Ann Oper Res (2006) 144:33–58
DOI 10.1007/s10479-006-0006-y

Using mixed integer programming for solving the capacitated arc routing problem with vehicle/site dependencies with an application to the routing of residential sanitation collection vehicles

John Sniezek · Lawrence Bodin

Published online: 24 May 2006

Abstract We present the Composite Approach for solving the Capacitated Arc Routing Problem with Vehicle/Site Dependencies (CARP-VSD). We also present two mixed integer programs, the Initial Fleet Mix Generator (IFM) and the Mathematical Programming Procedure (MPP), and a multi-criterion function called the Measure of Goodness. The IFM, MPP and Measure of Goodness are critical components of the Composite Approach. A key application area of the CARP-VSD is the routing of residential sanitation vehicles; throughout this paper, we derive parameters specific to this problem. In this paper, we describe the Composite Approach, the IFM, the MPP, the Measure of Goodness and work out several examples in detail.

Keywords Mixed integer programs · Heuristic methods

1. Introduction

A *vehicle class* is a set of vehicles with identical characteristics. A *vehicle/site dependency* on an arc exists if vehicles from one or more vehicle classes cannot service or traverse the arc while vehicles from other vehicle classes can service or traverse the arc. The Capacitated Arc Routing Problem with Vehicle/Site Dependencies (CARP-VSD) (Sniezek, 2000; Sniezek et al., 2002) is an arc routing problem where the vehicle fleet consists of at least two vehicle classes (fleet is non-homogeneous) and at least one arc in the network has a vehicle/site dependency. Vehicles from different vehicle classes can differ in capacity, daily capital cost, length of workday, crew size, or other characteristics. The vehicle/site dependency on an arc determines the vehicle classes that can (i) service the arc, (ii) travel along the arc but not

J. Sniezek
RouteSmart Technologies, 8550 Stanford Boulevard, Suite 2600, Columbia, MD 21044
e-mail: Jsniezek@routesmart.com

L. Bodin (✉)
R. H. Smith School of Business, University of Maryland, College Park, MD 20742
e-mail: Lbodin@rhsmith.umd.edu

service the arc, and (iii) neither travel along the arc nor service the arc. The CARP-VSD with no vehicle/site dependencies is called the Mixed Fleet Capacitated Arc Routing Problem (CARP-MF) (Dror, 2000).

In this paper, *an arc represents both sides of a street joining two intersections in a digital representation of a geographic area* and the algorithms assume that both sides of the street are serviced by the same vehicle, although not necessarily at the same time. A vehicle's travel path is the sequence of streets the vehicle would follow in its route. Each side of the street can be serviced by the vehicle at a different time if meandering of the street is not allowed. Both sides of the street can be serviced at the same time in a single traversal of the street if meandering is allowed. Our travel path generation (not described in this paper) takes these situations into account when forming the travel path for a route.

A major application of the CARP-VSD is the problem of routing residential solid waste collection vehicles. Vehicle/site dependencies occur in residential solid waste collection due to such factors as some streets (for example, alleys) being too narrow for vehicles from certain vehicle classes to traverse or service and some bridges or overpasses not able to support the weight of the vehicles from some vehicle classes. Another application of the CARP-VSD is the scheduling of field service or repair services. In this problem, the staff is broken down by skill level and some arcs can only be serviced by employees with certain skill levels.

1.1. Example of the CARP-VSD

Figure 1 illustrates a 95 arc example of the CARP-VSD. This network is extracted from a 1776 arc data set from Philadelphia, PA. In Fig. 1, both sides of the street are grouped together and drawn as a single arc. Vehicles come from 3 vehicle classes (called vehicle classes 1, 2, and 3) and there are 3 levels of vehicle/site dependencies on the arcs. The 10 dotted arcs must be serviced by a vehicle from vehicle class 1, the 14 dashed arcs must be serviced by a vehicle from vehicle classes 1 or 2, and the 71 solid arcs can be serviced by vehicles from any vehicle class. The subnetwork of dotted arcs and the subnetwork of dashed streets are disconnected whereas the subnetwork of solid arcs is connected. The disconnectivity of subnetworks is typical in the CARP-VSD.

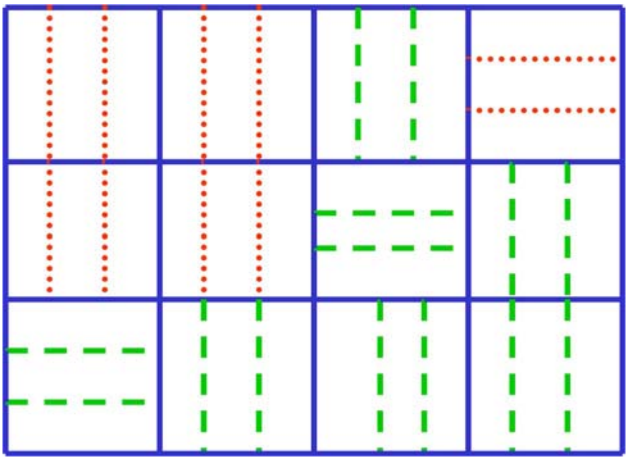

Fig. 1 Network for small example

Analysis of this network as well as other networks can be found in Sniezek (2001). This analysis includes plots of the networks color-coded by vehicle/site dependency, vehicle class or other characteristics. Upon request, Dr. Sniezek at jsniezek@routesmart.com will email a copy of Sniezek (2001) and color plots and digitized networks of some of the examples.

1.2. The capacitated arc routing problem

The CARP–VSD is a variant of the well known Capacitated Arc Routing Problem (CARP) (Assad and Golden, 1995; Bodin et al., 1983; Dror, 2000; Eiselt et al., 1995a, 1995b and many others) In a CARP, the following holds:

- The total travel time in any solution is the sum of the total service times on the arcs and the total deadhead time in the solution.
- The total service time on an arc is assumed to be a constant regardless of the solution.
- The total deadhead time in a solution is variable and generally much smaller than total service time.
- There is only one vehicle class; i.e. the vehicle fleet is homogeneous.

1.3. Glossary of terms

Some of the terms used in this paper are now briefly defined and these definitions are expanded in the denoted sections.

- The *Composite Algorithm (CA)*, outlined in Section 2, is our overall approach for solving the CARP-VSD.
- The *Initial Fleet Mix Generator (IFM)*, described in Section 4, is a mixed integer program whose solution gives an initial fleet mix for the CA. An example of the IFM, called IFM2, is given in Section 4.4. All instances of the IFM given in this paper were solved using LINDO.
- The *Measure of Goodness (MG)* is a function that allows the user to trade off between capital costs, operating costs, overtime costs, and a measure of partition compactness. The MG is a linear combination of the total cost of a solution, called the Total Solution Cost, and the measure of partition compactness for a solution, called GEOSPREAD. Variants of the MG are used throughout this paper. MG, *Total Solution Cost* and *GEOSPREAD* are described in Section 3.
- The *Mathematical Programming Problem (MPP)*, a mixed integer programming formulation of the CARP-VSD, is used as a route improvement procedure within the CA. A solution to the MPP over a subset of the partitions moves arcs between these partitions in order to achieve better balance, reduced cost and/or reduced route interlacing. MP1, MP2, and MP3 are three formulations of the MPP that arise when certain classes of variables in the MPP are fixed to a specific value. When MP1, MP2 and MP3 are solved, a new assignment of arcs to partitions and, possibly, a new fleet mix for MP2 and MP3 are found. The MPP is formulated in Section 5 and strategies for using the MP1, MP2 and MP3 as a between-route arc exchange procedure are described in Section 6. A detailed example of the overall approach is given in Section 7. In this analysis, AMPL was used to generate the problems, and CPLEX was used to solve the problems. The problems were solved on a UNIX SPARC workstation with 80 MB RAM and SUNOS 5.6.
- The original version of the *Vehicle Decomposition Algorithm (VDA)* (Sniezek et al., 2002) found an initial fleet mix, the breaking up of the area into partitions and a travel path for each partition. In the CA, the IFM finds the initial fleet mix and the VDA forms partitions

and travel paths. The original VDA is identical to the VDA used in the CA except that the original VDA determines an initial fleet mix concurrently with the partitions while the VDA used in the CA is given an initial fleet mix from either the IFM or some other source.

2. The composite approach (CA) for solving the CARP-VSD

It is difficult to develop an approach for solving the CARP-VSD that concurrently (i) determines a good fleet mix, (ii) forms reasonable partitions and travel paths and (iii) assigns a vehicle from the correct vehicle class to each partition. A problem that we have noted in reviewing papers for solving variants of the CARP is that the authors attempt to formulate mixed integer programs whose solution concurrently forms partitions and travel paths. These formulations either take an extraordinary amount of time to solve and can only solve very small problems (around 100 arcs).

The CA operates in a sequential manner. The CA first finds a very good fleet mix, then finds a set of partitions and travel paths for this fleet mix and then moves arcs between partitions in order to find an improved solution (possibly changing the fleet mix). The MG allows the user to tradeoff between capital costs, operating costs, overtime costs, and partition compactness and is necessary for finding good solutions to the CARP-VSD since vehicle deadheading is often only a minor component of cost in any solution to the CARP-VSD.

