

Improving collection efficiency through remote monitoring of charity assets

Fraser McLeod^a (*corresponding author*), Gunes Erdogan^b, Tom Cherrett^a, Tolga Bektas^b,
Nigel Davies^c, Duncan Shingleton^d, Chris Speed^d, Janet Dickinson^e, Sarah Norgate^f

^aTransportation Research Group, Engineering and the Environment, University of Southampton, Southampton.
SO17 1BJ. United Kingdom. F.N.McLeod@soton.ac.uk (Tel. 023 8059 3316) and T.J.Cherrett@soton.ac.uk

^bSouthampton Management School, University of Southampton, Southampton, SO17 1BJ, United Kingdom.
G.Erdogan@soton.ac.uk and T.Bektas@soton.ac.uk

^cComputing Department, University of Lancaster, LA1 4YR, United Kingdom. nigel@comp.lancs.ac.uk

^dEdinburgh College of Art, The University of Edinburgh, Edinburgh, EH3 9DF, United Kingdom.
C.Speed@ed.ac.uk and Duncan@shingleton.org

^eSchool of Services Management, Bournemouth University, Bournemouth, BH12 5BB, United Kingdom.
JDickinson@bournemouth.ac.uk

^fDirectorate of Psychology & Public Health, School of Health Sciences, University of Salford, Salford, M6
6PU, United Kingdom. S.H.Norgate@salford.ac.uk

ABSTRACT

Collection costs associated with servicing a major UK charity's donation banks and collecting unsold goods from their retail shops can account for up to 20% of the overall income gained. Bank and shop collections are commingled and are typically made on fixed days of the week irrespective of the amounts of materials waiting to be collected. . Using collection records from a major UK charity, this paper considers what vehicle routing and scheduling benefits could accrue if bank and shop servicing requirements were monitored, the former using remote sensing technology to allow more proactive collection scheduling. A vehicle routing and scheduling algorithm employing tabu search methods was developed, and suggested time and distance savings of up to 30% over the current fixed schedules when a minimum bank and shop fill level of between 50% and 60% was used as a collection trigger. For the case study investigated, this led to a potential revenue gain of 5% for the charity and estimated CO₂ savings of around 0.5 tonnes per week across the fleet of six heterogeneous vehicles.

KEYWORDS: Vehicle routing; remote monitoring; waste collection

1. INTRODUCTION

Donation banks (i.e. receptacles for receiving donated goods) (Figure 1) are often provided by charities at supermarkets, recycling centres and public car parks, as a convenient way for the general public to donate unwanted clothes, shoes, books and other materials. Alternatively, members of the public may prefer to deliver their unwanted goods directly to a charity shop. In both cases, the quality of the donated goods can vary widely, with some goods having no value (waste), other goods having some recycling value and the best quality goods being suitable for re-use, including items of high resale value such as wedding dresses and rare books. A 2009 survey of UK charity shops indicated that over 363,000 tonnes of textiles are sent on for reuse and recycling by charity shops every year and that around 2% of materials donated to charity shops end up in landfill (Charity Retail Association, 2013). The collection requirements associated with emptying donation banks and removing unsold/unsaleable items from shops can be considerable. This paper considers how a greater temporal visibility of both bank and shop fill rates, the former using remote monitoring sensors, could aid a major UK charity to more dynamically manage its collection schedules in order to reduce the CO₂ footprint, save time and increase revenue.



Figure 1 – Charity donation bank

A perceived problem with the fixed collection rounds often used in charity logistics, and a problem experienced in this case study, is that some banks may contain relatively few items when visited, while other banks may have been full for several days before a collection was made. In this case a more dynamic and flexible collection scheduling approach may be appropriate particularly where containers fill at highly variable and unpredictable rates, or, where unnecessarily long trips to only partially filled banks are being made (Johansson, 2006). This more dynamic and flexible approach can be enabled through remote monitoring of banks and other assets, Remote monitoring may also help in identifying incidences of theft from banks and could, in principle, affect the collection policy (e.g. visit more frequently), although this aspect has not been considered here.

Many charities also offer ad hoc services such as house clearances where a vehicle will be scheduled to visit a domestic dwelling or workplace to collect a quantity of items (books, clothes etc.). Such activity can be quite lucrative in terms of the potential quality of goods available and more informed logistics decisions could be made where remote monitoring identifies the relative worth of visiting certain banks in time over a request for a house clearance. The scale of the benefits depends on the flexibility in the collection schedule which in the case of some charities is limited by the need to service shops on a regular fixed interval basis, an activity undertaken in tandem with bank servicing. Aside from the scheduling benefits, remote monitoring of banks can also help charities quantify their performance in terms of average yields by area and time of year. This can also aid bank

placement strategy, which may need to be more temporal dependent on yields, theft rates and the changing characteristics of the local population.

