DYNAMIC COLLECTION SCHEDULING USING REMOTE ASSET MONITORING: A CASE STUDY IN THE CHARITY SECTOR

Fraser McLeod
Transportation Research Group, Engineering and the Environment, University of Southampton, Southampton. SO17 1BJ. United Kingdom. Phone: +44 (0)23 80593316 Fax: +44 (0)23 80593152
E-Mail: f.n.mcleod@soton.ac.uk

Gunes Erdogan
Southampton Management School, University of Southampton, Southampton. SO17 1BJ. United Kingdom. Phone: +44 (0)23 80598882 Fax: +44 (0)23 80593844 E-Mail: g.erdogan@soton.ac.uk

Tom Cherrett *(corresponding author)*
Transportation Research Group, Engineering and the Environment, University of Southampton, Southampton. SO17 1BJ. United Kingdom. Phone: +44 (0)23 80594657 Fax: +44 (0)23 80593152
E-Mail: t.j.cherrett@soton.ac.uk

Tolga Bektas
Southampton Management School, University of Southampton, Southampton. SO17 1BJ. United Kingdom. Phone: +44 (0)23 80598969 Fax: +44 (0)23 80593844 E-Mail: t.bektas@soton.ac.uk

Nigel Davies
Computing Department, Lancaster University. LA1 4YR. LA1 4YR. United Kingdom. Phone: +44 (0)1524 594337 Fax: +44 (0)1524 593608 E-Mail: nigel@comp.lancs.ac.uk

Chris Speed
Edinburgh College of Art, The University of Edinburgh, Edinburgh. EH3 9DF. United Kingdom. Phone: +44 (0)131 221 6000 Fax: +44 (0)131 651 4229 E-Mail: c.speed@ed.ac.uk

Janet Dickinson
School of Services Management, Bournemouth University, Bournemouth. BH12 5BB. United Kingdom. Phone: +44 (0)1202 965853 Fax: +44 (0)1202 515707 E-Mail: jdickinson@bournemouth.ac.uk

Sarah Norgate
Directorate of Psychology & Public Health, School of Health Sciences, University of Salford, Salford. M6 6PU. United Kingdom. Phone: +44 (0)161 2952324 Fax: +44 (0)161 295 5077
E-Mail: S.H.Norgate@salford.ac.uk

Number of words = 6720
Number of figures = 2 = 500 words
Number of tables = 1 = 250 words
Total = 7470 words

Submission Date: 15/11/12

ABSTRACT

In the waste collection sector, remote sensing technology is now coming onto the market, allowing waste and recycling receptacles to report their fill levels at regular intervals. This enables collection schedules to be dynamically optimised to better meet true servicing needs, so reducing transport costs and ensuring that visits to clients are made in a timely fashion. This paper describes a real-life logistics problem faced by a leading UK charity in servicing its textile and book donation banks and its High Street stores using a common fleet of vehicles with varying carrying capacities. This gives rise to a vehicle routing problem whereby visits to stores are on fixed days of the week, with time window constraints, and visits to banks (fitted with remote fill monitoring technology) are made in a timely fashion to avoid them becoming full before collection. A tabu search algorithm was developed to provide vehicles routes for the next day of operation, based on maximising profit. A longer look-ahead period was not considered on the basis that donation rates to banks are highly variable. The algorithm included parameters specifying the minimum fill level (e.g. 50%) required to allow a visit to a bank and a penalty function used to encourage visits to banks that are becoming full. The results showed that the algorithm significantly reduced visits to banks and increased profit by up to 2.4% with best performance obtained the more variable the donation rates.

INTRODUCTION

Many waste collection operations run on a traditional fixed round basis with collections typically made from individual locations on the same day(s) at approximately the same time(s) each week. With remote sensing technology now coming onto the market that allows waste receptacles to report their fill levels at regular intervals (*1*), collection schedules can be dynamically optimised to better meet their true servicing needs, so reducing transport costs and ensuring that visits are made in a timely fashion (e.g. to avoid overspilling). Dynamic optimisation of collection schedules will be particularly suitable where receptacles experience highly variable fill rates, from day to day and from week to week, making it difficult to accurately estimate the fill level and plan collection schedules without the use of remote monitoring data. Remote monitoring is also appropriate where the receptacles are situated in remote locations, and long distance trips to partly filled receptacles need to be avoided. Remote monitoring can also be used where the receptacles contain items of value, rather than waste, where the main objective is to maximise profit from sales of the materials collected. This is the case in the charity sector, where unwanted clothes, shoes, books and other materials are donated by members of the general public. This paper considers a real-life logistics problem currently being faced by a leading UK charity: transporting donated goods from ‘bring-banks’ (containers typically located in supermarket car parks) to a regional depot, acting as a transfer station, for onward transportation to a central recycling facility. A heterogeneous vehicle fleet is used and the drivers must also undertake collections of unsold goods from the charity’s High Street stores, in addition to the collections from the bring-banks.

Of growing interest to the charity is whether remote monitoring technology can be used to allow individual bring-banks to report their fill-level status at periodic intervals in order to measure bank performance, identify theft from banks and optimise vehicle collection schedules. In this scenario, the banks could dictate the daily/weekly collection schedule, responding to the ever fluctuating and dynamic nature of charity donations. This could help the charity reduce its considerable transport footprint and allow both logistics and store managers to visualise donation rates across the network in a more dynamic way. The optimisation problem behind this concept is being studied as part of an EU 7th Framework project called Straightsol (*2*), where infra-red sensors are being installed in bring-banks to report fill levels (e.g. 50% full) every 12 hours via GSM. The ways in which the sensor outputs and the recommendations from the optimisation process can be best visualised by the drivers, store managers and logistics controllers in space and time is being addressed as part of a UK Research Council (RCUK) funded project called 6th Sense Transport (*3*).