The CA is quite different from traditional approaches for solving vehicle routing problems. The CA uses good feasible solutions to mixed integer programs to derive an initial fleet mix and swap streets between components. The objective function in these mixed integer programs is the MG, a multiple criterion function. We found that minimizing travel time or travel distance did not generate solutions that could be used in practice. and that the CA will generate a solution to the CARP-VSD that is considered better even though there is more deadhead travel time and deadhead distance in the solution than in another feasible solution. Although the basic ideas in the CA are similar to other approaches to vehicle routing problems, the details in the CA are much different and challenging and offer insight into how aspects of the CA can be applied to other vehicle routing problems.

2.1. Steps in the CA

The CA consists of the following three steps.

Step a. The IFM determines an initial fleet mix that consists of vehicles from one or more vehicle classes.

Step b. The VDA determines partitions and travel paths for a fleet mix. Seed points are defined for each vehicle in the fleet mix. The arcs requiring service are divided into partitions in such a way that the vehicle/site dependency on each arc in each partition is satisfied. A few arcs (generally less than 5) are moved at a time between partitions in order to achieve better balance. Once the swapping is complete, a travel path that approximately minimizes deadhead time or distance is found for each partition. The fleet mix generally does not change when swapping is carried out. Once a solution is found, seed points can then be redefined for each partition and a new solution found. This process continues until the solutions on consecutive passes through Step b are identical or a fixed number of passes through Step b is carried out.

Step c. MP1, MP2 and MP3 are used to swap arcs between partitions while maintaining vehicle/site dependency feasibility and all other constraints in the problem. A solution to MP2 and MP3 can change the fleet mix.

2.2. Discussion

The VDA is fast computationally and can be used on large problems involving thousands of arcs and a large number of vehicles. Documented savings of over 12% have been realized at the Philadelphia Sanitation Department (Sniezek et al., 2002) using the original version of the VDA.

In its initial design, the VDA determined the initial fleet mix (Step a of the CA) at the same time that Step b was carried out. Since we found that integrating Step a and Step b of the CA generates inferior solutions, we developed the IFM to determine an initial fleet mix for the CARP-VSD without forming partitions and travel paths. We show in Section 4.3.6 that in 42 instances the estimated cost of a solution using the IFM without forming partitions and travel paths and the cost of the same solution using the VDA to develop the partitions and travel paths are very close. In most cases, the optimal solution to the IFM can be found in less than a second.

Since the IFM has the capabilities to find a very good fleet mix and an estimated cost that is close to the cost that will be realized when generating actual partitions with the VDA, the IFM can be used very effectively to estimate the final cost of the fleet mix that it determined. Multi-attribute ‘cost’ models are used in transportation applications for estimating the cost of a possible configuration without actually generating a feasible solution to the problem. Bodin et al. (1978, 1981, 1983), give examples for using functions like the MG for analyzing mass transit systems.

A solution to MP1, MP2 or MP3 can (i) improve an existing solution by exchanging many arcs between partitions at the same time, (ii) improve the compactness of the partitions, and (iii) possibly change the fleet mix. Between-route swapping procedures that only exchange a few arcs at a time have difficulty changing the fleet mix. Because we consider the capital costs of the vehicles, we want the ability to change the fleet mix.

We generally terminated runs of MP1, MP2 and MP3 after 30 min and used the best integer solution found. This approach is very reasonable since the applications of the CARP-VSD are only solved periodically—every 6 months or a year, for example. As such, the user can invest considerable time in determining the best solution. Computing time would become an issue if the problems were solved more frequently—daily or weekly for example.

3. Measure of goodness for the CARP-VSD

The MG serves as the basis for the cost models that evaluate and build solutions to the CARP-VSD. In the CARP-VSD, an objective that minimizes deadhead time or travel time alone is not sensitive enough to generate good solutions to the CARP-VSD since the costs associated with deadhead time and travel time are quite small when compared with the other costs involved in this problem.

3.1. Total cost of a solution

The *Total Solution Cost* is the total cost of a solution and is the sum of the Total Solution Cost(Partition p) taken over all partitions p; i.e.

$$\begin{aligned} \text{Total Solution Cost} &= \sum_p \text{Total Solution Cost (Partition } p) \\ &= \sum_p (\mathrm{FC}_p + \mathrm{VLC}_p + \mathrm{VRC}_\mathrm{p}) \end{aligned} \tag{1}$$

In (1), FC_p is the fixed cost of partition p, VLC_p is the variable labor cost of partition p, and VRC_p is the variable routing cost of partition p. These costs are defined as follows:

$$\text{Fixed Cost}: \quad \mathrm{FC}_\mathrm{p} = \$\mathrm{CAP}_k + \$\mathrm{SAL}_k \tag{2}$$

$$\begin{aligned} \text{Variable Labor Cost}: \quad \mathrm{VLC}_\mathrm{p} &= \$\mathrm{HOUR}_k * \mathrm{OT}_p \\ &= \$\mathrm{HOUR}_k * \mathrm{MAX}[(\mathrm{TIME}_p - \mathrm{TARGET}_k), 0] \end{aligned} \tag{3}$$

$$\text{Variable Routing Cost}: \quad \mathrm{VRC}_\mathrm{p} = \$\mathrm{MILE}_k * \mathrm{DIST}_\mathrm{p} + \$\mathrm{TIP}_k * Z_p \tag{4}$$

The various terms in (2), (3) and (4) are defined as follows:

- $\$CAP_k$ represent the *amortized daily capital cost* of a vehicle from vehicle class k and is generally vehicle class dependent since larger vehicles have a larger daily capital cost than smaller vehicles. The amortized daily capital cost can be determined from the vehicle purchase price, maintenance cost, insurance cost, etc, and the expected lifetime of the vehicle.
- $\$SAL_k$ represents the *daily salary cost of a crew* associated with vehicle class k working up to $TARGET_k$ hours where $TARGET_k$ is the length of the work of a crew from vehicle class k. Since the number of persons in a crew can be vehicle class dependent., $\$SAL_k$ is vehicle class dependent. This cost includes benefits and overhead. In sanitation routing, some vehicle classes require a one person crew while other vehicle classes require a two or three person crew. We have seen up to 10 person sanitation crews in some countries. When skill levels are involved, $\$SAL_k$ can be a function of the skill level as more skilled people earn more \$/hour.
- $\$HOUR_k$ represent the *hourly cost of overtime for a crew* from vehicle class k and represents the pay/hour that a crew from vehicle class k receives if the crew works longer than $TARGET_k$ hours.
- $TIME_p$ is the *total time* the crew requires to service partition p. This time includes the time to/from the disposal facility and the time to/from the depot.
- $OT_p = MAX\ [(TIME_p - TARGET_k), 0]$ is the *overtime* associated with partition p.
- $\$MILE_k$ represent this *cost per mile* for a vehicle from vehicle class k. $\$MILE_k$ is assumed to be the same whether the vehicle is performing a service or deadhead, and is vehicle class dependent.
- $DIST_p$ is the *total distance* in the travel path associated with covering the arcs in partition p and is the sum of the distance covered while servicing the streets, distance while deadheading, distance traversed to and from the disposal facility and the distance to and from the depot.

- $\$TIP_k$ is the *fixed cost portion of the tipping cost* for a vehicle from vehicle class k. The tipping cost of a vehicle at the disposal facility is a significant cost that is specific to sanitation VRPs. Since a vehicle from vehicle class k has a fixed capacity, M_k, the vehicle must be emptied at a disposal facility when the vehicle capacity is reached. A vehicle can make several trips to a disposal facility over a day and the vehicle must pay a tipping cost each time it visits the disposal facility. The tipping cost is typically measured as a fixed cost plus a variable cost measured in \$/unit of weight. Since the total weight collected over an entire solution is constant regardless of the fleet mix, the weight component of the tipping cost is a constant for any fleet mix.
- Z_p is the *total number of trips to the disposal facility* needed to service all arcs in partition p and includes one trip to the disposal facility at the end of the workday to unload any residual volume on the vehicle. In some unusual sanitation applications, the vehicles are not emptied at the disposal facility at the end of the day; instead, these vehicles return to the depot at the end of the day partially loaded.

If the fleet is homogeneous so that there is only one vehicle class, no overtime and one trip per route, then, for any partition p, $\$CAP_k$, $\$SAL_k$, $\$MILE_k$ and $\$TIP_k$ do not depend upon vehicle class k, $\$HOUR_k = 0$ and $Z_p = 1$. As such, the only variable in determining the cost of any partition p is DIST_p, the total distance of partition p. Thus, the total travel distance over all of the routes is $\sum_p \text{DIST}_p$. Minimizing $\sum_p \text{DIST}_p$ is the typical objective of most CARPs. In some cases, distance is converted to a cost and the objective of the CARP becomes to minimize cost

3.2. Partition compactness

A user will generally not accept a solution to a residential sanitation CARP-VSD unless the partitions are compact (interlacing of the routes is avoided as much as possible) and the solution has close to minimal cost. The following describes a methodology for measuring partition compactness.

