Ann Oper Res
DOI 10.1007/s10479-006-0057-0

The vehicle routing problem with demand range

Ann Melissa Campbell

Abstract We propose and formulate the vehicle routing problem with demand range (VRPDR), a new variation on the traditional vehicle routing problem. In the VRPDR, the delivery quantity for each customer i is allowed to vary from its original size d_i by an amount αd_i where $0 \leq \alpha < 1$. In adding this limited flexibility to the problem, there is potential to generate significant savings in the total distance traveled. We address issues such as bounding the impact of a given α on total distance and provide empirical results to illustrate "typical" behavior.

Keywords Vehicle routing problem · Delivery volume · Flexibility

In most real-life versions of the vehicle routing problem (VRP), the value used to represent the demand for a customer is based on an *estimate* of sales or product usage at that customer. Often this estimate is made by the customer himself at the time he places his order. Most routing algorithms treat these demand estimates as fixed values, though, and minimize the distribution costs assuming these exact amounts have to be delivered to the customers. This may result in unnecessarily high distribution costs that can be reduced by a small amount of collaboration between the vendor and the customers.

Consider, for example, a situation in which two customers far from the distribution center, but near each other, both place orders for 255 units of product while delivery truck capacity is 500. If demand is assumed to be fixed at 255, two trips have to be made to essentially the same location to satisfy the demand. On the other hand, if both customers had specified that they needed to receive a delivery in the range of 250 to 260, it would have been possible to make a single trip to satisfy demand (delivering 250 units of product to each of them). In many situations, this type of flexibility would be acceptable to the customers and may lead to significant savings in distribution costs for the vendors.

A. M. Campbell (✉)
Department of Management Sciences, Henry B. Tippie College of Business,
University of Iowa, Iowa City, IA 52242-1994
e-mail: ann-campbell@uiowa.edu

Because of such benefits, a recent trend has been the conversion to vendor managed inventory policies (VMI). This transfers inventory management and ordering responsibilities completely to the vendor. The vendor decides both the quantity and timing of customer deliveries. The customer is guaranteed he will not run out of product, and in return, the vendor gains a dramatic increase in flexibility that leads to a more efficient use of resources. The use of VMI policies has extended into a variety of industries from home furnishings (Brumback, 1995) to blue jeans (Mongelluzzo, 1998) to grocery stores (Purpura, 1997; Ross, 1998; Millstein, 1993).

Not all customers and industries that currently use a formal ordering system will be willing or able to change to such a model, though. A VMI policy requires remote monitoring of inventories either through remote telemetry or regular phone calls for readings. These readings are used to forecast usage and determine delivery times and quantities. If the product in question is not consumed at a steady, predictable rate or if the resources are not available, or are not economical, to provide these regular readings, a VMI policy will not work well for either the customer or the vendor. It can also be the case that the nature of the business is such that a customer is unwilling to transfer the delivery time and quantity decisions completely to the vendor. This can be motivated by a variety of reasons beyond a simple lack of trust in the vendor. In each of these situations, a strict VRP solution is currently the best option.

Through our experience with VMI policies, as described in Campbell et al. (1998), Campbell and Savelsbergh (2002, 2004), we observed the potential for improvement in situations where a full VMI policy can not be implemented and became interested in exploring the idea of adding limited delivery volume flexibility to the basic VRP model. More precisely, we want to evaluate the potential for improvement created by allowing a slight deviation, α, from the original delivery quantities and will refer to this as the vehicle routing problem with demand range (VRPDR). The primary contributions of this paper include the introduction and formulation of this new variation of the VRP, a discussion of the complexity of the problem, and computational results that illustrate the potential for large savings in distance even with small changes in total delivery volume. For the class of instances where all customers receive the same initial delivery quantity, we can establish lower bounds on the α required to decrease distance and describe a "stair step" behavior for the upper bound of savings as α increases.

The paper is organized as follows. We will discuss the related literature in Section 1, formally define and model the problem in Section 2, do some analysis of the problem in Section 3, describe a heuristic solution approach and provide computational results in Section 4, and discuss plans for future study in Section 5.

