Scheduling and Routing Grass Mowers Around Christchurch

James Tipping
Department of Management
University of Canterbury
New Zealand
jpt34@student.canterbury.ac.nz

Abstract

Creating daily routes for grass-mowers around a city's municipal parks can be modelled as an interesting variation of the Vehicle Routing and Scheduling Problem. In this case, complications mean that the typical VRSP solution methods must be extended. For example, not all of the grassed locations need to be mown every day, though there are soft constraints imposed on the length of the grass – the longer the grass at a location grows, the more urgently it needs to be mown. Further, many variables influence the time it takes to mow a location, and because of this uncertainty, the routes constructed must be long enough to ensure a full day's work. These complications result in a problem that is too large to be solved analytically on a daily basis. We report on the design, performance and suitability of a system containing a new simulated annealing search heuristic for the grass-mower scheduling and routing problem.

1 Introduction

City Care Limited, 100% owned by the Christchurch City Council, is responsible for the maintenance of over 1500 council-owned parks, cemeteries and roadside berms. Mowing the grass at these locations in Christchurch costs City Care approximately $2 million a year, and it is therefore vital that this operation be as efficient as possible.

We were approached by City Care in February 2002 and asked to provide a more efficient system for their grass-mowing operations than the manually-operated fixed-schedule system presently in use. In this environment, increasing efficiency is achieved through minimizing variable costs and increasing the quality of service provided. Our research identified two areas in which this could be achieved. Firstly, monitoring more closely the length of grass at each location would increase the chance of locations being mown only when necessary, thereby decreasing the cost of having them mown too often, and reducing complaints from mowing too infrequently. Secondly, the creation of an efficient routing system would reduce the travelling time of each mower operator, thereby reducing fuel costs and increasing the total time they are available to mow in a day.

Monitoring the approximate length of the grass at each location can be achieved by using a weather-based model for grass growth, and knowledge of the growth rate of the grass at each location. While the actual specifics of grass-growth modelling were ruled outside the scope of this project, combining the research of Moir *et al* (2000) and Scotter *et al* (2000) provided us with a model suitable for our purposes. Our system estimates daily grass growth at each location based on temperature, sunshine hours and rainfall.

The routing aspect of the project fell under the classical Operations Research label of a Vehicle Routing Problem (VRP). Cordeau *et al* (2002) and Laporte and Osman (1995) have compiled useful guides to the various classes of VRP that exist, and the methods available to solve them. However, whilst the City Care grass-mowing problem contained aspects similar to many of the classes described, the combination of these aspects and other features meant that in its entirety, the problem is unique. For example, the time-constrained nature of the routes was similar to the Vehicle Routing and Scheduling Problems (VRSP) extensively documented in the literature (see Desrosiers *et al* (1995)). A second feature of the grass-mower VRP was that, as not every location had to be included in the routes every day, allocation became an issue. This meant that the problem also fell into a more recently documented category, the Vehicle Routing-Allocation Problem (Beasley and Nascimento (1995)). Other aspects, such as fuzzy mowing times and flexible fleet configurations, further increased the problem's complexity. Due to the size and structure of the problem, it was obvious that a heuristic method would be required to obtain good routes for the grass-mowers. However, none of the models we found in the literature were suitable, and a new framework was required.

Section 2 of this paper details the scheduling and routing problem at City Care and their current approach. In Sections 3 and 4, we describe our model for the problem, along with details of the heuristic we developed to obtain daily routes, and the infrastructure supporting the heuristic. In the final sections of paper, examples of some of the solutions to the heuristic are provided to illustrate its performance, along with some suggestions for extensions to the program and implementation.

2 The City Care Grass Mower Routing and Scheduling Problem

City Care is under contract with the Christchurch City Council to maintain the length of grass at the locations on its asset list at between 20mm and 60mm at all times, and for the length of the grass at cricket grounds to be less than 40mm during the summer season. While these are not strictly "hard" constraints, if the grass grows much longer, complaints may be received from the public, or from the Council's auditors. A further complication to the requirements is the fact that during the cricket season, the outfields at cricket parks must be mown on both a Monday and a Thursday (or Tuesday and Friday if there is not enough time on the previous day), to ensure the grass is short enough for midweek training and weekend matches.