PROBLEM DESCRIPTION AND RESEARCH CHALLENGE

The problem concerns a part of the logistics operation undertaken by a leading UK charity which has around 650 High Street stores selling new and used goods and employs a network of around 1300 ‘bring-banks’ (subsequently referred to as ‘banks’) across the UK to receive donated goods from the general public. The charity operates a complex reverse logistics process (subsequently referred to as ‘take-back’ system) across several separate vehicle fleets, servicing these stores and banks. This enables the charity to transport goods, primarily second hand books and textiles, from banks to stores or processing centres, and to move goods between its stores for resale. The logistics operation also feeds recyclate generated by stores back into recognised commercial recycling streams and provides the take-back of low-grade clothing to a central sorting facility for separation and onward processing. The part of this take-back system being considered here concerns the collection of unsold stock from 50 stores and donated goods from 37 banks, undertaken by the same vehicle fleet, to a regional depot near Milton Keynes, UK. The collection region (Figure 1) used lies just north of London and comprises an area of approximately 8,000km^2 with a heterogeneous fleet of four vehicles, one van and three trucks, mainly working Monday to Friday with occasional collections made on Saturdays, if required.

The problem is formulated as the determination of schedules and routes over a rolling planning period (e.g. the next 5 working days) for a heterogeneous fleet of vehicles visiting clothing banks and High Street stores, in order to maximize profit. A requirement is that the stores should be visited on their existing specified days of the week, and that current, store specific time window constraints (related to access restrictions) are adhered to. Banks may be visited at any time on any day. The vehicle routes are subject to EU driving and working time constraints, which require the driver to take breaks at regular intervals (45 minutes for every 4.5 hours of driving) and to return to the depot before the end of the daily working time limit. Although profit is the primary goal, a secondary objective is not to allow banks to become full before a scheduled collection because this can create untidy sites that may discourage potential donors who could be lost to a neighbouring competitor's bank resulting in lost revenue. Profit is defined here as the value of the goods collected minus the associated transport costs minus penalty costs associated with banks overfilling.

The weights of materials to be collected from stores were assumed to accumulate at fixed store-specific average rates, derived from collection data recorded over a period of one year. The weights of materials contained in banks were assumed to be known at certain points in time, based on the availability and frequency of reliable and accurate remote monitoring data. [Note: these data were simulated here, as the remote sensors were being installed and tested, at the time of writing.]

In this paper, a daily update of bank fill levels was modelled, on the assumption that the banks would be monitored towards the end of one day in order to devise schedules for the subsequent day, with rounds typically starting early in the morning (e.g. 4.30am). Further donations into banks could be made between the time that schedules are devised and the times when banks are visited. In practice, this issue could be treated in various ways when devising schedules:

(i) Assume an average additional amount for each bank, based on the average accumulation rate and an estimate of how 'old' the data will be by the time the bank is visited. A difficulty with this approach is that the time of visit is not known until the schedule is devised, so an iterative method may be required.

(ii) Dynamically update the vehicle routes during the day based on updated remote monitoring data.

(iii) Either ignore the possibility of additional donations or build in a small amount of slack (i.e. spare capacity) in the schedules to allow an assumed overall average amount of added donations to be collected. If need be, some goods could be left in banks if vehicles did not have sufficient spare capacity available. This approach is the simplest of the three and the most likely to be adopted in practice by the charity.

RELATED WORK

The 'waste collection problem' typically involves using a fixed fleet of collection vehicles operating from a single depot to undertake collections at the minimum possible cost. Balanced vehicle rounds are usually desired to provide a fair distribution of work amongst the collection crews. The waste collection problem is usually modelled as a vehicle routing problem (VRP), or as a variant thereof, and the VRP is usually fundamental to all logistics planning (*4*,5). The best-known variant of VRP is the capacitated vehicle routing problem (CVRP), where a homogeneous fleet of vehicles with capacity constraints perform tours starting and ending at a depot, distributing (or collecting) a single commodity to a set of customers, all of which must be visited. Many variants of the VRP have been studied, involving features such as a heterogeneous fleet, multiple trips, split delivery, time windows, distance constraints, and profit collection with one of the best known exact methods being developed by Baldacci and Mingozzi (*6*). There have been several examples of heuristic algorithms that have been specifically developed for and applied to waste collection problems (*7-11*).

The heterogeneous vehicle routing problem (HVRP) is an extension of the CVRP where the fleet is composed of vehicles with varying, as opposed to uniform, capacity restrictions and is significantly harder to solve compared to the CVRP. A number of mathematical models and solution algorithms have been proposed (*6, 12, 13)*. The literature on the HVRP is relatively scarce and this is more so in the case for the variant of the HVRP with time windows (HVRPTW). Studies by Paraskevopoulos et al. (*14*) and by Ceschia et al (*15*) both describe variable neighbourhood tabu search algorithms to solve the HVRPTW, the objectives of which are to minimize the total costs, comprising fixed costs of using the vehicles and the travel cost incurred by visiting all customers.