Let C_p be the measure of *partition compactness* for partition p. If arc a is assigned to partition p, then D_{ap} is the distance from the closer of the two nodes making up the endpoints of arc a to the seed point of partition p.

Then, C_p is defined as follows:

$$C_p = \sum_a (D_{ap})^2 \text{ where the sum is over all arcs a assigned to partition } p \tag{5}$$

GEOSPREAD, the *measure of compactness for a solution*, is computed as follows:

$$\text{GEOSPREAD} = \sum_p C_p \text{ where the sum is taken over all partitions } p \tag{6}$$

3.3. Measure of goodness for a solution to the CARP-VSD

MG, the *Measure of Goodness for a solution* to the CARP-VSD, is the following:

$$\text{MG} = \alpha^* \text{Total Solution Cost} + \lambda^* \text{GEOSPREAD} \tag{7}$$

In (7), the Total Solution Cost is defined in (4), GEOSPREAD is defined in (6) and α and λ are parameters that allow for a tradeoff between the total cost of a solution and the measure of route compactness in a solution.

3.4. Discussion

(1)–(7) allows the user to estimate the total solution cost, measure the compactness of each route in a solution, and use the MG as a measure of the goodness of the solution. The parameters, α and λ, in (7) and the parameters used to compute (4) give the user great flexibility in trading off between the cost of a solution and the route compactness of the solution. *The above functions are developed specifically for residential sanitation CARP-VSD problems; similar functions can be developed for other vehicle routing applications*.

4. Using the MG to determine an initial fleet MIX for the CARP-VSD

The IFM, formulated in Section 4.3, can be used in Step a of the CA to find an initial fleet mix for the CARP-VSD. Input to the IFM are the costs identified with each vehicle class (Section 4.1) and the value of the *Vehicle Class Workload Matrix* V (Section 4.2).

4.1. Estimating the daily cost of each vehicle class

Table 1 gives the daily costs for four vehicle classes representing small, medium, large and very large vehicles for a typical vehicle fleet in residential sanitation collection and is formed under the following assumptions:

- Crew size for each vehicle class is specified.
- Each crew member costs \$150/day, including benefits.
- The capital cost of a vehicle from any vehicle class is amortized over 1000 workdays (approximately 4 years). The daily capital cost of a vehicle is the vehicle's capital cost divided by 1000. The vehicle capital costs by vehicle class is assumed to be the following: vehicle class 1 - \$100,000, vehicle class 2 - \$137,000, vehicle class 3 - \$156,000, vehicle class 4 - \$250,000.
- The fixed cost portion of the tipping fee is \$100. The average number of trips/day to the disposal facility for a partition as a function of vehicle class is estimated. Generally, an organization has a good estimate for average number of trips/day to the disposal facility by vehicle class.

Table 1 Estimated total daily vehicle cost by vehicle class

Vehicle Class	Daily capital cost	Crew size	Daily crew cost	**Daily fixed cost**	Disposal trips	Disposal cost	Route mileage	Mileage cost	**Daily routing cost**	**Total daily vehicle cost**
1	100	2	300	**400**	4	400	60	120	**520**	**920**
2	137	2	300	**437**	3	300	50	100	**400**	**837**
3	156	2	300	**456**	3	300	50	100	**400**	**856**
4	250	3	450	**700**	2	200	40	80	**280**	**980**

- The mileage on each route is estimated and includes the mileage between the routes and the disposal facility and the mileage to and from the depot. Each mile is assumed to cost \$2.
- The overtime associated with each route is zero.

The bold columns in Table 1 represented the total daily fixed costs, routing costs and vehicle cost by vehicle class.

4.2. Determining the *V* matrix

The entry, $V(i, j)$, in the *V* matrix represents the estimated number of vehicles from vehicle class j required to service all arcs whose it is assumed that vehicles from vehicle classes 1 through i are the only vehicles that can service these arcs. $V(i, j)$ does not consider any arcs considered in the estimates of $V(1, j)$ up to $V(i-1, j)$. In other words, the arcs used in the computation of the $V(i, j)$'s for any column j in the *V* matrix are mutually exclusive.

The following assumption is made in generating the *V* matrix and, thus, in the IFM and the VDA. *If a vehicle from vehicle class i can service or traverse an arc, then a vehicle from vehicles classes 1, 2, ..., $i-1$ can also service this arc.* This assumption is relaxed in the formulation of the MPP (Step c of the CA). The process for generating the *V* matrix can be found in Section 11.3.2.7 of Sniezek et al. (2002) (pages 298–300 of Toth and Vigo (2002)).

A sample *V* matrix based on the actual 1776 arc network from Philadelphia., PA is given in Fig. 2. This network has a significant number of site dependencies. The element $V(3, 1) = 3.604$ in V means that it is estimated that 3.604 vehicles from vehicle class 1 are needed to service all of the arcs who can be serviced by vehicles from vehicle classes 1, 2, and 3. Since the sum of the column 1 of V is 25.005, we estimate that 25.005 vehicle class 1 vehicles are needed to service the arcs of this network if only vehicle class 1 vehicles are used. The upper triangular portion of the *V* matrix is not defined since it is impossible to service an arc with a vehicle from vehicle class j if the service on the arc in constrained to vehicles from vehicle classes $1, 2, \ldots, i, i < j$. Many of the computational tests in this paper use this network and *V* matrix.

4.3. Mixed integer program to determine the initial fleet mix

The IFM formulated below determines an initial fleet mix for the CARP-VSD. The solution to the IFM allocates enough vehicles to satisfy the workload requirements of each row of the *V* matrix while minimizing the total cost of the fleet mix. As noted above, it is assumed in the IFM that if an arc can be serviced by a vehicle from vehicle class i, then this arc can be serviced by a vehicle from vehicle classes $1, 2, \ldots, i$.

Fig. 2 Sample *V* matrix

$$V = \begin{bmatrix} 1.987 & \text{------} & \text{------} & \text{------} \\ 0.958 & 0.756 & \text{------} & \text{------} \\ 3.604 & 2.807 & 2.807 & \text{------} \\ 18.456 & 15.380 & 13.445 & 11.530 \end{bmatrix}$$

4.3.1. Variables in the IFM

The variables in the IFM are as follows:

- X_j = Total number of vehicles in the fleet from vehicle class j.
- X_{ij} = Number of vehicles from vehicle class j to use to satisfy the requirements in row i of the V matrix.
- S_i = Excess number of vehicles (slack) available after satisfying the requirements for row i of the V matrix. This slack is present because the required number of vehicles needed to satisfy a row of the V matrix may not be integer while the number of vehicles allocated must be integer. This slack is used to meet the requirements of subsequent rows of the V matrix.

4.3.2. Data elements input to the IFM

The data elements input to the IFM are as follows:

- C_j = Estimated total daily cost of a vehicle from vehicle class j.
- *VCLASS* = Total number of vehicle classes considered in the IFM.
- $V = (V_{ij})$ matrix. The V matrix has VCLASS rows and columns.
- $QMAX_j$ = Maximum number of vehicles available from vehicle class j.
- $QMIN_j$ = Minimum number of vehicles from vehicle class j that must be used in the solution. Although this data element is typically zero, this lower bound allows an organization to ensure using QMIN_j vehicles of type j if that is a requirement of the organization such as an agreed-upon contract with the crews.

4.3.3. Formulation of IFM

The formulation of the IFM is the following:

$$\text{(IFM)} \quad \text{Min } Z = \sum_i \sum_{\text{j}} C_j X_{ji} \tag{8}$$

subject to the following conditions:

$$\left(\sum_{j=1}^{i} (V_{i1}/V_{ij}) X_{ji} \right) + S_{i-1} - S_i = V_{i1} \quad \forall i \tag{9}$$

$$\sum_{i=j}^{\text{VCLASS}} X_{ji} = X_j \quad \forall j \tag{10}$$

$$X_j \leq \text{QMAX}_j \quad \forall j \tag{11}$$

$$X_j \geq \text{QMIN}_j \quad \forall j \tag{12}$$

$$S_0 = 0 \tag{13}$$

$$\text{All } X_{ji} \text{ and } X_j \text{ variables are integer and nonnegative} \tag{14}$$

$$\text{All } S \text{ variables are nonnegative} \tag{15}$$

(8) is the MG where $\alpha = 1$ and $\lambda = 0$ or, equivalently, the Total Solution Cost defined in (1) and minimizes the estimated total cost of the fleet mix chosen. $\lambda = 0$ since there is no notion of partition compactness as no partitions have been formed as yet. (9) ensures that enough vehicles from the vehicle classes are used to satisfy the workload for each row of the V matrix. V_{i1}/V_{ij} is a straight-forward proration formula on workload. As an example, if $V_{i1} = 12$ and $V_{ij} = 3$, then it is assumed that a vehicle from vehicle class j can cover the workload of 4 vehicles from vehicle class i on a given day. V_{i1}/V_{ij} does not equal $V_{i+1,1}/V_{i+1,j}$ because the elements of the V matrix are formed with an estimate of the number of possible trips to the disposal facility.