2.1 The Current Situation at City Care

As of October 2002, City Care holds 29 mowers of seven different types, and has 22 permanent mowing staff to whom they allocate one mower each per day. There is a normal configuration of mowers used every day, and City Care has the option of hiring contract mowers during the busy spring and early summer months. At the start of a week, each operator is given a week's worth of mowing in the form of a single route. Their task is to get as far as they can along the route in a day, return to the depot after their 8 hours of regular time, then start mowing the following day from the next location on their route. The routes for each operator are found using a single-route TSP among all the locations in a certain sector of Christchurch.

The major shortcoming of the current route-construction method is that it takes no account of the length of grass at each location, with fixed routes and intervals between mowing at each location. The grass at each of the locations around Christchurch grows at different rates, without any discernible geographic pattern. This combination of factors causes losses in efficiency if, for example, the grass at two adjacent locations grows at completely different rates. At times when there is no water in the root zones of the grass at non-irrigated parks, an irrigated park may still grow up to 7 mm per day, but an adjacent park without irrigation may have no growth at all. With the period between mowing at each location roughly fixed (to ensure that the faster growing locations are kept below the required length), the slower growing locations will be mown more often than they need to be. This creates the greatest loss of efficiency in the current system.

2.2 The Project Requirements

The problem we faced was to decide each day whether each location should be mown (based on how urgently it needed mowing), and, given the single type of mower required for each location, what route each mower operator should take around Christchurch that day. The optimal assignment of mower types to individual locations was ruled outside the scope of this project.

The estimated set-up and mowing times for each location supplied were fairly fuzzy, and most likely to be overestimated. A separate study (again outside the scope of this project) would be required to estimate the relationship between grass length, the weather, and mowing times. For this reason, City Care required that the daily routes constructed be a couple of hours longer than the eight hours the mower operators actually have available. If the mowing times were in fact accurate, then, in the interests of efficiency, it would be important that the operator had a relatively short trip back to the depot. Therefore, it was necessary that for the final two hours of each route, the later a location's mowing would be finished, the closer that location must be to the depot.

Due to this uncertainty in the mowing times, the later in the day a location is scheduled to be mown, the less likely it is actually to be mown that day. Therefore, the longer the grass is at a location, the more urgently it requires mowing, and the earlier in the day its mowing should be scheduled. In addition, as not all locations selected for the routing heuristic may be included in the final routes, the more urgently a location requires mowing the more its inclusion in a route should be encouraged.

Obviously, if City Care were to implement the results of this project, the routing and scheduling system would have to be packaged inside an easy-to-use interface, and come with a clear and concise user manual. The program would have to store data on each of the different locations, house the grass-growth model for estimating the length of grass at each location, contain the heuristic and allocation methods, and illustrate the daily routes on a map. We decided that this complete program and user manual, ready for immediate implementation, would be the major deliverables to our client.

3 The Routing and Scheduling Model

As part of the problem involves allocating which locations will be mown in a day and which will not, the first decision required is to decide which locations will be selected for possible inclusion in a route. From this selection, all the locations not in routes are placed in the "pool of excess locations", from which locations can be swapped into and out of routes. Locations in the final excess pool will not be mown that day, but are likely to be included in a route the following day, as their grass length (and corresponding urgency) will have increased.

The mower routing and scheduling model follows most of the usual conventions of a typical time-constrained VRP, with the use of penalty functions as an extension to the usual distance-based objective function. The aim of the model is to construct routes that cover as many of the selected locations as possible, and to ensure that the mowing of the most urgent locations is scheduled within the expected finish time of the rounds. The lower the urgency of a location, and the closer it is to the depot, the later in the route it can be scheduled. Note, however, that it must always be preferable for a location to be placed in a route rather than in the pool of excess locations. The descriptive model for a single mower type takes the following form:

Sets

L = Locations that have been assigned the particular mower type
I = Locations urgent enough to be selected for the model ($I \in L$)
C = Cricket grounds selected for the model ($C \in I$)
R = Locations scheduled in a route ($R \in I$)
E = Locations in the pool of excess locations ($E = I - R$)

Parameters

r = The number of routes constructed
u_i = The mowing urgency of location i: an exponential function of tomorrow's forecasted grass length at location i
t_i = The estimated travel time from location i to the depot
p_x = The constant term in the function for the penalty on x (x ={urgency, distance, travel time from depot, pool of excess locations})

