A Location Based Heuristic for General Routing Problems*

Julien Bramel†and David Simchi-Levi‡

February 1992

Revised: March 1993

Abstract

We present a general framework for modeling routing problems based on formulating them as a traditional location problem called the Capacitated Concentrator Location Problem. We apply this framework to two classical routing problems: the Capacitated Vehicle Routing Problem and Inventory-Routing Problem. In the former case, the heuristic is proven to be asymptotically optimal for any distribution of customer demands and locations. Computational experiments show that the heuristic performs well for both problems and in most cases outperforms all published heuristics on a set of standard test problems.

1 Introduction

Vehicle routing problems have received much attention in recent years due to the increased importance of determining efficient distribution strategies to reduce operational costs in distribution systems. A typical routing problem consists of a fleet of vehicles located at a central depot or warehouse that must be scheduled to provide some type of service to customers geographically dispersed in a service region. The service may involve the delivery of goods to retailers from a central warehouse, the pick-up and delivery of children in school buses, or the pick-up of packages for express mail delivery, just to name a few of the possible applications.

In this paper, we present a general framework for solving several different routing problems. We apply the algorithm to two classical problems: the Capacitated Vehicle Routing Problem (CVRP)

*Research supported in part by ONR contract N00014-90-J-1649 and NSF contracts DDM-8922712 and DDM-91-01179

†Graduate School of Business, Columbia University, NY, NY, 10027.

‡Dept. of Industrial Engineering and Operations Research, Columbia University, NY, NY, 10027.

and the Inventory-Routing Problem (IRP), also known as the One Warehouse Multi-Retailer Distribution Problem. In the CVRP, a fleet of vehicles of fixed capacity are initially located at a central depot. A number of items must be delivered by the vehicles to each of the customers. We may consider the problem of delivering the goods from the central depot to satisfy customer demands, or the problem of picking up the loads at the customers to be brought to the depot. For the sake of consistency, we address only the former since these two cases are mathematically equivalent. The objective is to deliver the items to the customers such that each customer receives its demand, the vehicle capacity is not exceeded and the total distance traveled is minimized.

In the Inventory-Routing Problem, a central warehouse with an unlimited supply of items serves a set of retailers distributed in a given area. The retailers experience a fixed demand per unit of time for the items, and vehicles of limited capacity must be dispatched to replenish the retailer inventories. Each retailer incurs a holding cost per item per unit of time and a fixed cost per order placed. The objective is to schedule the vehicle departures and specify the loads destined for each retailer such that total cost per unit of time is minimized. This cost includes transportation cost, fixed order cost and inventory holding cost at the retailers. Examples of systems that can be modeled in this way occur when the warehouse is an outside supplier or when the depot is a manufacturing facility producing just to meet demand; see Anily and Federgruen (1990) and Gallego and Simchi-Levi (1990) for a more detailed description.

Since all non-trivial routing problems are NP-hard, much of the research has focused on finding heuristics that give good solutions, but not necessarily optimal ones. Most routing heuristics fall into the class called, by Christofides (1985), two-phase methods. These heuristics are of two types: (i) cluster first-route second, or (ii) route first-cluster second. In the first category, one clusters customers into groups (phase I) and then designs efficient routes for each cluster (phase II). In the second category, one constructs a traveling salesman tour through all the customers (phase I) and then partitions the tour into segments (phase II). One vehicle is assigned to each segment and visits the customers according to their appearance on the traveling salesman tour. The distinction between these two categories of heuristics on the quality of their solutions is very important, as demonstrated in Bienstock, Bramel and Simchi-Levi (1991). They show that no heuristic in the *route first-cluster second* class can be asymptotically optimal for the CVRP. A heuristic is asymptotically optimal if the relative error between the cost of the solution provided by the heuristic and the cost of the optimal solution decreases to zero as the number of customers increases.

We introduce here a new heuristic for general routing problems. This heuristic, called the Location Based Heuristic (LBH), is based on formulating the routing problem as a location problem commonly called the Capacitated Concentrator Location Problem (CCLP). This location problem is subsequently solved and the solution is transformed back into a solution to the routing problem. The method enables us to incorporate many different routing features into the model, and hence it is possible to apply the technique to many different problems.

In Section 2, we provide some motivation for the Location Based Heuristic that stems from recent results on the probabilistic analysis of the CVRP.

In Section 3, we present the Location Based Heuristic. We also formulate the Capacitated Concentrator Location Problem and present solution techniques for it.

In Section 4, we apply the heuristic to the CVRP. We present some enhancements to the LBH that we have found to work well for this problem. In addition, we prove that the LBH is asymptotically optimal. That is, the solution produced by the heuristic tends to the optimal solution value as the number of customers increases. To assess the quality of the solution on realistic size problems, we have performed computational experiments on a set of standard test problems.

In Section 5, we describe the IRP in more detail and apply our algorithm to it. To evaluate the quality of our solutions we develop a new lower bound on the cost of any policy that belongs to a specific subset of policies, called fixed partition policies.

In Section 6, we present some concluding remarks, and in particular we point out that the general framework can handle several other types of combinatorial problems.

2 Preliminaries

The Location Based Heuristic is motivated by some recent probabilistic results on the CVRP performed in Simchi-Levi and Bramel (1991) (see also Bramel at el. (1992)). To describe the results, we first present some notation. Let $N = \{x_1, x_2, \cdots, x_n\}$ be the set of customers served by the common depot x_0, w_k the demand of customer x_k, d_k the distance from customer x_k to the depot, and d_{kl} the distance between customers x_k and x_l and Q the vehicle capacity. Let $L_0(S)$ be the length of the optimal traveling salesman tour through the customers of a set $S \subseteq N$ and the depot. We denote by Z^* the value of the optimal solution to the CVRP, and by Z^{H} the value of the solution produced by heuristic H.