1. Literature review

There is little in the routing literature about adding flexibility to the VRP short of a full vendor managed inventory policy. The closest is the study of the vehicle routing problem with split deliveries (VRPSD) which will be discussed first. In some recent studies of ship routing problems, decisions concerning which harbors each ship should visit and how much to pickup/drop off at each are modeled with a delivery quantity range at the various harbors. Therefore, we also review the relevant literature in ship routing.

1.1. Vehicle routing with split deliveries

This variation of the VRP removes the restriction that each customer can receive a delivery from only one vehicle. This relaxation is needed if the total volume delivered to a customer exceeds a vehicle's capacity, but it also can be used so that better combinations of deliveries can occur. The total quantity delivered to each customer must still equal a specified quantity, though how it is parsed is not restricted. Since the size of each delivery to a customer is now a variable, the problem formulation for the VRPDR is very similar to the VRPSD. In Dror, Laporte, and Trudeau (1994), the authors propose and describe valid inequalities that strengthen the integer programming formulation for the VRPSD. Other work on split deliveries includes their earlier work in Dror and Trudeau (1989) and Dror and Trudeau (1990) and adaptations using grid distances such as in Frizzell and Giffin (1995). Belenguer, Martinez, and Mota (2000) have found new valid inequalities and have used these to establish a lower bound for the VRPSD.

1.2. Ship routing

In Christiansen and Nygreen (1998), the described ship routing problem is a multi-vehicle pickup and delivery problem with inventory constraints. The pickup and delivery points are harbors, and each harbor has a predetermined quantity interval and set of time windows. The problem is solved over a rolling horizon with an objective each period of minimizing the cost of ship travel while maintaining harbor inventories within these quantity intervals. The authors discretize the load quantities to reduce the number of potential options. They give good ideas for reducing the problem size when dealing with varying delivery volumes, but do not actually explore how this flexibility affects the resulting solutions.

2. Problem definition and formulation

2.1. Definition

Next we will formally define the Vehicle Routing Problem with Demand Range (VRPDR). The problem is clearly very similar to the traditional vehicle routing problem except for the variability in delivery quantity. The problem is defined by the following parameters:

- m = number of identical vehicles in the fleet.
- Q = the capacity for each vehicle.
- n = number of customers to which a delivery must be made. Customers are numbered 1 to n with 0 representing the depot. We will alternately refer to the set of all customers plus the depot as N.
- d_i = the original delivery quantity to customer i.
- α = the percentage of variability in the delivery quantity to each customer. In other words, the delivery quantity for each customer i will range from $(1-\alpha)d_i$ to $(1+\alpha)d_i$ so that all customers will have the same amount of "flexibility."
- C_{ij} = cost of direct travel between customers i and j. Unless otherwise specified, the cost matrix C is symmetric and the triangle inequality applies.

The VRPDR is to determine m vehicle routes such that each customer is assigned to exactly one route, and the sum of delivery quantities assigned to each route is less than or equal to Q. The primary objective that we will focus on in this paper is to minimize the total distance traveled over all m routes, but we can alternately consider an objective that considers the tradeoff between distance and total volume delivered.

2.2. Integer programming formulation

We can model the above problem with the following integer programming formulation. Variables represent decisions about which arcs are used and the quantity delivered to each customer by each vehicle:

$$x_{ikv} = 1 \quad \text{if vehicle } v \text{ travels directly from customer } i \text{ to customer } k \tag{1}$$

$$= 0 \quad \text{otherwise} \tag{2}$$

$$y_{iv} = \text{quantity delivered to customer } i \text{ from vehicle } v \tag{3}$$

The simplest objective function for the VRPDR is defined by

$$\text{minimize} \quad \sum_{i=0}^{n} \sum_{j=0, j\neq i}^{n} \sum_{v=1}^{m} C_{ij} x_{ijv}. \tag{4}$$

As discussed earlier, we may also want to consider an alternate objective function where delivery costs are minimized and delivery quantities are maximized. This requires the use of a penalty factor p to balance the dual objectives appropriately:

$$\text{minimize} \sum_{i=0}^{n} \sum_{j=0, j\neq i}^{n} \sum_{v=1}^{m} C_{ij} x_{ijv} - p \sum_{i=1}^{n} \sum_{v=1}^{m} y_{iv}. \tag{5}$$