The UK charity studied in this research was interested in understanding how the adoption of new remote fill rate monitoring sensors fitted into its network of textile donation banks could help optimise its collection strategy on a day-to-day basis. This paper uses real and simulated donation data to quantify the impacts of various collection strategies assuming different minimum fill level collection triggers and maximum fill level penalties. In reality, estimating bank fill level can be difficult given that bags of clothes may not lie flat in the bank or may shift as new donations are added. In this application, an infra-red sensor was mounted on the underside of the roof of each bank, measuring the distance to the top of the pile of materials lying below, translated into a percentage fill estimate for the bank. The fill level was transmitted via GSM (Global System for Mobile Communications) to a webserver every 12 hours. Initial indications from a trial of the sensors suggested that the fill estimates were accurate to within 20%. Another practical concern is the battery life of the sensor and the need to replace or recharge batteries. Other remote sensing technologies are potentially suitable for waste and recycling applications: for example, Rovetta et al (2009) reported the use of ultra-sonic, light emitting diode (LED) and pressure sensors installed in waste containers in Shanghai, China

1.1 Background to the routing problem

Devising a service schedule where a collection vehicle has to visit specific points on fixed days and others, depending on dynamic fill rate information, is an extension of the well-studied capacitated vehicle routing problem (CVRP), with a heterogeneous vehicle fleet (HVRP) and time window constraints (HVRPTW). This problem is considerably harder to solve compared to the CVRP and a number of mathematical models and solution algorithms have been proposed (Yaman, 2006; Baldacci and Mingozzi, 2009; Brandão, 2011). The literature on the HVRP is relatively scarce and this is more so where time windows are included (HVRPTW). Studies by Paraskevopoulos et al (2008) and by Ceschia et al (2011) both describe variable neighbourhood tabu search algorithms to solve the HVRPTW, the objectives of which are to minimize total cost. Vehicle routing and scheduling algorithms have been applied to a wide variety of waste collection problems: for example, Kim et al (2006) considered time windows associated with commercial waste collections; Angelelli and Speranza (2002) compared alternative collection methods, with varying types of containers and vehicles used; and Amponsah

and Salhi (2004) focused on waste collection problems faced in developing countries. Practical applications of remote monitoring in the waste/recycling sector have included cardboard collection in Sweden (Johansson, 2006) and the collection of disassembled materials from cars (Krikke et al, 2008) where, in both studies, the authors estimated potential vehicle mileage savings of 26% through dynamic scheduling.

2. DEFINING THE CASE STUDY

The case study concerns a subset of the logistics operation undertaken by a leading UK charity operating a network of around 650 shops selling new and used goods, and approximately 1300 donation banks across the UK. The charity operates a complex reverse logistics process across several separate vehicle fleets, servicing these shops and banks. This enables the charity to transport goods, primarily second-hand books and textiles, from banks to shops or processing centres, and to move goods between its shops for resale. The logistics operation also feeds recycle generated by shops back into recognised commercial recycling streams and provides the return of low-grade clothing to a central sorting facility for separation and onward processing. The part of this reverse logistics system being considered here concerns the collection of rejected or unsold goods from 75 shops and donated goods from 58 textile bank sites, undertaken by the same vehicle fleet, to a regional depot near Milton Keynes, UK (Figure 2). The collection region covers an area of approximately 11,000km², including the towns/cities of Oxford, Cambridge, Northampton and Peterborough.