Remote monitoring of 3300 banks collecting cardboard waste in Sweden, with hourly updates of fill levels (*16*) showed that the remote monitoring data were rarely used by the waste contractors, due to either a lack of time and/or expertise in using them effectively for scheduling and routing purposes, or a reluctance to change working practices. The author estimated potential vehicle mileage savings of 26% and cost savings of 6%, for the city of Malmo through using dynamic schedules. A similar application was reported by Krikke et al (*17*) for the remote monitoring of 267 car dismantling sites in the Netherlands. Disassembled materials, both solids and liquids, were placed in containers and monitored using sensing equipment. As collection frequencies were relatively low (some materials only being collected annually), weekly sensor updates were used. The authors estimated vehicle mileage savings of 26% and cost savings of 19% through making more informed collection decisions. A more complex system whereby both containers and vehicles are equipped with remote monitoring sensors has also been proposed (*18*).

Various remote sensing technologies have been used in practice: Rovetta et al (*19*) reported the use of various types of sensor (ultra-sonic, LED, pressure) in domestic waste bins in Pudong, Shanghai, China; a specially designed electro-mechanical device for domestic waste collection was used in a province of Italy (*20*); pressure sensors tend to be used for measurement of liquids and gases (e.g. for fuel deliveries); while infra-red sensors have been employed in the application described in this paper. Practical difficulties associated with the use of remote monitoring equipment and data relate to: the ability of the sensor to accurately measure fill-level, bearing in mind that solid materials (e.g. bags of clothes) may not lie flat in the container; battery life of the sensor; operational methods and the flexibility for change; and lack of expertise in the use of such data. Initial evidence here suggests that although the technology providers claim sensor accuracy of +/- 5%, measurements can be +/- 20% due to the materials not lying flat in the container.

Dynamic scheduling is an inherently difficult problem, particularly when various operational constraints have to be considered. There is no ‘standard’ optimisation technique for this type of problem so different heuristics have been devised, often based on the concepts of containers that must be visited (e.g. >75% full), those that may be visited (e.g. > 50% full) and those that should not be visited yet because they are considered not to contain sufficient stock to warrant a visit (e.g. < 50% full). Faccio et al (*18*) used an ‘oversize risk’ parameter specifying an upper limit on the percentage of containers that are allowed to become full at any given time. They suggested that this parameter should be set close to zero, although the values they used in testing ranged from 5% to 15%. Dynamic routing, on the day of operation, may also be envisaged for this type of problem, although not considered here, with the remote monitoring data being used to dynamically reschedule collections during the day based on updated information received. Dynamic routing is more typically used in courier services, ambulance and dial-a-ride applications (*21-23*).

A similar type of problem to that considered here is the so-called ‘team orienteering problem’ (TOP), whereby the locations to be visited are chosen to maximise a stated objective function. Variations of the TOP have included consideration of time windows and of capacity constraints (*24*). Our problem is similar in concept to the TOP with time window and vehicle capacity constraints but is unique, as far as the authors are aware, in its consideration of a heterogeneous vehicle fleet, fixed collection days for shops and capacity constraints associated with the banks.

PROBLEM SPECIFICATION

The problem is specified here in terms of the objectives, constraints and the parameters used; a mathematical formulation is not given here but the interested reader is referred to a formulation given for the related ‘team orienteering problem’ (*24*), discussed above.

Objectives

Vehicle routes are required for a series of N consecutive working days (Monday to Friday), where N = 20 was considered here, to make collections from 37 bank sites and 50 High Street stores, subject to time window, vehicle capacity and driving constraints. Each vehicle may undertake only one route each day, starting and ending at a single depot, and the same heterogeneous vehicle fleet is available each day, although all vehicles need not necessarily be used on any particular day. Each location may only be visited by one vehicle on any given day. The objective is to maximise profit, defined as the value of goods collected minus the associated transport costs and minus any penalty costs imposed as a result of banks overfilling.

Constraints

Time Windows

Collections from stores have to be made on specified fixed days of the week; each store also has its own associated time window, the same each day, when a collection can be made. Collections from banks may be made at any time. A vehicle may wait at a location until the time window opens.

Vehicle capacity

A heterogeneous vehicle fleet of one van, with a carrying capacity of 1400kg and three trucks, each with a carrying capacity of 2500kg, were available to be used each day. Vehicle capacities were specified by payload and not by volume.

Driving constraints

All rounds were assumed to start at the same time (0430 hours) with a maximum round time of 11 hours normally specified although this was extended to 12 hours for some of the Monday rounds where 11 hours proved to be insufficient time to permit a feasible set of routes to be calculated due to the additional weights of materials which gathered over the weekend. Break periods were modelled in line with UK and EU regulations: a 45-minute break was required for every 270 minutes of accumulated driving time (excluding servicing time).

Parameters

Distances and travel times between sites

Driving distances and driving times between all sites (banks, stores and the vehicle depot) were initially obtained using commercial software; driving times were then calibrated with reference to recorded driver logs, as it was considered that the commercial software significantly underestimated travel times. Here, there were 88 sites (1 depot+37 banks+50 stores), giving rise to 3828 (= 88 x 87/2) pairs of postcodes for which times and distances were calculated.

Demand

The weight of materials (kg) available to be collected from a site on any given day is assumed to be known, or, at least, capable of being estimated well. In reality, exact weights will never be known, however, the effects of estimation errors were not considered here. It was assumed that all collections would be full (i.e. 100% of the demand) with the exception that for the last collection on any round a partial collection was permitted in the case that the vehicle would otherwise go over capacity. For stores, fixed daily accumulation rates were assumed based on average values recorded over a period of 10 months (April 1, 2011 to March 31, 2012); similar data were available for banks and were used to specify the average and standard deviations used to randomly generate bank donations. For both shops and banks, the amount accumulated between the Friday and Monday collections was assumed to be equal to two days' worth of materials.