Variables

T = Total scheduled travel time on all the routes for that mower type
D_i = The difference between the scheduled mowing completion time for location i and the prescribed mowing round time (PMRT). The mowing round time is the time it takes to travel from the depot to each location on the route, mow each location, and travel back to the depot.
P_i^u = $p_u u_i \exp^{D_i}, \forall D_i > 0$ = Urgency penalty on a location i
P_i^d = $p_d t_i \exp^{D_i}, \forall D_i > 0$ = Travel time from depot penalty on a location i
P_i^c = $p_c \exp^{D_i}, \forall D_i > 0$ = Cricket penalty on a location i
P_i^e = $p_e u_i$ = Penalty on a location i in the pool of excess locations

Model GRASS

1) Minimise $T + \sum_{i \in R} P_i^u + \sum_{i \in R} P_i^d + \sum_{i \in C \in R} P_i^c + \sum_{i \in E} P_i^e$
2) All routes must take less than the PMRT + the assumed buffer time for each route, which allows for early completion.
3) Every location selected for the VRP heuristic must either be in the pool of excess locations or in one of the routes

4 The Routing and Scheduling Heuristic

Due to the sheer size and nature of the problems we had to solve, it was obvious we would have to use a heuristic rather than an optimisation technique, with a higher-level procedure to determine the best daily allocation. We based our initial heuristic on a tabu search for simple VRPs proposed by Barbarosoğlu and Özgür (1999). Despite the popular idea that tabu search is better suited for VRPs, our tests found that solution times could be reduced significantly by using a simulated annealing move acceptance criterion rather than using a tabu list, without any noticeable loss in solution quality.

By the time we had finished adapting heuristic to create mowing rounds for City Care, only some of its original structure remained. Descriptions of the modified versions of their construction and search heuristics are in Sections 4.1 and 4.2. Section 4.3 contains information on the higher-level allocation procedure, which calls the construction and search methods many times, and the overall framework of our heuristic.

4.1 The Construction Procedures

The first step was to construct a list of locations from which the heuristic would create its routes. If we were creating r routes, we ranked all the locations for a particular mower type in decreasing order of urgency, then selected the $(8+2)r$ hours of most urgent mowing, and the next c most urgent locations, to be on this list.

Barbarosoğlu and Özgür's (1999) paper described two construction procedures, one random and one a nearest neighbour-type method, both of which we kept. Once the required r routes were constructed, the rest of the locations on the list were placed in the pool of excess locations. Please refer to the paper for details of these two methods.

4.2 The Search Procedure

The search procedure, which we labelled the "Route Search Method" (RSM), involves iterating through two procedures and using a simulated annealing approach to neighbour acceptance. These procedures, the Basic Improvement Method (BIM) and Relative Improvement Method (RIM), incorporate different approaches to generating neighbouring solutions to the current solution. At the end of each iteration of RSM, the best neighbours found by BIM and RIM are compared, and the one with the lower of the two objective function values (OFVs) is considered for acceptance. If the neighbour has an improved OFV over the current solution, then it automatically becomes the current solution. Otherwise, the neighbour may be accepted if the following condition on the increase in OFV is met:

$$\text{Uniform }[0,1] < e^{-\left(\frac{B-C}{t}\right)}$$

where: B = the OFV of the neighbouring solution
C = the OFV of the current solution
t = the current temperature of the SA system.

4.2.1 Basic Improvement Method (BIM)

A single iteration of BIM involves swapping randomly-selected locations between two randomly-selected routes (or one randomly-selected route and the pool of excess locations), inserting the swapped locations into their destination routes, and performing a fully enumerated 2-opt procedure until no more improvements can be found. The insertion procedure finds the best feasible position in which to insert the swapped location into the destination route, according to the penalty- and distance-based objective function. However, if one of the 'routes' selected is the pool of excess locations, whereby order is irrelevant, the 2-opt procedure is not called after the locations are added to the pool. Multiple iterations of BIM are performed on the provided solution, with the resulting solution that has the lowest objective function value proposed as the neighbouring solution for consideration.

4.2.2 Relative Improvement Method (RIM)

The key distinction between an iteration of RIM and an iteration of BIM lies in the way that RIM selects the locations to consider for swapping between the two chosen routes. Rather than randomly selecting these locations, RIM uses intelligent methods to make the selections, based on the position of each location relative to the centroid of its own route and that of the other route. The goal of this process is to remove locations from a route when they are relatively far from the other locations in their current route and relatively close to the locations in the destination route.