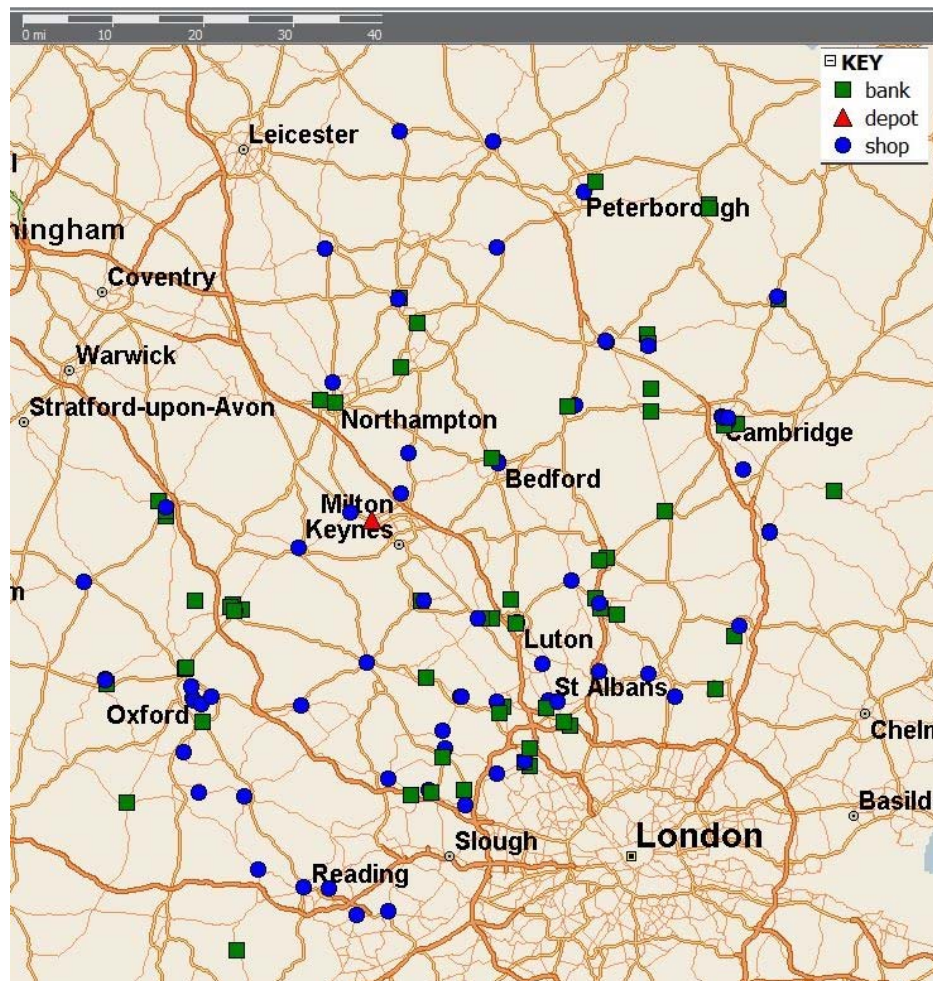


FIGURE 2 Case study area showing depot, banks and shops

Previous research (McLeod et al, 2013) found that the charity's existing requirement to maintain fixed collection days for the shops severely constrained the vehicle schedules and that remote monitoring of banks under such a scenario returned small benefits in overall round time and distance savings due to the primary requirement to service shops. Of increasing interest is whether the traditional fixed-interval shop collection schedule is too onerous and whether a more dynamic approach could be adopted. This concept would work in a similar way to the remote bank monitoring, with shop managers reporting their actual collection requirements each day (or potentially via remote monitoring the stock room) to better understand actual servicing needs. In this research, the collection requirements for the shops are relaxed by treating them in a similar way to banks, only visiting them as and when required, based on how near capacity they are. This temporal relaxation provides much greater flexibility in when both shops and banks can be visited and should improve logistical performance. Since shops can accumulate rejected/unsold goods at varying rates, it is envisaged that, in practice, shops managers would report their collection requirements on a daily basis, or as and when they would like a

collection to be made. A smart phone application has been developed to provide the asset fill level visibility and facilitate communications between the shops and the transport manager as part of this research (www.sixthsensetransport.com). In terms of the practical operation, the charity's transport manager would produce a vehicle schedule for the next working day of operation based on the latest fill readings from the bank sensors and shops, aided by a vehicle scheduling algorithm, specifically developed for this application and described in this paper.

The problem is formulated as the determination of vehicle tours for the given capacitated heterogeneous vehicle fleet for the next day of operation over a series of consecutive working days (20 days considered in our computational study). The vehicles may not return to the depot until the end of the working day and not all vehicles need be used every day. The weights of materials available to be collected from shops and bank sites on any given day were assumed to be known. Shops and banks may be visited on any day (i.e. there are no pre-specified collection days) and each site may only be visited by one vehicle on any given day. All sites should ideally be visited before their storage capacity is reached and, as this is a soft constraint, a penalty function is appropriate to avoid, although not necessarily prevent, instances of banks or shops becoming full. The objective is to maximise revenue, calculated as the estimated value of goods collected (£ per kg), where this value can vary between sites, minus the estimated collection costs (£ per mile, dependent on vehicle type). In reality, such an approach also reduces collection time which equates to cost savings as remote bank sites might not need visiting so frequently if a reasonable collection demand is not identified via the remote monitoring. This saved time could be utilised by the fleet manager in other ways, potentially to undertake specialist servicing.

