterials (Zone 3 – 88A, 88B, 88C), 10 cm in Zone 2 (67A, 67B, 67C), and all roads to and from these regions. The snow storm is assumed to last for 2 hours with its snowfall intensity (rate) increasing gradually during the first 50 minutes and decreasing then after, as described in Table 2, and illustrated in Fig. 4.

The following weather forecasting scenarios are considered, each assuming the same total amount of snowfall over the analysis area but a different level of forecasting accuracy:

1. *Time and spatially accurate*: We have perfect information on the snowfall rates for each road segment in the treatment area at every time interval in the simulation period. This corresponds to a perfect forecast of the storm in both a time and spatial sense.
2. *Time averaged – Spatially accurate*: An average snowfall rate is assumed for each zone over the entire simulation period. Assuming the total amount of snowfall is the same, the snowfall rates, averaged over the simulation period (2 hours), are therefore 0.4, 0.8 and 1.3 cm/time interval, for Zone 1, Zone 2 and Zone 3 respectively.
3. *Time accurate – Spatially averaged*: We have an exact record of snowfall rates for each time interval in the simulation period,

Table 1
Summary of routes used in case study

Route number	Total length (time periods)	Sequence								
R1	14	PY ^a -88B	88B-88C	88C-67B	67B-67B	67B-88C	88C-88B	88B-PY		
R2	6	PY-88B	88B-88B	88B-PY						
R3	11	PY-88B	88B-67A	67A-67A	67A-88B	88B-PY				
R4	13	PY-88B	88B-88C	88C-88C	88C-88B	88B-PY				
R5	14	PY-88B	88B-98C	98C-98B	98B-98A	98A-98B	98B-98C	98C-88B	88B-PY	
R6	8	PY-88B	88B-88C	88C-67A	67A-67B	67B-67A	67A-88C	88C-88B	88B-PY	
R7	12	PY-88B	88B-98C	98C-98C	98C-88B	88B-PY				
R8	14	PY-88B	88B-88A	88A-98B	98B-98B	98B-88A	88A-88B	88B-PY		
R9	18	PY-88B	88B-88C	88C-67B	67B-67C	67C-67C	67C-67B	67B-88C	88C-88B	88B-PY
R10	14	PY-88B	88B-88A	88A-88A	88A-88B	88B-PY				
R11	8	PY-88B	88B-98C	98C-98A	98A-98A	98A-98C	98C-88B	88B-PY		

^a PY – patrol yard.

Table 2
Summary of snowfall by time period and total

Start of time period	10 cm total, accurate	10 cm total, averaged	20 cm total, accurate	20 cm total, averaged	30 cm total, accurate	30 cm total, averaged
0:00	0.0	0.4	0.0	0.8	0.0	1.3
0:05	0.0	0.4	0.0	0.8	0.0	1.3
0:10	0.0	0.4	0.1	0.8	0.1	1.3
0:15	0.1	0.4	0.2	0.8	0.3	1.3
0:20	0.2	0.4	0.4	0.8	0.5	1.3
0:25	0.3	0.4	0.7	0.8	1.0	1.3
0:30	0.5	0.4	1.1	0.8	1.6	1.3
0:35	0.8	0.4	1.6	0.8	2.4	1.3
0:40	1.1	0.4	2.1	0.8	3.2	1.3
0:45	1.3	0.4	2.5	0.8	3.8	1.3
0:50	1.3	0.4	2.7	0.8	4.0	1.3
0:55	1.3	0.4	2.5	0.8	3.8	1.3
1:00	1.1	0.4	2.1	0.8	3.2	1.3
1:05	0.8	0.4	1.6	0.8	2.4	1.3
1:10	0.5	0.4	1.1	0.8	1.6	1.3
1:15	0.3	0.4	0.7	0.8	1.0	1.3
1:20	0.2	0.4	0.4	0.8	0.5	1.3
1:25	0.1	0.4	0.2	0.8	0.3	1.3
1:30	0.0	0.4	0.1	0.8	0.1	1.3
1:35	0.0	0.4	0.0	0.8	0.0	1.3
1:40	0.0	0.4	0.0	0.8	0.0	1.3
1:45	0.0	0.4	0.0	0.8	0.0	1.3
1:50	0.0	0.4	0.0	0.8	0.0	1.3
1:55	0.0	0.4	0.0	0.8	0.0	1.3
Total	10.0	10.0	20.0	20.0	30.0	30.0