4.3 The Structure of the Heuristic

Figure 1 shows the overall structure of the heuristic. As explained above, both BIM and RIM incorporate the 2-opt and Insertion methods, each performed for a number of iterations. After constructing several solutions from each of the two construction methods, the First Improvement method develops each of them using RSM, to enable an educated decision as to which of the constructed solutions most warrants further consideration. The SASearch procedure brings together the whole process of finding routes for a single mower type, by incorporating the two construction methods, the First Improvement stage and the Major Improvement stage (which involves many iterations of RSM on a single solution in order to find the best routes possible).

The highest level of the heuristic involves the overall allocation of mower operators to types of mower. This process begins by running a preliminary SASearch process for each mower type, over a range of numbers of routes. For each of these solutions, a measure of their goodness is obtained by summing the urgencies of all locations for the particular mower type that have been left out of the routes (i.e. in (L – R)). The allocations are then determined by considering first the standard configuration of mowers, and then investigating the marginal benefits of slight variations in these allocations. Following the allocation procedure, the cooling parameters are reset, and the solutions for the resulting allocation selected undertake a Final Intensification stage of the RSM procedure.

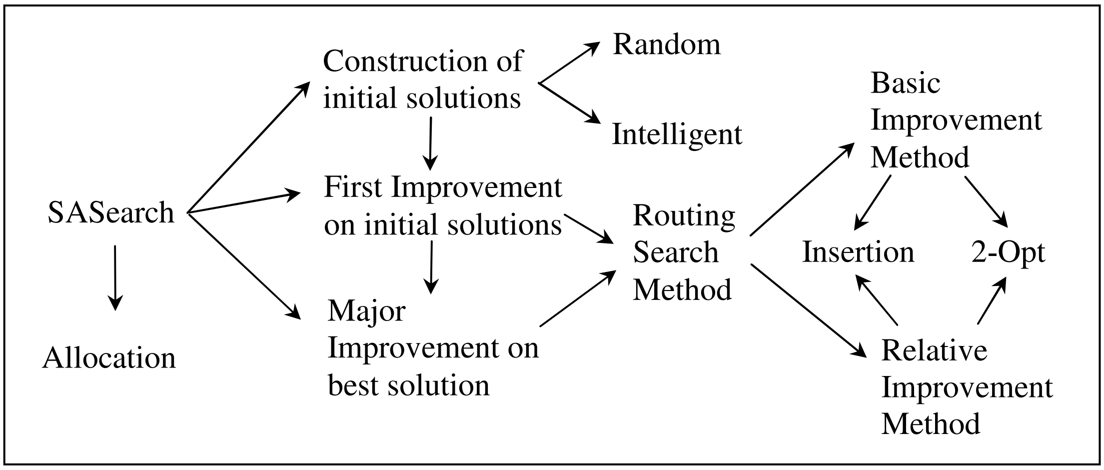

Figure 1. The structure of the routing and scheduling heuristic

5 Results

When testing our cut-down heuristic (with the S.A. acceptance criterion) from the original paper on the test problems in Cordeau *et al* (2002), our best solution values were between 0 and 3% worse than the best-known results, and solution times were significantly faster than those of the five metaheuristics listed.

After coding the heuristic described in Section 4, we conducted exhaustive testing to find an appropriate cooling schedule. We found that any schedule that had a high and very slowly decreasing temperature provided good solutions, due to the nature of the acceptance criterion, and the scale of the changes in the objective function value.

5.1 The Structure of the Routes Constructed and the Overall Solutions

The heuristic performed extremely well both on the set of data for which it was built, and when the data set was manipulated to test for over-fitting. Figure 2 below illustrates the influence of the time of day / travel-time penalty on the heuristic. Note that the locations scheduled to be completed after eight hours are much closer to the depot than most of the other locations on the route.