Note that with $p = 0$, we have the first objective function. Thus, we will use the second objective in the formulation:

$$\text{minimize} \sum_{i=0}^{n} \sum_{j=0, j\neq i}^{n} \sum_{v=1}^{m} C_{ij} x_{ijv} - p \sum_{i=1}^{n} \sum_{v=1}^{m} y_{iv}. \tag{6}$$

subject to

$$\sum_{k=0, k\neq i}^{n} x_{kiv} - \sum_{j=0, j\neq i}^{n} x_{ijv} = 0 \quad i = 0, \ldots, n \quad v = 1, \ldots, m \tag{7}$$

$$\sum_{v=1}^{m} y_{iv} \geq (1-\alpha) d_i \quad i = 1, \ldots, n \tag{8}$$

$$\sum_{j=0, j\neq i}^{n} (1+\alpha) d_i x_{ijv} \geq y_{iv} \quad i = 1, \ldots, n \quad v = 1, \ldots, m \tag{9}$$

$$\sum_{i=1}^{n} y_{iv} \leq Q \quad v = 1, \ldots, m \tag{10}$$

$$\sum_{j=0, j\neq i}^{n} \sum_{v=1}^{m} x_{ijv} \leq 1 \quad i = 1, \ldots, n \tag{11}$$

$$\sum_{v=1}^{m} \sum_{i\in S, j\in \bar{S}} x_{ijv} \geq V(S) \quad S \subseteq N \backslash \{0\}, \quad |\, S \,| \geq 2. \tag{12}$$

$$x_{ijv} \in \{0, 1\} \quad i, j = 0, \ldots, n \quad j \neq i \quad v = 1, \ldots, m \tag{13}$$

$$0 \leq y_{iv} \leq (1+\alpha)d_i \quad i = 1, \ldots, n \quad v = 1, \ldots, m \tag{14}$$

In the above, the flow balance constraint Eq. (7) is the same as in a typical VRP formulation. The total delivery volume for each customer is forced above the new minimum level in Eq. (8), and the x and y values are connected in Eq. (9). The truck capacity for each vehicle is imposed via the y variables in Eq. (10). The formulation, thus far, theoretically allows a delivery to be split and served by more than one vehicle if it is cheaper, so Eq. (11) restricts service to one vehicle. Subtour elimination constraints for the VRPDR are similar to corresponding constraints for the VRP and VRPSD. Set notation is often used to express such constraints compactly, as done here in Eq. (12). We let S represent each subset of N containing at least two customers, where $V(S)$ is the minimum number of vehicles required to serve the customers in each subset S feasibly. The value of $V(S)$ can be found by solving a bin-packing problem (exactly or heuristically). It is straightforward to prove that the subtour elimination constraint in Eq. (12) which is valid for the VRPSD (Dror, Laporte, and Trudeau, 1994) is also valid for the VRPDR and is sufficient to remove any subtours. Since $V(S)$ is based on minimum delivery volumes and deliveries in the VRPDR are allowed to exceed $(1-\alpha)d_i$, there can potentially be a big gap between $V(S)$ and the actual number of vehicles required for a set of customers. Equation (15) is a relaxation of Eq. (12) when delivery volumes are fixed, but with variable delivery quantities, Eq. (15) can actually create a stronger cut depending on the values involved.

$$\sum_{v=1}^{m} \sum_{i \in S, j \in \bar{S}} x_{ijv} \geq \sum_{i \in S} \sum_{v=1}^{m} \frac{y_{iv}}{Q} \quad S \subseteq N \backslash \{0\}, \quad | S | \geq 2 \tag{15}$$

Thus, in a solution scheme where violated subtour elimination constraints are added to a reduced formulation, it might be helpful to add both types of constraints. Variable definitions for x and y variables are included in Eq. (13) and Eq. (14).

3. Analysis

In this initial study, we are interested in evaluating the impact of α on the total distance ($p = 0$) required to serve a set of customers. By focusing strictly on distance, we are effectively conducting a study of the sensitivity of the vehicle routing problem to an added amount of flexibility. In this section, we identify some of the new questions related to this sensitivity and determine some bounds based on α for a class of instances.