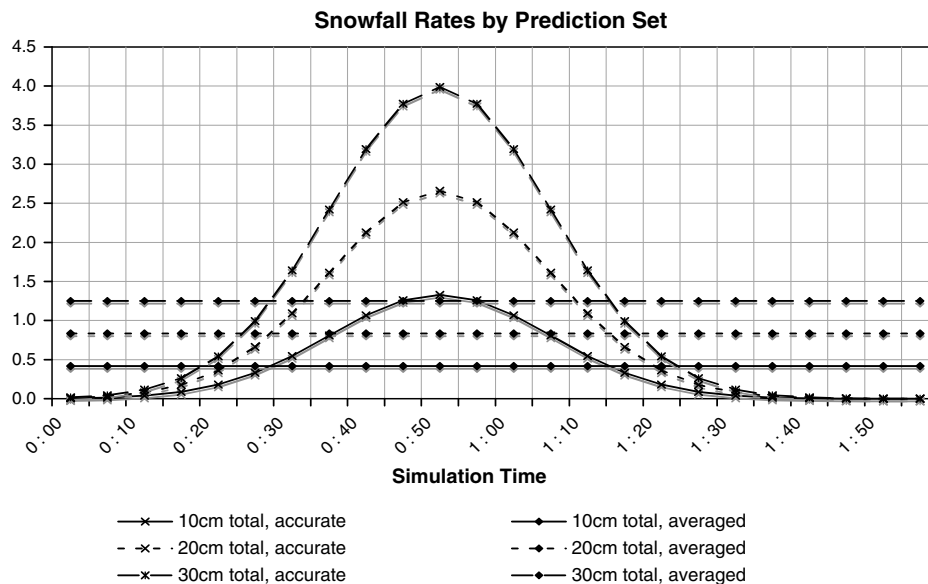


Fig. 4. Graph of snowfall values by time period and total.

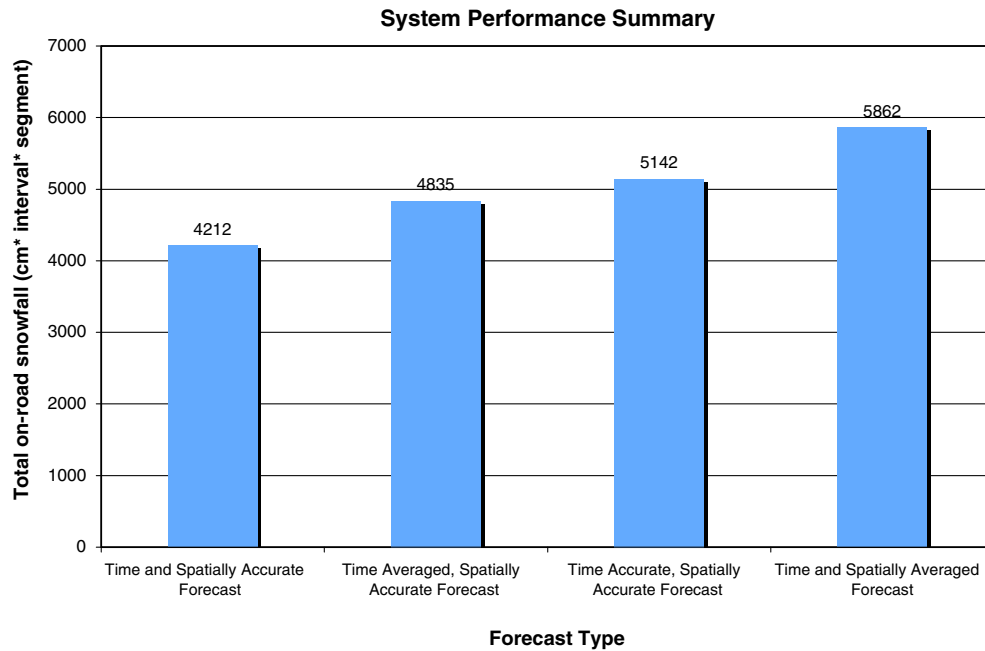


Fig. 5. Summary of case study results.

but that value is averaged over zones, and represents the average value of snowfall for the given time interval over all the road segments.