In their work, Simchi-Levi and Bramel relate the asymptotic optimal solution value of the CVRP to the asymptotic optimal solution of the bin-packing problem defined by the customer demands with bins of size equal to the vehicle capacity. To present their result, let b_n^* be the minimum number of bins of capacity Q needed to pack n demands drawn from some (general) distribution Φ. Results on the bin-packing problem tells us that there exists a constant γ such that $\lim_{n\to\infty} b_n^*/n = \gamma$, (a.s.). This means that for large n, the minimum number of bins required (b_n^*) is very well approximated by γn, where γ depends only on the distribution Φ. They prove the following.

Theorem 2.1 *Let the customers be independently and identically distributed in a compact region of $\Re^2$ with expected distance $E(d)$ to the depot. Let the demands (w_i/Q) be independently and identically distributed according to a probability measure Φ with support on $[0,1]$. Then,*

$$\lim_{n\to\infty} \frac{1}{n} Z_n^* = 2\gamma E(d) \qquad \text{(a.s.)}. \tag{1}$$

That is, for large n, the cost of the optimal solution to the CVRP can be very well approximated by the value $2n\gamma E(d)$.

The proof of the above result is based on constructing upper and lower bounds on Z_n^* that converge to the desired value as n tends to infinity. The structure of the upper bound is of special interest to us since it provides a method to construct a feasible solution which is asymptotically optimal. This upper bound, which provides the motivation for the Location Based Heuristic, is based on the following procedure.

Superimpose a grid of squares with side $\epsilon > 0$ on the area where the customers are located. For each square induced by the grid, solve the bin-packing problem defined by the demands of customers in the square and bins of capacity Q. For each bin in the solution to the bin-packing problem, send one vehicle to serve the customers assigned to the bin. By definition, the total load in a bin will not violate the vehicle capacity. The actual sequence or tour taken by each vehicle can be found by solving a traveling salesman problem on the customers in the bin and the depot. However, for the purpose of constructing an asymptotically optimal heuristic, Simchi-Levi and Bramel show that the following tour, asymptotically, is good enough. The tour starts at the depot, goes to one particular customer on its route, called the *seed point of the route*, and then proceeds to go back and forth from this customer to all the other customers on the route, and then back to the depot, see Figure 1.

This heuristic is very nearly asymptotically optimal; that is, as the number of customers increases this method will provide a solution whose relative error decreases to ϵ. Since ϵ can conceivably be picked as small as we like, we can ensure an arbitrarily small error. At a first glance, one might be tempted to use a similar heuristic in practice, by choosing ϵ very small. The problem is that one needs to weigh the advantages of a small ϵ, which will give a small error, and a large ϵ which will ensure enough points in each grid to be able to pack customers efficiently. To overcome these difficulties, we must turn to methods that do not use this type of region partitioning, but nevertheless have the same structure as the above described upper bound.

To do that, observe that the cost of each route in the above upper bound can be decomposed into two parts. The *first* is the cost of the *simple tour* that starts at the depot goes to the seed point and back to the depot; the *second* is the sum of the costs associated with having the vehicle travel to and from each customer to the seed point. It is therefore appropriate to construct a heuristic that clusters customers together so as to minimize the sum of the lengths of simple tours plus the total insertion cost of customers into simple tours. This can be achieved by approximating the CVRP with another combinatorial problem called the *Capacitated Concentrator Location Problem* (CCLP). This problem has applications in telecommunications network design.

The CCLP:

The Capacitated Concentrator Location Problem can be described as follows: given m possible sites for concentrators of fixed capacity Q_j, $j = 1, 2, \ldots, m$, we would like to locate concentrators at a subset of these m sites and connect n terminals, where terminal i uses w_i units of a concentrator's capacity, in such a way that each terminal is connected to *exactly one* concentrator, the concentrator capacity is not exceeded and the total cost is minimized. A site-dependent cost is incurred for locating each concentrator; that is, if a concentrator is located at site j, the *set-up* cost is v_j, for $j = 1, 2, \ldots, m$. The cost of connecting terminal i to concentrator j is c_{ij} (the *connection* cost), for $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, m$.

In formulating an instance of the CVRP as an instance of the Capacitated Concentrator Location Problem, we make every customer (in the CVRP) a possible site for a concentrator in the CCLP. We want to make the concentrator selection problem in the CCLP correspond to the seed selection problem in the CVRP. Therefore, the set-up cost for locating a concentrator at site j corresponds to the cost of choosing customer j as a seed customer. This cost is simply the cost of sending the vehicle to the seed customer (customer j) and back; that is, the length of the simple tour through the depot and customer j. Each customer (in the CVRP) is also made a terminal in the CCLP. The cost of connecting terminal i to a concentrator at site j is exactly the cost of inserting customer i into a simple tour through seed customer j and the depot.

In the next section, we use this insight to construct an effective method for solving general routing problems, not just the CVRP.

3 The Framework of the Location Based Heuristic

In this section we formulate a general routing problem and present the LBH. We then formulate the CCLP and discuss an effective technique for solving it.