3.1. Complexity issues

It is well known that the traditional VRP is an NP-Hard problem. It should be quite obvious that the VRPDR is also NP-Hard since the VRP is a special case of the VRPDR with $\alpha = 0$. We are interested in whether we can still say something, though, about the value of α needed to create some change in total distance or if we can find a limit on how a given value of α may reduce the total distance. Such information would be helpful in determining whether it is worthwhile for a vendor to negotiate such a change in policy with customers.

We will refer to the problem of finding the minimum α required to reduced the distance objective from the value associated with the original VRP as *MINALPHA*. A polynomial algorithm for this optimization problem is unlikely to exist because the following related decision problem is NP-Hard:

FIND-ALPHA: Is there an α, $0 < \alpha \leq a$, such that there is a feasible routing solution for a set of n customers with delivery quantities $(1-\alpha)d_i$, m vehicles of capacity Q, and cost matrix C of total cost at most k?

Theorem 1. *The FIND-ALPHA problem is NP-Hard.*

Proof: This problem is NP-Hard with a reduction from the traveling salesman problem. To determine if a TSP with cost matrix C for n cities has a solution with cost at most l, we can create an instance of FIND-ALPHA with delivery volumes for all customers equal to $\frac{Q}{n}$, $k = l$, $m = 1$, and $a = .5$ (a can actually have any value between 0 and 1). This choice of delivery quantity is sufficient for all customers to fit on one tour, so α will not change the possible solutions. If we can find the solution to FIND-ALPHA, then we will have determined whether or not there is a TSP tour with length at most l. Restricting distances in FIND-ALPHA to be symmetric and obey the triangle inequality does not change the complexity since the TSP is well known to be hard under these conditions. □

Even if it is hard to determine the minimum α required to yield improvement, it is polynomial to determine if improvement is possible in several circumstances including:

Theorem 2. *If the triangle inequality holds strictly and $\sum_{i=1}^{n} d_i > Q$, an improving α exists with value $\leq .5$.*

Proof: If $\sum_{i=1}^{n} d_i > Q$, then there are at least two routes in the optimal solution to the original VRP. If $\alpha = .5$, then any two routes can be combined ($\alpha = .5$) for a savings in distance because of the strict triangle inequality assumption. □

We can make more definitive statements about the impact of α when we consider the special case where all deliveries (d_i) are the same size before flexibility is added.

3.2. All deliveries the same size

3.2.1. Minimum α required

When all deliveries are the same size, we can create meaningful lower bounds for the size of α required to create a reduction in the distance required. No improvement in total distance can occur until the number of deliveries on a route can increase. A value of α less than this will not allow any exchanges of customers between routes that were not already feasible. This does not mean we need to change the number of routes, just the number on the routes for an improvement to be possible.

We can determine the minimum α required for improvement once we compute the maximum number R on a route currently:

$$R = \left\lfloor \frac{Q}{d_i} \right\rfloor \tag{16}$$

Fig. 1 Three routes better than two

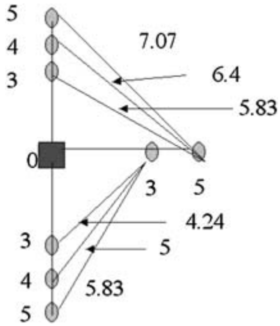

Set $\hat{R} = R + 1$ and then we can derive the minimum α:

$$\alpha = 1 - \frac{\frac{Q}{\hat{R}}}{d_i} \tag{17}$$

Note that the computation of α is dependent only on the truck capacity and the customer delivery volume and not dependent on the number of customers.