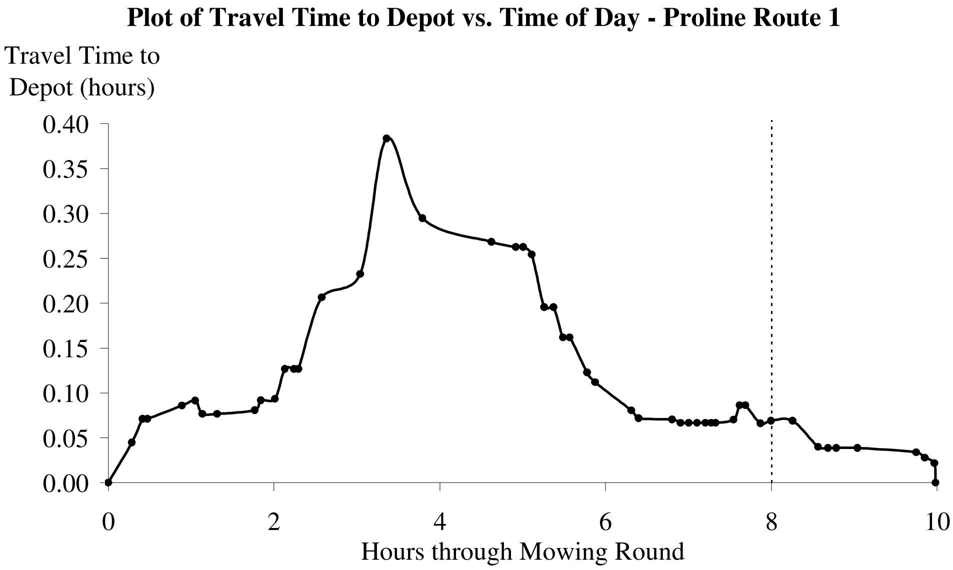

Figure 2. The structure of a route in terms of travel time to depot

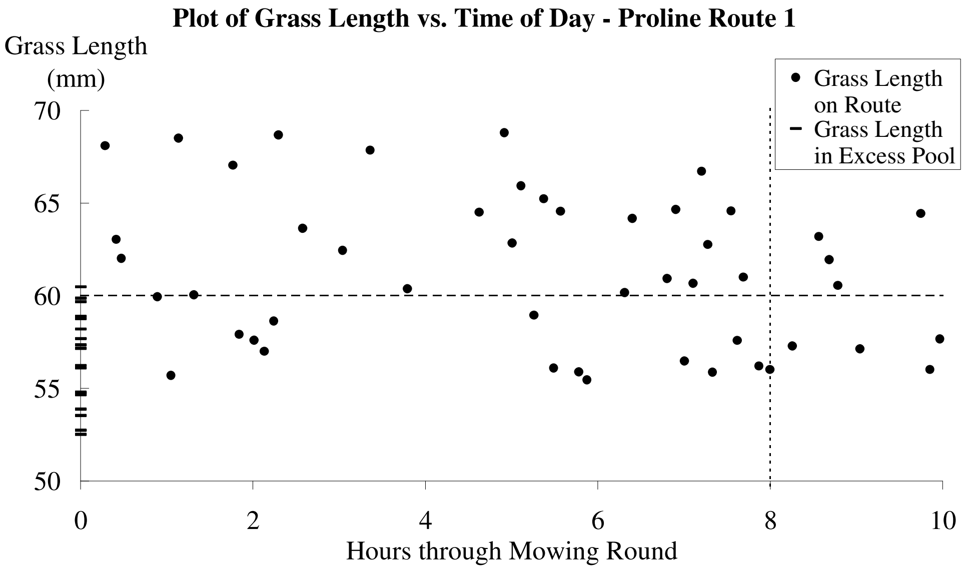

Figure 3. The structure of a route in terms of grass length / urgency

Figure 3 above shows a typical graph of location grass length against scheduled location finishing time, calculated for one of the mowers' routes. Very seldom were any truly urgent locations scheduled in the last couple of hours of a route, and there were never any very urgent locations placed in the pool of excess locations either.

The routes of an optimal solution to a distance only-based VRP or TSP will have no crossovers, and the routes will look like petals of a flower centred at the depot. As minimising total distance travelled was not the only aim of this heuristic, we cannot say with certainty what the optimal solution will look like – crossovers inside a route may indeed be best. However, due to the use of the modified 2-opt procedure, crossovers in the routes are rare, and the flower structure is certainly evident from the map plots of all the routes for a particular type of mower. The daily rounds created for any type of mower can be printed out by our program, as shown in Figure 4 below.