Bank donations

The amount of potential donations, X kg, for a particular bank on a particular day, was determined as X = Y.Z, where Y is a Bernoulli variable with a P value representing the probability of someone making a donation to the bank on that day and Z being a Normal random variable, representing the amount donated. The P value is a model parameter; values of 0.2, 0.5 and 0.8 were used to test model sensitivity, with the lower value representing relatively infrequent donations and the higher value representing relatively frequent donations. The Z value was bounded at [0, bank capacity], with the mean equal to the average daily donation amount, excluding days where no donations are made, and the standard deviation estimated from collection data. To allow fair comparisons between different collection strategies, a linear congruential random number generator was used to ensure that the same seeds were generated. Ten different starting seed values were used to increase the sample size.

Lost donations

In reality, a potential donation (or part of it) may be lost if the bank is full. In this case, the person may decide to take their goods elsewhere or they may decide to leave bags outside the bank. The latter practice is discouraged, using notices on the banks, but can happen and could be modelled by assuming a slightly greater bank capacity. However, it was assumed here that any potential donations modelled to exceed the bank capacity would be considered as 'lost donations', not collected. Lost donations were output by the model to give an indication of lost profit due to allowing banks to become full before collection.

Servicing times

The times needed to collect materials from sites were assumed to be known, fixed values and were derived from driver logs recorded during one week in January 2012. For banks, these times ranged from 7 to 65 minutes, with the longer times being for sites with two or more banks; the range for stores was similar, from 7 to 67 minutes. It is recognised that, for robustness, considerably more collection time data should be obtained prior to any implementation in the field. Also, it seems likely that collection times will be positively correlated with weights collected. With this in mind, the model developed here allows collection times to be specified as the sum of a fixed part and a variable part, depending on the weight, that is:

$$collection\ time_i = a_i + b_i X_i$$

where i specifies the site, X_i the demand (kg) and a_i and b_i are constants; however, this feature has yet to be used due to the lack of sufficient collection time data.

Transport costs

Transport costs were estimated to be £1.50/mile (~$2.35 per mile) by the charity's transport manager, taking into account all vehicle fixed and variable costs including overheads.

Values of goods

The value of goods per unit weight was modelled to be site-specific, giving scope to vary values between sites according to the type of site (e.g. some stores may specialise in books, whereas others may focus on clothes) or to the geographical area (some areas may donate better quality goods than others). In the absence of such detailed valuations, all banks were assumed to generate the same value (£0.80/kg), while unsold goods from stores were valued lower (£0.50/kg) by the charity donation banks manager.

Bank capacity and overfilling penalty

Each bank site contains one, two or three containers. Most of the containers are a standard size (1000L) and are estimated to contain 270kg of clothes, when full; some other larger containers are also used. A site with more than one bank is modelled as a single entity here: for example, a site with three banks is modelled as a single bank with a capacity of 810kg. A penalty function is used to discourage banks overfilling and a financial penalty (£/ kg) is levied on any bank not visited whose demand is modelled to exceed a specified percentage fill level. The size of the penalty and the fill level at which it applies are both modelled as variable parameters.

Minimum percentage fill level to be collected

To inhibit collections from banks with low levels of fill, a minimum percentage fill level parameter is used.

SOLUTION APPROACH

An important consideration when planning vehicle routes for a series of days is the 'look-ahead period', that is, the number of days in advance for which one is producing routes. At present, the charity uses the same fixed schedules and routes from week to week, only adjusting them in response to changes in the collection network (e.g. a new bank) or when new vehicles with different carrying capacities are purchased. The fixed schedules see the vehicles visiting both banks and stores on set days of the week irrespective of how full the banks might be. The introduction of remote monitoring data has highlighted the fact that donations can be highly variable in terms of their amounts and how frequently they occur. With this in mind, the solution approach adopted here was the use of only a one-day look-ahead period, that is, calculating routes for the next working day based on maximising profit for the next day. It is recognised, however, that this approach is rather myopic, as the resulting model has no concept of being able to increase profit by delaying collections. To prevent collections taking place too soon, a minimum collection amount, expressed as a proportion of a full bank (e.g. 0.5 = half full) was used, with values of 0.5 and 0.7 being considered. To discourage collections from taking place too late, a high penalty value of £10/kg was applied to banks not visited and whose fill level was modelled to exceed penalty levels of 0.75 and 0.95, respectively. The assumed probability of a donation taking place on any particular day was also varied, with values of 0.2, 0.5 and 0.8 being tested; however, total donations were kept approximately the same between sets of runs, subject to some rounding errors due to an assumption of the mean donation rate being an integer variable.