Though the above defines how to compute the lowest value of α such that improvement is possible, it unfortunately does not guarantee that the objective value will decrease with that particular α choice. This is not surprising since FIND-ALPHA is still NP-Hard with identical delivery quantities. We were surprised, though, at the difficulty of finding a value of α that was sufficient to guarantee a decrease in total distance. The results were rather non-intuitive due to the following:

Theorem 3. *If all customers receive deliveries of the same size, it is not always optimal to assign the customers to the fewest number of routes.*

Proof: The proof is by example, specifically Fig. 1. The numbers on each axis indicate the distance from the origin/depot, and these values were used to compute Euclidean distances between the customers. Assume that all customers have a delivery quantity of 1, and truck capacity equals 4. Since there are 8 customers, 2 routes with 4 stops are feasible with an minimum distance objective of 30.90($(5 + 5.83 + 3) + (5 + 7.07 + 5)$). Using 3 routes, though, yields a lower distance cost of 30 because no diagonal arcs, which are expensive, are needed. □

Due to this theorem and the above example, it should be clear that choosing α sufficient not only to increase the number on a route but to reduce the minimum number of routes by one does not guarantee a reduction in distance traveled. The more surprising result is that choosing an α that reduces the number of routes by more than one also does not guarantee a reduction in distance, as in the case of Fig. 2. There are 2 customers at each of 6 points in the plane, with a distance of 0 between each pair. Distance from each pair to its closest neighbor pair is 6. If each customer initially demands 1 unit and trucks have a capacity of 2, we need 6 routes, each costing 12, for a total cost of 72. Let α be $\frac{1}{3}$ such that 3 deliveries can now fit on a single route and not violate capacity restrictions. Now a total of 4 routes (12/3) are required instead of 6 (12/2). An optimal set of routes with 4 trucks is indicated by the darkened lines in Fig. 3. Each route costs 18 summing to a total cost of 72, representing no

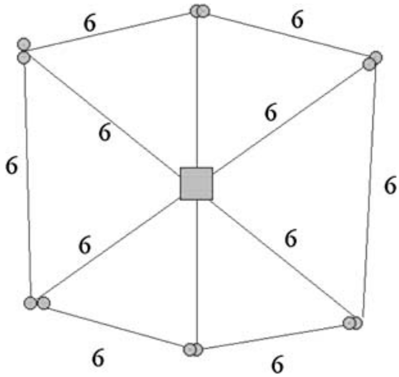

Fig. 2 Six routes

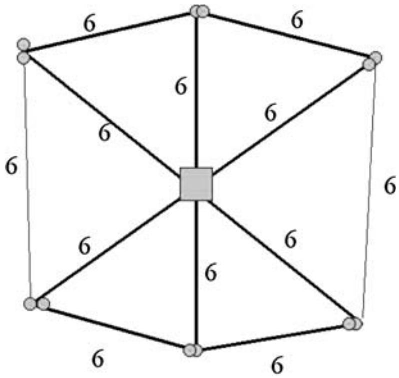

Fig. 3 Six routes same as four routes

improvement in distance even though 2 fewer trucks are required. Based on the geometry of this problem, we make the following conjecture which we have not yet been able to prove:

Conjecture 1. With Euclidean distances, if all customers receive deliveries of the same size, a decrease in total distance is guaranteed when the choice of α allows the minimum number of routes to be decreased by 3.

Even if the conjecture can be proved, the fact that it may take such a potentially significant reduction in delivery quantity to insure improvement indicates that there may be many instances where additional flexibility is not as valuable as we had expected. Note, though, it is likely that instances of the VRPDR with varying delivery quantities are more likely to be responsive to low values of α because small changes in volume can make certain beneficial swaps feasible without necessarily changing the number on the routes. The computational experiments will test this theory.

3.2.2. Upper bound on the impact of α

Even though an α sufficient to decrease the number of routes may not guarantee a reduction in costs, in many instances it will decrease the distance required. In fact, computing the possible reduction in the number of routes required along with the maximum percent of the total cost occupied by these routes is the basis for the computation of the upper bound on improvement. The simplest example is the case with $d_i = Q\ \forall i$. No more than one customer can fit on a route until $\alpha = 50\%$. At that value, the best scenario is to divide the number of routes by 2 (and costs by 2) which can result in a corresponding distance savings of at most 50%. Each time α is increased to a value that can further reduce the number of routes, we can compute a new maximum savings relative to the case with $\alpha = 0$. The computation of savings is based on the extreme scenario that all customers are co-located.