4. *Time and spatially averaged*: A single snowfall rate is assumed for the entire storm duration for all roads in the treatment area. Assuming the total snowfall is the same, the average snowfall rate for this area over the 2-hour period is 0.8/time interval. This scenario represents the least accurate forecast about the weather.

Each of these four data sets is used to generate a dispatch schedule that is optimal *with respect to the forecast data at hand*. The resulting four dispatch schedules are then evaluated to see how effectively they service the planning region, by evaluating them against the (known) actual weather record (Case 1). In this way, we can quantitatively determine the benefit to be obtained from progressively more detailed weather forecasts.

Aside from the differing weather forecasts, all other data is kept consistent between the four test cases. Specifically, we use the set of routes described above for all four trials, and assume the same fleet size (5 plows), initial weather conditions (0 cm of snow accumulation), road classes (all roads were considered to have the same importance), and service level restrictions (30 cm maximum accumulation) for all four cases. For the objective function weighting factor (β), we used a relatively small β value (0.003) to emphasize the importance of finding the most effective service strategy. As a result, all four test cases output the same cost service strategy (they all output relatively 'expensive' strategies, each with 12 plow dispatches over the 2-hour simulation period). Note that a sensitivity analysis could be conducted to obtain a β value that best replicates local conditions and service policy.

4. Results

The four scenarios were carried out in identical conditions, save for the differing weather sets. The resulting dispatch schedule was obtained for each scenario, and each schedule was evaluated against the actual weather data to determine the schedule's realized performance. We expect that a more detailed weather forecast

will allow more targeted vehicle deployment, and thus improve overall system performance.

System performance is measured by comparing the overall snow accumulation on the streets of the treatment area that results from each of the four computed schedules. In this way, we can determine how effective each schedule was in light of the actual observed snowfall values, rather than in light of the average predicted values (as would be the case if we simply compared the objective values obtained by each test case). Lower values in this respect are better.

The results of this case study are presented in Fig. 5. We observe a 28% predicted reduction in overall snow accumulation throughout the test period when the time and spatially accurate case is compared with the time and spatially averaged forecast. Between the time accurate – spatially averaged, and time averaged – spatially accurate scenarios, we observe that the spatially accurate scenario predicts slightly (6%) improved results, but that both scenarios predict better results than the simple time and spatially averaged case (by a margin of 12% and 18%, respectively). Thus, we expect that any improvement in the resolution of forecast data will yield an improved ability to target vehicle deployments, and thus to reduce overall snowfall accumulation in a planning region. Given the choice between obtaining spatially or temporally improved forecasts, the data suggests that an improvement in the spatial resolution of forecast data will yield superior results. This follows from logic, as snowfall will generally still be on the ground even if a forecast is off by a period of time, but fallen snow will not generally move from area to area over time.

5. Conclusions

We have attempted to provide both a method for producing optimal deployment schedules for winter road maintenance planning in an automated manner, as well as a consistent framework to carry out and compare current and future research work in the field. The underlying winter road maintenance scheduling problem (WRMSP) was formulated as an integer programming problem and solved using a standard algorithm. The scheduling model considers both operating costs and quality of service requirements, as well as

several key road and weather condition factors such as road network topology and class and weather forecasts.

A case study is conducted to assess the benefit of highly accurate real-time weather information on a deployment schedule, and thus on realized treatment levels. A substantial benefit is predicted to be obtained from better resolved weather forecasts, with improved spatial resolution providing a greater expected benefit to service levels than temporal resolution improvements.

The framework as presented here is still in a very early stage. Research is ongoing to improve the fidelity of the model to real life applications, primarily through improving the service and weather accumulation function and treatment models. Research into better route selection algorithms is also ongoing, as a careful mix of route length and segment coverage seems to have a sizeable effect on the effectiveness of solutions proposed by the WRMSP. The user interface of the WRMSP is evolving as well, in order to improve the usefulness of the WRMSP in operational field environments. Further applications of the WRMSP in garbage and recycling vehicle routing, transit routing, and emergency vehicle dispatching are possible.

Appendix A. Deriving location information from departure information

The principal difficulty in determining the effect of a deployment schedule from a central depot is in correlating the effects of a given route departure to the treatment times of roads in the network. A given road segment may be part of many routes, and will typically only be reached after a known delay from the departure time along the route. We need to construct a function that can take as input a set of route departures, determine all of the effects of those departures, and output how many service vehicles will be on each road segment in each time interval. Fig. A1 illustrates this transformation.