The problem considered here is NP-hard (*25*) as it includes the CVRP as a special case, suggesting that an optimal method for the size of the considered problem will be difficult to obtain within short computation times. For this reason, a tabu search algorithm was developed to solve the problem. Tabu search is a metaheuristic algorithm used to solve difficult optimization problems which employs a tabu list to prevent the search from being trapped in local optima (*26*,*27*). The implementation here employs three local search operators that work on a given incumbent solution: customer addition, customer removal, and customer swap, where the 'customers' are the sites to be visited (i.e. banks and stores). The customer addition operator determines customers not yet visited in the incumbent given solution, the inclusion of which results in a maximal increase (or a minimal decrease) of the profit collected. The customer removal operator determines the customer already visited in the incumbent given solution, the removal of which results in a maximal increase (or a minimal decrease) of the profit collected. Finally, the customer swap operator determines two customers already visited in the incumbent given solution, where the swapping of the two results in a maximal increase (or a minimal decrease) of the profit collected. Details of the tabu search algorithm are:

- **Step 1 (Initialization):** Set the incumbent solution and the best known solution to the empty solution in which no customers are visited. Initialize the tabu list to an empty list. Set the iteration counter to 1. Set the tabu tenure of all customers to 0.

- **Step 2 (Stopping condition):** If the iteration counter is greater than a pre-specified iteration limit, stop and report the best known solution.

- **Step 3 (Local search):** Tentatively apply the customer addition, customer removal, and customer swap operators on the current solution, ignoring moves that involve customers in the tabu list as well as moves that result in infeasible solutions. Apply the operator that yields the best possible improvement and update the incumbent solution.

- **Step 4 (Best solution update):** If the incumbent solution is better than the best known solution, update the best known solution with the incumbent solution and set the iteration counter to 1. Else, increase the iteration counter by 1

- **Step 5 (Tabu list update):** Add the customer(s) in the selected operator to the tabu list and increase their tenure by 1. Remove customers with a tabu tenure that is greater than a pre-specified tenure limit. Go to Step 2.

The computational experiments conducted on a computer with an Intel Core i7 2.50 GHz CPU have not required more than 15 CPU seconds per daily set of routes generated.

RESULTS

Model results were obtained for a series of 20 consecutive working days (Monday to Friday) and for various combinations of parameter values; results shown in bold (Table 1) represent the base case for comparison. The numbers shown are totals over the 20 days that have been averaged over ten sets of runs using different random seeds for the generation of donations data. The following observations can be made:

- The numbers of visits to banks were drastically reduced with the use of remote monitoring data. This is especially the case where donations are highly sporadic (e.g. P=0.2) and where the minimum collection level is set quite high (e.g. 0.7).
- Despite the number of visits to banks being substantially reduced, profit gains were modest: up to 2.5%, in the case where (P, Minimum Level, Penalty Level) = (0.2, 0.5, 0.75). Small profit losses were observed in all three cases where a minimum collection value of 0.7 was used, suggesting that the lower value of 0.5 should be preferred. Further sets of runs considering other minimum collection levels are required.
- The desired profit gains did not materialise as driving distances were only reduced in four of the cases, between 1.2% and 2.4%, and increased by up to 4% in other cases. The reason why driving distance can increase despite bank visits being reduced is due to the limiting constraint of having to visit stores on their existing fixed days of the week, combined with the fact that bank visits are now constrained by their fill levels. The ‘myopic’ one-day look-ahead approach is also a factor here as there may be situations where it would be better to delay a collection, for example, to wait until the following day when a collection vehicle may be better placed to visit a particular bank. This phenomenon will be investigated in further research.
- Time taken was modelled to reduce in all cases by up to 7.6% due to the reduced numbers of bank visits; however, the results slightly overestimate the benefits here since fixed collection times were assumed rather than weight-dependent collection times.
- The amounts of materials modelled to be lost, due to banks filling up before the collection, were slightly reduced through the use of remote monitoring. For example, in the base cases, the ratio of materials lost to materials collected was (3.0%, 2.3%,1.5%) for P = (0.2, 0.5, 0.8), respectively, which reduced to (1.9%,1.0%,0.5%) respectively in the case where the minimum collection level of 0.5 and the penalty level was 0.75.
- The total number of vehicle rounds was modelled to reduce in all cases with up to 4.6 fewer rounds required, on average, during the modelled four-week period.
- With the minimum collection level set to 0.5, there was little difference in results between penalty levels of 0.75 and 0.95, suggesting that visits to banks over 75% full would have been scheduled in any case, without being forced to do so via the use of the penalty.

DISCUSSION AND CONCLUSIONS

A tabu search algorithm has been developed to solve a heterogeneous vehicle routing problem (HVRP) with time windows, working time restrictions and customer selection, relevant to remotely monitored collections. The objective was to maximize profit from collections made from bank sites and stores, which involves a trade-off between the value obtained from the goods collected and the transport costs associated with collecting them. A penalty function was also employed to reduce incidences of banks becoming full before collection.

The effectiveness of using remote monitoring data to influence vehicle schedules was shown to be dependent on the variability of demand, in this case equating to variability of donations made by the general public. The results showed that the adopted solution approach, which involved only a one-day look-ahead period, gave greatest improved performance compared to existing fixed schedules, the more variable the donation rates. If donation rates are fairly predictable then there is less need for remote monitoring at collection sites and the further ‘look-ahead’ one can contemplate when developing the vehicle schedules. The results here have been based on varying levels of simulated donations some of which may overestimate levels of variability; the real ‘proof of the pudding’ will come once the live remote monitoring data become available.

Further work required prior to implementation in the field includes: (i) more extensive testing of different combinations of parameter values; (ii) rigorous testing of the remote monitoring sensors and analysis of the data received in order to develop procedures for dealing with missing and/or suspect data due to equipment failure; (iii) developing a method to run the model with live bank fill rate data and a mechanism to enable the outputs of the model to be effectively visualised.