Theorem 4. *If all customers i have identical* $d_i = Q$*, then the maximum savings from using* $\frac{k-1}{k} \le \alpha < \frac{k}{k+1}$ $(k = 2, \dots, n)$ *is* $\frac{k-1}{k}100\%$.

These savings represent the largest change in distance resulting from increasing the maximum number on a route from 1 to k.

We can generalize the above result for any initial delivery volume, still assuming the initial delivery volume is the same for all customers. If the initial volume d^* for each customer is such that $\frac{Q}{b} < d^* \le \frac{Q}{b-1}$ for any integer $b \ge 2$ the following holds:

Theorem 5. *For* $\frac{k-1}{b+k-2} < \alpha \le \frac{k}{b+k-1}$ $(k = 1, \dots, n - b + 1)$*, the maximum savings is* $\frac{k}{b+k-1}100\%$.

This represents the savings from increasing the maximum number on a route from $b - 1$ to $b + k - 1$. Once $k = n - b + 1$, no more savings are possible since all n customers are on the same route.

Both of these equations demonstrate that the upper bound on the change in distance has a "stair step" behavior with respect to the choice of α. The maximum savings stays the same between ranges of α values and increases when the maximum number on a route, k, is incremented. Note that these formulas indicate the potential for a one-to-one relationship between the reduction in delivery volume and the resulting percentage change in distance for boundary values of α. In the computational experiments, we will try to identify how the average improvement compares to this.

4. Computational experiments

As discussed in Section 3, certain instances provide little to no improvement with high values of α where other instances, such as the ones used in creating the upper bounds, respond well to α. To get an idea of the average savings resulting from such a change in policy and see how these savings change with increasing α, we generated several instances of the VRPDR with different characteristics. Before presenting the results, we will briefly present the methodology used to solve these instances.

4.1. Insertion heuristic

We can solve the VRPDR, as defined in Section 2.2, exactly with tools such as CPLEX that specialize in solving integer programs. To solve large instances of the problem in seconds, we will instead use modified heuristics developed for the VRP. We will use insertion algorithms, similar to the ones used in our earlier work (Campbell and Savelsbergh, 2004), because their virtues include speed and efficiency, as well as the quality of the solutions constructed. Without increasing the algorithmic complexity of the basic insertion heuristic ($O(n^3)$), we can modify it to keep track of the maximum delivery quantity for each vehicle in addition to the current delivery sequence on each vehicle. This will allow us to compare how much of the total possible discount in delivery volume is required to achieve the corresponding reduction in distance even with $p = 0$. Our hypothesis is that a reduction in delivery volume for certain key customers creates most of the savings in distance, so comparing these values should help in evaluating this.

In the insertion algorithm, the m routes are created in parallel. The cost of an inserting j between $i-1$ and i is approximated by the standard insertion criterion in 18:

$$\text{insertion cost} = C_{i-1,j} + C_{j,i} - C_{i-1,i} \tag{18}$$

At each iteration, we evaluate whether any of the customers can not be inserted feasibly into one of the routes that already has at least one customer assigned to it or if it is cheaper for any of them to be visited by one of the as of yet unused m vehicles. If any customers satisfy this criteria, we assign one of them to an unused vehicle to initiate a new route. Otherwise, the customer with the cheapest insertion cost is inserted into the current schedule. To improve the solution, a 2-OPT procedure, which was first introduced in Croes (1958), is used to locally improve the solution after the initial route construction is complete.

4.2. Results

To develop insights into the typical impact resulting from increasing flexibility by α, we tested increasing α values in a variety of contexts. The tables presented here are representative of a large number of experiments. In each table, we provide the following information:

- n = Number of customers in each data set
- α = discount in delivery volume allowed
- Q = vehicle capacity
- *AvgVeh* = average number of vehicles required
- *AvgDel* = average total quantity delivered
- *AvgObj* = average final objective value (distance)
- *%Decrease* = average percent decrease in objective from same experiment but with $\alpha = 0$

Each table provides the average results from 10 data sets of 100 customers spread uniformly over a grid of dimension 50 by 50 units for each experiment.

