An Efficient Approximation Algorithm for the Fixed Routes Problem

M. W. P. Savelsbergh

Eindhoven University of Technology
Georgia Institute of Technology

M. Goetschalckx

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205

Version 2.0, September 29, 1992

This research was supported in part by the member companies of the Material Handling Research Center of the Georgia Institute of Technology.

ABSTRACT

The Fixed Routes Problem is a variation of the Vehicle Routing Problem in which the routes that have to be constructed will be operated unchanged for an extended period of time while the customer demands within that period will vary. If, in a delivery scenario, a vehicle cannot satisfy the demand of a customer it musts return to the depot for replenishment before continuing its route. The fixed routes problem can be modeled as a vehicle routing problem with stochastic demands, which in turn can be solved with a stochastic programming model with recourse.

An effective and efficient approximate algorithm to generate fixed routes is presented. It is based on the generalized assignment model for the standard vehicle routing problem. The algorithm consists of two phases. In the first phase, customer demands are considered deterministic and a standard vehicle routing problem is solved to construct an initial set of routes. In the second phase, customer demands are again stochastic and local search methods are applied to find an improved set of routes.

Computational experiments show that the algorithm is able to solve realistic problem sizes with acceptable computational times. They also indicate that the travel distance of the fixed routes being generated increases only modestly when the variance of the customer demands is increased. Furthermore, the algorithm compares favorably to the two methods for the vehicle routing with stochastic demands published in the literature.

1. INTRODUCTION

In the standard Vehicle Routing Problem (VRP) a number of vehicles are located at a single depot and must serve a number of geographically dispersed customers. Each vehicle has a given capacity and each customer has a given demand. The objective is to minimize the total cost of travel for servicing the customers. The routes are recomputed for each instance of the customer demand data. The vehicle routing problem is an important problem in the area of distribution management. Many distribution managers spend hours on the day-to-day routing and scheduling of delivery vehicles. The problem has also received extensive attention in the research literature. A recent survey of VRP literature can be found in Golden and Assad [1988].

A common phenomenon in practical distribution management is the fact that problem instances for different periods show certain similarities. In various situations, the same set of customers has to be visited every day, and it may even be possible to detect a weekly pattern. In these circumstances the question arises whether designing a set of routes to be operated unchanged over a given period of time is a viable alternative to constructing a set of routes every day based on the particular instance of that day. The problem of designing a set of routes to be operated unchanged over a given period of time is known as the Fixed Routes Problem (FRP) and has not received the same amount of attention in the vehicle routing literature as the standard VRP.

One of the primary situations that leads to fixed routes is the case where actual customer demands only become known at the customers' premises. This often occurs in the petrochemical industry and in convenience and grocery store deliveries. Another application of fixed routes arises when in part collection for just-in-time manufacturing operations.

Advantages of using fixed routes include reduced management cost, since the need to optimize routes on daily bases is eliminated; individual and predictable service towards customers; increased performance by the drivers, since they familiarize themselves with the routes and the delivery area; and the possibility to increase the efficiency at the depot, since standard depot procedures can be developed. A disadvantage of using fixed routes is the lack of flexibility in the use of the resources (vehicles and drivers) which may cause a possible increase in the delivery cost and distance compared to daily routing.

The main difficulty in the design of fixed routes is the need to address demand fluctuations, i.e., demands are no longer deterministic but have become stochastic. Contrary to the standard VRP, where all demands are known in advance, it is possible for a particular realization of customer demands that a vehicle arrives at a customer and does not carry enough goods to satisfy its demand. One possible recourse action is for the vehicle to return to the depot, replenish its load and then return to the customer and continue its route. A different scenario occurs when the demands are known before the vehicle leaves the depot. In this case the fixed route that would violate the vehicle capacity is split up between two vehicles while the sequence of customers remains unchanged. The discussion above shows that the fixed routes problem is a Vehicle Routing Problem with stochastic demands and recourse action.

Only recently, researchers have started to investigate Stochastic Vehicle Routing Problems (SVRP). A comprehensive overview of stochastic vehicle routing problems and a survey of existing solution methodologies can be found in Dror, Laporte, and Trudeau [1989].

This paper describes our efforts to develop an effective and efficient algorithm for the Fixed Routes Problem based on a stochastic programming model with recourse. We assume that customer demands are independently and normally distributed and that the mean and variance of the customer demands are known. The algorithm is designed in such a way that if the customer demand variance approaches zero, i.e., the demand becomes deterministic, the algorithm is competitive with existing solution methods for the standard vehicle routing problem.