A *General Routing Problem* is presented as follows. Given a set of customers N, define the collection of *servable sets* (denoted $\mathcal{C}$) to be those subsets of N that can be served by one vehicle. The term "servable" means that the set can be served by one vehicle without violating any of the constraints of the routing problem. The cost of serving a set $S \subseteq N$ is given by a real-valued *routing function* $\phi(S)$ and is defined for all subsets of N, even those that are not servable.

Define a *partition* of a set N to be a collection of disjoint non-empty sets $S_1, S_2, \ldots, S_m$ such that $\cup_{i=1}^{m} S_i = N$. Define a *feasible partition* to be a partition made up of only servable sets, say, $\{S_i\}_{i=1}^{r}$ such that $S_i \in \mathcal{C}$, for $i = 1, 2, \ldots, r$. The objective is the following:

$$\min_{\text{all feasible partitions: } S_1, \ldots, S_r \in \mathcal{C}} \sum_{i=1}^{r} \phi(S_i).$$

3.1 The Heuristic

In its most general form, the Location Based Heuristic consists of the following three phases:

Phase I: For an integer m, choose m non-empty subsets of N, say $T_1, T_2, \ldots, T_m$, called *seed sets.* These are just generalizations of *seed points.* These sets may overlap, and their union may not even cover all of N. Calculate the *set-up* costs $v_j = \phi(T_j)$, for each $j = 1, 2, \ldots, m$. Moreover calculate the *connection costs* $c_{ij} = \phi(T_j \cup \{x_i\}) - \phi(T_j)$, for each $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, m$.

Phase II: Solve the CCLP with the data defined in Phase I. The CCLP becomes the problem of choosing some of the seed sets, and "connecting" nodes to these sets, such that the total set-up cost of the seed sets chosen plus the sum of connection costs is as small as possible.

Phase III: Transform the solution to the CCLP into a solution to the routing problem.

In the above formulation of the LBH, the sets T_j, $j = 1, 2, \ldots, m$, correspond to sets of customers, that, if selected, are served together. Therefore, the set-up cost v_j represents the cost of selecting the set T_j, i.e., the cost of serving this set of customers. The connection cost c_{ij}, on the other hand, represents the added cost of serving customer x_i with the set T_j.

3.2 A Solution Method for CCLP

Phase II of the LBH requires a solution method for CCLP. We first formulated the CCLP as the following integer linear program. Let

$$y_j = \begin{cases} 1, & \text{if a concentrator is located at site } j, \\ 0, & \text{otherwise,} \end{cases}$$

and let

$$x_{ij} = \begin{cases} 1, & \text{if terminal } i \text{ is connected to concentrator } j, \\ 0, & \text{otherwise.} \end{cases}$$

Then CCLP is:

$$\text{Problem P} : Min \quad \sum_{i=1}^{n}\sum_{j=1}^{m} c_{ij}x_{ij} + \sum_{j=1}^{m} v_j y_j$$

$$s.t. \quad \sum_{j=1}^{m} x_{ij} = 1 \qquad \forall i, \tag{2}$$

$$\sum_{i=1}^{n} w_i x_{ij} \le Q_j \qquad \forall j, \tag{3}$$

$$x_{ij} \le y_j \qquad \forall i, j, \tag{4}$$

$$x_{ij} \in \{0, 1\} \qquad \forall i, j, \tag{5}$$

$$y_j \in \{0, 1\} \qquad \forall j. \tag{6}$$

Constraints (2) ensure that each terminal is connected to exactly one concentrator, and constraints (3) ensure that the concentrator's capacity constraint is not violated. Constraints (4) guarantee that if a terminal is connected to site j, then a concentrator is located at that site. Constraints (5) and (6) ensure the integrality of the variables.

Unfortunately, CCLP is NP-hard which indicates that the existence of a polynomial time algorithm for its optimal solution is unlikely. Hence, at a first glance it seems that we have not gained much; we have transformed one NP-hard problem (the routing problem) into another NP-hard problem (the CCLP). The advantage, however, is that, while both are NP-hard, the CCLP is considerably easier to solve in the sense of finding a "good" solution in a "reasonable" amount of time. One reason is that the constraints of the CCLP are simple compared to the constraints that appear in the routing problem, namely the subtour elimination constraints. In addition, the structure of the objective function in the CCLP is substantially simpler than the cost structure in the general routing problem.

Several algorithms have been proposed to solve the CCLP in the literature; all are based on the celebrated Lagrangian relaxation technique. This includes Neebe and Rao (1983), Barcelo and Casanovas (1984), Klincewicz and Luss (1986), and Pirkul (1987). The one we use is derived in a similar fashion as Pirkul (1987) which seems to be the most effective.

This solution method concentrates on relaxing a set of constraints, bringing them into the objective function with a multiplier vector giving a lower bound, then using a subgradient search method to find the best lower bound. At each step of the subgradient procedure (i.e., for each set of multipliers) we try to make use of the information given by the multipliers to find a feasible solution to the location problem. This step consists of a simple and efficient subroutine. After a prespecified number of iterations the algorithm is terminated.

More specifically, we relax the problem by including constraints (2) in the objective function. For any vector $\lambda \in \Re^n$, consider the following problem P_λ:

$$\text{Problem } \mathrm{P}_\lambda : Min \quad \sum_{i=1}^{n}\sum_{j=1}^{m} c_{ij}x_{ij} + \sum_{j=1}^{m} v_j y_j + \sum_{i=1}^{n} \lambda_i \Big(\sum_{j=1}^{m} x_{ij} - 1 \Big)$$

subject to (3)-(6). Let Z_λ be its optimal solution with $\{\overline{y}, \overline{x}\}$ its optimal variables.

