I. COMBINATORIAL OPTIMIZATION

A COMPUTATIONAL COMPARISON OF ALGORITHMS FOR THE INVENTORY ROUTING PROBLEM

M. DROR

Ben-Gurion University of the Negev, Israel

M. BALL*

College of Business and Management, University of Maryland, College Park, MD 20742, USA

and

B. GOLDEN

College of Business and Management, University of Maryland, College Park, Maryland 20742, USA

Abstract

The inventory routing problem is a distribution problem in which each customer maintains a local inventory of a product such as heating oil and consumes a certain amount of that product each day. Each day a fleet of trucks is dispatched over a set of routes to resupply a subset of the customers. In this paper, we describe and compare algorithms for this problem defined over a short planning period, e.g. one week. These algorithms define the set of customers to be serviced each day and produce routes for a fleet of vehicles to service those customers. Two algorithms are compared in detail, one which first allocates deliveries to days and then solves a vehicle routing problem and a second which treats the multi-day problem as a modified vehicle routing problem. The comparison is based on a set of real data obtained from a propane distribution firm in Pennsylvania. The solutions obtained by both procedures compare quite favorably with those in use by the firm.

Keywords and phrases

Vehicle routing, inventory, graph algorithms, heuristics.

*Part of this work was performed while this author was visiting the University of Waterloo.

1. Introduction

In this paper, we define an *inventory routing problem* (IRP) and present and test solution schemes for it. The IRP involves a set of customers, where each customer has a different demand on each day. For example, each customer uses a commodity such as heating oil or methane at an estimated consumption rate. Each customer possesses a known capacity (for example, the size of his tank to hold home heating oil). The objective is to minimize the annual delivery costs, while attempting to ensure that no customer runs out of the commodity at any time. The impetus behind this distribution system is the importance of maintaining a sufficient supply of inventory at a customer's location. The optimization problem which is actually solved, is the problem of minimizing a 'cost' for assigning customers to specific days for servicing, plus the delivery cost for a period of, say, a week or two weeks. We refer to the actual time period for which the delivery problem is solved as the 'planning period'. The problem which we formulate and solve reflects an attempt to minimize the annual delivery costs, which may not be the same as minimizing the delivery costs over each planning period separately. Throughout the analysis of this problem, we assume that if a customer runs out of inventory, he is replenished immediately by emergency service. In addition, we assume that each customer consumes more than his fixed tank capacity can hold during the course of a year; thus, in order not to run out of the commodity, a customer requires a number of replenishments annually. On a given day, the customer's demand is defined as the capacity of the tank minus the stock on hand. This implies a replenishment policy under which a customer's tank is always filled up when he is serviced. Consequently, if the cutomer's consumption rate is known, the amount delivered is determined by the day of delivery.

All the customers are served from one central depot at which the vehicles start and end all their routes. This central depot serves both as a garage and as a refilling facility for the fleet of vehicles. The staffing of drivers and number and type of vehicles to be used on a given day are given parameters in our problem formulation, and their cost is considered as constant. We note that while the IRP is an operational problem, directly addressing the day-to-day activities, the questions pertaining to staffing and composition of the fleet are medium-range tactical decisions. We do not address the tactical questions here, but rather the day-to-day operation of an inventory routing system. We assume that no delivery time windows are specified; thus, the replenishment can be performed at any time during the workday.

There are two major steps in the development of a detailed solution procedure for this annual horizon optimization problem: In the first step, we reduce a problem of minimizing annual distribution costs to a problem of optimizing over the planning period. A procedure is developed to select the customers for the distribution problem which is actually solved for each planning period. In the second step, the limited time horizon distribution problem is modeled and the algorithmic procedures that solve this model are developed.

The problem reduction for the IRP (step 1) is described in detail in [9]. The principal content of that paper is an analysis leading to a single period problem formulation and the definition of the data required for the formulation. In this paper, we address the issue of solving the model. In particular, we define and compare approximate algorithms for this problem.

In sect. 2 of this paper, we describe the single planning period optimization problem. Based on two formulations, we outline two solution strategies. The first involves an allocation/routing approach where, in step 1, customers are assigned to days and, in step 2, a routing problem is solved for each day. The second strategy involves the enhancement of a vehicle routing algorithm to solve the allocation and routing problems simultaneously. Sections 3 and 4 describe the two solution schemes and present computational results for them. All computational studies make use of data derived from a real problem environment. Section 5 examines the sensitivity of the approach to parameter changes, and sect. 6 compares these approaches to other approaches and to solutions used in a real problem setting.

2. Problem description and algorithmic approaches

The problem we are attacking is a multi-period problem. The objective is to minimize the long-term distribution costs. For several reasons, most notably (i) the stochastic nature of the demand data, and (ii) the great size of long-term problems, we repeatedly solve a problem over a short planning period, e.g. one week. In another paper [9], we address the problem of reducing a long-term problem to the single time period. Reference [9] gives a refinement of the model presented in [8]. The values used for the computational results presented here are those defined in [8]. In this paper, we use certain parameters and conditions defined in [8] and [9]. The parameters and conditions are only briefly discussed here so that we may concentrate on solving the problem over the planning period.

We start with a set of customers, where each customer has an initial inventory level. Two assumptions are that when we visit a customer, we will always fill that customer's tank to capacity and we only visit customers once during a planning period. We now define:

M = the customer set,

ND = the number of days in a planning period,

d_{ik} = the demand of customer i on day k.