A further dimension to this work will be to treat the High Street stores in a similar way to banks, in that they have variable amounts of stock donated into them by members of the public, and therefore have similar, variable collection requirements. The current requirement to service stores on fixed days severely constrains performance and may not be appropriate if stores were to report the number of bags of stock that required collection at the end of each working day. The developed algorithm has been designed to allow this possibility to be readily modelled. Alternative solutions to this problem could be: (i) to solve a periodic vehicle routing problem for the stores only and once these are determined, incorporate the ‘fixed’ locations into the tabu search model before re-solving for each day. Such a strategy would ensure that stores are visited periodically whilst the dynamic aspect is still kept for the banks; (ii) to solve a dynamic periodic vehicle routing problem to take into account stores and banks simultaneously. Another interesting dimension not studied here is to increase the rate at which the banks are remotely monitored (e.g. hourly).

Of wider interest is the way in which the individual shop managers’ needs can be catered for. The charity is keen, wherever possible, for shops to ‘adopt’ banks, as the stock will then be sorted locally and greater value obtained, as opposed to the textiles being sold on as a bulk concern. To this end, there are interesting interactions at the local level that need to be considered in a practical application of the model where a shop manager might need an urgent top-up of stock from a local bank as soon as possible, regardless of the suggested collection day, and this would need to be fed into the system. Smartphone technology could provide this enhanced visibility and provide a link between the model outputs and the various parties implementing the strategies on the ground. This is being investigated as part of the 6th Sense Transport project *(3)* where drivers’ current and projected locations in relations to banks and stores can be viewed along with the bank yield and fill rate history (Figure 2). This could provide an interface for the HVRP algorithm and allow proposed new next-day or longer ‘look-ahead’ schedules to be viewed by the logistics scheduler, drivers and store managers.

REFERENCES

1. Smartbin. *Remote monitoring of bin fill level*. www.smartbin.com. Accessed July 26, 2012.
2. Straightsol. *Strategies and measures for smarter urban freight solutions.* http://www.straightsol.eu/index.htm. Accessed July 26, 2012.
3. 6th Sense Transport. www.sixthsensetransport.com. Accessed July 26, 2012.
4. Golden, B., S. Raghavan, and E. Wasil. *The vehicle routing problem: latest advances and new challenges.* Springer, New York, 2008.
5. Laporte, G. Fifty Years of Vehicle Routing. *Transportation Science,* Vol. 43, 2009, pp. 408-416.
6. Baldacci, R. and A. Mingozzi. A unified exact method for solving different classes of vehicle routing problems. *Mathematical Programming,* Vol. 120, 2009, pp. 347-380.

7. De Rosa, B., G. Improta, G. Ghiani, and R. Musmanno. The arc routeing and scheduling problem with transshipment. *Transportation Science*, Vol. 36, No. 3, 2002, pp. 301-313.
8. Viotti, P., A. Polettini, R. Porni, and C. Innocenti. Genetic algorithms as a promising tool for optimization of the MSW collection routes. *Waste Management and Research,* Vol. 21, No. 4, 2003, pp. 292-298.
9. Bautista, J. and J. Pereira. *Ant algorithms for urban waste collection routing.* vol. 3172, Springer, Berlin, 2004.
10. Kim, B.-I., S. Kim, and S. Sahoo. Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, Vol. 33, No. 12, 2006, pp. 3624-3642.
11. Nuortio, T., J. Kytöjoki, H. Niska, and O. Bräysy. Improved route planning and scheduling of waste collection and transport. *Expert Systems with Applications,* Vol. 30, No. 2, 2006, pp. 223-232.
12. Yaman, H. Formulations and valid inequalities for the Heterogeneous Vehicle Routing Problem. *Mathematical Programming,* Vol. 106, No. 2, 2006, pp. 365–390.
13. Brandão, J. A tabu search algorithm for the heterogeneous fixed fleet vehicle routing problem. *Computers & Operations Research,* Vol. 38, No. 1, 2011, pp. 140–151.
14. Paraskevopoulos, D., P. P. Repoussis, C. D. Tarantilis, G. Ioannou, and G. P. Prastacos. A reactive variable neighbourhood tabu search for the heterogeneous fleet vehicle routing problem with time windows. *Journal of Heuristics,* Vol. 14, No. 5, 2008, pp. 425–455.
15. Ceschia, S., L. Gaspero, and A. Schaerf. Tabu search techniques for the heterogeneous vehicle routing problem with time windows and carrier-dependent costs. *Journal of Scheduling,* Vol. 14, No. 6, 2011, pp. 601-615.
16. Johansson, O. M. The effect of dynamic scheduling and routing in a solid waste management system. *Waste Management,* Vol. 26, No. 8, 2006, pp. 875-885.
17. Krikke, H., I. le Blanc, M. van Krieken, and H. Fleuren. Low-frequency collection of materials disassembled from end-of-life vehicles: On the value of on-line monitoring in optimizing route planning *International Journal of Production Economics*, Vol. 111, No. 2, 2008, pp. 209-228.
18. Faccio, M., A. Persona, and G. Zanin. Waste collection multi objective model with real time traceability data. *Waste Management,* Vol. 31, No. 12, 2011, pp. 2391-2405.
19. Rovetta, A., F. Xiumin, F. Vicentini, Z. Minghua, A. Giusti, and H. Qichang. Early detection and evaluation of waste through sensorized containers for a collection monitoring application. *Waste Management*, Vol. 29, No. 12, 2009, pp. 2939-2949.
20. Waste Management World. *Bin Monitoring Allows for 'Smart' Refuse Collection*, http://www.waste-management-world.com/index/display/article-display.articles.waste-management-world.collection-transport.2011.12.Bin_Monitoring_Allows_for__Smart__Refuse_Collection.QP129867.dcmp=rss.page=1.html. Accessed 12 July 2012.
21. Angelelli, E., N. Bianchessi, R. Mansini, and M. G. Speranza. Short term strategies for a dynamic multi-period routing problem. *Transportation Research Part C: Emerging Technologies,* Vol. 17, No. 2, 2009, pp. 106-119.
22. Ichoua, S., M. Gendreau, and J. Potvin. Diversion Issues in Real-Time Vehicle Dispatching. *Transportation Science*, Vol. 34, No. 4, 2000, pp. 426-438.
23. Khouadjia, M. R., B. Sarasola, E. Alba, L. Jourdan, and E. Talbi. A comparative study between dynamic adapted PSO and VNS for the vehicle routing problem with dynamic requests. *Applied Soft Computing*, Vol. 12, 2012, pp. 1426-1439.
24. Souffriau, W., Vansteenwegen, P., Vanden Berghe, G. and D. Van Oudheusden. The Multiconstraint Team Orienteering Problem with Multiple Time Windows. *Transportation Science,* published online before print, October 5, 2011, doi: 10.1287/trsc.1110.0377
25. Garey, M.R. and D.S. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. W.H. Freeman, New York, 1979.