In Section 2 the problem and its assumptions are defined and in Section 3 the current literature is reviewed. In Section 4 the stochastic programming model with recourse is derived. In Section 5 an approximation algorithm is proposed and in Section 6 its performance is compared with other algorithms. Finally, Section 7 discuses possible extensions and Section 8 contains the conclusions.

2. THE PROBLEM

The specific single-depot Fixed Routes Problem studied in this paper has the following characteristics. The set of vehicles is homogeneous and each vehicle has a capacity Q. The location of the depot and the customers are known and fixed. The demand of each customer i is stochastic and modeled by an independent random variable ξ_i that is normally distributed with finite mean μ_i and finite variance σ_i^2. The objective is to plan a set of routes in anticipation of customer demands such that all customers are visited and fully serviced and the expected total cost of travel is minimized. The exact customer demand becomes only known during the execution of the planned routes.

The most important characteristic of the VRP with stochastic demands is that it is no longer possible to assume, contrary to the deterministic VRP, that routes can be followed as planned. A service failure may occur on a planned route because at some point along the route customer demand cannot be met. In that case some recourse action has to be taken.

Various recourse strategies, depending on the acceptable service policy and the amount of available information on customer demands, are possible. We have adopted the following recourse action. If a failure occurs, the vehicle will make a round trip to the depot for replenishment and resumes the planned route at the point of failure. Note that if the depot is denoted by 0 and if c_{ij} denotes the cost of traveling from i to j, a service failure at customer i will incur an additional cost of $c_{i0} + c_{0i}$. This recourse action is equivalent to the situation where customer demands are known before the vehicle leaves the depot and a fixed route that violates the vehicle capacity is split over two vehicles, while the customer sequence remains unchanged and no planning of the route break is allowed.

In addition, we assume that the parameters of the demand distributions are such that

$$\mathbf{P}[\xi_i < 0] = 0 \qquad (1)$$
$$\mathbf{P}[\xi_i > Q] = 0$$

for all customers i on every route R_k.

Finally, our algorithm will not construct any routes on which the demand ***on a single route*** can exceed two vehicle capacities, i.e.

$$\mathbf{P}[\sum_{i \varepsilon R_k} \xi_i > 2Q] = 0. \qquad (2)$$

The third restriction limits the possible number of service failures on a single route to at most one, since two full vehicle loads will always be able to satisfy all demands on a single route. This restriction is not unreasonable from a practical point of view if there are many customers with a limited demand variance (e.g. a coefficient of variation of the customer demand distribution smaller than or equal to 0.3).

3. LITERATURE REVIEW

Christofides [1971] was the first one to investigate the Fixed Routes Problem. He proposes the following solution approach. Each instance in a sample set, based for example on historic data, is solved as a deterministics VRP. The inter customer links are then sorted in decreasing order of frequency with which they appeared in any solution . Start the problem afresh and add links from the top of the list downwards provided they do not make the route infeasible on any of the sample days. Continue this process until every customer is routed. The resulting set of routes is further improved by the application of the 3-opt exchange improvement method of Lin [1965]. Christofides also showed by example that increasing the maximum allowed service break probability reduces the number of vehicles.

Beasley [1984] modifies the savings algorithm by Clarke and Wright [1964] and the k-exchange procedure by Lin [1965] so that they can be applied to the fixed routes problem. In the savings algorithm links are added only if the probability that the resulting route becomes infeasible remains below a amximum allowed value. Similarly, in the k-exchange procudures exchanges are alonly accepted if the probability that the resulting route becomes infeasible remains below a maximum allowed values. Computational experiments indicate that for a demand variation of ± 10 % the increase in travel distances is less than 2 % compared to deterministic routing for each instance. Furthermore, Beasley observes that his methode and the one proposed by Christofides [1971] perform equivalently.

Stewart and Golden [1983] prospose two models for the vehicle routing problem with stochastic demands. The first one is a chance-constraint model:

$$\begin{aligned} &\min z = cx \\ &\text{s.t. } \mathbf{P}[q_x \le Q] \ge 1-\alpha \\ &\qquad x \in T_m \end{aligned} \tag{3}$$

were cx is the deterministic route distance, q_x is the stochastic route demand, Q the vehicle capacity, and T_m is the set of all possible combinations of m routes.

The second models is a stochastic programming model with recourse:

$$\begin{aligned} &\min z = cx + \sum_{k=1}^{m} \lambda_k \mathbf{E}[l_k] \\ &\text{s.t. } x \in T_m \end{aligned} \tag{4}$$

where cx is the deterministic route distance, λ_k is the penalty per unit excess demand on route *k*, $E[l_k]$ is the expected excess demand of route *k*, and T_m is the set of all possible combinations of *m* routes.