The d_{ik}'s are simply the difference between the customer's tank size and inventory level. Thus, they depend on the customer's usage rate, tank size, and initial inventory level.

In this paper, we treat the d_{ik}'s as deterministic values. In reality, of course, they are stochastic. If we do not schedule a customer for delivery on day k and that customer stocks out on that day, then a special delivery is made to that customer at a very high cost. Thus, the longer we wait to make a delivery, the greater the probability of a stockout and the higher the expected stockout cost. On the other hand, the earlier we deliver to a customer, the more often we would deliver to that customer over the course of a year. Thus, the stockout cost increases with k, whereas the 'long-term delivery' cost decreases with k. In [9] we give an explicit description of these costs, sum them and compute a minimum. We define the customer inventory level on the day that achieves this minimum as our 'safety stock' level, i.e. treating the d_{ik}'s as deterministic, we always deliver to a customer before that customer reaches his safety stock level. Thus we define

$\bar{d}_i$ = (customer i's tank size)
− (the safety stock level for customer i).

$\bar{M}$ = $\{i : d_{ik} \geqslant \bar{d}_i$ for some k, $1 \leqslant k \leqslant \mathrm{ND}\}$
= the set of customers we must deliver to during the planning period,

n_i = $\max\{k : d_{ik-1} < \bar{d}_i\}$ for $i \in \bar{M}$
= the last day on which we can deliver to a customer in set $\bar{M}$.

It is entirely possible that we may wish to deliver to customers not in $\bar{M}$. For example, if customer $i \notin \bar{M}$ was a neighbor of customer $j \in \bar{M}$ and customer i's tank was 1/4 full, then it would probably make economic sense to deliver to i. For similar reasons, it could easily make economic sense to deliver to a customer $i \in \bar{M}$ prior to the day n_i. In [9], to quantify these effects, we associate a cost with each customer's ending inventory level. That cost reflects the future costs associated with the inventory, i.e. if the closing inventory was high, then the future costs would be low since we could wait a long time to deliver to that customer, whereas if the closing inventory was low, then the future costs would be high. Two parameters were defined:

c_{ik} = (customer i's 'future costs' if a delivery is made on day k)
− (customer i's 'future costs' if a delivery is made on day n_i) for $i \in \bar{M}$,

g_{ik} = (customer i's 'future costs' if no delivery is made)
− (customer i's 'future costs' if a delivery is made on day k) for $i \in M - \bar{M}$.

It is clear that if g_{ik} is very small, then it would never be worthwhile to deliver to a customer during the current planning period. In [8], based on a comparison between g_{ik} and an estimate of the minimum marginal delivery cost, we define

M' = those customers in $M - \bar{M}$ that could possibly be included in a solution.

The preceding discussion gives the necessary inputs to a single period problem that reflects long-term considerations. We now can define our single period routing problem. Let

a_{ij} = the cost of servicing customer i and then traveling to customer j,

W = the number of vehicles,

q = the vehicle capacity,

i^* = the depot.

Before defining the IRP, we define two well-known problems. The *traveling salesman problem* over customer set $S \subseteq M$ is the problem of finding a least-cost tour through all members of S. Its cost function is (see [13]):

$$\mathrm{TSP}(S) = \Big\{ \min a_{i^*\sigma(1)} + \sum_{j=2}^{|S|} a_{\sigma(j-1)\sigma(j)} + a_{\sigma(|S|)i^*}, \quad \sigma(j) \text{ is an ordering of the members of } S \Big\}.$$

The day k *vehicle routing problem* over customer set S [7] is

$$\mathrm{VRP}(S,k) \equiv \text{minimize} \quad \sum_{w=1}^{W} \mathrm{TSP}(S_w)$$

$$\text{subject to} \quad \sum_{w=1}^{W} x_{iw} = 1 \qquad \text{for } i \in S, \tag{1}$$

$$\sum_{i \in S} d_{ik} x_{iw} \leqslant q \qquad \text{for } w = 1, 2, \ldots, W, \tag{2}$$

$$S_w = \{i : x_{iw} = 1\} \qquad \text{for } w = 1, 2, \ldots, W, \tag{3}$$

$$x_{iw} \in \{0, 1\} \qquad \text{for all } i, w. \tag{4}$$

The assignment variable x_{iw} indicates whether or not customer i is assigned to vehicle w. Constraint (1) ensures that each customer is assigned to exactly one vehicle, and constraint set (2) ensures that the vehicle capacity q is met. This is essentially the formulation used by Fisher and Jaikumar [12] to motivate their generalized assignment heuristic.