26. Glover, F. Tabu Search - Part 1. *ORSA Journal on Computing,* Vol. 1, No. 2, 1989, pp. 190–206.
27. Glover, F. Tabu Search - Part 2. *ORSA Journal on Computing* Vol. 2, No. 1, 1990, pp. 4–32.

LIST OF FIGURES AND TABLES

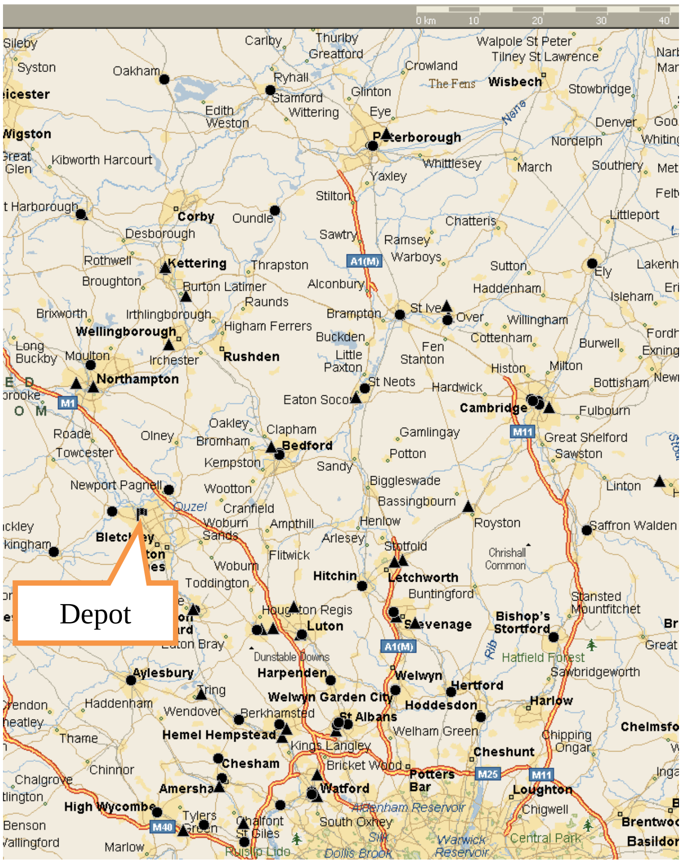

Key: ▲ Banks; ● Stores

FIGURE 2 Case study area showing depot, banks and stores.