3. METHODOLOGY

The developed solution was based on a 'one-day look-ahead' approach, where only the revenue for the next day of operation was considered at any point in time. Vehicle rounds were calculated on a day-by-day basis over a series of N consecutive working days, where $N=20$ was used, representing a period of approximately one month of operation. The one-day look-ahead is somewhat short-sighted but may be justified where the demands for collections from shops and banks are highly variable from day to day, making it difficult to predict demand accurately over a longer look-ahead period where the impact of season can also be significant (McLeod et al, 2011). Observations from remote monitoring sensors installed in the banks indicated substantial day-to-day

variability in donation rates at some sites in the case study area. It was hoped that the remote monitoring data from banks would have been available for use in this case study, however, significant reliability issues meant that simulated values were used in this analysis.

The solution algorithm was based on tabu search techniques (Glover, 1989; Glover, 1990) as the problem being considered is essentially NP-hard (Garey & Johnson, 1979) since it includes the CVRP as a special case, suggesting that an optimal method for the size of the considered problem would be difficult to obtain within short computation times. 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. 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 shops). The customer addition operator determines the customer not yet visited in the incumbent solution, the inclusion of which results in a maximal increase (or a minimal decrease) of the revenue collected. The customer removal operator determines the customer already visited in the incumbent solution, the removal of which results in a maximal increase (or a minimal decrease) of the revenue collected. Finally, the customer swap operator determines two customers already visited in the incumbent solution, where the swapping of the two results in a maximal decrease (or a minimal increase) of the distance travelled (McLeod et al., 2013). The algorithm design is summarised by the high-level pseudo-code given in section 3.1, while the details of the algorithm input and output parameters are given in descriptive terms in section 3.2.

3.1 Algorithm pseudo-code

The design of the implemented tabu search algorithm is now described, in pseudo-code format, including a summary of the main parameters, operators and functions used.

Parameters:

k_{max} : the limit on the number of iterations without improving the best solution

t_{max} : the limit on the number of iterations a customer (=bank or shop) can stay in the tabu list

Sets:

S : incumbent solution, set of sequences of customers for each vehicle

B : best known solution, chosen among S , encountered during the search

T : tabu list, set of two tuples (vertex, tenure)

Operators:

Customer_Addition(S, T) : Returns the best feasible solution after adding a vertex to S that is not in T

Customer_Removal(S, T) : Returns the best feasible solution after removing a vertex from S that is not in T

Customer_Swap(S, T): Returns the best feasible solution after swapping two vertices in S that are not in T

Functions:

$z(S)$: returns the revenue obtained from solution S i.e. $z(S) = \sum_{i \in C(S)} p_i - \sum_{(i,j) \in A(S)} c_{ij}$, where $C(S)$ is the set of customers visited in S , p_i is the value of goods collected from customer i , $A(S)$ is the set of trips between customers in S , and c_{ij} is the cost of travelling from customer i to j .

Tabu Search algorithm

```

 $S = \emptyset ; B = \emptyset ; T = \emptyset$  // Initialization
Set  $k = 1$  // iteration counter
While  $k < k_{max}$ 
    // operator selection
     $o^* = \operatorname{argmax}_{o \in \{Customer\_Addition, Customer\_Removal, Customer\_Swap\}} \{z(o(S, T))\}$ 
    If  $z(o^*(S, T)) > z(B)$ 
        Set  $B = o^*(S, T)$  // best solution update
        Set  $k = 1$  // reset the iteration counter
     $T = T \cup \{(o^*(S, T) \setminus S, 0)\}$  // tabu list update
     $S = o^*(S, T)$  // apply the operator
    For all  $(j, t) \in T$ , set  $T = (T \setminus \{(j, t)\}) \cup \{(j, t + 1)\}$  // tenure increase
    For all  $(j, t) \in T$ , if  $t > t_{max}$ , set  $T = T \setminus \{(j, t)\}$  // tenure limit
     $k = k + 1$  // increment the counter
End While

```

3.2 Algorithm inputs and outputs

The algorithm runs were characterised by the following key input and output data:

Distances and travel times between sites: Driving distances and driving times between all 8911 (= 134 x 133/2) pairs of the 134 sites (58 bank sites, 75 shops and the vehicle depot) were obtained using commercial vehicle routing software and calibrated against recorded driver logs. The driving times between sites corresponded to

off-peak travelling conditions and the algorithm did not take traffic congestion into account. In practice, it is known that drivers avoid visiting the busiest towns and cities (e.g. Cambridge and Oxford) during the morning peak period; this was considered through the use of time window restrictions to ensure that these locations were visited early.