In Table 1, all customers have the same initial delivery volume of 10, and α is increased along with the vehicle capacity Q. As predicted by Theorem 5 we see some of the "stair step" behavior in the objective values. The total distance values remain stable before changing and then may change significantly, as exemplified by the average results when $Q = 100$. We also see that even when selecting customers to insert based only on distance and not volume, we are able to realize the improvements in distance without the full decrease in volume. For example, with vehicle capacity 300, distance reduces from 537.33 to 485.46 with an α of .20, but total volume reduces only 10%. Not surprisingly, the delivery volume "stair steps" to match the change in total distance. The changes in the objective function are significant, but definitely not one-to-one with α. Thus, an analysis of the revenues and all relevant costs would be needed to determine it if it worthwhile, for example, to exchange a 9% decrease in total delivery quantity to get an average savings of 8.22% in distance cost (as found with $Q = 300$). Note that with a larger vehicle capacity (e.g. $Q = 300, \alpha = .10$ vs. $Q = 100, \alpha = .10$), changes in the objective occur sooner, but the savings are not always larger than with a smaller vehicle capacity (e.g. $Q = 100, \alpha = .20$ vs. $Q = 300, \alpha = .20$).

If we let delivery quantities among the 100 customers vary randomly between 1 and 20, as in Table 2, we see less "stair stepping" in the total distance and delivery quantity. There is much greater improvement from using small values of α with small vehicle capacities than with uniform initial delivery quantities. We can also see that many of these distance savings

Table 1 50 by 50 grid, identical initial delivery quantities ($d = 10$)

n	α	Q	Avg Veh	Avg Del	Avg Obj	% Decrease
100	0.00	50.00	20.00	1000.00	1271.56	0.00
100	0.05	50.00	20.00	1000.00	1271.56	0.00
100	0.10	50.00	20.00	1000.00	1271.56	0.00
100	0.15	50.00	20.00	1000.00	1271.56	0.00
100	0.20	50.00	20.00	848.00	1118.59	12.03
100	0.00	100.00	10.00	1000.00	798.20	0.00
100	0.05	100.00	10.00	1000.00	798.20	0.00
100	0.10	100.00	10.00	911.00	763.32	4.37
100	0.15	100.00	10.00	911.50	763.32	4.37
100	0.20	100.00	10.00	848.00	747.08	6.40
100	0.00	300.00	4.00	1000.00	537.33	0.00
100	0.05	300.00	4.00	974.55	535.29	0.38
100	0.10	300.00	4.00	911.00	493.16	8.22
100	0.15	300.00	4.00	900.00	486.60	9.44
100	0.20	300.00	4.00	900.00	485.46	9.65

Table 2 50 by 50 grid, varying initial delivery quantities ($1 \leq d \leq 20$)

n	α	Q	Avg Veh	Avg Del	Avg Obj	% Decrease
100	0.00	50.00	23.00	1048.51	1388.73	0.00
100	0.05	50.00	21.80	1041.58	1324.92	4.60
100	0.10	50.00	21.60	994.38	1263.12	9.05
100	0.15	50.00	21.60	948.93	1228.23	11.56
100	0.20	50.00	21.60	882.78	1164.59	16.14
100	0.00	100.00	11.70	1048.51	844.46	0.00
100	0.05	100.00	11.00	1029.29	828.14	1.93
100	0.10	100.00	11.00	986.67	791.90	6.23
100	0.15	100.00	11.00	922.33	767.10	9.16
100	0.20	100.00	11.00	871.80	739.72	12.40
100	0.00	300.00	4.00	1048.51	534.24	0.00
100	0.05	300.00	4.00	1022.17	528.56	1.06
100	0.10	300.00	4.00	974.56	516.58	3.31
100	0.15	300.00	4.00	926.78	500.90	6.24
100	0.20	300.00	4.00	892.86	491.60	7.98

are achievable with only minor changes in delivery quantity. For example, the average 9.05% change in delivery costs associated with $\alpha = .10$ and vehicle capacity of 50 requires only an average 5% reduction in total delivery volume. This indicates that our hypothesis appears to be correct in that the savings occurs from a small number of key changes in the routing.