One can see that P_λ separates into m easily solvable subproblems. For a given $j = 1, 2, \ldots, m$, define the following:

$$\begin{aligned} \text{Problem } \mathrm{P}_\lambda^j : Min \quad & \sum_{i=1}^{n} \overline{c}_{ij}x_{ij} + v_j y_j \\ s.t. \quad & \sum_{i=1}^{n} w_i x_{ij} \le Q_j, \\ & x_{ij} \le y_j \qquad \forall i = 1, 2, \ldots, n, \\ & x_{ij} \in \{0, 1\} \qquad \forall i = 1, 2, \ldots, n, \\ & y_j \in \{0, 1\} \end{aligned}$$

where $\overline{c}_{ij} \equiv c_{ij} + \lambda_i$, for all i, j.

Clearly, problem P_λ^j is no more difficult than a single constraint 0-1 knapsack problem, for which efficient algorithms exist; see, e.g., Nauss (1976). If the optimal knapsack solution is less than $-v_j$, then the corresponding optimal solution to P_λ^j is found by setting $\overline{y}_j = 1$ and $\overline{x}_{ij}$ according to the knapsack solution, indicating whether or not terminal i is connected to concentrator j. If the optimal knapsack solution is more than $-v_j$, then the optimal solution to P_λ^j is found by setting $\overline{y}_j = 0$ and $\overline{x}_{ij} = 0$ for all $i = 1, 2, \ldots, n$. Let Z_λ^j be the optimal solution value of P_λ^j.

The solution to P_λ, the lower bound on the optimal solution to CCLP, is therefore easily found. To find the best possible lower bound, we use a subgradient procedure.

Using an initial vector $\lambda^{(0)}$, we solve the m knapsack problems and get a solution $\{\overline{y}^{(0)}, \overline{x}^{(0)}\}$. This solution in most cases is not a feasible solution to P, since the values $\overline{x}^{(0)}$ do not necessarily satisfy constraints (2). We generate new multipliers using the following formula:

$$\lambda_i^{(k+1)} = \lambda_i^{(k)} + t_k\Big(\sum_{j=1}^{m} \overline{x}_{ij}^{(k)} - 1\Big), \ \forall i = 1, 2, \ldots, m.$$

The step size t_k is determined by

$$t_k = \alpha \cdot \frac{(\overline{Z} - Z_{\lambda^{(k)}})}{\sum_{i=1}^{n}(\sum_{j=1}^{m} \overline{x}_{ij}^{(k)} - 1)^2},$$

where α is a scalar and $\overline{Z}$ is an upper bound on the optimal solution to P (see Held, Wolfe and Crowder (1974) for a justification of this formula). The scalar α is initially set to 2 and halved after the bound has not improved in a prespecified number of iterations. When α reaches some lower bound fixed beforehand, the algorithm is terminated.

For a given set of multipliers, if the values $\overline{x}^{(k)}$ satisfy (2), then we have an optimal solution to problem P, and we stop. Otherwise, we perform a quick subroutine to find a feasible solution to P. This procedure is based on the observation that the knapsack solutions found in the lower bound give us some information concerning the benefit of setting up a concentrator at a site (relative to the current multipliers $\lambda^{(k)}$). If, for example, the knapsack solution corresponding to a given concentrator is 0, i.e., the optimal knapsack is empty, then this is most likely not a "good" concentrator to select at this time. In contrast, if the knapsack solution has a very negative cost, then this is a "good" concentrator. In this sense, the multipliers and the knapsack solutions tell us which concentrator sites are the best ones to select. Given the values $Z_{\lambda^{(k)}}^j$ $(j = 1, 2, \ldots, m)$, renumber the concentrators so that

$$Z_{\lambda^{(k)}}^1 \leq Z_{\lambda^{(k)}}^2 \leq \cdots \leq Z_{\lambda^{(k)}}^m.$$

The procedure we perform is called GREEDY, since it allocates terminals to concentrators in a myopic fashion. Let M be the minimum possible number of concentrators used in the optimal solution to CCLP. This can be found by solving the bin-packing problem defined on the values w_i with bin capacities Q_j; see Johnson et al. (1974). Starting with the "best" concentrator, in this case concentrator 1, connect the terminals in its optimal knapsack to this concentrator. Then, following the order of the renumbered knapsack solutions, take the next "best" concentrator (say concentrator j) and solve a new knapsack problem: one defined with costs $\overline{c}_{ij} = c_{ij} + \lambda_i^{(k)}$ for each terminal i still unconnected. Connect all terminals in this knapsack solution to concentrator j. If this optimal knapsack is empty, then a concentrator is not located at that site, and we go on to the next concentrator. Continue in this manner until M concentrators are located. Let $\{y', x'\}$ be the resulting solution.

The solution $\{y', x'\}$ may still not be a feasible solution to P since some terminals may not be connected to a facility. In this case, unconnected terminals are connected to facilities in use where they fit with minimum additional cost. If needed, additional facilities may be opened following the ordering of the renumbered knapsack solutions. A local improvement heuristic is then performed to improve on this location solution, using simple interchanges between terminals, and the best solution is kept as the upper bound to P.