We now present two formulations of the single period IRP. Each of these suggest different algorithmic approaches for which we will present computational results. The first we call the *vehicle assignment formulation:*

$$\text{minimize} \sum_{k=1}^{\text{ND}} \sum_{w=1}^{W} \left(\text{TSP}(S_{wk}) + \sum_{i \in \bar{M}} c_{ik} y_{iwk} - \sum_{i \in M'} g_{ik} y_{iwk}\right)$$

$$\text{subject to} \sum_{w=1}^{W} \sum_{k=1}^{n_i} y_{iwk} = 1 \quad \text{for all } i \in \bar{M}, \tag{5}$$

$$\sum_{w=1}^{W} \sum_{k=1}^{\text{ND}} y_{iwk} \leq 1 \quad \text{for all } i \in M', \tag{6}$$

$$\sum_{i \in \bar{M} \cup M'} d_{ik} y_{iwk} \leq q \quad \text{for } w = 1, 2, \ldots, W, \quad k = 1, 2, \ldots, \text{ND}, \tag{7}$$

$$S_{wk} = \{i : y_{iwk} = 1\} \quad \text{for } w = 1, 2, \ldots, W, \quad k = 1, 2, \ldots, \text{ND}, \tag{8}$$

$$y_{iwk} \in \{0, 1\} \quad \text{for all } i, w, k. \tag{9}$$

In comparing this formulation to the previous one, note that an index k corresponding to days has been added. Furthermore, the customer set is divided into $\bar{M}$ and M'. Each customer $i \in \bar{M}$ must be served by day n, and customers in M' need not be served at all. The incentive to serve customers in M' is the future savings g_{ik}. The c_{ik}'s encourage future savings as well by pushing deliveries for customers in $\bar{M}$ as late as possible.

The natural algorithmic decomposition suggested by this formulation is an assignment of each customer to a vehicle on a particular day, followed by the solution of a traveling salesman problem for each vehicle and day. An alternative decomposition would be the assignment of customers to days of the week, followed by the solution of a single day vehicle routing problem. The corresponding *day assignment formulation* is given by:

$$\text{minimize} \quad \sum_{k=1}^{ND} \left(\mathrm{VRP}(S_k) + \sum_{i \in \bar{M}} c_{ik} z_{ik} - \sum_{i \in M'} g_{ik} z_{ik}\right)$$

$$\text{subject to} \quad \sum_{k=1}^{n_i} z_{ik} = 1 \qquad \text{for all } i \in \bar{M}, \tag{10}$$

$$\sum_{k=1}^{ND} z_{ik} \leqslant 1 \qquad \text{for all } i \in M', \tag{11}$$

$$\sum_{i \in \bar{M} \cup M'} d_{ik} z_{ik} \leqslant q^* W, \quad \text{for } k = 1, 2, \ldots, \mathrm{ND}, \tag{12}$$

$$S_k = \{i : z_{ik} = 1\} \qquad \text{for } k = 1, 2, \ldots, \mathrm{ND}, \tag{13}$$

$$z_{ik} \in \{0, 1\} \qquad \text{for all } i, k. \tag{14}$$

In this formulation, z_{ik} variables assign customers i to days k. Constraint (12) is obtained by summing constraint (7) over w.

As is illustrated in fig. 1, these two formulations suggest two different 'generalized-assignment' approaches [12] to the IRP. One approach initially assigns customers

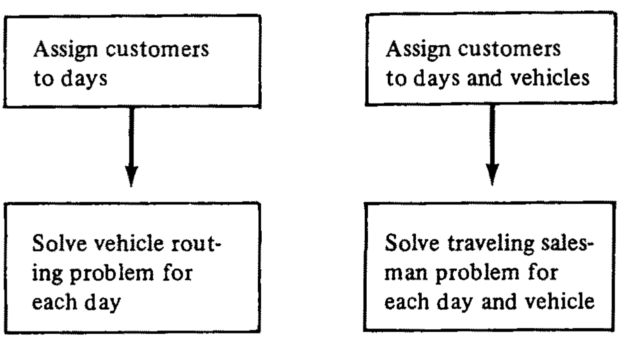

Fig. 1. Generalized assignment approaches.

to days of the week and then solves a VRP for each day; the second initially assigns customers to a vehicle on a particular day and then solves a TSP for each vehicle and day. The constraint set for the first approach is obtained by dropping constraint (13) from the day assignment formulation. The constraint set for the second approach is obtained by dropping constraint (8) from the vehicle assignment formulation. A key aspect of the generalized assignment approach is to construct a surrogate objective to replace the objective component corresponding to the constraint set dropped, i.e. $\mathrm{VRP}(S_k)$ or $\mathrm{TSP}(S_{wk})$. The approach used by Fisher and Jaikumar [12] involves the definition of seed locations placed in areas where vehicles are likely to travel. The cost of assigning a customer to a vehicle would be a function of the distance between the customer and the seed.

We chose to implement a day assignment approach. Two factors led us to this decision. First, for the day assignment approach, the generalized assignment problem is considerably easier. In particular, our solution approach involved perturbing a solution to the linear programming relaxation of the generalized assignment problem. It is well known (see Lasdon [15]) that the maximum number of non-integer values in the linear programming solution to the generalized assignment problem is twice the number of knapsack-type constraints [constraint set (7) or (12)]. For the day assignment formulation this is $2 \times \mathrm{ND}$, but for the vehicle assignment formulation it is $2 \times \mathrm{ND} \times W$.

The second factor was that for the vehicle assignment formulation there did not appear to be a natural surrogate objective other than using the same set of seeds for each day. Such a surrogate, which effectively treats each day equally, eliminates the advantages of the vehicle assignment approach. For the day assignment approach, we simply ignored the $\mathrm{VRP}(S_k)$ cost component. The effect of this was that in the assignment phase the solution procedure ignored any spatial considerations. We felt this was reasonable since on any given day, the fleet of vehicles would tend to cover a fairly wide geographic area. In addition, we applied an exchange procedure at the very end which considered exchanges between days and vehicles.