These formulations can be solved with adapted heuristic algorithms for the VRP such as Clarke and Wright [1964] for model (3) and with a multiplier adjustment algorithm for formulation (4). They report that an algorithm based on generalized Lagrangean multipliers consistently outperforms Clarke and Wright but a higher computational cost. For more complex situations, such as correlated demands, they report that the simple Clarke and Wright algorithm seems to strike the best tradeoff between effectiveness and efficiency.

The above model exhibits two deficiencies as first observed by Dror and Trudeau [1986]. First, the recourse cost does not reflect the true cost incurred by a failure, since it does not take the location of a failure into account. Consider the two routes depicted in Figure 1.

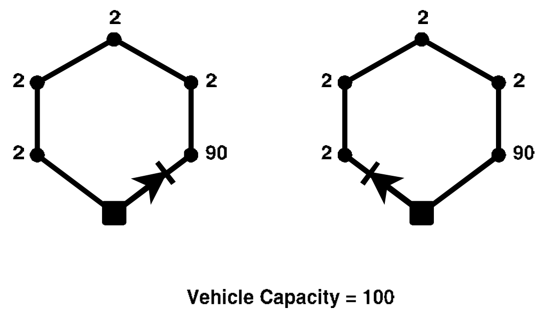

Figure 1. Recourse Cost Depending on Route Direction

The only difference between the two routes is their direction. Therefore, the expected excess demand, and thus the recourse cost as defined in Stewart and Golden, is the same for both routes. However, it is clear that the second route is preferred since a failure is only likely to occur at the last customer, which is close to the depot.

The second deficiency of the model lies in the fact that the recourse cost is linear in the amount of excess demand. Most practical recourse actions lead to a nonlinear recourse cost, since the cost of returning to the depot is independent of whether the vehicle goes back to pick up one unit or to pick up a full truck load.

Dror and Trudeau [1986] formulate the stochastic vehicle routing problem (SVRP) where only the customer demand is a random variable as a chance-constrained programming problem and a stochastic optimization problem with recourse. They identify the importance of the location of the service break and the orientation of the route. They propose a savings algorithm adapted from Clarke and Wright [1964] with as recourse individual (out-and-back) delivery to the customers at and following the service break. When computing the savings both route orientations are examined. They conclude that their algorithm performs better than the one by Stewart and Golden [1983] but a higher computational cost.

Dror, Laporte, and Trudeau [1989] suggest two solution frameworks for the SVRP with stochastic demands: They argue that Markov decision processes are well suited to model the problem in which reoptimization of the routes is allowed after more information becomes available, and that stochastic programming is the best alternative if route reoptimization is not allowed. In stochastic optimization, chance-constrained programming controls the probability of service failure, but does not take the cost of such failure into consideration. They show that the SVRP does not posses the classical VRP properties of non-intersection, convex hull sequence, and the dynamic programming principle of optimality. The paper provide a literature review and establishes a framework and classification, but does not contain any numerical results.

Stochastic programming with recourse seems the best alternative to model the real life fixed routes problem It should be noted that various recourse actoions that could be considered. Our solution approach falls in this category of stochastic programming with recourse.

4. THE MODEL

Our stochastic programming model with recourse takes the location of a failure into account and also uses a non-linear recourse cost. Since the variables ξ_i are assumed to be independently and normally distributed with mean μ_i and variance σ_i^2, their cumulative demand distribution is also normal with mean $\sum_{1 \le i \le n} \mu_i$ and variance $\sum_{1 \le i \le n} \sigma_i^2$. Therefore, given a route (0,1,2,...,*n*,*n*+1) where for convenience we have split the depot in an `origin' 0 and `destination' *n*+1, the probability P_i that the cumulative demand up to customer *i* is greater than the vehicle capacity given by

$$P_i = \mathbf{P}[\sum_{1 \le j \le i} \xi_j > Q] = 1 - \mathbf{P}[\sum_{1 \le j \le i} \xi_j \le Q] = 1 - \Phi\left(\frac{Q - \sum_{1 \le j \le i} \mu_j}{\sqrt{\sum_{1 \le j \le i} \sigma_j^2}} \right) \tag{5}$$

where Φ denotes the standard normal distribution function. As a consequence, the probability that a failure occurs at customer i, i.e. $\mathbf{P}[\sum_{1 \le j \le i-1} \xi_j \le Q, \sum_{1 \le j \le i} \xi_j > Q]$, can now be computed, since

$$\mathbf{P}[\sum_{1 \le j \le i-1} \xi_j \le Q, \sum_{1 \le j \le i} \xi_j > Q] = P_i - P_{i-1}. \tag{6}$$