Upon termination of this algorithm, if the relative error between the best upper bound and the best lower bound is more than a threshold value (typically 0.5 percent), we start a branch and bound algorithm to reduce this gap. The branching is done by fixing the values $\overline{y}_j$ to either 0 or 1. The procedure described above is repeated at each node of the branch and bound tree, until the relative error is reduced below the threshold value. If the best upper bound does not decrease in the search of a number of consecutive nodes of the tree (typically 15), the branching is terminated.

4 The Capacitated Vehicle Routing Problem

In this section, we describe the application of the LBH to the CVRP. This problem has been analyzed extensively in the literature in the last three decades. For a survey, see Christofides (1985).

4.1 Formulation

The Capacitated Vehicle Routing Problem can be stated as follows: a set of n geographically distributed customers needs to be served by a fleet of identical vehicles of fixed capacity Q. Associated with customer x_k is a positive demand $w_k \leq Q$ which is the amount of load that needs to be delivered to that customer. The objective is to design efficient routes to serve the customers at minimum cost, where cost is proportional to distance traveled. We concentrate here on the case where the vehicles have identical capacities since all of the benchmark problems from the literature have this property. However, the adaptation of the Location Based Heuristic to the different capacity case is straightforward.

In the CVRP, the collection of servable sets is

$$\mathcal{C} = \Big\{S \subseteq N \Big| \sum_{i \in S} w_i \leq Q\}$$

and the routing function is given by

$$\phi(S) = L_0(S).$$

In Phase I, for a given number m, we choose seed sets, $T_1, T_2, \ldots, T_m$; each set being a subset of the customers. In Section 4.2 we present the types of seed sets that we have found to work well in practice. The cost of selecting set T_j, or setting up a concentrator at site j, is then $v_j = \phi(T_j)$, for $j = 1, 2, \ldots, m$. For $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, m$, the connection costs c_{ij} is a measure of the cost of inserting node x_i into set T_j, i.e.,

$$c_{ij} = \phi(T_j \cup \{x_i\}) - \phi(T_j).$$

Since finding the exact values of c_{ij} can be quite time consuming, in Section 4.2 we present what we have found to be satisfactory approximations.

In Phase II, we solve the CCLP with the data in this form. The solution to CCLP specifies which terminals (customers) to connect to which concentrators (seed sets).

In Phase III, we transform the location solution provided in Phase II into a feasible routing solution. Let $\{y^*, x^*\}$ be the best solution found for P and for each j with $y_j^* = 1$ define $S_j = \{1 \leq i \leq n | x_{ij}^* = 1\}$. Assume, that $S_1, S_2, \ldots, S_r$ are the non-empty sets after renumbering. Each S_j is a set of customers that can be served by one vehicle since they represent feasible connections in the CCLP (since x^* satisfies (3)). The cost of the Location Based Heuristic solution to the CVRP is then:

$$Z^{\text{LBH}} = \sum_{j=1}^{r} \phi(S_j).$$

4.2 Selection of Seed Sets and Connection Costs

It is clear that many possible variations of the LBH can be implemented depending on two decisions: first, the types of seed sets chosen, and second, the connection cost approximations used.

The selection of seed sets provide much flexibility in the implementation. If two or more customers must be served together, for reasons inherent in the particular application, then they can be inputed as a seed set, which will ensure that they are served together in the final routing solution. Also, if some routes are known to be good routes by an experienced dispatcher, they can be inputed at this phase and will be in the final solution.

The choice of connection costs also provides much flexibility. Let the optimal traveling salesman tour through a seed set T_j be the cycle $\{x_0 = x_{j_0}, x_{j_1}, x_{j_2}, \ldots, x_{j_p}, x_{j_{p+1}} = x_0\}$. There are many possible connection costs, of which we have used the following:

$$\textit{direct cost}\text{: } c_{ij} = \min_{l=0,\ldots,p} \{2d_{ij_l}\}, \quad \text{or}$$

$$\textit{nearest insertion cost}\text{: } c_{ij} = \min_{l=0,\ldots,p} \{d_{j_l i} + d_{ij_{l+1}} - d_{j_l j_{l+1}}\}.$$

Direct cost has the advantage that, when added to $\phi(T_j)$, it provides an upper bound on the routing cost, while the nearest insertion cost works well because it is accurate for small sets T_j.

We have implemented several different versions of the LBH. Each one starts with the seed sets $T_j = \{x_j\}$ for $j = 1, 2, \ldots, n$ with $m = n$. In this case $\phi(T_j) = 2d_j$. This seems to work well for the CVRP.

The heuristics differ in the types of connection costs. In the first implementation, the connection costs are determined by the nearest insertion cost, i.e., $c_{ij} = \phi(\{x_j\} \cup \{x_i\}) - \phi(\{x_j\}) = d_i + d_{ij} - d_j$. We call this version the *Seed-Tours Heuristic* (ST).

Another implementation has connection costs determined by the direct cost $c_{ij} = 2d_{ij}$. We call this version the *Star-Connection Heuristic* (SC), since connections are made in the form of stars.

In both cases, for each $j = 1, 2, \ldots, m$, the customer that defines the seed set T_j, i.e. x_j, is called the *seed* customer for that seed set. Note that when the seed sets have only one customer, all calculations of v_j and c_{ij} are trivial.

One can note the relationship between the ST Heuristic and the Generalized Assignment Heuristic due to Fisher and Jaikumar (1981). In their heuristic, Fisher and Jaikumar choose an initial set of m *seed* customers, say $\{x_{j_1}, x_{j_2}, \ldots, x_{j_m}\}$. For each seed customer, say x_{j_k}, they determine the cost of inserting a non-seed customer x_i into the tour containing only customer x_{j_k}, i.e., their cost is exactly $d_i + d_{ij_k} - d_{j_k}$. The problem then is to "add" the customers to "tours" at minimum cost. To do this they solve a generalized assignment problem. The solution is a partition of m sets, all containing at most a total demand of Q and each containing one seed customer.