(10) tallies the number of vehicles used from each vehicle class. (11) and (12) are the upper and lower bounds on the total number of vehicles from each vehicle class that can be used in the fleet mix. (13) ensures that there is no slack from any other vehicle class available towards satisfying the demand in row 1 of the V matrix. The slack variables account for an integer number of vehicles being used to satisfy a non-integer demand in a row of the V matrix. S_0 is set to zero since there can be no slack yet realized from a previous row to row 1 of the V matrix. (14) ensures that the number of vehicles by vehicle class used in the fleet mix are both nonnegative and integer and (15) ensures the slack variables are nonnegative.

4.3.4. Example

An example of the IFM is as follows. The cost estimates are defined in Table 1, the V matrix is defined in Fig. 2 and all vehicles have to be purchased. Further, we assume than QMAX_j = infinity and $\mathrm{QMIN}_j = 0$ for all vehicle classes j. The IFM becomes the following mixed integer program (IFM2):

$$\begin{aligned}\text{(IFM2) MIN } Z = {} & 920X_{11} + 920X_{12} + 920X_{13} + 920X_{14} + 837X_{22} + 837X_{23} + 837X_{24} \\ & + 856X_{33} + 856X_{34} + 980X_{44}\end{aligned}$$

SUBJECT TO

$$\begin{aligned}& X_{11} + S_0 - S_1 = 1.987 \\ & X_{12} + 1.27X_{22} + S_1 - S_2 = 0.958 \\ & X_{13} + 1.28X_{23} + 1.28X_{33} + S_2 - S_3 = 3.604 \\ & X_{14} + 1.20X_{24} + 1.37X_{34} + 1.60X_{44} + S_3 - S_4 = 18.456 \\ & X_1 - X_{11} - X_{12} - X_{13} - X_{14} = 0 \\ & X_2 - X_{22} - X_{23} - X_{24} = 0 \\ & X_3 - X_{33} - X_{34} = 0 \\ & X_4 - X_{44} = 0 \\ & S_0 = 0 \\ & \text{All variables nonnegative} \\ & \text{All } X_{ij} \text{ and } X \text{ variables are integer.}\end{aligned}$$

The solution to the IFM2 gives the following fleet mix: two vehicles from vehicle class 1 ($X_1 = 2$), six vehicles from vehicle class 2 ($X_2 = 6$), three vehicles from vehicle class 3 ($X_3 = 3$), and seven vehicles from vehicle class 4 ($X_4 = 4$). The total cost of this solution is \$16,290/day.

4.3.5. Discussion

If the company presently owns 25 class 1 vehicles and wishes to use these 25 vehicles to provide its service, then 25 class 1 vehicles are needed to cover the demand since the sum of the entries in column 1 in the *V* matrix of Fig. 2 is 25.005. As there is no capital cost associated with these class 1 vehicles, the daily cost of a Class 1 vehicle as given in Table 1 is $820 = $920–$100 where the $100 is the fixed cost for purchasing a class 1 vehicle. The total cost of this fleet of 25 class 1 vehicles is $20,500/day ($820 ∗ 25).

The IFM2 solution costs $16,290/day and the 25 class 1 vehicle solution costs $20,500/day. It appears that the company would save $4,210/day (over 20%) by purchasing the non-homogeneous fleet dictated by the IFM2 solution and incurring additional capital costs rather than using the 25 class 1 vehicles that they already own. The cost savings occur since there are only 13 crews (33 crew members) required in the new fleet versus 50 crew members in the original fleet. Further savings occur when the salvage value of the company owned vehicles is taken into account and the company already owns 2 class 1 vehicles so these vehicles would not need to be purchased.

In a public sanitation service, implementing the IFM2 solution generally does not result in jobs being lost. The sanitation service will use these personnel on other activities, such as street sweeping, that are not normally carried out due to budgetary restrictions.

4.3.6. Using the IFM as a cost estimation tool

If the IFM is a good cost estimation tool, we only need to execute the CA once or twice, using the solution from the IFM as an initial solution and be assured that the solution is close to minimum cost and have a reasonable amount of interlacing. In Sniezek (2001), we assumed 42 different fleet mixes for the 1776 arc network from Philadelphia and ensured that each fleet mix was feasible to the IFM. We then developed partitions and travel paths for these fleet mixes using the VDA. In each case, we compared the total cost of the solution from the objective function of the IFM with the final cost of the solution found using the VDA. The results were very encouraging. In 8 out of 42 instances, the two costs were within 1% of each other. In 12 out of 42 instances, the two costs were between 1% and 2% of each other. In 8 out of 42 instances, the two costs were between 2% and 3% of each other. In 13 out of 42 instances, the two costs were between 3% and 6% of each other. In one instance, the two costs were between 6% and 7% of each other. Average deviation of these two costs was 2.31% over all 42 scenarios and the IFM tended to overestimate the cost found by the VDA in 38 of the 42 scenarios.

5. A mixed integer programming formulation of the CARP-VSD

The MPP is a mixed integer programming formulation of the CARP-VSD. A solution to the MPP breaks the arcs in a network into partitions, assigns a vehicle from a vehicle class to each of the partitions and ensures that the assignment of arcs to partitions is feasible with respect to vehicle/site dependency on the arcs. The objective function of the MPP is the MG given in (7). A difficulty with the MPP is that it gets large rapidly (in terms of number of integer variables) as the size of the network grows. It appears that the MPP, in its complete generality, can solve networks having about 500 arcs, 10 vehicles and 4 vehicle classes.

Our experiments have shown that the MPP has great promise as a route improvement procedure for the CARP-VSD (see Sections 6 and 7). Using the parameters that are available makes the MPP extremely flexibility in performing between-route swaps. Our strategy for solving large problems is outlined as follows (details are given in Section 6). We extract a set of routes from a solution and solve a variant of the MPP over this region. We then replace the old routes by the new routes and repeat the process. We continue until we are satisfied with the solution.

As an example, suppose we can have a region of 2000 arcs with vehicle/site dependencies that is serviced by a heterogeneous fleet of 40 vehicles. We select five routes containing approximately 200 arcs and resolve that portion of the network using the MPP while keeping the remainder of the solution fixed. It is possible that a solution to the MPP finds partitions for these 200 arcs with a different fleet mix and partitions where both the cost and interlacing are reduced from the previous solution. We then repeat with another set of routes. An example that illustrates the use of the MPP in this context is given in Section 7.

A major difficulty in solving mixed fleet vehicle routing problems is to develop an approach that both performs between-route exchanges and changes the fleet mix. Most between-route swapping procedures only move one or two arcs at a time between routes. These procedures do not alter the fleet mix since the fleet mix gets locked in as the procedure is carried out and these procedures are not global enough to generate a fleet mix change. Throughout this discussion, we used *seed point* to be a node around which arcs are accumulated and partitions formed. Each partition will represent a route for a vehicle, once a travel path is found.

5.1. Variables in the MPP

The variables in the formulation of the MPP are the following:

- $X_{ikp} = 1$ if arc i is assigned to vehicle class k at seed point p and 0 otherwise.
- $Y_{kp} = 1$ if vehicle class k is assigned to the partition being build around seed point p and 0 otherwise.
- Z_{kp} = Number of trips to the disposal facility that vehicle class k makes while servicing partition p when vehicle class k is assigned to service partition p.
- $TIME_p$ = Total time associated with partition p.
- $DIST_p$ = Total distance associated with partition p.
- OT_p = Total overtime associated with partition p.
- $GEOSPREAD$ = Measure of geographic spread in a solution.
- $COSTVALUE$ = Total cost of a solution.

5.2. Data input and parameters for the MPP

The following data inputs and parameters are needed in the MPP.

- P = Number of candidate seed points.
- R = Maximum number of seed points to be used. The user specifies whether $R = P$ or $R \leq P$. (See discussion of (20) in 5.5.)
- L_i = Length of arc i, $i = 1,2,\ldots$ no. of arcs in the network.
- D_{ikp} = Distance from arc i to seed point p for vehicle class k. The index $i = 0$ or $p = 0$ denotes the depot. The index $i = 0'$ or $p = 0'$ denotes the disposal facility. D_{ikp} can vary by vehicle class because there may be certain arcs that can be traversed by vehicles from some vehicle classes and not traversed by vehicles from other vehicle classes. To compute D_{ikp}, we first compute the shortest path from seed point p to all nodes for vehicle class k. D_{ikp} is then set to the minimum distance to the two nodes that make up the endpoints of arc i.