5. Conclusions and future research

From this initial study of the impact of added flexibility in demand quantities, we find many interesting questions and promising computational results. Flexibility appears to on average have a better and faster payoff when there is a wide variety in the initial delivery quantities requested by customers. Across the scenarios, the savings in distance costs are almost always less than the associated α value but still represent significant improvements. In many cases, only a fraction of the potential reduction in total delivery volume is required to create the improvements. An interesting further study would be to identify which customers regularly incur the largest delivery reduction in these solutions. If the same customers fall into this classification over and over, it might be worthwhile to try to characterize these customers and evaluate the impact of negotiating agreements for added flexibility for just this subset of customers. As relationships between vendors and customers become more collaborative, the opportunities for such flexibility are increasing. It is important to study how and where these opportunities can best be exploited.

Acknowledgments This research was partially funded by the National Science Foundation under Award Number DMI 02-37726.

References

Belenguer, J.M., M.C. Martinez, and E. Mota. (2000). “A Lower Bound for the Split Delivery Vehicle Routing Problem.” *Operations Research*, 48(5), 801–810.

Brumback, N. (1995). “Rubbermaid’s New System Scores.” *Home Furnishing Network*, 69, 11.

Campbell, A., L. Clarke, A. Kleywegt, and M. Savelsbergh. (1998). “Inventory Routing.” In: T. Crainic and G. Laporte (Eds.), *Fleet Management and Logistics*, pp. 95–112. Kluwer Academic Publishers.

Campbell, A. and M. Savelsbergh. (2004). “A Decomposition Approach for the Inventory Routing Problem.” *Transportation Science*, 38, 488–502.

Campbell, A. and M. Savelsbergh. (2004). “Efficient Insertion Heuristics for Vehicle Routing and Scheduling Problems.” *Transportation Science*, 38, 369–378.

Campbell, A. and M. Savelsbergh. (2002). “Inventory Routing in Practice.” In: P. Toth and D. Vigo (Eds.), *The Vehicle Routing Problem*, pp. 309–329. SIAM Monographs on Discrete Mathematics and Applications.

Christiansen, M. and B. Nygreen. (1998a). “A Method for Solving Ship Routing Problems with Inventory Constraints.” *Annals of Operations Research*, 81, 357–378.

Christiansen, M. and B. Nygreen. (1998b). “Modeling Path Flows for a Combined Ship Routing and Inventory Management Problem.” *Annals of Operations Research*, 82, 391–412.

Croes, G.A. (1958). “A Method for Solving Traveling Salesman Problems.” *Operations Research*, 6, 791–812.

Dror, M., G. Laporte, and P. Trudeau. (1994). “Vehicle Routing with Split Deliveries.” *Discrete Applied Mathematics*, 50, 239–254.

Dror, M. and P. Trudeau. (1989). “Savings by Split Delivery Routing.” *Transportation Science*, 23, 141–145.

Dror, M. and P. Trudeau. (1990). “Split Delivery Routing.” *Naval Research Logistics*, 27, 383–402.

Frizzell, P.W. and J.W. Giffin. (1995). “The Split Delivery Vehicle Scheduling Problem with Time Windows and Grid Network Distance.” *Computers and Operations Research*, 22, 655–667.

Laporte, G. and Y. Nobert. (1985). “A Branch and Bound Algorithm for the Capacitated Vehicle Routing Problem.” *OR Spektrum*, 5, 77–85.

Millstein, M. (1993). “Giant Smoothing out Supply Flow.” *Supermarket News*, 43, 15.

Mongelluzzo, B. (1998). “Shippers let Vendors Manage the Stock: Wal-Mart’s Suppliers Share in Databases.” *Journal of Commerce and Commercial*, 417, 12A.

Purpura, L. (1997). “Vendor-Run Inventory: Are its Benefits Exaggerated?.” *Supermarket News*, 47, 59–60.

Ross, J. (1998). “HEB Project Leads Expansion of Vendor Managed Inventory Programs.” *Stores*, 80, 46–47.

Sierksma, G. and G.A. Tijssen. (1998). “Routing Helicopters for Crew Exchanges on Off-Shore Locations.” *Annals of Operations Research*, 76, 261–286.