To accomplish this translation, we use a four-dimensional binary matrix Q , which we derive directly from l and R . Q contains entries of the form $q_{k,i,t,t'}$, where k, i and t all range over their standard domains, and t' ranges over the same domain as t . It is easiest to understand the use of Q by considering the entries as two-dimensional 'slices' formed by fixing k and i values. For a fixed k and i , we have an array slice $q_{k,i}$ that is indexed by t and t' . Entries of this array are set to 1 if and only if a departure of a vehicle on route k at time t' will result in a vehicle on road segment i at time t . For example, if we consider a road segment i that takes 3 time intervals to traverse, and that is reached on route k at the start of the second time interval after a vehicle departs on route k , then $q_{k,i}$ looks like Fig. A2. As noted, Q is obtained directly from l and R . In the sample implementation of the WRMSP, this is accomplished by defining Q as a derived variable, using the following function to define its values automatically. Formally, the function used to define the entries in Q is as follows:

$$q_{k,i,t,t'} = \begin{cases} 1 & \text{if } t \geq t' + \sum_{\forall t': r_{k,t'} > 0 \text{ and } r_{k,t'} \leq r_{k,i}-1} l_{t'} \\ \text{and } t < t' + \sum_{\forall t': r_{k,t'} > 0 \text{ and } r_{k,t'} \leq r_{k,i}} l_{t'} \\ 0 & \text{otherwise.} \end{cases} \tag{A.1}$$

A.1. Using Q to relate X and Y

Having Q at hand, we need to use it to relate X to Y , as that is the main purpose behind Q 's existence. Recall that X is a two-dimensional matrix of values $x_{k,t}$, where a value is set to 1 if and only if a vehicle departs on route k at time t . Thus, a single row of X represents all of the departures of vehicles on the corresponding route, laid out in time increasing order. Notice also that the '1' entries in a row of a slice of Q represent the time intervals at which a vehicle would have had to have left the service yard on the corresponding route to be on the corresponding road segment at the corresponding time. Thus, the dot product of these two vectors is exactly the number of vehicles from route k that will be on road segment i at time t . Thus, for a fixed k and i , the matrix product

$$z_{k,i} = q_{k,i} \cdot X_k^T \tag{A.2}$$

Let $Z = \{z_{k,i,t}\}$ where $z_{k,i,t}$ represent the number of vehicles on road segment i at time t due to route k . Note that the T superscript above refers to the transpose matrix operation, and not to the number of time periods in the model. A simple summation would produce the values of Y as a function of X, l , and R , as shown in Eq. (A3)

$$y_{i,t} = \sum_{k=0}^{m-1} z_{k,i,t} \tag{A.3}$$

A.2. Treatment time modeling

The values of Y are crucial to derive and maintain fleet sizing restrictions. However, to correctly model the treatment of a road segment, we need to assume a single time of treatment for an entire road segment. This is due to the fact that a single service vehicle can only be at one place on a road segment at a time. Thus, if we consider the road as being 'treated' at all times that a vehicle is on it, we end up having the first physical section of a road segment being considered as 'treated' when the service vehicle is at the other end of the road segment.

For this purpose, we derive a table Y' which we will use for treatment time decisions. This table is based directly on $S = \{s_{k,i,t,t'}\}$, a modified version of Q in which only one entry per column in a slice of s is set to 1. This means that in the resultant table Y' , for every vehicle service event on a road segment, only

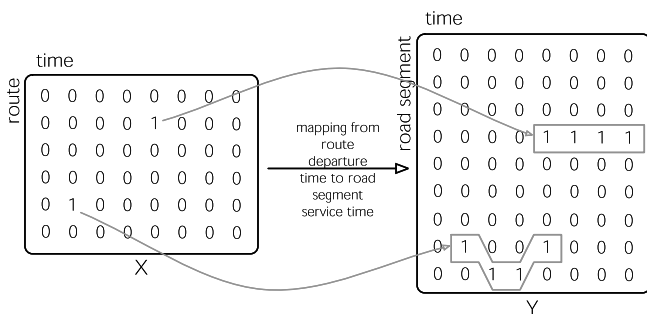


Fig. A1. Mapping from route departures to service times.