Since we have assumed that the parameters of the demand distributions and the generated routes are such that at most one failure is likely to occur in a route, the expected recourse cost of the route *k* is equal to

$$\mathbf{E}[c_k] = \sum_{1 \le i \le n} (P_i - P_{i-1})(c_{i0} - c_{0i}) \tag{7}$$

The complete stochastic programming model with recourse that forms the basis for the algorithm to be described in the next section thus looks as follows:

$$\begin{aligned} &\min z = cx + \sum_{1 \le k \le m} \mathbf{E}[c_k] \\ &s.t.\ x \in T_m \end{aligned} \tag{8}$$

where $E[c_k]$ is the expected recourse cost of route *k*, and T_m is as before.

Note that the only place where the uncertainty of the customer demand explicitly occurs is in the objective function. The structure of the set of feasible solutions for the VRP with stochastic demands does not differ from the structure of the set of feasible solutions for the deterministic VRP even though the set of feasible routes is larger.

5. THE ALGORITHM

Since the uncertainty of the demand only appears in the objective function of the stochastic programming model with recourse, an obvious solution approach is to use any of the existing algorithms for the deterministic VRP to obtain a feasible set of routes that minimizes the total distance traveled, with respect to the mean customer demands. Given this set of routes, the expected total cost of travel using the objective function of the stochastic programming model with recourse can then be computed. Observe that if the variances are relatively small, or if the number of customers per route is relatively large, this should give a reasonable answer.

Another approach is to use an existing algorithm for the deterministic VRP and apply it using a modified cost function, anticipating the fact that we are not interested in minimizing the total distance traveled, but in minimizing the total cost of travel as specified by the objective function of the stochastic programming model with recourse. This approach is taken in Dror and Trudeau [1986], where the savings algorithm of Clarke and Wright [1964] is adapted. The `savings' term is modified to account for the location of a potential failure, the direction of the route, and the customers already supplied on the route.

The major concern in modifying an algorithm for the deterministic VRP for use in the context of the VRP with stochastic demands is the increase in computation time due to the introduction of a more complicated and computationally intensive cost function. As the cost of a route is related to the location of possible failures, a minor change to a route, such as the addition or deletion of a customer, forces a recomputation from scratch of the cost of the route.

Our two-phase method implements the underlying ideas of both approaches sketched above and does not require prohibitive computation times. In the first phase, a deterministic VRP algorithm is applied to construct an initial set of routes that minimizes the total distance traveled with respect to the mean customer demands. In the second phase, local search procedures are applied to improve the current set of routes with respect to the objective function of the stochastic programming model with recourse. This has the clear computational advantage that only in the second phase we have to work with a complicated cost function.

Observe that by the nature of a deterministic VRP algorithm the mean demand of each route of the set of initial routes does not exceed the vehicle capacity. However, in the context of the VRP with stochastic demands it is sufficient that the maximum demand of each route is less than or equal to two vehicle capacities so that the number of failures is at most one.

5.1. The Generalized Assignment Heuristic

The deterministic algorithm used in the first phase is our implementation of the generalized assignment heuristic proposed by Fisher and Jaikumar [1981].

The generalized assignment heuristic is usually presented as a two-phase method. In the first phase, an assignment of customers to vehicles is obtained by solving a generalized assignment problem with an objective function that linearizes and thereby approximates the cost of the traveling salesman tours of the vehicles through the customers. In the second phase, once the assignment decision has been made, a routing of each vehicle through its set of customers is obtained by solving a traveling salesman problem. The approximation of the delivery cost is obtained by constructing seed routes and considering the cost of inserting customers into these seed routes. A seed route is an artificial route consisting of the depot and a seed point, which is supposed to indicate an area that is expected to be visited by a single vehicle.

The choice of a good set of seed points is of crucial importance for the performance of the method. In our opinion, it is therefore better to consider the generalized assignment heuristic as a three-phase method, to emphasize the importance of seed selection. In this perspective, a set of seed points is chosen in the first phase, an assignment of customers to seed points is determined in the second phase, and routes are constructed for each of the obtained clusters in the third phase.

Since the algorithms we have implemented for the solution of the generalized assignment problem and the traveling salesman problems are well known, we will not describe them in detail. However, the method used to generate seed points is original and we will elaborate on it below.