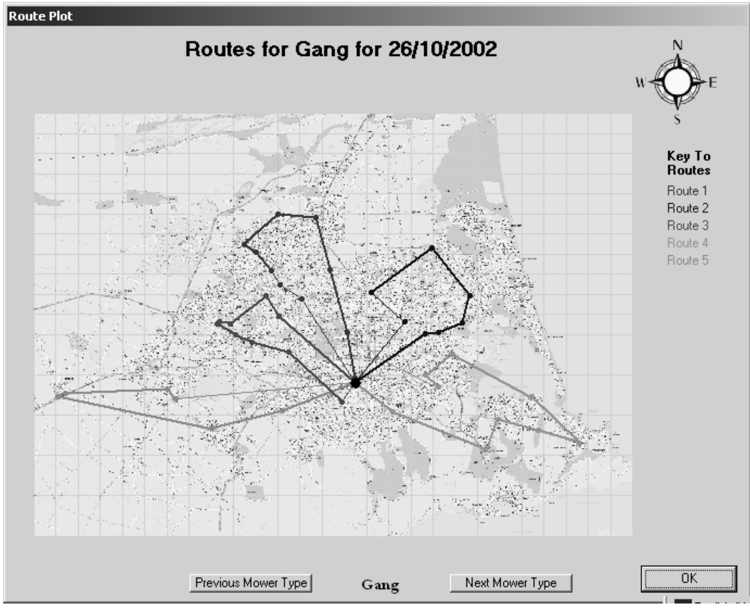

Figure 4. Map of Christchurch showing daily rounds for one type of mower

Using the pool of excess locations and multiple selection criteria over and above travel time certainly aids in constructing intelligent routes, as Figures 2, 3 and 4 above show. For example, if a location were far away from most of the other locations mown by the same type of mower, including it in a route as soon as its length reaches 60mm would cost a great deal in terms of travelling time. However, our penalties are tuned to allow it to be left in the excess pool for a few days until either its urgency becomes too great, or other locations nearby become urgent and are put in a route together.

5.2 Solution Time

After observing progress graphs of the search, and taking into account the testing of the cooling schedule, we decided upon the most suitable schedule of iterations for the heuristic: 120 for each First Improvement search, and 1000 for each of the Major Improvement and Final Intensification searches. With this schedule of iterations, the entire procedure takes on average 140 minutes to run for the seven types of mowers currently held[1]. This solution time is perfectly acceptable to City Care, as the entire set of routes can be created overnight. The program offers a contingency for breakdowns, creating a single route of the most urgent locations, with a much faster solution time.

5.3 Numerical Results

A measure of performance can be drawn from the heuristic's ability to schedule the most urgent locations into the first eight hours of a route rather than in the last two

[1] The computers used to test the program were P3 663 MHz.

hours, or in the pool of excess locations. The first two columns of Table 1 below attest to the performance of the heuristic in this regard.

We were able to compare the relative performance of our 8-hour daily routes with a single 40-hour week-long route, and it showed that having daily routes could reduce total travel by up to 12%, a significant saving in terms of fuel costs. City Care can compare the last three columns of Table 1 to their current cost figures; however, as the current mowing rounds are calculated solely on distance, comparisons are limited.

Mower	Avg urgency in final two hrs	Avg urgency in excess pool	Avg % Travel Time in Round	Avg Round Length	Avg Locations per Round
72 Inch	47.2%	26.5%	9.3%	9.78 hrs	12.0
Gang	44.2%	15.3%	9.7%	9.86 hrs	10.0
Proline	48.9%	12.6%	19.7%	9.97 hrs	59.9
Walker	51.9%	6.6%	17.0%	9.98 hrs	71.8
Z-Master	35.0%	0.2%	6.7%	9.88 hrs	13.0
11ft Rotary	60.3%	9.7%	17.3%	9.95 hrs	12.5

Table 1. Average urgencies as a percentage of average urgency in first eight hours, and route statistics for individual mower types

6 The Benefits to City Care of Implementation

Implementation of our system would obviously lead to more accurate monitoring of grass lengths and more efficient routing and scheduling of grass mowers. This achieves the overall aims specified in Section 1. Aside from the calculable benefits to City Care, we foresee a reduction in the amount of labour-intensive and skilled work required to create routes for the grass mowers. The only human activity required to run our program is in the form of data entry, i.e. ticking boxes at the end of each round when locations are mown, and entering mowing times if desired.