- T_{0kp} = Time from the depot to seed point p for vehicle class k.
- $T_{0'kp}$ = Time from the disposal to seed point p for vehicle class k.
- $T_{0'k0}$ = Time from the disposal to the depot for vehicle class k.
- C_{ik} = Volume on arc i for vehicle class k. Although the compaction ratio for different vehicle classes can vary based upon the type of garbage being picked up, we assume that $C_{ik} = C_i$ independent of k.
- S_{ik} = Service time of arc i for vehicle class k. In some residential sanitation collection problems, the service time for the arc can vary by vehicle class because of the type of vehicle (i.e. rear loader vs. side loader). The MPP allows for service time of an arc to vary by vehicle class. Since the VDA assumed a service time for an arc to be independent of vehicle class, we will assume in the MPP that the service time on an arc is independent of vehicle class. The MPP does not require this assumption.
- $a(i)$ = Set of vehicle classes that can service arc i; i.e., $a(i) = \{k|$ vehicle class k can service arc i and arc i requires service. If $a(i) = \Phi$, the empty set, arc i does not require service. The MPP does not require the hierarchical constraint assumption on site dependencies by vehicle class that is needed by the VDA. Thus, a more general CARP-VSD can be solved by the MPP than the VDA.
- M_k = Capacity of the vehicles from vehicle class k and expressed in the same units as C_{ik}
- DT_k = Time to unload a vehicle from vehicle class k at the disposal facility. We are assuming a vehicle is filled when going to the disposal facility, except at the end of the day. Further, for a given vehicle class, the majority of the time spent at the disposal facility is time in queue before unloading, and the actual time unloading is relatively independent of the fullness of the vehicle. We could assume DT to be independent of k, without much loss of generality, however, using DT_k rather than DT independent of k, adds more realism to the model.
- Q_k = Number of vehicles from vehicle class k, $k = 1, \ldots,$VCLASS, that can be used in the solution to the MPP.
- $Target_k$ = Upper bound on the duration of the workday of a crew from vehicle class k without paying overtime. Overtime must be paid to the crew who service a partition with duration longer than TARGET_k.
- Z'_k = Maximum number of trips to a disposal facility allowed in any partition for vehicle class k. The default value of Z'_k is a large number.
- $\$CAP_k$ = Daily capital cost for a vehicle from vehicle class k. Details in 3.1.
- $\$SAL_k$ = Daily salary cost of a crew associated with a vehicle from vehicle class k. Details in 3.1.
- $\$HOUR_k$ = Hourly cost of overtime for a crew from vehicle class k. In the formulation of the MPP, the hourly cost of overtime is assumed to be independent of vehicle class; i.e., for each value of k \$HOUR = $\$HOUR_k$. Details in 3.1.
- $\$MILE_k$ = Cost per mile traveled by a vehicle from vehicle class k. In the MPP, the cost per mile is assumed to be independent of vehicle class; i.e., for each value of k, \$MILE = $\$MILE_k$. Details in 3.1.
- $\$TIP_k$ = Fixed cost portion of the tipping cost. Details in 3.1.
- $\$FC_k$ = Fixed cost of a vehicle from vehicle class k. $\$FC_k = \$CAP_k + \$SAL_k$. Details in 3.1
- *MAX_GEOSPREAD_ALLOWED* = Upper bound on the GEOSPREAD component of the objective function. MAX_GEOSPREAD_ALLOWED can be set to infinity if no upper bound is known.

- *MAX_COSTVALUE_ALLOWED* = Upper bound on the COSTVALUE component of the objective function. MAX_GEOSPREAD_ALLOWED can be set to infinity if no upper bound is known.
- *MAX_PARTITION_OVERTIME* = Upper bound on the overtime allowed in a partition. MAX_PARTITION_OVERTIME can be set to infinity if no upper bound is known.
- $\phi = 1+$ *Estimated* percentage of deadhead time in a solution to the CARP-VSD. ϕ is ≥ 1.
- $\theta = 1+$ *Estimated* percentage of deadhead distance in a solution to the CARP-VSD. ϕ is ≥ 1.
- λ = Scaling factor applied to measure of geographic spread (GEOSPREAD) in the MPP objective function.
- α = Scaling factor applied to measurement of partition cost (COSTVALUE) in the MPP objective function.

ϕ and θ can be determined from a feasible solution to the CARP-VSD and this estimate is quite consistent from one solution to the next. The user-defined scaling factors, α and λ, ensure that both terms in the objective function defined below have the same units. As illustrated in the example in Section 7, λ and α gives the MPP great flexibility to generate solutions that trade off between geographic spread and cost. In a commercial implementation of the system, examples would be given to illustrate to the user how to set these parameters.

5.3. Formulation of the MPP

$$\text{MIN } Z = \lambda(\text{GEOSPREAD}) + \alpha(\text{COSTVALUE}) \tag{16}$$

Subject to

$$\sum_{k\in a(i)} \sum_{p} X_{ikp} = 1 \quad \forall i \text{ where } a(i) \text{ is not empty} \tag{17}$$

$$\sum_{k} Y_{kp} \leq 1 \quad \forall p \tag{18}$$

$$X_{ikp} \leq Y_{kp} \quad \forall i, k, p \tag{19}$$

$$\sum_{k} \sum_{p} Y_{kp} \leq R \tag{20a}$$

$$\sum_{k} \sum_{p} Y_{kp} = R \tag{20b}$$

$$\sum_{p} Y_{kp} \leq Q_k \quad \forall k \tag{21}$$

$$Z_{kp} \leq Y_{kp}{}^{*} Z'_k \quad \forall k, p \tag{22}$$

$$\sum_{i} \sum_{k\in a(i)} C_{ik} X_{ikp} \leq \sum_{k} M_k Z_{kp} \quad \forall p \tag{23}$$

$$\text{DIST}_p = \theta \sum_{i} \sum_{k\in a(i)} L_i\, X_{ikp} + \sum_{k} [(D_{0kp} - D_{0'kp} + D_{0'k0}) * Y_{kp} + 2D_{0'kp} * Z_{kp}] \quad \forall p \tag{24}$$

$$\text{TIME}_p = \phi \sum_{i} \sum_{k\in a(i)} S_{ik} X_{ikp}$$

$$+\sum_{k}[(T_{0kp} - T_{0'kp} + T_{0'k0})*Y_{kp}\ 2T_{0'kp}*Z_{kp} + DT_k Z_{kp}] \quad \forall p \tag{25}$$

$$\text{OT}_p \geq \text{TIME}_p - \sum_{k} \text{TARGET}_k * Y_{kp} \quad \forall p \tag{26}$$

$$\text{OT}_p \leq \text{MAX\_PARTITION\_OVERTIME} \quad \forall p \tag{27}$$

$$\text{GEOSPREAD} = \sum_{i}\sum_{k\in a(i)}\sum_{p}(D'_{ikp})^2 X_{ikp} \tag{28}$$

$$\text{COSTVALUE} = \sum_{k}\sum_{p}(\$\text{FC}_k * Y_{kp} + \$\text{TIP}_k Z_{kp})$$

$$+\sum_{p}(\$\text{HOUR} * \text{OT}_p + \$\text{MILE} * \text{DIST}_p) \tag{29}$$

$$\text{GEOSPREAD} \leq \text{MAX\_GEOSPREAD\_ALLOWED} \tag{30}$$

$$\text{COSTVALUE} \leq \text{MAX\_COSTVALUE\_ALLOWED} \tag{31}$$

$$\text{All Variables} \geq 0 \tag{32}$$

$$X_{ikp} = 0 \text{ or } 1\ \forall(i, k, p) \tag{33}$$

$$Y_{kp} = 0 \text{ or } 1\ \forall(k, p) \tag{34}$$

$$Z_{kp} \geq 0 \text{ and integer } \forall(k, p) \tag{35}$$

5.4. Explanation of the objective function of the MPP

The objective function measures the total compactness of the partitions in the solution as represented by GEOSPREAD and the total cost of the solution as represented by COSTVALUE. By appropriately setting the values of λ and α, the objective function allows the user to trade off between GEOSPREAD and COSTVALUE. Setting $\lambda = 0$ allows the MPP to minimize cost and setting $\alpha = 0$ allows the MPP to minimize interlacing. When set appropriately, (30) and (31) ensure that the existing solution is feasible to the MPP; i.e., the new solution has values of GEOSPREAD and COSTVALUE no worse than the old solution and the value of (16) is improved.

5.5. Explanation of the constraints of the MPP

(17) ensures that every arc requiring service is assigned to a partition. (18) ensures that at most one vehicle class is assigned to a seed point, (19) ensures that an arc is assigned to a vehicle/seed point combination if that vehicle class is assigned to service the arcs associated with that seed point.

(20a) ensures that at most R partitions can be created from P candidate seed points. (20b) ensures that exactly R partitions can be created from P candidate seed points. The user can select whether to use (20a) or (20b). These constraints play a role in the solutions strategies proposed for this problem in Section 6.

(21) ensures that the number of vehicles used from a vehicle class does not exceed the number of available vehicles in the vehicle class. (22) and (23) are coupled together and determine the number of trips to the disposal facility for partition p. (22) ensures that no more than the maximum number of disposal facility trips allowed are used and (23) ensures

that the total capacity assigned to partition p is no greater than the vehicle capacity of the vehicle assigned to the partition times the number of trips to the disposal facility the vehicle makes.