As seed points are supposed to indicate areas that are expected to be visited by a single vehicle, it is clear that customer demand will play an important role in the determination of seed points. The method originally proposed by Fisher and Jaikumar divides the area around the depot in cones, each representing a total demand that is close to vehicle capacity, and locates the seed points on the rays that bisect the cones and at such a distance from the depot that a fixed percentage of the total demand in the cone is closer to the depot.

The method we propose is not only based on customer demand but also on `customer proximity'. Consider the two possible solutions for the same situation depicted in Figure 2.

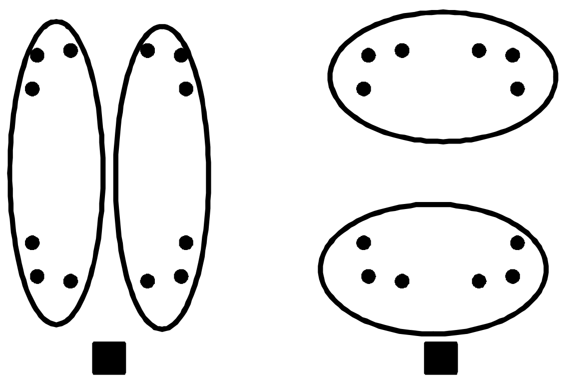

Figure 2. Cone Covering versus Circle Covering

The first solution results if the *cone covering method* as proposed by Fisher and Jaikumar is applied. Instead of the cone covering method we propose what might be called the *circle covering method*. The basic idea is the following. For each customer, determine the smallest circle with the center at this customer such that the total demand covered by the circle is less than or equal to the vehicle capacity, but adding the next closest customer would violate that property. Next, order the customers by increasing radius of their associated circles. Finally, cover all customers by circles as follows. Iteratively take the first not yet covered customer on the list and add its associated circle to the covering until all customers are covered. Take the set of seed points equal to the set of centers of the circles that constitute the covering. Note that this approach is independent of the location of the depot.

The generalized assignment problem is solved with a savings-regret heuristic [Martello and Toth 1981] complemented with various improvement methods.

The heart of the algorithm used to solve the traveling salesman problems is a dual ascent method based on the 1-tree relaxation, as proposed by Held and Karp [1971]. It is extended with a modified Christofides heuristic [Christofides 1975] and 2-opt and Or-opt techniques [Croes 1958, Or 1976].

5.2. Iterative Improvement Procedures

Iterative improvement procedures are based on what is perhaps the oldest optimization principle: neighborhood search. It is a simple and natural idea, which has proven to be surprisingly successful on a variety of problems. The general iterative improvement procedure proceeds as follows. We start at some initial feasible solution and search in its neighborhood for a better (cheaper) one. As long as an improved solution exists, we adopt it and repeat the neighborhood search from the new solution. Finally, we will reach a local optimum and stop.

The most often used neighborhood for vehicle routing problems is the k-exchange neighborhood. A k-exchange is a substitution of k arcs of a route with k others. Since the computational requirement of k-exchanges increases rapidly with k one usually only considers the cases k=2 and k=3.

To improve a single route, we have implemented a 2-exchange procedure [Croes 1958] and an Or-exchange procedure [Or 1976]. In addition, we designed two new k-exchange neighborhoods that aim at improving two routes by swapping customers between them. All neighborhoods are such that testing for optimality over the neighborhood requires $O(n^2)$ time.

In our description of the two new k-exchange neighborhoods, we will refer to the route that currently contains the customers we want to relocate as the *origin* route and the other as the *destination* route. Furthermore, a vertex i will always refer to a vertex from the origin route and pre_i and suc_i will denote its predecessor and successor, and a vertex j will always refer to a vertex from the destination route and pre_j and suc_j will denote its predecessor and successor.

The *relocate* neighborhood contains all the sets of at most two routes that can be obtained by removing the arcs (pre_i, i), (i, suc_i), and (h, suc_h) and replacing them with the arcs (pre_i, suc_i), (h, i) and (i, suc_h), i.e., we try to insert a vertex from the origin route into the destination route. A relocation is pictured in Figure 3.