time to depart on route k (t')

0	0	0	0	0	...	0	0
1	0	0	0	0	...	0	0
1	1	0	0	0	...	0	0
1	1	1	0	0	...	0	0
0	1	1	1	0	...	0	0
...	0	0
0	0	0	1	1	1	0	0
0	0	0	0	1	1	1	0

$Q_{k,i}$

Fig. A2. Sample slice of Q .

one entry of Y' is set to 1, and the corresponding time interval is used as the representative treatment time for this service event on this road segment. Specifically, entries in Y' are set to 1 if and only if they represent the first time interval during which a vehicle services a road segment.

The derivation of S proceeds from that of Q as follows:

$$s_{k,i,t,t'} = \begin{cases} 1 & \text{if } q_{k,i,t,t'} = 1 \text{ and } \sum_{u=0}^t q_{k,i,u,t'} = 1, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A.4})$$

From S , we proceed to define Y' (and thus Z') as follows:

$$z'_{k,i} = s_{k,i} \cdot x_k^T. \quad (\text{A.5})$$

Similar to Z , Z' is a three-dimensional matrix. However, in contrast to Z , the entries $z'_{k,i,t}$ represent the number of vehicles that are considered to be treating road segment i at time t due to route k . Again, note that the T superscript above refers to the transpose matrix operation, and not to the number of time periods in the model. A simple summation of the following form produces the required values of Y' :

$$y'_{i,t} = \sum_{k=0}^{m-1} z'_{k,i,t}. \quad (\text{A.6})$$

References

- Andrey, J., Mills, B., Jones, B., Hass, R., Hamlin, W., 1999. Adaptation to climate change in the Canadian Transportation Sector. Natural Resources Canada, Adaptation Liaison Office, Ottawa.
- Fourer, R., Gay, D.M., Kernighan, B.W., 2003. AMPL: A modeling language for mathematical programming. Duxbury Press & Brooks/Cole Publishing Company.
- GNU Software Project, 2006. GNU Linear Programming Kit. Available from <<http://www.gnu.org/software/glpk/>> (cited 01.09.06) (online).
- Jordan, R., 1991. A one-dimensional temperature model for a snowcover: Technical documentation for sntherm.89. Technical Report CRREL Special Report 91-16, US Army Cold Regions Research and Engineering Laboratory, Hanover, NH.
- Keyser, J.H., 1981. Chemicals and abrasives for snow and ice control. In: Gray, D.M., Male, D.H. (Eds.), Handbook of Snow. Pergamon Press, pp. 580–612.
- MDSS National Laboratory Consortium, 2004. The maintenance decision support system (MDSS) project: MDSS prototype release-3 technical description. Technical report, MDSS National Laboratory Consortium.
- Ontario Ministry of Transportation, 2004. Maintenance Technology Project 2003–2004 Overview.
- Pell, K., 1994. An improved displacement snowplow. Technical report, Department of Mechanical Engineering, University of Wyoming.
- Perchanok, M., 2002. Decision Support System for Winter Maintenance: The DART Database. Ontario Ministry of Transportation Report.
- Perrier, N., Langevin, A., Campbell, J.F., 2006a. A survey of models and algorithms for winter road maintenance. Part I: System design for spreading. Computers and Operations Research 33, 209–238.
- Perrier, N., Langevin, A., Campbell, J.F., 2006b. A survey of models and algorithms for winter road maintenance. Part II: System design for snow disposal. Computers and Operations Research 33, 239–262.
- Perrier, N., Langevin, A., Campbell, J.F., 2007a. A survey of models and algorithms for winter road maintenance. Part III: Vehicle routing and depot location for spreading. Computers and Operations Research 34, 211–257.
- Perrier, N., Langevin, A., Campbell, J.F., 2007b. A survey of models and algorithms for winter road maintenance. Part IV: Vehicle routing and fleet sizing for plowing and snow disposal. Computers and Operations Research 34, 258–294.
- SHRP Anti-Icing/RWIS Lead States Team, 2000. 1999–2000 technology transfer and usage survey results. Technical report, Strategic Highway Research Program.
- Trudel, M.R., 2005. A Mathematical Model for Winter Maintenance Operations Management. Master's Thesis, Department of Civil Engineering, University of Waterloo.