Time windows: Some of the shops may only be visited at certain times of day relating to whether access during shop opening hours is required, prohibited or not desired. All time-of-day restrictions on shop visits were provided by the charity and used as input to the algorithm. The algorithm permitted a vehicle to wait outside a shop until the specified time window opened, although, in practice, this option was rarely used due to the flexibility available in scheduling the rounds. It was assumed that banks could be visited at any time of day. Opening and closing times were also specified for the vehicle depot (4.30am and 5pm), respectively, to ensure that all vehicle rounds were no longer than 12.5 hours in duration. In practice, the majority of drivers start their rounds at 4.30am to try to avoid peak traffic as far as possible.

Demand: The weight of materials available to be collected from a site on any given day (referred to as the 'demand') was assumed to be known, or, at least, capable of being accurately estimated from the available bank and shop collection records. The daily total weights of textile donations made into banks and of unsold/rejected goods set aside for collection by shop managers were calculated as the product of a Bernoulli random variable and a Normal random variable, with the former representing whether any amount had accumulated (0=No, 1=Yes) and the latter, the amount, based on mean and standard deviations derived from the charity's collection data obtained for a period of one year. If the sum of the donation amount and amount of textiles already at the bank site exceeded the site capacity, then the excess amount was considered to be lost (e.g. if a bank was full, the person donating would take the excess amount elsewhere). For banks, the assumed probability value for the Bernoulli variable was 0.8, (i.e. on average, on 20% of days it would be assumed that no donations were made to the bank). This was a 'best guess' assumption based on known trends and would need to be revisited in future research when sufficient remote monitoring data from banks becomes available. For shops, it was assumed that rejected and unsold goods would accumulate on all days.

To avoid 'end effects', the amounts of materials waiting to be collected on the first day were assumed to be equal to the average daily accumulation rate for the site, multiplied by the number of days since the previous

scheduled collection day (from the existing schedules); similarly, the amounts of materials that were left uncollected at the end of the 20-day period were considered when evaluating performance. Amounts accumulated between the Friday and Monday collections were assumed to be equal to two days' worth of materials. Ten different starting seed values for the random sampling were used for each set of runs to improve the robustness of the results, with the randomly sampled values being identical across the different collection scenarios to prevent the introduction of any random bias.

Collection times: The amount of time (minutes) needed to collect materials from sites was assumed to be a linear function of the weight collected: $collection\ time_i = a_i + b_i X_i$, where i specifies the site, X_i the demand (kg) and a_i (mins) and b_i (mins/kg) are constant site-specific parameter values. It was assumed that all collections would be made in their entirety 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. In previous research (McLeod et al, 2013), fixed site-specific collection times (i.e. $b_i = 0$) were assumed. Here, detailed collection round data comprising weights collected and times taken at over 4,000 site visits during a four-month period (August to November, 2012), routinely recorded by drivers in the course of their work, were used to set the parameter values. Although the algorithm was designed for site-specific parameter values, a single common formula for all banks (1), derived from linear regression, and similarly for all shops (2) was used due to some small sample sizes recorded at certain sites. The associated R^2 values for the linear regressions (1) and (2) were 0.47 and 0.31, respectively, indicating unexplained variation in collection times which, particularly for the shops, could be due to other factors yet to be explored such as likelihood of finding a nearby parking space, whether stairs have to be negotiated in shops and whether the driver is offered a cup of coffee.

$$bank\ collection\ time\ (mins) = 0.05 \times weight\ collected\ (kg) + 7 \quad (1)$$

$$shop\ collection\ time\ (mins) = 0.05 \times weight\ collected\ (kg) + 12 \quad (2)$$

Vehicle fleet operating constraints and costs: The algorithm allows any vehicle fleet to be specified in terms of the number of vehicles available of each type, where vehicle type is specified by:

- (i) carrying capacity (kg)
- (ii) driving time limit (hours)
- (iii) working time limit (hours), where this includes driving and loading time but excludes any breaks

- (iv) cost per unit distance (£/km), taking into account all vehicle and crew fixed and variable costs (e.g. labour, fuel, vehicle maintenance, insurance, depreciation) .

Two vehicle types were used here: a transit van (gross vehicle weight = 3.5t), with a carrying capacity of 1400kg; and a lorry (gross vehicle weight = 12t), with a carrying capacity of 6000kg. One transit van and five lorries were available for use, all operating Monday to Friday only. As vans and lorries are subject to different regulations, the driving time constraints were specified as 10 hours for the van and 9 hours for the lorries with drivers also required, by law, to take breaks at regular intervals: a minimum 45-minute break being taken for every 4.5 hours of accumulated driving. Working time constraints were specified as 11 hours for the van and 12.5 hours for the lorries. Transport costs were assumed to be £0.93/km (=£1.50/mile) for the van and £1.09/km (=£1.75/mile) for the five larger vehicles, based on estimated figures provided by the charity. It should be noted that the algorithm does not explicitly consider any costs associated with working time, as these are subsumed in the mileage costs, nor are any cost savings considered due to the possibility of using a reduced size of vehicle fleet and crew on some days.