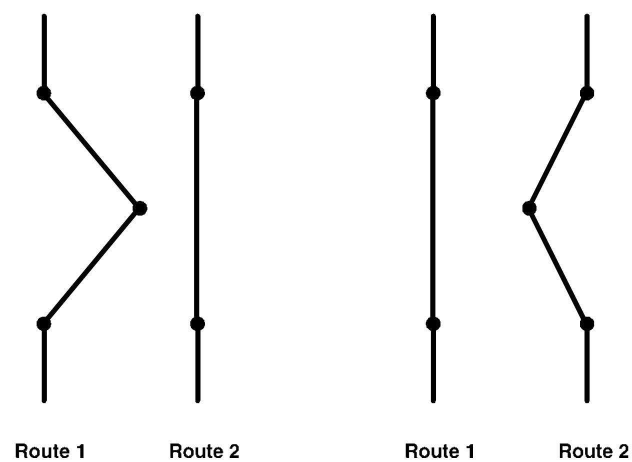

Figure 3. A Relocation

The *cross* neighborhood contains all the sets of at most two routes that can be obtained by removing the arcs (i, suc_i) and (h, suc_h) and replacing them with the arcs (i, suc_h) and (h, suc_i), i.e., we try to remove crossing arcs. Computational experience indicates that iterative improvement based on this neighborhood is very powerful. Note that replacing (0,suc_0) and (pre_{n+1}, n+1) with (0, n+1) and (pre_{n+1}, suc_0) combines the two routes. One of the significant drawbacks of variants of the Clarke and Wright algorithm is the fact that once customers are joined on a route they are never split again. Our algorithm does not have this deficiency. A cross-exchange is pictured in Figure 4.

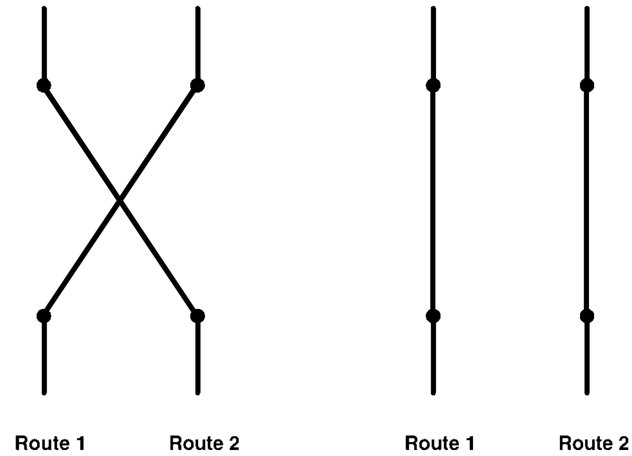

Figure 4. A Cross-Exchange

The above described iterative improvement procedures can easily be extended to larger neighborhoods by the introduction of paths, i.e. sequences of vertices instead of a single vertex, that are relocated.

6. COMPUTATIONAL RESULTS

The algorithm is implemented on an IBM PS/2 Model 70 using Microsoft C version 5.1. However, since the code does not use any operating system specific functions, it should be easy to port it to other environments.

The basic question, raised in the introduction, related to the fixed routes problem is whether or not a set of routes designed to be operated unchanged over a period of time provides a viable alternative to constructing a set of routes every day based on the particular instance of that day.

To be able to provide, at least an empirical, answer to this question, we conducted the following experiment. For several randomly generated instances of the VRP with stochastic demands, we compared the average cost that results if the fixed routes produced by our algorithm are used for twenty realizations of the customer demands with the average cost of the solutions produced by a deterministic VRP algorithm for the same twenty realizations of the customer demands. The deterministic VRP algorithm was also restricted so that it should never produce a solution that requires more vehicles than are needed for the fixed routes. The realization of the demand of customer *i* is obtained by setting q_i to $\max(0,\mu_i+RN(\cdot)\sigma_i)$, Random instances were generated for two different ratios of maximum mean demand over vehicle capacity, namely 0.3 and 0.15, and for four different values of the coefficient of variation σ/μ, namely 0.05, 0.1, 0.3, and 0.5. A summary of the results is given in Tables 1 and 2.

Table 1. Fixed Routes versus Variable Routes

σ/μ	E[cost]	fixed routes	variable routes	% cost increase
0.05	945.81	942.00	928.30	1.45
0.10	1182.06	1185.30	1121.65	5.37
0.30	1382.62	1365.70	1229.80	9.95
0.50	1713.48	1756.00	1444.30	17.75

Table 2. Fixed Routes versus Variable Routes

σ/μ	E[cost]	fixed routes	variable routes	% cost increase
0.05	784.11	780.90	802.25	-2.73
0.10	818.10	822.10	808.10	1.70
0.30	836.19	833.20	798.55	4.15
0.50	823.77	816.70	785.65	3.80

First of all, a comparison of the expected cost and the actual fixed routes cost shows that the expected cost is a good predictor of the average fixed routes length. But the most striking, and counter intuitive, phenomenon is that the average cost of the solutions produced by the deterministic VRP algorithm is not always smaller that the average cost that results if the fixed routes produced by our algorithm are used. The explanation for this phenomenon is that approximation algorithms are used for the solution of both problems and that the solution space of the vehicle routing problem with stochastic demands is larger than the solution space of the vehicle routing problem with deterministic demands, since multiple trips per vehicle are allowed.