We also implemented a second algorithm which was not based on a generalized assignment approach. The second procedure viewed the IRP as a modified version of the VRP (MVRP). It involved an enhanced version of the Clarke and Wright [6] algorithm. In avoiding a day-of-the-week decomposition, it was able to consider simultaneously spatial issues, the $\mathrm{TSP}(V_{kw})$ cost, and temporal issues, the g_{ik} and c_{ik} costs, in putting together vehicle routes. The disadvantage of such an approach is that it must solve a much larger routing problem. Figure 2 illustrates the two competing approaches. Note that both are followed by a solution improvement algorithm.

We should mention that for simplicity of presentation, we left out of the preceding model a key problem characteristic. That characteristic is that individual vehicles typically handle more than one route per day. The limitation on the number of routes handled involves a constraint on the time a vehicle can be out during a day and an overtime penalty. In many problem settings, these issues could cause major modeling and algorithmic complications. Due to the characteristics of our data, we were able to handle them rather simply as follows. In nearly all cases, vehicles handled

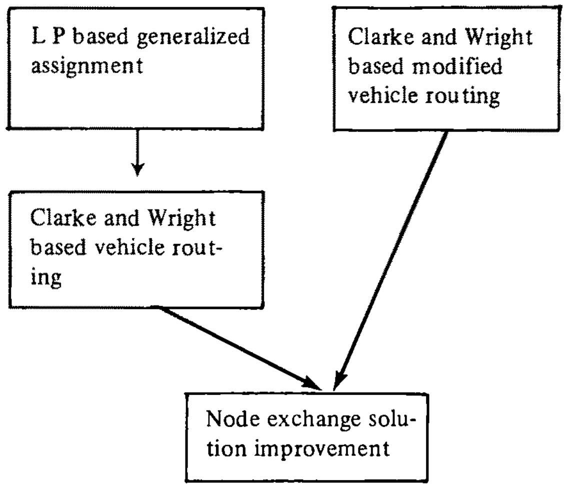

Fig. 2. Algorithms implemented.

two routes per day. Thus, in the vehicle allocation model, the number of routes handled can be effectively doubled and in the day allocation model, the capacity of a route can be doubled. We describe algorithmic implications in the next section. A second issue involves the stochastic nature of the demand and the resultant possibility of 'route failure', i.e. the vehicle runs out of product before delivering to all customers. This issue is treated using the capacity adjustment methods described in [17] and [18].

We refer the reader to [3,5,11,16] for the definition of other inventory/routing and allocation/routing problems. The problem studied by Bell et al. [3] is the closest of the three to ours. There are, however, some fundamental differences between the two problems. First, our problem has a major routing component and theirs does not. Second, the objective functions are different. Third, the two problems handle stochastic demands in dissimilar ways. The problem studied by Federgruen and Zipkin [11] is similar to but simpler than the one addresses by Bell et al.

The problem addressed by Christofides and Beasley [5] is nearly identical to the one studied by Russell and Igo [16]. In this assignment/routing problem, demand per week is deterministic and known in advance. The objective is to assign days of the week to each customer in order to minimize distribution costs over the weekly planning period.

3. Solution algorithms

As fig. 2 illustrates, two different algorithmic approaches were compared. Note that four algorithmic modules are involved. The problems they solve are:

I. Generalized assignment,
II. VRP,
III. MVRP,
IV. Improvement of IRP solution.

In subsects. 3.1, 3.2, 3.3 and 3.4, we describe the algorithm used for problems I, II, III and IV. We should note that the MVRP algorithm uses the VRP algorithm as a subroutine.

3.1. THE GENERALIZED ASSIGNMENT ALGORITHM

The generalized assignment problem we solved is

$$\text{minimize} \quad \sum_{i \in \bar{M}} c_{ik} z_{ik} \phi - \sum_{i \in M'} g_{ik} z_{ik}$$

subject to (10), (12), (14).

Note that it is a little more complicated than the simplest version of the generalized assignment problem in that the sum in (10) stops at n_i; (11) is an inequality constraint and the coefficients in (12) depend on k. This last fact implies that the linear programming relaxation can not be solved as a network flow problem. Instead, simplex-based or generalized network codes must be used.

Our approach to solving this problem was to solve the linear programming relaxation and then to apply various rounding techniques. We felt this approach would provide good solutions due to the fact that the linear programming relaxation will contain a limited number of fractions (see Lasdon [15] page 173). For our problem, the number is 2 × ND. In this case, ND is usually small, e.g. 5 or 10, compared to $|\bar{M}| + |M'|$, e.g. 75 to 200.

The rounding methods were rather complicated, so we do not describe them here but rather refer the reader to [8]. We do present, in table 1, some computational results indicating the accuracy of the solutions obtained. That is, we present the value of both the linear programming relaxation and the feasible solution obtained by rounding. We should note that due to the presence of the c_{ik}'s and $-g_{ik}$'s in the objective function, the value can be either positive or negative. These results indicate that this approach is very accurate in all but a small number of cases. In sect. 4, we present more details on the data used in this computational experiment.

3.2. THE VRP ALGORITHM

The VRP solved is more general than the one explicitly defined in sect. 2. In particular, each vehicle handles one or more routes per day. There is an explicit

Table 1

Computational results for the rounding approach to the generalized assignment problem

Explanation: this table gives four sets of sequences of twelve weekly runs; these were generated in conjunction with the other algorithms so that the solution from one week generated the problem in the next. For this data, ND = 5 and $w = 4$.