The remaining constraints in the problem represent accumulation variables or bounds. An *accumulation variable* computes the value of a variable as the (weighted) sum of other variables in the problem. (24) computes DIST_p, the estimated total distance for partition p, as the sum of the following two components:

a. $\theta \sum_i \sum_k L_i X_{ikp}$ where $\sum_i \sum_k L_i\, X_{ikp}$ represents the total service distance and θ is a multiplier that estimates within-route deadhead distance as a function of total service distance. In residential sanitation collection, θ is generally no larger than 1.10 as within-route deadheading is generally no more than 10% of the route service distance.
b. $\sum_k [(D_{0kp} - D_{0'kp} + D_{0'k0})^* Y_{kp} + 2D_{0'kp}{}^* Z_{kp}]$ is the estimated initial distance from the depot to the partition, the estimated distance to/from the disposal facility, and the exact distance between disposal facility and depot.

(25) computes TIME_p, an estimate of the total time of partition p, as the sum of the following two components:

a. $\phi \sum_i \sum_k S_{ik} X_{ikp}$ where $\sum_i \sum_k S_{ik}\, X_{ikp}$ represents the total service time and ϕ is a multiplier that estimates within-route deadhead time as a function of total service time. While θ and ϕ are highly correlated, since vehicles tend to drive faster when deadheading, the percentage of deadhead time in a route is generally less than the percentage of deadhead distance in the route. Thus, ϕ is generally smaller than θ. Since deadhead time in residential sanitation collection is no more than 5% of total service time, ϕ is generally set no larger than 1.05.
b. $\sum_k [(T_{0kp} - T_{0'kp} + T_{0'k0})^* Y_{kp} + 2T_{0'kp}{}^* Z_{kp} + DT_k\, Z_{kp}]$ represents the estimated initial time from the depot to the partition, the estimated time to/from the disposal facility, the dump time at the disposal facility, and the estimated time from the disposal facility back to the depot.

(26) computes OT_p, an estimate of the overtime associated with partition p, and (27) limits the total amount of overtime in a partition. OT_p cannot be negative due to (32).

(28) computes GEOSPREAD as the sum of the square of the distances from each arc requiring service to the seed point that the arc is assigned to. As noted before, it is generally accepted in residential sanitation collection, that more compact partitions are preferred because they are more efficient and easier to manage (although they may not have less deadheading).

(29) computes COSTVALUE, a measurement of the total cost over all the partitions. In order to keep the MPP formulation linear, \$HOUR, overtime cost/hour, and \$MILE, cost/mile, are assumed constant across all vehicle classes and not dependent on each individual vehicle class. (30) is an upper bound on the value of GEOSPREAD and (31) is an upper bound on the total cost, COSTVALUE. MAX_GEOSPREAD_ALLOWED in (30) and MAX_COSTVALUE_ALLOWED in (31) can be set from a solution to the VDA, some other solution, user defined, or equal to ∞ if no upper bound is known.

(32) stipulates that all variables are non-negative. (33) stipulates that the arc assignments to a vehicle class and a seed point (X_{ikp}) are binary. (34) stipulates that the assignment of a vehicle class to a seed point (Y_{kp}) is binary.

(35) stipulates that the number of trips to the disposal facility for a vehicle (Z_{kp}) is integer and (35) forces a trip to the disposal facility at the end of the day with a partially filled vehicle. This constraint can be relaxed if the model does not need to force a trip to the disposal facility

at the end of the workday. As noted earlier, not emptying a vehicle at the end of the day is not a common situation.

6. Solution strategies for the CARP-VSD using the MPP and the VDA

Since CARP-VSD problems in practice can involve thousands of arcs and it is difficult to solve the CARP-VSD over the entire region as a single MPP, the MPP can be used strategically as a between-route improvement procedure. A typical residential sanitation route involves around 50 arcs. Under all of the strategies proposed below, we have been successful in deriving feasible solutions to the MPP involving 500 arcs, 10 vehicles, and 4 vehicle classes in under 30 min of computer time.

As noted previously, the variables in the MPP can be fixed to represent a known feasible solution. This feasible solution can come from a solution to the steps a and b of the CA, a previous run of the MPP, the existing solution, or a user defined solution. The right hand sides of (27), (30), and (31) are computed from this solution to ensure that this solution is feasible to the MPP. The value of the objective function for the feasible solution is an upper bound to the optimal solution to the MPP since any feasible solution to the MPP defined in this manner can be no worse than the initial feasible solution. Then, some of the variables in the MPP are relaxed.

We now propose three different solution strategies called MP1, MP2 and MP3 for using the MPP as a between-route swapping procedure that improves a feasible solution to the CARP-VSD. Each of these strategies has the right hand sides of (27), (30) and (31) set to the appropriate values as dictated by the existing solution and then some of the variables are relaxed. Figure 3 displays the various components of the MPP that can be opened up to allow the MPP to improve the feasible solution. In the VDA column in Fig. 3, all variables in the mathematical formulation of the MPP are fixed to represent the VDA solution

6.1. MP1

In solution strategy MP1, the vehicle type for each seed point in the MP1 is assumed to be the same as the vehicle type used for each seed point in the existing solution. The decision variables Y_{kp} are set to 0 or 1 to represent this existing solution. The decision variables, X_{ikp} and Z_{kp}, are freed up. MP1 acts as a giant between-partition arc swapper. It appears that integer solutions to MP1 involving several thousand arcs can be found.

We always solve MP1 when examining a feasible solution from an outside source such as the VDA or manually generated. A solution to the MP1 can reduce overtime and/or improve the geographic layout of the partitions depending upon the weights used

Decision Variables	**VDA**	**MP1**	**MP2**	**MP3**
Fleet Mix Used	Fixed	Fixed	Fixed	Open
Assignment of Fleet to Seed Points	Fixed	Fixed	Open	Open
Geography	Fixed	Open	Open	Open
Number of Trips	Fixed	Open	Open	Open

Fig. 3 Variable options in the MPP

in the objective function (16). Using MP1, we found integer solutions to the 1776 arc network in a few minutes. Proving optimality was difficult because of the degeneracy in MP1.

6.2. MP2

In solution strategy MP2, a subset of the routes are extracted, the Y_{kp} variables are freed up in such as a way that more than one vehicle class can be assigned to each seed point, the seed points are the same as in the existing solution and the fleet mix in the new solution is the same as in the existing solution. For example, if 5 routes are extracted and contain 3 vehicle class 3 vehicles and 2 vehicle class 2 vehicles, then the new solution will have the same fleet mix but the vehicle class assigned to each seed point and resulting partition could be different. The assignment of arcs to partitions and the number of trips to the disposal facility are made variables by freeing up the decision variables, X_{ikp} and Z_{kp}. By setting the right hand sides of (27), (30) and (31) equal to the values in the given solution, the solution to MP2 can be no worse than the feasible solution.

A solution to MP2 can lead to a global change in the geographic assignment of arcs to partitions for all arcs being considered in MP2 and a change in the vehicle class assigned to a partition. Since MP2 is more difficult to solve than MP1, MP2 can only be used on problems involving several hundred arcs. Concurrently determining the vehicle class to assign to the seed points and the arcs to associate with each seed point complicate finding a solution.

6.3. MP3

The MP3 solution strategy is similar to the MP2 solution strategy. We separate MP2 from MP3 because MP2 only uses the existing seed points and existing fleet mix where MP3 allows for more seed points than partitions and allows for an altering of the existing fleet mix. The solution to MP3 gives the seed points to use, the vehicle class to assign to that seed point and the arcs assigned to each partition. The computational difficulty is solving MP3 and MP2 appear the same even though MP3 has more 0-1 variables.

6.4. Discussion

In many cases, in using MP1, we try to minimize GEOSPREAD by setting λ large and α small. A solution to MP1 reduces the interlacing among the partitions at no additional cost. The estimates of ϕ and θ obtained from the feasible solution may be higher than the values determined by the MP1 solution or additional cost may be incurred as the within-route deadhead time and distance may increase after the travel path for each partition is determined. We always run MP1 after an initial feasible solution is determined and after several runs of MP2 and MP3 are carried out. MP1 is very effective.

After running MP1 for the first time, MP2 or MP3 are applied to repartition a few routes at a time. We assume that the partitions not considered in this repartitioning are frozen. For example, if there is a small geographic region containing partitions with high GEOSPREAD, we can extract the subnetwork for these partitions and neighboring partitions and run MP2 or MP3 to find new partitions. We have limited MP2 and MP3 to no more than 500 arcs involving 5–15 routes. Since a feasible solution to MP2 and MP3 can be found in a few minutes of computer time (optimality is difficult to prove), we generally terminate a run after about 30 min and use the best integer solution found. To reduce the number of variables in

MP3, we only consider assigning every arc to the seed point it is currently assigned to and the 3 closest candidate seed points to the arc.

The variety of solution strategies that can be used is endless. In most cases, we use the last feasible solution to set the upper bounds in (27), (30) and (31). In many cases, the solution to MP2 and MP3 finds a different fleet mix with improved cost and better GEOSPREAD. Because the fleet mix may change in the solution to MP2 and MP3, care must be taken to ensure that the upper bound on the number of vehicles in each vehicle class is satisfied. In some cases, we relax (30) and or (31) slightly in order to reduce cost (or interlacing) while assuming slightly more interlacing (or cost). As will be seen in Section 7, this strategy is often very effective.