It is clear that the performance of the Generalized Assignment Heuristic depends highly on the initial set of seed customers. For this purpose, Fisher and Jaikumar suggest several methods including an interactive approach (leaving the decision to the scheduler) or an automatic approach (based on a region partitioning scheme). Using the terminology of Fisher and Jaikumar, the ST heuristic chooses simultaneously the best m seeds (out of a possible n) and the best way to assign the customers to these seeds. That is, it combines the seed selection problem with the problem of assigning customers to the selected seeds by solving the CCLP.

4.3 An Asymptotically Optimal Heuristic

In this section we show that the Star-Connection Heuristic (SC) is asymptotically optimal. This means that the relative error between the solution it produces and the optimal solution decreases to zero as the number of customers increases. We prove this result by showing that the solution to the CCLP defined by the parameters in the implementation of the Star-Connection Heuristic can be transformed into a routing solution which is asymptotically optimal to the CVRP.

The specific set-up and connection costs used in the Star-Connection Heuristic imply that the cost of the solution to the CCLP (at the end of Phase II of the LBH) is an upper bound on the cost of the routing solution produced in Phase III. That is, the cost of the routing solution generated by the SC Heuristic is bounded from above by the cost of the solution to the CCLP. This is true since the Star-Connection Heuristic approximates the routing cost by having the vehicle travel back and forth to and from the seed point to each customer. This provides an upper bound on any efficient routing of the customers. All that needs to be shown, therefore, is that there is a solution to CCLP whose cost asymptotically approaches the value on the right-hand side of (1). We do this in the following theorem.

Theorem 4.1 *Let the customers be independently and identically distributed in a compact region of $\Re^2$ with expected distance $E(d)$ to the depot. Let the demands (w_i/Q) be independently and identically distributed according to a probability measure Φ with support on $[0,1]$. Then, the Star-Connection Heuristic is asymptotically optimal, i.e.,*

$$\lim_{n\to\infty} \frac{1}{n} Z_n^{\mathrm{SC}} = 2\gamma E(d) \qquad \text{(a.s.)}.$$

Proof. We start with an upper bound on the cost of the solution produced by the LBH, that is, we construct a feasible solution to the CCLP. Pick a fixed $\epsilon > 0$ and let $G(\epsilon)$ be an infinite grid of $\epsilon \times \epsilon$ squares. Let A be the compact support of the distribution μ. Let $A_1, A_2, \ldots, A_{t(\epsilon)}$ be the subregions of $G(\epsilon)$ that intersect A and have $\int_{A_i} d\mu > 0$. Let $n(i)$ be the number of customers located in subregion A_i.

For a given subregion A_i, find an optimal bin-packing of customer demands in the subregion and bin capacity Q. Let $b^*(i)$ be the number of bins used in this optimal packing, and let $B_j(i)$ be the set of customers in the j^{th} bin of this packing. Now arbitrarily select one customer from each bin; say $x_{l_1}, x_{l_2}, \ldots, x_{l_{b^*(i)}}$. Each of these customers is a "seed" customer, that is, they correspond to the selection of the seed sets $T_{l_1}, T_{l_2}, \ldots, T_{l_{b^*(i)}}$. Now connect each terminal (or customer) to the concentrator corresponding to the seed customer in its bin. Repeating this for each subregion defines a solution to the CCLP with value Z_L.

Then,

$$\begin{aligned} Z_n^{\text{SC}} \leq Z_L &= \sum_{i=1}^{t(\epsilon)} \sum_{j=1}^{b^*(i)} \Big\{ v_{l_j} + \sum_{x_k \in B_j(i)} c_{kl_j} \Big\} \\ &= \sum_{i=1}^{t(\epsilon)} \sum_{j=1}^{b^*(i)} \Big\{ 2d_{l_j} + \sum_{x_k \in B_j(i)} 2d_{kl_j} \Big\}. \end{aligned}$$

Clearly, $d_{kl_j} \leq \epsilon\sqrt{2}$ for each $x_k \in B_j(i)$. Hence,

$$Z_n^{\text{SC}} \leq Z_L \leq \sum_{i=1}^{t(\epsilon)} \sum_{j=1}^{b^*(i)} \Big\{ 2d_{l_j} + 2(|B_j(i)| - 1)\epsilon\sqrt{2} \Big\}.$$

Let $\underline{d}(i)$ be the distance from the depot to the nearest point in subregion A_i. Then $d_j \leq \underline{d}(i) + \epsilon\sqrt{2}$ for each customer x_j in region A_i. Hence,

$$\begin{aligned} Z_n^{\text{SC}} \leq Z_L &\leq \sum_{i=1}^{t(\epsilon)} \Big\{ 2b^*(i)\underline{d}(i) + 2n(i)\epsilon\sqrt{2} \Big\}, \\ &\leq 2\sum_{i=1}^{t(\epsilon)} b^*(i)\underline{d}(i) + 2n\epsilon\sqrt{2}. \end{aligned}$$