Run No.	Value of LP objective	Solution value of objective	No. of customers available in ICP formulation	No. of customers served in ICP solution
1	− 81.6	− 80.9	190	185
2	− 71.8	− 71.0	75	70
3	− 24.6	− 19.6	108	[illegible]
4	− 73.9	− 68.3	105	[illegible]
5	− 72.0	− 70.1	104	[illegible]
6	188.2	189.8	92	[illegible]
7	− 18.4	− 15.1	106	102
8	− 40.4	− 35.4	119	[illegible]
9	− 34.3	− 28.9	106	[illegible]
10	− 88.5	− 82.9	102	[illegible]
11	− 148.7	− 145.7	109	[illegible]
12	− 104.0	− 63.5	93	[illegible]
1	− 81.6	− 80.9	190	[illegible]
2	− 67.5	− 66.6	74	[illegible]
3	− 15.1	− 9.4	100	[illegible]
4	− 81.1	− 75.5	109	[illegible]
5	− 15.6	− 11.4	107	104
6	186.0	206.3	97	74
7	− 49.6	− 48.2	103	[illegible]
8	− 12.2	− 9.3	121	[illegible]
9	− 47.5	− 41.4	116	[illegible]
10	− 127.7	− 126.2	103	[illegible]
11	− 115	− 112.4	106	[illegible]
12	− 104.4	− 51.4	95	[illegible]
1	29.8	32.9	188	[illegible]
2	− 25.4	− 22.9	76	68
3	31.8	35.7	106	99
4	− 39.7	− 37.0	113	108
5	− 26.6	− 16.9	117	108
6	227.4	230.2	84	77
7	6.3	11.0	114	108
8	51.6	59.5	128	[illegible]
9	− 0.13	4.4	123	115
10	− 41.7	− 38.1	115	109
11	− 98.5	− 94.8	111	103
12	− 53.2	− 21.9	93	70

Table 1 (continued)

Run No.	Value of LP objective	Solution value of objective	No. of customers available in ICP formulation	No. of customers served in ICP solution
1	- 199.1	- 194.7	195	191
2	- 122.5	- 117.6	71	70
3	- 75.5	- 68.5	103	97
4	- 101.9	- 96.9	101	98
5	- 117.2	- 109.3	101	96
6	149.8	175.7	85	67
7	- 94.6	- 91.2	117	114
8	- 78.6	- 64.1	119	108
9	- 102.7	- 97.8	106	104
10	- 170.8	- 168.9	99	94
11	- 216.3	- 210.2	102	100
12	- 191.9	- 188.1	99	93

maximum time for each vehicle as well as a smaller preferred time beyond which overtime is paid. We solved this problem in two stages. The first stage used a modified Clarke and Wright algorithm [6,13] to generate routes and the second stage used a variant of the first fit decreasing bin packing algorithm to pack routes into vehicles.

3.3. THE MVRP ALGORITHM

In contrast to the hierarchical optimization approach described earlier, the solution procedure for the MVRP attempts to construct a 'good' solution to the inventory routing problem (IRP) without partitioning the problem into separate optimization stages. The criteria used throughout the search for a solution is the total objective function of the single period IRP. Both the cost for assigning the customer to a specific day and the cost for routing are included.

Before describing the solution procedure, it is instructive to compare the VRP and MVRP. In particular, one could view the MVRP as a VRP with $ND \times W$ vehicles. Starting with this point of view, we point out three key differences:

(1) All customers need not be served.

In the VRP, all customers are included in the equality constraints (1), whereas in the MVRP some are included in equality constraints (5) and some in inequality constraints (6).

(2) Customer demand varies with 'vehicle'.

In the VRP, the customer demand is independent of the vehicle, whereas in the MVRP it depends on which day the vehicle serves the customer.

(3) Costs depend on which vehicle services the customer.

In the VRP, costs only depend on the sequence of customers on a route, whereas in the MVRP the c_{ik} and g_{ik} terms imply a vehicle/customer cost.

The solution procedure for the MVRP, which is displayed next, uses as a sub-procedure the VRP algorithm described earlier. It initially assigns each customer for delivery on every day of the week (step 1). Then it reassigns customers to a single delivery day based on the costs from the initial solution (steps 2 and 3). Finally, it performs additional reassignments to ensure feasibility with respect to the vehicles on hand (step 4). Steps 0 through 3 are self-explanatory. Step 4 involves a rather complicated interchange procedure. For details, see [8].

MVRP algorithm

Step 0

For $k = 1, 2, \ldots, \mathrm{ND}$, set $S_k = \bar{M} \cup M' - \{i: n_i < k\}$.

Set $M^\star = \phi$.

Step 1

For $k = 1, 2, \ldots, \mathrm{ND}$,

solve $\mathrm{VRP}(S_k, k)$ using a number of vehicles (W) sufficiently large to satisfy all demand.

(We note that at the end of this step, each customer has been assigned for delivery on several days.)