The overall results indicate that even for relatively large variability, i.e., $\sigma/\mu = 0.3$, the increase in cost is less than 10 percent. Therefore fixed routes may indeed provide a viable alternative to constructing routes on a daily basis.

Another interesting observation is that an algorithms for the vehicle routing problem with stochastic demands tend to spread the work load more evenly over the vehicles than an algorithm for the vehicle routing problem with deterministic demands, because it tries to avoid routes that are filled to near capacity. Solutions with that characteristic are often preferred in practice.

Stewart and Golden [1983] as well as Dror and Trudeau [1986] apply a modified savings algorithm to a stochastic extension of a 75 customer problem that appeared in Christofides and Eilon [1969]. To turn the deterministic instance into a stochastic instance, the mean of the demand is taken to be the original demand quantity and a standard deviation, between zero and one-third of the mean demand, is generated using a uniform random generator.

To assess the performance of our algorithm, we have also applied it to the instance described above. Table 3 summarizes the results obtained by the three algorithms and Table 4 presents the solution constructed by our algorithm in more detail. Since both modified savings algorithms work with a different recourse action which in turn is different from our recourse action comparisons have to be made carefully. Table 3 presents the expected cost and deterministic distance computed based on our definition of the recourse action. Table 4 presents two values for the expected cost. The first shows the expected cost for the set of routes based on the recourse cost function where each customer at or after the service break is service on an individual out-and-back route as specified by Dror and Trudeau (1986). The second shows the expected cost for the set of routes based on the recourse cost function associated with our recourse action mentioned above. We emphasize that optimization in our algorithm is done with respect to the objective function of the stochastic programming model with recourse presented in this paper and that the other expected cost is computed only from the final set of routes. Also note that mean demand of route 7 exceeds the vehicle capacity of 160.

Table 3. Summary of the Results Obtained by the Three Algorithms

	Number of routes	Distance	E[cost]	Fullest route			Emptiest Route		
				μ	σ^2	P[d>Q]	μ	σ^2	P[d>Q]
Stewart & Golden	11	1016.34	1019.98	145	99.95	0.067	30	3.38	0.000
Dror & Trudeau	10	856.16	884.91	166	167.31	0.679	53	5.63	0.000
Savelsbergh & Goetschalckx	9	805.42	855.55	164	145.75	0.629	139	97.105	0.016

Table 4. Solution Obtained for the 75 Customer Test Problem

	μ	σ^2	*P*[*d*>*Q*]	*E*[*cost*]$_1$	E[cost]$_2$	Distance	Route sequence
1	142	50.021	0.000	39.26	39.26	39.26	4,27,52,46,34,67
2	155	176.746	0.353	85.09	85.27	82.30	45,48,47,21,61,74,30,75
3	158	102.201	0.421	91.29	91.51	85.66	12,40,9,39,31,58,26
4	149	14.547	0.152	89.23	89.23	84.94	72,10,38,65,66,11,7
5	153	158.871	0.289	72.76	74.04	67.28	6,33,1,73,62,28,2,68
6	139	97.105	0.016	115.37	115.37	114.76	57,15,37,20,70,60,71,69,36,5,29
7	164	45.750	0.629	127.93	129.77	115.62	49,24,18,55,25,50,32,44,3,17
8	146	77.837	0.056	104.07	104.07	102.30	13,54,19,59,14,53,35,8
9	158	141.104	0.433	124.67	127.03	113.30	22,64,42,43,41,56,23,63,16,51
	1364			849.72	855.55	805.42	

As anticipated, the use of a more sophisticated deterministic VRP algorithm to obtain an initial set of routes, i.e. the use of the generalized assignment heuristic versus the savings method, was clearly beneficial as it provided a solution with nine vehicles.

7. EXTENSIONS

Observe that the only crucial assumption in our approach is that the cumulative demand distribution can be computed. Therefore, it is relatively easy to extend the approach to handle some side constraints and to allow other than normal distributions for customer demands.

As an example, consider route duration constraints. Let t_{ij} denote the travel time from customer i to j, and T a universal upper bound on the duration of the routes. Furthermore, let the unloading time at customer i be given by $u_i = s_i + r_i q_i$, where s_i is a fixed part related to the customers premises and $r_i q_i$ a variable part depending on the actual demand q_i. In that case, the cumulative distribution of the travel time is again normal with parameters $\sum_{1 \le i \le n} (t_{i,i+1} + s_i) + \sum_{1 \le i \le n} r_i \mu_i$ and $\sum_{1 \le i \le n} (r_i \sigma_i)^2$. Results of incorporating such duration of tour and length of tour constraints will be reported on in a later paper.