In the example in Section 7, we illustrate the following effective variant to MP1, MP2 and MP3. We compute the seed point of each partition in the existing solution as the centroid node of the partition. These seed points become the new candidate seed points and we run MP1 using these seed points.

As implemented, the CA is not totally automated since we manually identify the partitions that make up the network that MP2 or MP3 is going to repartition. To facilitate the visual inspection of the partitions and travel paths in a solution and to decide upon and set up the scenario to be carried out, the VDA, MP1, MP2 and MP3 have been imbedded within a Geographic Information System (GIS).

7. Example of the composite approach

In this section, the Composite Approach is applied to the 1776 arc network described earlier. The fleet mix consists of 4 vehicle classes. Site dependencies are defined for each arc in the network that has site dependencies. The Composite Approach begins by solving the IFM and the VDA to get a feasible solution and MP1 in order to swap arcs between routes. Then, MP1, MP2 and MP3 are run as needed to improve the solution where MP2 and MP3 are run on manually selected partitions from the network.

The overall objective is to reduce route interlacing significantly while not incurring any additional cost. We set λ =10,000 and α =.00001 in the objective function for all runs of MP1, MP2 and MP3, and fix the right hand sides of (27), (30), and (31) to carry out each scenario. We present the results of our analysis in Table 2 and give details in Sections 7.1–7.8. In Sniezek (2001), other examples of the CA are given. In some of these examples, we used different values of λ and α to get different tradeoffs between route interlacing and total solution cost.

The quantities in Table 2 are defined as follows:

- *Scenario* defines the scenario that was run.
- *Fleet Mix* gives the number of vehicles by vehicle class in the fleet mix in the scenario.
- *Fixed Cost* is the fixed costs in the solution.
- *# Trips* is the number of trips to the disposal facility.
- *Tip Cost* is the total cost for disposing of the refuse (i.e. tipping cost) at the disposal facility.
- *Xtra Cap* is the excess capacity or empty space in all vehicles in the solution at the end of the day.
- *Total Time* is the total time to cover all of the routes, including travel to and from the disposal facility and the depot and the within-route deadhead time (in hours). The within-route deadhead time used in computing Total Time (and total deadhead distance used in computing Total Dist) in a solution is determined by solving the Rural Postman Problem

Table 2 Summary statistics for the large problem

Scenario	Fleet mix	Fixed cost	# Trips	Tip cost	Xtra cap	Total time	OT cost	Total dist	Dist cost	Total cost	Actual GEO	Marginal GEO	Within deadhead
VDA	4-9-3-19	20201	96	9600	66.70	317.98	633.04	826.92	1653.8	32087.89	93.838	36.747	30.25
MP1	4-9-3-19	20201	94	9400	28.70	315.79	834.51	825.70	1651.4	32086.91	73.834	26.751	27.79
MP3	4-9-4-18	19957	95	9500	26.70	316.85	828.51	825.24	1650.5	31935.98	70.217	23.972	27.05
MP2	4-9-4-18	19957	95	9500	26.70	316.95	834.14	825.67	1651.3	31942.48	67.317	21.823	27.48
MP3	4-9-6-16	19469	98	9800	32.70	320.44	934.66	840.80	1681.6	31885.26	64.408	20.059	26.64
MP1	4-9-6-16	19469	97	9700	18.70	319.42	1050.46	839.16	1678.3	31897.79	57.321	17.203	27.76
MP1	4-9-6-16	19469	97	9700	20.70	319.28	1019.30	834.77	1669.5	31857.85	55.327	16.171	26.00

over each of the partitions in the solution (see Dror, 2000 for a discussion of procedures for solving the Rural Postman Problem).
- *OT Cost* is the overtime cost for all the routes (in $).
- *Total Dist* is the total distance over all routes, including travel to and from the disposal facility and the depot and the within-route deadhead distance (in miles).
- *Dist Cost* is the cost (at $2/mile) to cover the routes.
- *Total Cost.* = *Fixed Cost* + *Tipping Cost* + *OT Cost* + *Dist Cost* is the total cost of the solution.
- *Actual GEO* is a measure of the geographic spread of the route and is computed as the sum of the square of the distances from the closest node of each arc requiring service to the seed point of that arc where the seed point represents the partition to which that arc is assigned.
- *Marginal GEO* is the sum of the square of the difference in distance between each arc and the seed point it is assigned and the distance between that arc and the closest seed point of a partition where the vehicle class associated with that partition can service the arc. An arc increases the marginal GEO of the route to which it is assigned if it is closer to a feasible seed point of a different route.
- *Within Deadhead* is the total within-route deadhead mileage for the route and does not account for the deadhead distance travelling to or from the disposal facility or the depot.

7.1. VDA solution

The VDA solution gives a fleet mix of four vehicles from vehicle class 1 (capacity = 5), nine vehicles from vehicle class 2 (capacity = 8), three vehicles from vehicle class 3 (capacity = 10) and nineteen vehicles from vehicle class 4 (capacity = 16). Row 1 of Table 2 shows statistics for this solution.

7.2. First MP1 solution

The between-route swapping procedure, MP1, was run on the entire 1776 arc network. MP1 used the 35 seed points and vehicle class assigned to each of the seed points as determined by the VDA. MAX_GEOSPREAD_ ALLOWED was set to 93.838 in (30), and MAX_COSTVALUE_ALLOWED was set to 32,087.89 in (31).

CPLEX was terminated after 30 min and the results displayed in row 2 of Table 2. The best integer solution found was 1.13% from the *lower bound to the optimal solution* to MP1 (i.e. the *MIP gap* was 1.13%). This lower bound, computed by CPLEX while solving the mixed integer programming, is used throughout this section to represents the best value of the objective function that could be obtained by a feasible solution to the mixed integer program. The first integer solution to MP1 was found after 76 seconds and was 3.10% from the MP1 lower bound. For the best integer solution found, GEOSPREAD decreased from 93.838 to 73.834, total cost was reduced by $.98, an additional $201.47 in overtime was incurred, and there were 2 less trips to the disposal facility. This step shows how the strategic use of overtime and the different cost factors in the Measure of Goodness work together to produce a better solution.

7.3. First MP3 solution

We applied MP3 to two routes extracted from the solution found in 7.2. One route used a class 3 vehicle and the other route used a class 4 vehicle. These two routes consisted of 105 arcs, had severe interlacing and one of the routes had a high GEOSPREAD value of 8.22. MP3 solved this problem to optimality in 12 seconds.

The MP3 solution changed the fleet mix to two class 3 vehicles, reduced interlacing from 9.937 to 6.320, within-route deadhead from 2.79 miles to 2.05 miles, and cost by $150.93. This cost reduction was due to lower capital costs since a class 3 vehicle is less expensive than a class 4 vehicle. Variable cost increased slightly due to an additional trip to the disposal facility. These results are displayed in row 3 of Table 2.

7.4. First MP2 solution

MP2 was then applied to the 220 arcs making up the 4 class 3 vehicle routes in the MP3 solution analyzed in 7.3. The optimal solution was found in 70 seconds. GEOSPREAD dropped from 15.441 to 12.524, Marginal GEOSPREAD decreased from 7.603 to 5.454, within-route deadhead distance increased from 5.67 miles to 6.1 miles (after generating the travel paths in order to compute within-route deadhead distance), total cost increased from $3,416.28 to $3,422.77 (despite the improved geographic compactness), and the fleet mix was unchanged. The results for this solution are given in row 4 of Table 2.

This solution illustrates the situation noted earlier that cost can increase when the travel paths are formed since within-route deadhead time and distance are estimated in the solutions to MP1, MP2 and MP3. The user can make a decision at this point to disregard this solution and revert to the previous solution. We decided to continue this analysis using the MP2 solution because of the substantial decrease in GEOSPREAD.

7.5. Second MP3 solution

MP3 was run on a 219 arc subnetwork. This subnetwork consisted of the arcs from 2 class 3 vehicle partitions and 2 class 4 vehicle partitions in the current solution. These partitions were selected because the class 3 vehicle partitions had considerable geographic disconnectivity and interlaced with the arcs from the 2 class 4 vehicle partitions.

The optimal solution, found in 60 s, changed the fleet mix. The two class 4 vehicle partitions became class 3 vehicle partitions, GEOSPREAD and marginal GEOSPREAD decreased from 8.921 to 5.809 and 3.23 to 1.243 respectively, within-route deadhead decreased from 3.06 miles to 2.22 miles, and the total cost for the four routes decreased from $3,619.38 to $3,562.17. The reduction is cost was due to the reduction in capital cost because the two class 4 vehicles were replaced by two class 3 vehicles. The number of trips to the disposal facility and overtime increased.

The results of this MP3 solution are impressive. GEOSPREAD was reduced by over 30% and the cost was reduced by $60. These results are given in row 5 of Table 2.

7.6. Second MP1 solution

MP1, the global arc swapping procedure, was next applied to the current solution over the entire network to try to reduce GEOSPREAD further. MP1 was terminated after 30 min. The best integer solution, found in 5.05 min, was 7% from the lower bound. This solution is displayed in the row 6 of Table 2.