Step 2

While $M^\star \neq \bar{M} \cup M'$,

For $i \in \bar{M} \cup M' - M^\star$,

compute $\mathrm{SAV}(i, k)$ = the decrease in the objective function obtained by deleting customer i from the VRP solution for day k;

Find $i^\star \in \bar{M} \cup M' - M^\star$ such that $\sum_{k=1}^{\mathrm{ND}} \mathrm{SAV}(i^\star, k)/\mathrm{ND}$ is maximum;

Find $k^\star$ such that $\mathrm{SAV}(i^\star, k^\star) = \min_{k=1,\ldots,\mathrm{ND}} \{\mathrm{SAV}(i^\star, k)\}$;

Set $S_k = S_k - \{i^*\}$ for all $k \neq k^*$; adjust the cost of the current VRP solutions accordingly;

Set $M^* = M^* \cup \{i^*\}$;

Step 3

For $k = 1, 2, \ldots, ND$,

resolve $VRP(S_k, k)$ using a number of vehicles (W) sufficiently large to satisfy all demand.

(We note that at the end of this step, each customer has been assigned for delivery on exactly one day.)

Step 4

Using a set of node interchanges/insertions/deletions, adjust the solution obtained in step 3 to make it feasible with respect to the set of vehicles on hand. During this process, periodically resolve step 3.

IMPROVEMENT ALGORITHM

After a feasible solution to the IRP is obtained, using either the assignment/routing approach or the MVRP approach, we then apply an improvement algorithm. The improvement procedure developed is categorized as a node interchange heuristic and is designed to examine a given feasible solution and to search for favorable trade-offs when interchanging customers' positions on a single route or between different routes. It attempts to exploit the costs related to the two characteristics of a feasible IRP solution:

(a) temporal considerations (i.e. customer's demand level depends on what day it is scheduled for replenishment).

(b) spatial considerations (i.e. the customer's location relative to other customers in the solution).

In addition, we form one artificial route that includes all the customers from the subset M' that are not assigned for replenishment during the planning period. A customer that belongs to the artificial route can be inserted into any other route, but only customers from M' can be interchanged with a customer on the artificial route.

The goal of this third stage is to reduce the total distribution cost for the planning period as reflected by our objective function, while maintaining route capacity constraints. The node interchange procedure used for this purpose considered moving single customers or exchanging the position of two customers either on the same route or between two different routes. The pair of routes could be on different days. Details can be found in [8] and [10].

4. Computational comparison of algorithms

The test runs were based on twelve successive weekly problems, with the output of one week serving as the input to the next week. Thus, the only data set which is the same for all eight 12-week runs is the first week's data set. The number of days in the planning period ND was five. The number of vehicles was two, but since each vehicle performed two routes per day, the value of w was set to four. The total customer data base contained several thousand entries, but the number delivered to in a week was between 75 and 200. Typical routes had between five and twenty customers. The number of days between deliveries to individual customers varied between 5 and 60. There were a few customers who required almost daily deliveries. These were not included in the assignment problems, but were manually assigned to particular days. The measure used in order to evaluate the efficiency of the distribution solution provided for each planning period is expressed in the number of inventory units delivered per hour of distribution operation. (This includes the routing times, the unloading times, and the fixed times per delivery.) This measure for the efficiency of the distribution operation is widely accepted in the heating oil delivery industry. It is considered a more objective measure than a cost-based measure because cost figures usually depend on many local and accounting factors. In addition, a solution to the IRP is not examined only for a single planning period, but for twelve dependent successively-solved weekly problems. Different solutions to this problem start with the same initial data input; however, for each period it is most probable that different specific inventory problems will be solved. Consequently, a direct cost comparison would no longer be appropriate. Nonetheless, the results obtained using this measure should be viewed with some degree of caution since it does not express the amount (or the cost) of drivers' slack time in the system. In particular, it implicitly assumes that if driver time is reduced, then costs can be reduced, either by reducing total driver pay hours or by employing drivers in other productive ways.

Essentially, the cost function that we so elaborately developed serves as a surrogate objective for our true measure for the efficiency of distribution operation as expressed in units delivered per hour.

In table 2, we present the performance of the resulting weekly solutions after the first two stages of the three-stage hierarchical approach and the MVRP approach before the improvement stage is implemented.

As table 2 makes clear, when we measure the general quality of the feasible solutions obtained by the assignment/routing approach in units per hour, the results are superior in quality to the feasible solutions obtained using the MVRP procedure.

In table 3, we present the improvement (as measured in percentage) and the final weekly units per hour which occur when we apply the improvement procedure.

As we can clearly observe from table 3, the assignment/routing approach outperforms the modified vehicle routing approach. The first solution is superior to the

Table 2

Week	Assignment/routing approach-feasible solution (units/hour)	MVRP approach-feasible solution (units/hour)
1	4410	3704
2	5452	4384
3	4899	3950
4	4441	3623
5	5050	4622
6	5408	4144
7	5311	4474
8	5248	3913
9	4617	4058
10	4845	4679
11	4852	4614
12	4861	4011

Table 3

	Assignment/routing approach		MVRP approach	
Week	% improvement	units/hour	% improvement	units/hour
1	46.6	6463	21.3	4494
2	32.3	7215	39.0	6092
3	35.6	6645	41.3	5689
4	32.8	5896	33.9	4852
5	32.4	6687	21.5	5516
6	29.4	6996	34.8	5588
7	36.5	7251	12.4	5031
8	30.8	6864	44.4	5650
9	39.3	6432	22.8	4982
10	25.7	6088	8.1	5058
11	31.3	6372	17.4	5415
12	34.0	6512	29.5	5051
entire period	33.9	6606	26.9	5278

second solution by over 25.1%. We should note that the number of replenishments performed implementing the first solution is 1104 and in the second solution this number is only 1078. As far as the performance of the improvement stage on the two different feasible solutions, we note that in seven out of twelve cases greater improvement is accomplished on the first solution procedure, and in the other five cases it is the other way around.