Values of goods: The value of goods per unit weight collected was assumed to be site-specific, allowing scope to vary values between sites according to the type of site (e.g. some shops 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 rejected goods from shops were estimated by the charity to have a lower value (£0.50/kg).

Site capacities and overfilling penalties: Each bank site contained between one and four containers. Most of the containers were a standard size (1000L) and were estimated to contain 270kg of textiles when full; some other larger containers were also used. A site with more than one bank was considered as a single entity: e.g. a site with three standard banks was specified in the algorithm's input as single bank with a capacity of 810kg. The capacity of each shop (kg) was assumed to be equal to the maximum weight of goods that had been collected from the shop during a four-month period (August to November, 2012), as the charity's transport manager stated that most shops had little spare storage capacity available by the time of the existing scheduled collections.

Relatively large overfilling penalty values of £5/kg and £10/kg for shops and banks, respectively, were used in the algorithm to strongly encourage collection before a site reached capacity. These were levied on any site not visited whose demand exceeded a specified percentage fill level (50%, 75% and 95% fill levels were compared in this paper).

Minimum collection: A one-day look-ahead approach, as adopted here, means that the resulting algorithm has no concept of being able to increase revenue by delaying collections. To counteract this, a minimum collection parameter, specifying the minimum percentage fill level to be considered for collection, was used to prevent visits to sites containing relatively small proportions of goods, relative to the capacity of the site. This approach is in common with similar waste collection applications reported by Johansson (2006) and Krikke et al (2008). This parameter was specified as a site-specific variable, although, in the results presented, variation between sites has not been considered. The algorithm runs tested the sensitivity of the results to changing the minimum collection parameter value across all sites simultaneously.

4. RESULTS AND DISCUSSION

The two main algorithm parameters, affecting how collections are scheduled, are the minimum permitted collection, expressed as a proportion of a full bank, and the maximum fill level, beyond which penalties apply if a collection is not scheduled. The sensitivity of the results to these parameter values is discussed.

4.1 Effect of the minimum collection parameter value

The minimum collection (MC) parameter is the key input here, as it affects when collections can be made. MC values of 25%, 40%, 50%, 60% and 75% were considered alongside a penalty threshold value of 75%. In the latter case, this meant that a collection could not be made until the site was at least 75% full but once this level was reached, it then became an urgent service priority. Results for these five cases are shown (Table 1) alongside the existing base case situation where all collection days for banks and shops were fixed to those used in practice by the charity, with percentage savings illustrated in Figure 3. In each case, the results were totalled over a period of 20 working days and were averaged over the ten sets of runs made for each scenario. Performance was assessed in terms of distance travelled, time taken, goods lifted and remaining at sites for later collection, numbers of site visits and estimated revenue. Two values of estimated revenue are shown: the first

based on the materials actually collected and the second, including a consideration of the materials left over for collection until the next day of operation (i.e. beyond the time period being considered). This future revenue from the leftover materials was estimated to be 80% of the estimated value of the goods, on the basis that collection costs were up to 20% of the value of the goods, depending on the efficiency of the collection regime.

TABLE 1 Influence of minimum collection parameter value

MC	Rounds (#)	Distance (km)	Time (hrs)	Goods (tonnes)		Revenue (£K)		#Site visits	
				Collected	Left over	Collected	+Left over	Banks	Shops
Existing	99	22953	923	225.2	26.8	107.3	120.0	444	656
25%	96	25303	964	231.7	25.0	108.9	120.8	379	752
40%	80	20805	791	226.2	30.3	110.2	124.6	262	541
50%	72	18834	718	222.8	33.2	110.0	126.0	216	450
60%	69	17867	672	218.9	36.0	108.6	126.0	183	387
75%	70	17831	668	216.8	37.1	107.6	125.4	182	384

Note: Existing collection strategy is fixed and not based on minimum collection (MC)




Figure 3 - Relative savings (compared to existing situation) for different minimum collection (MC) values

Comparing the existing situation with the results for the dynamic scheduling (Table 1 and Figure 3) indicated the following:

- *Number of vehicle rounds* – the total number of vehicle rounds required during the 4-week period reduced in all cases, with the greatest reduction (from 99 to 69 rounds) observed when MC=60%. In this case, on some days, as few as two vehicles were utilised although, on others, typically Mondays, all six vehicles were required. Although not included in the revenue estimates here, a 30% reduction in the total number of vehicle rounds required would significantly reduce vehicle operating costs, or could release ‘spare’ vehicles and staff to undertake other, potentially lucrative activities. It should be remembered however that practical difficulties could result in managing such highly dynamic workloads, where the true number of vehicles required is not known much in advance, if the fleet is stripped down to a bare minimum, allowing no spare capacity..
- *Distance travelled* –reduced by up to 22%, as the minimum collection parameter value was increased, but increased by 10% where the minimum collection parameter was set at 25%. The latter corresponded to a 15% increase in the number of shop visits (from 656 to 752) combined with a 15% reduction in bank visits (from 444 to 379). The reduction in distance travelled translated into CO₂ emissions savings of around 0.5 tonnes per week, assuming an average emissions factor of 400g/km, based on the vehicle types used and average emissions factors derived by den Boer et al (2011). [Note: the emissions estimate here was simplistic and did not consider engine performance, driving cycles and other details that would be needed to gain a more accurate estimate.]
- *Time taken* –reduced by up to 28% but increased by 4.4% where the minimum collection parameter was 25%, suggesting that this parameter setting is too low. These savings were a combination of driving time savings (up to 24%), corresponding to the reduced distances travelled, and collection time savings (of up to 32%) due to fewer site visits being made.
- *Goods collected and left over* – as the minimum collection parameter increased, the amount of goods collected over the 4-week period tended to reduce, as more goods were left over for collection on the next working day. It should be noted that ‘goods collected’ plus ‘goods left over’ are not identical between the sets of runs due to the possibility of lost donations resulting from sites reaching capacity before the collection is made. Here, in the existing situation, some 4.7 tonnes of goods were estimated to be lost due to the fixed collections being more geared towards shop servicing requirements rather than bank servicing needs. The lost donations for the dynamic schedules were reduced to 0.12t, 0.30t, 0.77t, 1.9t and 2.8t for the MC values of 25%, 40%, 50%, 60% and 75%, respectively, with the extra losses being due to the move towards ‘just-in-time’ collections.

- *Revenue* – Taking the value of left over goods into consideration, revenue gains were highest (5%) when MC = 50% or 60%. The revenue dropped slightly (to 4.6%) when MC = 75% was used due to the assumed lost donations. These savings may initially appear to be fairly modest but do not include possible further revenue gains that could result from reducing the vehicle and crew requirements or utilising the time gained to undertake further profitable activities such as house clearances.
- *Site visits* – In all cases, visits to bank sites were substantially reduced in number and similarly for shops, except in the case MCP = 25%, suggesting that many sites are currently being visited too frequently, although it could be argued that such frequent collections are needed at some bank sites to avoid the chances of valuable materials being stolen from banks.

4.2 Effect of the penalty fill level parameter value

Whereas the minimum collection parameter prevents collections being made too early, the penalty fill level parameter acts as a ‘hurry-up’ call to try to prevent collections being made too late. Penalty fill levels of 50%, 75% and 95% were considered, in association with a minimum collection of 50%, to test the sensitivity of the results to this parameter. The first of these values (50%) meant that any site reporting a 50% fill became an immediate priority for collection, whereas the 95% value meant that sites did not become urgent until later. It can be seen from the results (Table 2) that the three sets of runs recorded very similar performance, indicating that the hurry-up call was not necessary in this instance, as most, if not all of the collections would have been scheduled in any case before the site reached capacity.

TABLE 2 Influence of penalty fill level parameter value

Penalty fill level	Rounds (#)	Distance (km)	Time (hrs)	Goods (tonnes)		Revenue (£K)		Site visits	
				Collected	Left over	Collected	+Left over	Banks	Shops
Existing	99	22953	923	225.2	26.8	107.3	120.0	444	656
50%	73	19124	722	222.7	33.3	109.7	125.7	216	451
75%	72	18834	718	222.8	33.2	110.0	126.0	216	450
95%	71	18888	717	222.8	33.3	109.9	125.9	216	450

5. CONCLUSIONS

While previous research showed that remote monitoring of charity donation banks was of negligible benefit where the collection rounds were constrained by fixed shop collections, relaxation of this constraint has resulted in estimated time and distance savings of around 30%, comparable to performance suggested by other authors in

similar related waste-recycling applications (Johansson, 2006; Krikke et al, 2008). The results are based on the assumption that bank and shop demands for the next working day can be obtained to a reasonable degree of accuracy. For banks, this may be done through the use of remote monitoring sensors reporting fill levels at regular intervals, with procedures in place to be able to cope with any missing data, if, for example, a sensor fails to report. No system is currently in place to allow the charity's shop managers to report their collection requirements, although a smartphone application has been developed for this study to demonstrate how this may be done (www.sixthsensetransport.com). This application also provides visibility of the outputs from the remote monitoring sensors installed in banks. Areas for further research, if such a dynamic collection strategy was to be considered for implementation, include: investigating how shop managers would respond, in terms of sorting out goods for collection; understanding how driver activity schedules might change as a result of the time gained through remote monitoring; and further investigation of the reasons for collection time variability.