Dividing by n and taking the limit gives

$$\overline{\lim_{n\to\infty}} \frac{1}{n} Z_n^{\mathrm{SC}} \leq \overline{\lim_{n\to\infty}} \frac{1}{n} Z_L \leq 2 \overline{\lim_{n\to\infty}} \frac{1}{n} \sum_{i=1}^{t(\epsilon)} b^*(i)\underline{d}(i) + 2\epsilon\sqrt{2}.$$

In Simchi-Levi and Bramel (1990) (see also Bramel et al. (1991)) there is a simple proof of the almost sure result

$$\forall \epsilon > 0, \quad \overline{\lim_{n\to\infty}} \frac{1}{n} \sum_{i=1}^{t(\epsilon)} b^*(i)\underline{d}(i) \leq \gamma E(d).$$

Then,

$$\overline{\lim_{n\to\infty}} \frac{1}{n} Z_n^{\mathrm{SC}} \leq \overline{\lim_{n\to\infty}} \frac{1}{n} Z_L \leq 2\gamma E(d) + 2\epsilon\sqrt{2}.$$

Since ϵ was arbitrary and with the lower bound of equation (1), this proves that the Star-Connection Heuristic is asymptotically optimal for the CVRP. ∎

4.3.1 Computational Issues

To solve the CVRP, we perform an enhancement phase in parallel with the GREEDY procedure presented in Section 3.2. The GREEDY procedure constructs solutions to CCLP at each iteration of the subgradient procedure while this procedure at the same time constructs solutions to the CVRP.

The connection costs used in the CCLP only approximate the real cost of adding a customer to a tour. Therefore, to get a better approximation we try to update the connection costs as we add terminals to concentrators. Each time we connect a terminal to a concentrator we update the connection costs to take into account this new customer. Specifically, for each set of multipliers, we perform the following procedure. Select the M "best" concentrators according to the current knapsack solutions; these are concentrators $1, 2, \ldots, M$ after renumbering. Consider the set of terminals that are connected to only one concentrator in the m knapsack solutions, i.e., that appear in only one "knapsack". The subset of these terminals that are connected to one of the M concentrators selected are each connected to the concentrator whose knapsack they appear in. For each concentrator, determine the tour through the terminals connected to that concentrator (and the concentrator itself) using the nearest insertion method (see, Rosenkrantz, Stern and Lewis (1977)); a customer is inserted into a tour without changing the orientation of the tour, but simply by inserting the customer in the cheapest way between two other customers.

Then determine for each unconnected x_i and each j $(1 \leq j \leq M)$, the costs $\hat{c}_{ij}$, which represents the cost of inserting node x_i into the tour associated with concentrator j, using the nearest insertion cost. If x_i does not fit in tour j (because of the capacity constraint) then let $\hat{c}_{ij} = +\infty$. The value of $\hat{c}_{ij}$ represents the "closeness" of terminal i to the tour associated with concentrator j. Next, determine the penalty associated with inserting customer x_i into its second "closest" tour instead of into its "closest" tour. Let x_{i^*} be the customer with the largest such penalty. Insert node x_{i^*} into its "closest" tour, say tour j^*, using the nearest insertion method. Update the insertion costs $\{\hat{c}_{ij}\}$, (in fact, only $\hat{c}_{ij^*}$ needs to be updated) and continue in this manner until all terminals are in tours. The resulting routing solution is then compared with the best solution found so far and the better one is kept.

4.4 Computational Results

In this section we report on computational experiments with the Location Based Heuristic on a set of 7 standard test problems from the literature. The problems vary in size from 50 to 199 customers as reported in Table 1. The problems are from Christofides, Mingozzi and Toth (1979). We compare the performance of the LBH to the performance of the following 9 published heuristics:

- SAV = Clarke and Wright's Savings Algorithm (1964),
- M&J = Mole and Jameson (1976),
- PSA = Altinkemer and Gavish's Parallel Savings Algorithm (1985),
- MBS = Desrochers and Verhoog's MBS Algorithm (1989),
- SWP = Gillett and Miller's Sweep Algorithm (1974),
- TPM = Christofides, Mingozzi and Toth's Two-Phase Method (1979),
- F&J = Fisher and Jaikumar's Generalized Assignment Heuristic (1981),
- TRE = Christofides, Mingozzi and Toth's Incomplete Tree Search Algorithm (1979),
- P&F = Pureza and Franca's Tabu Search Algorithm (1991),

The CPU time of the LBH (in seconds) is based on running the algorithm on an RS6000 Model 550.

Table 1

Computational Results on the Standard Test Problems

Prob. No.	Size	SAV	M&J	PSA	MBS	SWP	TPM	F&J	TRE	P&F	Alg. ST[1]	CPU Time	Alg. SC[2]	CPU Time
		Published Heuristics												
[1]	50	585	575	564	586	532	547	524	534	536	524.6	68	524.9	240
[2]	75	900	910	878	885	874	883	857	871	842	848.2	406	884.3	656
[3]	100	886	882	868	889	851	851	833	851	851	832.9	890	894.9	1237
[4]	100	831	879	845	828	937	827	824	816	—	826.1	400	828.9	110
[5]	120	1079	1100	1066	1058	1266	1066	—	1092	—	1051.5	1303	1051.2	2570
[6]	150	1204	1259	1104	1133	1079	1093	1014	1064	1081	1088.6	2552	1123.7	3412
[7]	199	1540	1545	1370	1424	1389	1418	1386	1386	—	1461.2	4142	1438.2	8021

([1]) Seed-Tours Heuristic.

([2]) Star-Connection Heuristic.