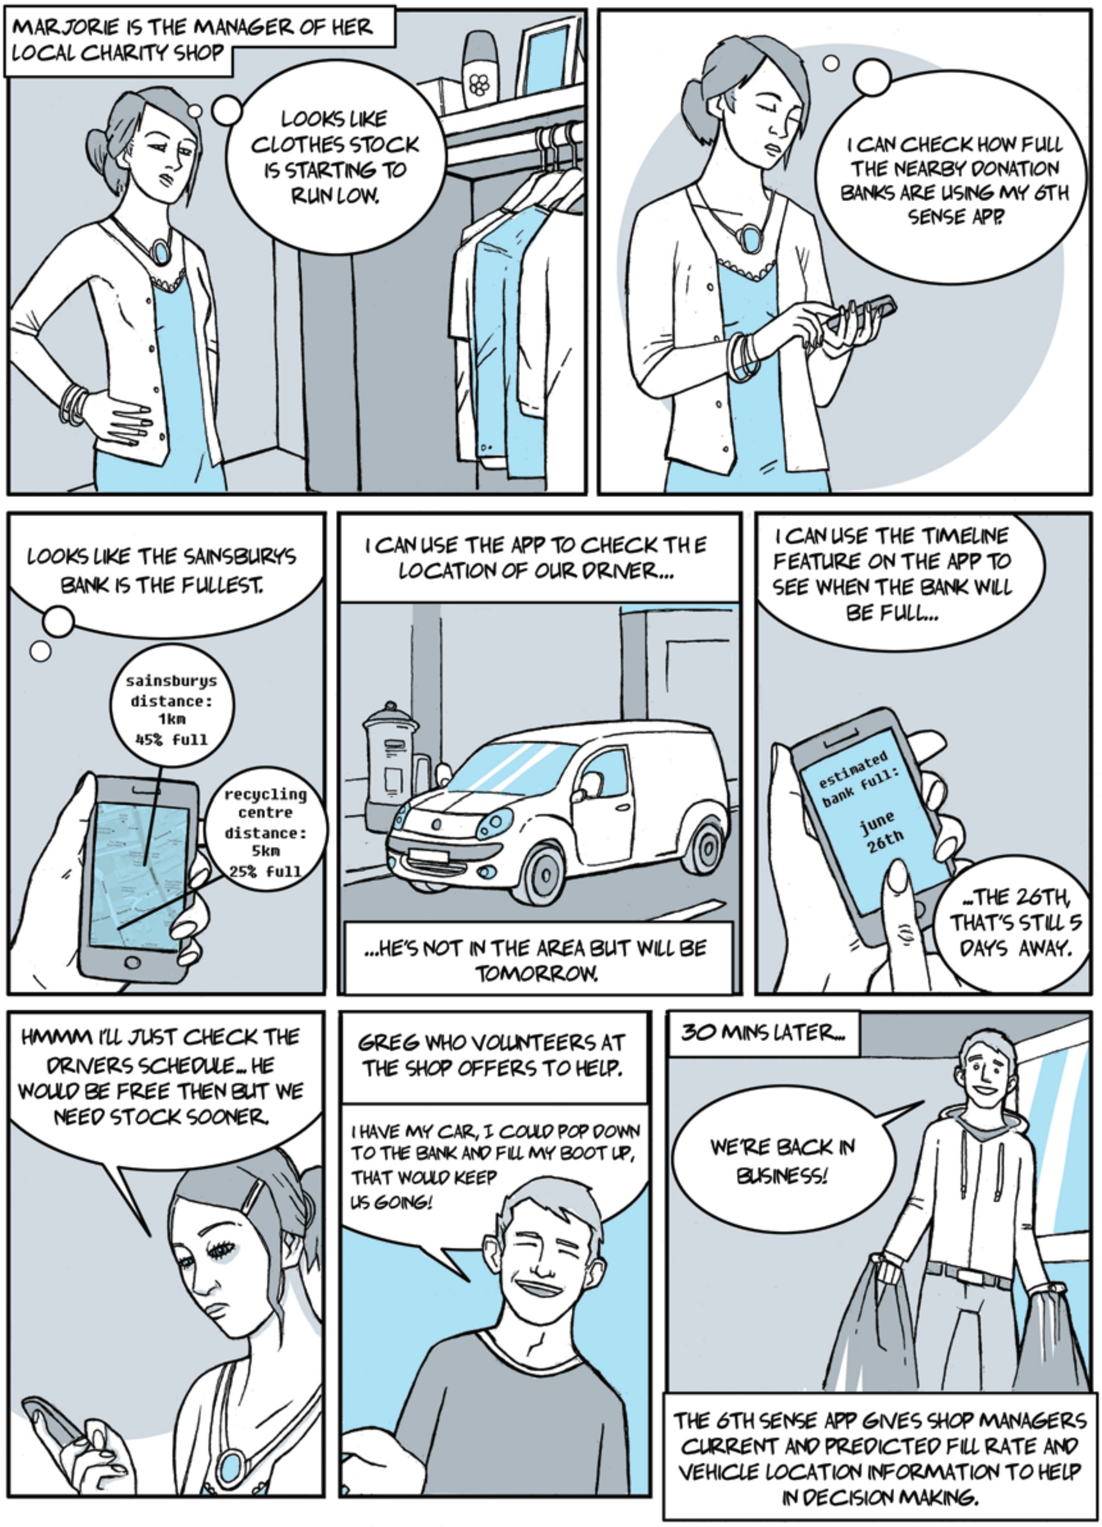

FIGURE 2 Prototype app showing driver locations and bank fill levels to aid decision making.

TABLE 1 Model Results

P	Min. Level	Penalty Level	Rounds (#)	Distance (km)	Time (hours)	Weight collected (kg)	Amount lost (kg)	Profit ($)	Bank Visits
0.2	**-**	**-**	**76.0**	**13651**	**737**	**127340**	**3813**	**88606**	**240**
0.2	0.5	0.75	72.3	13434	691	128954	2463	90813	86
0.2	0.5	0.95	72.6	13493	692	128929	2455	90710	86
0.2	0.7	0.95	71.4	13329	681	126451	3848	87964	69
0.5	**-**	**-**	**76.0**	**13770**	**743**	**133793**	**3104**	**96204**	**240**
0.5	0.5	0.75	74.0	14159	716	135162	1384	97332	124
0.5	0.5	0.95	73.7	14142	719	135113	1420	97275	124
0.5	0.7	0.95	73.0	13541	695	132945	2745	95457	94
0.8	**-**	**-**	**76.0**	**13854**	**743**	**135394**	**1997**	**98132**	**240**
0.8	0.5	0.75	74.5	14426	731	136351	704	98409	137
0.8	0.5	0.95	74.4	14369	729	136388	754	98486	136
0.8	0.7	0.95	73.4	13861	708	134113	1667	96501	102

Key to TABLE 1: P = probability of a donation(s) on any particular day; results in bold represent the current situation using fixed schedules.