In this solution, GEOSPREAD dropped from 64.408 to 57.321, Marginal GEOSPREAD dropped from 20.059 to 17.203, within-route deadheading increased from 26.64 miles to 27.76 miles and total cost increased from $31,885.26 to $31,897.79. This run reduced interlacing by about 10% and total cost increased by about \$12.50 (out of over \$31,885). The total cost increased slightly because total mileage increased after forming the travel paths.

7.7. Third MP1 solution

MP1 was run with new seed points where the seed point for each partition was computed as the centroid node of that partition. The vehicle class for each new seed was the same as the vehicle class for the partition found in 7.6. The results were very good and can be found in row 7 of Table 2. The run was terminated after 30 min. The best integer solution, found in 6 min, was 5% from the lower bound. In the solution, GEOSPREAD was reduced from 57.321 to 55.327, marginal GEOSPREAD dropped from 17.203 to 16.171, within-route deadhead decreased from 27.76 to 26.00 miles and total cost decreased from \$31,897.79 to \$31,857.85.

7.8. Summary of the results for the problem

The results for this example can be summarized as follows.

- The fleet mix changed twice when running the MP3 between-route arc swapping procedure.
- Total GEOSPREAD over all partitions was reduced on each step from an initial value of 93.838 to a final value of 55.327 (a reduction of 41 percent). As GEOSPREAD measures the distance from each arc to its seed point, this 41% reduction in GEOSPREAD indicates that the partitions are more compact.
- Marginal GEOSPREAD over all the partitions was reduced on each step from an initial value of 36.747 to a final value of 16.171 (a reduction of 56 percent). As Marginal GEOSPREAD measures the penalty in distance from each arc to its seed point as compared to its distance to the closest seed point, this 56% reduction in Marginal GEOSPREAD indicates that the partitions are more compact.
- Within-route deadhead was reduced from 30.25 miles to 26.00 miles (a reduction of 14 percent where .14= (30.25–26)/30.25).
- The total cost of the solution was reduced about \$230 from \$32,087.89 to \$31,857.85 (a reduction of 0.7 percent).
- Overtime cost increased but was offset by savings in other costs in this example. This example illustrates the benefits of the selective use of overtime in generating a solution to the CARP-VSD.

8. Conclusions

In this paper, we present the successful use of the Composite Approach for solving the CARP-VSD. Imbedded within the Composite Approach are the following:

- Two mixed integer programs—the MPP, a between-route route improvement procedure for any feasible solution to the CARP-VSD, and the IFM, a generator of an initial fleet mix for the CARP-VSD.
- The Measure of Goodness, a multiple criterion function that allows us to generate feasible solutions to the CARP-VSD based on the daily fixed capital costs, variable costs, overtime costs and a measure of route interlacing.

The successful use of mathematical programs in the solution process is a particularly appealing aspect of the Composite Approach.

This paper addresses issues rarely discussed when analyzing vehicle routing problems. Few vehicle routing procedures consider multiple criteria in their objective functions. In the case of mixed fleet problems, however, we contend that the ordinary objective of minimizing total travel time or distance (or total deadhead time or distance) is not sensitive enough to consider all of the factors that make up a good solution. Further, the 'best' solution to virtually any mixed fleet vehicle routing problem, and, in some cases, homogeneous fleet problems, has to consider the tradeoffs between several factors such as the factors we consider in the Measure of Goodness.

The following anecdote illustrates the importance of reducing interlacing. In studying the street sweeper problem for the New York City sanitation department, Bodin and Kursh, 1978 divided an area in New York City (one of about 50 sanitation districts) into 7 sweeper routes. This solution contained considerable interlacing and little deadheading. At that time, the New York City sanitation department used 8 sweeper routes to service the area. These 8 routes had virtually no interlacing but with some deadheading. The routes found by Bodin and Kursh generated a cost savings of about 1 route/area. Bodin and Kursh estimated that a citywide implementation would save about 50 routes (or about $2,000,000 a year). Despite the savings, the NYC sanitation department rejected the solution because they believed it was easier to administer the routes that contained no interlacing.

A key problem in mixed fleet routing problems is determining an initial fleet mix. The IFM gives a very good fleet mix estimate for the CARP-VSD. This estimate is used to initiate the Composite Approach and can generally be found in less than a second. Any algorithm designed for solving mixed fleet problems requires an approach such as the IFM to find an initial fleet mix.

The between-route swapping possibilities of the MPP offer the opportunity to derive a solution that not only improves the solution but changes the fleet mix. The success of the between-route swapping procedures comprising the MPP contrasts the situation in many vehicle routing problems where the between-route swapping procedures are not effective, computationally prohibitive and almost always unsuccessful in altering the fleet mix. Altering the fleet mix while doing route improvement is critical in mixed fleet problems because the initial guess of the fleet mix may be incorrect. The MPP as applied to the CARP-VSD and its variants such as the CARP-MF may be somewhat unique in its ability to carry out between-route swapping at the level proposed in this paper.

An issue that is rarely discussed in the vehicle routing literature is the judicious use of overtime in solving a vehicle routing problem. We believe that in vehicle routing problems, as well as in other fleet planning problems, a little overtime, used very carefully, can be a good thing. A little overtime can often save a vehicle and crew. Overtime can be eliminated in the Composite Approach if overtime is not permitted.

Some of the questions that remain unanswered in determining the effectiveness of the Composite Approach in getting reasonable solutions to the CARP-VSD and its variants are as follows:

- How effective is the Composite Approach in solving the CARP-MF?
- Can the Composite Approach be used to solve the CARP-VSD with skill levels of the employees replacing vehicle/site dependencies?
- Can the Composite Approach be adapted to solve the mixed fleet node routing problem?
- Is the Composite Approach effective when the cost attributed to within-route deadheading is a large percentage of the overall cost?
- Can the MPP improve any reasonable solution to the CARP-VSD? In other words, can we modify the CA to operate as follows. The IFM generates an initial fleet mix. Steps b

of the Composite Approach is skipped. Instead, partitions are arbitrarily formed (perhaps, formed manually) for the fleet mix. Some of these partitions may not be feasible and other partitions may be underutilized. Then, the MPP is employed to attempt to improve the solution and make it feasible.

All of these questions (and others) offer opportunities for both the researcher and the user.

Acknowledgments We wish to thank Dr. Laurence Levy, Professor Michael Ball, Professor M. Grazia Speranza and the referees for their advice, encouragement and insights.

References

Assad, A. and B. Golden. (1995). "Arc Routing Methods and Applications." In M.O. Ball et al. (eds.), *Handbooks in Operations Research and Management Science*, Vol. 8, pp. 375–483.

Bodin, L., B. Golden, A. Assad, and M. Ball. (1983). "Routing and Scheduling of Vehicles and Crews: The State of the Art." *Computers and Operations Research*, 10, 65–211.

Bodin, L. and S. Kursh. (1978). "A Computer-Assisted System for the Routing and Scheduling of Street Sweepers." *Operations Research*, 26, 525–537.

Bodin, L., D. Rosenfield, and A. Kydes. (1978). "UCOST, A Micro Approach to the Transit Planning Problem." *Journal of Urban Analysis*, 15, 47–69.

Bodin, L., D. Rosenfield, and A. Kydes. (1981). "Scheduling and Estimation Techniques for Transportation Planning." *Computers and Operations Research*, 8, 25–38.

Dror, M, (ed.)., (2000). *ARC ROUTING: Theory, Solutions, and Applications*. Kluwer Academic Publishers.

Eiselt, H., M. Gendreau, and G. Laporte. (1995a). "Arc Routing Problems, Part I: The Chinese Postman Problem." *Operations Research*, 43, 231–242.

Eiselt, H., M. Gendreau, and G. Laporte. (1995b). "Arc Routing Problems, Part II: The Rural Postman Problem." *Operations Research*, 43, 399–414.

Golden, B., A. Assad, and E. Wasil. (2002). "Routing Vehicles in the Real World: Applications in the Solid Waste, Beverage, Food, Dairy and Newspaper Industries." *Chapter 10 of The Vehicle Routing Problem* In P. Toth and D. Vigo (eds.), *SIAM Monograph on Discrete Mathematics and Applications* pp. 243–286.

Sniezek, J. (2001). "The Capacitated Arc Routing Problem with Vehicle/Site Dependencies: An Application of Arc Routing and Partitioning." Ph.D. Dissertation, University of Maryland, College Park, Md. 20742.

Sniezek, J., L. Bodin, L. Levy, and M. Ball. (2002). "Capacitated Arc Routing Problem with Vehicle-Site Dependencies: The Philadelphia Experience." *Chapter 11 of The Vehicle Routing Problem*, In P. Toth and D. Vigo (eds.), *SIAM Monograph on Discrete Mathematics and Applications*, pp. 287–308.

Toth, P., D. Vigo (eds.) (2002). "The Vehicle Routing Problem." SIAM Monograph on Discrete Mathematics and Applications.