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 (i.e. take on the servicing of them), as the goods 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 would need to be considered in a practical application of the algorithm where a shop manager might need an urgent top-up of goods from a local bank as soon as possible, regardless of the suggested collection day. Smartphone technology could provide this enhanced visibility and provide a link between the algorithm outputs and the various parties implementing the strategies on the ground. This could provide an interface for the algorithm developed here and allow proposed new next-day or longer 'look-ahead' schedules to be viewed by the logistics scheduler, drivers and shop managers.

The research described here was limited to considering a one-day look-ahead period only (i.e. planning for tomorrow), justified by the fact that donations rates in this case were highly variable. For applications where container fill rates are less variable, there is scope for considering a longer planning period of two or more days, which would likely improve vehicle rounds. However, the more predictable the fill rate the less need there is for the remote monitoring of containers.

6. REFERENCES

Amponsah, S. K. and Salhi, S., 2004. The investigation of a class of capacitated arc routing problems: the collection of garbage in developing countries. *Waste Management*, Vol. 24, No. 7, pp. 711-721.

Angelelli, E. and Speranza, M. G., 2002. The application of a vehicle routing model to a waste-collection problem: two case studies. *Journal of the Operational Research Society*, Vol. 53, No. 9, pp. 944-952.

Baldacci, R. and Mingozzi, A., 2009. A unified exact method for solving different classes of vehicle routing problems. *Mathematical Programming*, Vol. 120, pp. 347-380.

Brandão, J., 2011. A tabu search algorithm for the heterogeneous fixed fleet vehicle routing problem. *Computers & Operations Research*, Vol. 38, No. 1, pp. 140–151.

Ceschia, S., Gaspero, L., and Schaerf, A., 2011. Tabu search techniques for the heterogeneous vehicle routing problem with time windows and carrier-dependent costs. *Journal of Scheduling*, Vol. 14, No. 6, pp. 601-615.

Charity Retail Association, 2013. Charity shops, re-use and recycling.

<http://www.charityretail.org.uk/reuse.html> (Accessed 7/11/2013)

den Boer, E., Otten, M. and van Essen, H., 2011. Comparison of various transport modes on a EU scale with the STREAM database. http://www.cedelft.eu/publicatie/stream_international_freight_2011/1174 (Accessed 7/11/2013)

Garey, M.R. and Johnson D.S., 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W.H. Freeman, New York.

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

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

Johansson, O. M., 2006. The effect of dynamic scheduling and routing in a solid waste management system. *Waste Management*, Vol. 26, No. 8, pp. 875-885.

Kim, B.-I., Kim, S., and Sahoo, S., 2006. Waste collection vehicle routing problem with time windows. *Computers & Operations Research*, Vol. 33, No. 12, pp. 3624-3642.

Krikke, H., le Blanc, I., van Krieken, M., and Fleuren, H., 2008. 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, pp. 209-228.

McLeod, F. N., Cherrett, T. J., and Waterson, B. J., 2011. The scope for joint household/commercial collections: a case study. *International Journal of Logistics: Research and Applications*, Vol. 14, No. 6, pp. 399-411.

McLeod, F. N., Erdogan, G., Cherrett, T. J., and Bektas, T., 2013. Dynamic collection scheduling using remote asset monitoring: A case study in the charity sector, *Transportation Research Board 92nd Annual Meeting*, Washington DC.

Paraskevopoulos, D., Repoussis, P. P., Tarantilis, C. D., Ioannou, G., and Prastacos, G. P., 2008. A reactive variable neighbourhood tabu search for the heterogeneous fleet vehicle routing problem with time windows. *Journal of Heuristics*, Vol. 14, No. 5, pp. 425–455.

Rovetta, A., Xiumin, F., Vicentini, F., Minghua, Z., Giusti, A., and Qichang, H., 2009. Early detection and evaluation of waste through sensorized containers for a collection monitoring application. *Waste Management*, Vol. 29, No. 12, pp. 2939-2949.

Yaman, H., 2006. Formulations and valid inequalities for the Heterogeneous Vehicle Routing Problem. *Mathematical Programming*, Vol. 106, No. 2, pp. 365–390.