In developing the different procedures implemented in this paper, we did not emphasize their performance with regard to running time. Our main objective has been to test concepts without investing heavily in algorithmic efficiency as measured by CPU time on a UNIVAC 1100/82. With that in mind, one should view the running times for the procedures as a rough upper bound on their time performance. Nevertheless, in our opinion the running times for the present algorithms are quite good, and certainly their cost is affordable by most moderately-sized distribution companies. If greater emphasis were placed on coding for fast execution, these times could be cut substantially.

In table 4 we present the CPU times in seconds for two different instances of the IRP; the first week, in which we selected 190 customers, and the second week, in which we selected 78 customers in total for the sets M and M'.

Table 4

Running time in CPU seconds for the different procedures

Procedure	Week 1	Week 2
Customer selection procedure	19.70	16.28
Procedure for obtaining feasible solution for the ICP	1.35	1.23
Modified Clarke and Wright procedure	1.32	0.26
Improvement procedure	19.68	8.37
MVRP procedure	21.11	7.24

5. Sensitivity of the computational results

5.1. ALLOCATION/ROUTING APPROACH – TESTING THE TIME CONSTRAINT

When designing the solution scheme using the allocation/routing approach, we claimed that the time constraint inherent in the feasible length of a driver's workday would take care of itself when a solution to the problem was computed. Under that assumption, we did not test whether the lengths of the routes scheduled for a given day could be 'packed' into the drivers' workday. In fact, we were always able to do so. All of the solutions obtained were feasible when the daily routes were 'packed' into the drivers' workday. In a few cases, some overtime driving was required; in those cases, the overtime limit of two hours was never violated.

5.2. SENSITIVITY OF ALLOCATION/ROUTING APPROACH TO CHANGES IN CAPACITY

In this section, we present the changes in the results which are obtained when we restrict the vehicle capacity available on each day to 1.5 vehicles or 2.5 vehicles; then we compare these results to the base case of two vehicles. In the allocation/routing approach we restrict the assignment of customers on specific days by imposing a certain capacity constraint in the ICP formulation. In our base case of two vehicles available daily, we assume that each of the vehicles will be able to complete at least two routes daily; consequently, the daily fleet capacity available equals four times the capacity of one vehicle.

In order to take into account the stochastic nature of the demand, 'artificial' capacities were used. Recall that when the routes are run, only an estimate of the customers' demand is known. Thus it is possible that the vehicle would be unable to satisfy all demand. We call this situation 'route failure'. To take this possibility into account, we limit the daily capacity to 90% of the actual capacity and call it the 'artificial capacity'. Thus, in the base case, this artificial capacity equals 79 520 units. In the VRP, we use 95% of the vehicle's tank size as its capacity. This artificial capacity provides a supply buffer which with certain probability guards against route failure. For more details about artificial capacity and its relation to probability of route failure, see [17,18].

When testing the case of 1.5 vehicles available daily, we reduced the artificial capacity to 59 640 units; in the case of 2.5 vehicles available daily, we increased this capacity to 99 400 units. We call the run with 59 640 units of daily capacity run (1); run (2) is our base case and run (3) is the run with 99 400 units of daily capacity available. Run (2) and run (3) provide a solution for the entire twelve weeks. Run (1) could not solve the IRP for twelve consecutive weeks, but only for five consecutive weeks, after which it could not solve the problem for week 6. Run (1) could not find a feasible assignment for customers selected to be replenished in week 6 because the demand of the customers to be replenished on that Monday exceeded the capacity of 59 640 units. According to our rule, those customers could not be shifted for delivery to any day later than Monday. The excess demand on that Monday is 1 510 units.

For the five weekly runs generated with run (1), we tested whether the daily routes obtained can fit into a work day of 1.5 drivers, where the 0.5 driver available can work only four hours a day. The answer is yes; however, the first driver is required to work some amount of overtime.

In table 5 we present the results of runs (1)–(3).

We can not compare the results from run (1) (which only covered five weeks) with the results of run (2) and run (3). As for comparing runs (2) and (3), we can see that by increasing the available daily capacity by one vehicle route, we improved the measure of units per hour from 6 606 to 6 632.

Table 5

Week	Run (1) (units/hour)	Run (2) (units/hour)	Run (3) (units/hour)
1	6 221	6 463	5 938
2	6 972	7 215	7 352
3	6 450	6 645	7 248
4	5 733	5 896	5 801
5	6 368	6 687	6 873
6	-	6 996	7 321
7	-	7 251	6 826
8	-	6 864	6 988
9	-	6 432	6 617
10	-	6 088	6 008
11	-	6 372	6 085
12	-	6 512	6 870
entire period	6 318	6 606	6 632

6. Interpretation of computational results

Throughout this paper, we have been concentrating on minimizing the objective to the single period IRP. In the previous sections we introduced the concept of units per hour to evaluate solutions over a multi-period time horizon. One crucial solution parameter that has been ignored is the number of stockouts. We did not treat this issue explicitly in this paper, since the characteristics of stockouts were the primary input in the definition of the safety stock levels (see [8] and [9]). These safety stock levels in turn determined the n_i's which were a primary input to our model. In order to compare our solutions with solutions used in practice and solutions used by other methods, we must get an accurate estimate or bound on the number of stockouts.