Many of the activities currently performed in a haphazard manner at City Care are automated in our program. For example, if a mower breaks down in the middle of the day, the operator usually returns to the depot, picks up a new mower (often of a different type) and returns to complete their original round. Our system can create a new round for the number of hours left in the day, and will assign the type of mower that can most reduce total urgency. Further, if a complaint is received about the grass at a location, its length in the system can be greatly increased, to ensure that it is mown the following day. Similarly, if a particular location must be kept strictly below the upper length limit, all that is required for it to be mown more frequently is to increase its growth rate – the system will always think the grass at the location is longer than it actually is.

7 Shortcomings, Possible Extensions and Conclusions

Using real distances by road, instead of scaled Euclidean distances, would create routes that are more logical than some of the current ones, particularly given Christchurch's geographical features (e.g. the rivers and estuary). Further, using a time- or road-dependant travel speed instead of a single average speed would have been desirable.

On a wet day, the routes created are likely to be far too long, due to an up-to-fivefold increase in mowing times. Applying a user-specified scale factor to all the mowing times would be one way of increasing them on a wet day, but it would be preferable instead to use some link between the weather data and the mowing times. Our program does use a weighted average of the most recent mowing times for each location (as entered by the system user) to calculate that location's expected mowing time, and we would consider that estimates should improve with time.

Incorporating the efficient assignment of mowers to locations into the heuristic would be a very large (and possibly intractable) extension to the current model, but would allow a more detailed vehicle requirements study.

One benefit of the weekly routes created currently is that the operators know in advance which locations they are going to be visiting every day of the week, which may suit them better for planning their routes. However, sharing locations around the operators is a better practice for auditing mowing times. Further, daily schedules offer the best solution to breakdowns and non-completion of rounds.

Overall, we believe that implementing our routing and scheduling system would be of great benefit to City Care, not only from a financial point of view but also in terms of increasing organisational efficiency. Without any suitable methods from previous studies available, the heuristic we have designed creates good routes that have the desired structure. The integration with the grass growth model for scheduling purposes is one of the key features of this system, which differentiates it from any previous studies we have found.

8 Acknowledgements

This paper was produced as the result of a project conducted in conjunction with Paul Stewart.

Dominic Heritage, the Operations Researcher for City Care, compiled most of the data for this project, which saved us many hours of sifting through timesheets. Along with his colleagues, Steve Chandler and Mark Aydon, he also gave several hours' assistance to us in explaining the requirements of the project and City Care's current mowing operations.

Our project supervisors, Dr Shane Dye and Dr John Giffin, provided invaluable expertise and assisted us in many of areas of the project. Dr Ross James also provided significant help with our software and ideas for the application of the heuristic.

9 References

Barbarosoğlu, G. and D. Özgür. 1999. "A tabu search algorithm for the vehicle routing problem." *Computers & Operations Research*, **26**: 255-270.

Beasley, J. E. and E.M. Nascimento. 1995. "The vehicle routing-allocation problem: A unifying framework." Working Paper, Imperial College, London, England.

Cordeau, J-F., M. Gendreau, G. Laporte, J-Y. Potvin and F. Semet. 2002. "A guide to vehicle routing heuristics." *Journal of the Operational Research,* **53**: 512-522.

Desrosiers, J., Y. Dumas, M. Solomon and F. Soumis. 1995. "Time constrained routing and scheduling." In Ball, M.O., M.L. Magnanti and G.L. Nemhauser (Eds). *Handbooks in Operations Research and Management Science: Network Routing,* Elsevier, Amsterdam, **8**: 35-139

Laporte, G. and I.H. Osman. 1995. "Routing Problems: A bibliography." *Annals of Operations Research*, **61**: 227-265.

Moir, J.L., D.R. Scotter, M.J. Hedley and A.D. Mackay. 2000. "A climate-driven, soil fertility dependent, pasture production model." *New Zealand Journal of Agricultural Research*, **43**: 491-500.

Scotter, D.R., D.J. Horne and S.R. Green. 2000. "A procedure for generating daily rainfall and evaporation data: an evaluation and some applications." *Journal of Hydrology (NZ)*, **39**(1): 65-82.