8. CONCLUSIONS

In this paper, we have discussed the development of an approximation algorithm for the fixed routes problem modeled as a vehicle routing problem with stochastic demands. Computational experiments indicate that its performance compares favorably to other existing algorithms. The algorithm reflects our initial design goals: effectiveness and efficiency, i.e., a high quality solution in a reasonable amount of time. The fact that the cost functions in a stochastic environment are very complicated and computationally intensive makes it hard to accomplish an acceptable balance between effectiveness and efficiency. We have designed a two-phase approach. In the first phase, some effectiveness is sacrificed for efficiency. Using a simplified cost function we obtain a reasonable solution relatively fast. In the second phase, some efficiency is sacrificed for effectiveness. Using the actual cost function and the initial solution we obtain a high quality solution reasonably fast. In fact, the approach and the algorithm can be used for any variant of the vehicle routing problem that calls for a complicated computationally intensive cost function.

Our algorithm compares favorably with other fixed routes algorithms, primarily, we believe, because of the superiority of the generalized assignment procedure to get initial good routes as compared to algorithm based on modifications of the Clarke and Wright savings heuristic. Our new circle covering method to generate seed routes should generated good seeds both for uniformly distributed and clustered customers. The routes generated by our algorithm were within 10 % for coefficients of variations smaller than 0.3 of the customer demand. This indicates that fixed routes provide a viable alternative to daily recomputed routes.

ACKNOWLEDGMENT

This research was supported in part by the member companies of the Material Handling Research Center of the Georgia Institute of Technology. The authors thank Gideon Weiss for stimulating discussions and useful suggestions and Jeff Dillon for running the computational experiments.

REFERENCES

1. Beasley, J. E., 1984. "Fixed Routes," *Journal of Operational Research*, Vol. 35, No. 1, pp. 49-55.

2. Christofides, N., 1971. "Fixed Routes and Areas for Delivery," *International Journal of Physical Distribution*, Vol. 1, pp. 87-92.

3. Christofides, N., "Worst-Case Analysis of a New Heuristic for the Traveling Salesman Problem", Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh, PA, 1975.

4. Christofides, N. and S. Eilon, "An Algorithm for the Vehicle Dispatch Problem," *Operational Research* Quarterly 20, pp. 309-318, 1969.

5. Clarke, C., and J.W. Wright, "Scheduling of Vehicles from a Central Depot to a Number of Delivery Points," *Oper. Res.* 12, pp. 568-581, 1964.

6. Croes, A., "A Method for Solving Traveling Salesman Problems," *Oper. Res*. 5, pp. 791-812, 1958.

7. Dror, M. and P. Trudeau, "Stochastic Vehicle Routing with Modified Savings Algorithm," *European J. Oper. Res.* 23, pp. 228-235, 1989.

8. Dror, M., G. Laporte, and P. Trudeau, "Vehicle Routing with Stochastic Demands: Properties and Solution Frameworks," *Transportation Science,* 23, pp. 166-176, 1989.

9. Fisher, M. L. and J. Jaikumar, "A Generalized Assignment Heuristic for Vehicle Routing," *Networks* 11, pp. 109-124, 1981.

10. Golden, B. L. and A. A. Assad (Eds.), 1988. **Vehicle Routing: Methods and Studies**. Elsevier Science Publishers, New York, New York.

11. Held, M. and R. M. Karp, "The Traveling Salesman Problem and Minimal Spanning Trees, Part II," *Math. Programming* 1, pp. 6-25, 1971.

12. Laporte G., F Louveaux, and H. Mercure, "Models and exact solutions for a class of stochastic location-routing problems," *European Journal of Operational Research* 39, pp. 71-78, 1989.

13. Lin, S., "Computer Solution to the Traveling Salesman Problem," *Bell System Tech. J*. 44, pp. 2245-2269, 1965.

14. Martello, S. and P. Toth, "An Algorithm for the Generalized Assignment Problem," J. P. Brans (ed.), *Operations Research* 81, North-Holland, Amsterdam, pp. 589-603, 1981.

15. Or, I., "Traveling Salesman-Type Combinatorial Problems and Their Relation to the Logistics of Blood Banking," Ph.D. Thesis, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL, 1976.

16. Stewart, W. R. and B. L. Golden, "Stochastic Vehicle Routing: A Comprehensive Approach," *European J. Oper. Res.* 14, pp. 371-385, 1983.