In [9], the safety stock levels were derived based on (i) knowledge of the mean and variance of each customer's daily consumption, (ii) the assumption that consumption rate was normally distributed, and (iii) the assumption that consumption rates were independent between customers and days. Using these assumptions and estimates of the mean and variance based on the problem data, we obtained an upper bound on the expected number of stockouts for the entire twelve week period of 28. One could argue that none of these assumptions would be valid in realistic problem environments. However, if assumptions (ii) or (iii) were changed, then the analysis could be changed in a like manner. That is, if the consumption rate was not normally distributed or the rates were not independent, then the analysis could be appropriately adjusted. The real key to the results we obtained is explicit knowledge of the customers consumption characteristics.

As we observe in table 3, the solution scheme which has the best results for the twelve weeks of inventory deliveries is the one based on the allocation routing approach. The number of units per hour exceeds 6 600. When we examine the results obtained in the studies reported in [1,2,14], we observe that the quality of our solution far surpasses the quality of the solution reported in those studies.

In order to avoid unfair comparisons, we should point out the differences between our approach and the one used for the earlier studies. We observe that the earlier solution is based on a simulation approach with an optimization component in which the vehicle routes are constructed. The methodology differs considerably from the one proposed here. While both studies use the same forecasting procedure to estimate the distribution parameters, they differ drastically in the use they make of these parameters once they are obtained. In the earlier studies, the standard deviation of a customer's demand is used only for simulating the actual demand and for counting the number of stockouts. That is, the system does not know the form of the demand distribution, the mean, or the standard deviation for each customer when it constructs routes. In the system described in this paper, however, this detailed information is actually used in generating the daily routes. Thus we believe that the earlier results are much more robust in terms of their sensitivity to parameter changes than are the results reported here. In addition, the two studies do not react the same way to stockouts and route failures.

We should also note that while the earlier studies have substantially overlapping data bases, the data bases are not identical.

For the twelve-week period analyzed in the earlier studies, the solution proposed has 5 453 units per hour, with 42 stockouts. These results are compared with the simulated performance of the distribution company, which delivers 4 799 units per hour, with 97 stockouts. When the same runs were performed using the data base we analyzed, the results were somewhat reduced in quality; 4 620 units delivered per hour, with 54 stockouts, for the solution proposed, and 4 101 units per hour, with 60 stockouts, for the simulated performance of the distribution company.

We believe that this comparison indicates that our approach is quite effective and that explicit knowledge of customer consumption characteristics can be extremely valuable in structuring algorithmic approaches.

References

[1] A. Assad, B. Golden, R. Dahl and M. Dror, Design of an inventory routing system for a large propane-distribution firm, *Proc. Southeastern TIMS Conf.*, Myrtle Beach (1982) p. 315.

[2] A. Assad, B. Golden, R. Dahl and M. Dror, Evaluation of the effectiveness of an integrated system for fuel delivery, *Proc. 1983 Southeastern TIMS Conf.*, ed. J. Eastman, Myrtle Beach (1983) p. 153.

[3] W. Bell, L. Dalberto, M. Fisher, A. Greenfield, R. Jaikumar, P. Kedia, R. Mack and P. Prutzman, Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer, Interfaces 13, 6(1983)4.

[4] L. Bodin, B. Golden, A. Assad and M. Ball, Routing and scheduling of vehicles and crews: The state of the art, Computers & Operations Research 10, 2(1983)62.

[5] N. Christofides and J. Beasley, The period routing problem, Networks 14, 2(1984)237.

[6] G. Clarke and J. Wright, Scheduling of vehicles from a general depot to a number of delivery points, Oper. Res. 12(1964)568.

[7] G. Dantzig and J. Ramser, The truck dispatching problem, Management Science 6(1959)81.

[8] M. Dror, The inventory routing problem, Ph.D. Diss., College of Business and Management, University of Maryland (1983).

[9] M. Dror and M. Ball, Inventory routing: Reduction from an annual to short period problem, Working Paper No. 167, Dept. of Management Sciences, University of Waterloo, Ontario, Canada (1985).

[10] M. Dror and L. Levy, A vehicle routing improvement algorithm-comparison of a 'greedy' and a 'matching' implementation for inventory routing, Centre Recherche sur les Transports, Pub. No. 353 (1984).

[11] A. Federgruen and P. Zipkin, A combined vehicle routing and inventory allocation problem, Oper. Res. 32(1984)1019.

[12] M. Fisher and R. Jaikumar, A generalized assignment heuristic for vehicle routing, Networks 11, 2(1981)109.

[13] B. Golden, T. Magnanti and H. Nguyen, Implementing vehicle routing algorithms, Networks 7, 2(1977)113.

[14] B. Golden, A. Assad, M. Dror and R. Dahl, Planning a distribution system in the presence of routing and inventory subproblems, presented at ORSA/TIMS National Meeting, San Diego (1982).

[15] L. Lasdon, *Optimizing Theory for Large Systems* (McMillan, 1970).

[16] R. Russell and W. Igo, An assignment routing problem, Networks 9, 1(1979)1.

[17] W.R. Stewart, New algorithms for deterministic and stochastic vehicle routing problems, Ph.D. Diss., College of Business and Management, University of Maryland (1981).

[18] W.R. Stewart and B. Golden, Stochastic vehicle routing: A comprehensive approach, Eur. J. Oper. Res. 14(1983)371.