We observe that the Seed-Tours Heuristic finds solutions better than most of the other published heuristics. The running time is comparable to the running time of many heuristics, including the recently published, Parallel Saving Algorithm; see Altinkemer and Gavish (1991).

5 The Inventory-Routing Problem

We now turn our attention to another routing problem that involves a more complex cost structure, but can however be handled by the Location Based Heuristic.

Consider the problem where n retailers are geographically dispersed in a given area. A central warehouse has an unlimited supply of items. Retailer i faces a deterministic demand of D_i items per unit of time, a fixed cost K_i for each order placed, and an inventory holding cost of h_i per item per unit of time. We assume an unlimited amount of inventory can be kept at each of the retailers. We seek a dispatching and routing strategy that delivers items to retailers from the central warehouse such that total inventory holding cost, order cost and transportation cost per unit of time is minimized. We assume all demands must be met without backlogging, that is, shortages are not allowed. The problem is called the One-Warehouse Multi-Retailer Distribution Problem, or the Inventory-Routing Problem (IRP).

The problem is clearly difficult since the set-up cost for each order is very complicated. It consists of the fixed cost plus the cost of sending a vehicle to serve a set of customers, which is proportional to the total distance traveled by the vehicle. This set-up cost is not separable and this is what makes the problem drastically more difficult to solve than the CVRP.

As is pointed out by Anily and Federgruen (1990), optimal policies for this problem may be very complicated and in addition characterizing them mathematically may not be easy. Moreover, in practice, policies that are not easy to implement are not often used. For example, a policy where a retailer receives orders at very irregular intervals would not be easy to implement. Therefore, much of the research on this problem concentrates on policies that are, in some sense, simple. The totality of possible transportation and scheduling policies is much too large, therefore in general researchers concentrate their efforts on studying subsets of policies.

Many approaches have been used to attempt to tackle this problem. Gallego and Simchi-Levi (1990) prove that a direct shipping policy, a policy where each vehicle serves only one customer, is within 6% of optimality under certain conditions. Herer and Roundy (1990) restrict their attention to power of two policies, and show some good empirical results when vehicles have unlimited capacity. Anily and Federgruen (1990) suggest region partitioning strategies that are asymptotically optimal within a specific class of policies.

Consider the following set of policies, which we call *Fixed Partition Policies*. The set of customers is partitioned into m disjoint sets, $S_1, S_2, \ldots, S_m$ and each set is served separately. That is, whenever a customer in a set is served, all the customers in the set are served. What is the justification for this subset of policies? Clearly, these types of policies are easy to implement. Each set has its own cycle time and all retailers get orders at constant regular intervals. In addition, drivers need only learn a small number of possible routes.

It is clear that if a given set of customers are always served together, then the set-up cost for ordering is just the cost of the optimal traveling salesman tour through the customers of the set and the depot plus the fixed order costs. In this case, it is well known that optimal deliveries occur at regular fixed intervals. Since the set-up cost is known, the optimal *cycle time*, the time between deliveries, can be found using the traditional Economic Order Quantity formula. Let S be a set of customers served every $t(S)$ units of time and define $K(S) = \sum_{i \in S} K_i$ and $D(S) = \sum_{i \in S} D_i$. Then, the cost per unit time for serving the set S is

$$\frac{1}{t(S)}\Big(L_0(S) + K(S) + \sum_{i \in S} \frac{1}{2} h_i t^2(S) D_i\Big)$$

$$= \frac{L_0(S) + K(S)}{t(S)} + t(S) \sum_{i \in S} \frac{1}{2} h_i D_i. \tag{7}$$

If the vehicles have unlimited capacity, the optimal cycle time, denoted by $t^*(S)$, can be found by minimizing on $t(S)$:

$$t^*(S) \equiv \Big(\frac{2(L_0(S)+K(S))}{\sum_{i\in S} h_i D_i}\Big)^{1/2}.$$

Note that the vehicle capacity restriction disallows us from always choosing this minimum, and in fact, the cycle time $t(S)$ must satisfy

$$t(S)D(S) \leq Q.$$

The best feasible cycle time for a set S is therefore given by:

$$t(S) = \min\Big\{t^*(S), \frac{Q}{D(S)}\Big\}.$$

Hence, in this problem, the routing function is

$$\phi(S) = \frac{L_0(S)+K(S)}{t(S)} + \frac{1}{2}t(S)\sum_{i\in S} h_i D_i, \tag{8}$$

and the collection of servable sets is simply $\mathcal{C} = \{S | S \subseteq N\}$.

The Location Based Heuristic can now be implemented. In Phase I, we select m seed sets and calculate

$$v_j = \phi(T_j),\ \forall j = 1, 2, \ldots, m$$

and

$$c_{ij} = \phi(T_j \cup \{x_i\}) - \phi(T_j),\ \forall i = 1, 2, \ldots, n,\ j = 1, 2, \ldots, m.$$

In Phase II, since in the IRP any subset of N is a servable set, there is no need to have capacities on the concentrators. Hence we use the formulation of the CCLP without constraints (3), in this case the location problem is simply a facility location problem. The solution method described in Section 3.2 can still be used, and in fact it runs more efficiently since no knapsack algorithm is needed.

In Phase III, the solution to the CCLP corresponds to a partition of the customers into disjoint sets, say sets $S_1, S_2, \ldots, S_r$. These sets correspond to a feasible solution to the IRP: a fixed partition policy. The cost of the Location Based Heuristic solution to the IRP is then:

$$Z^{\text{LBH}} = \sum_{j=1}^{r} \phi(S_j).$$

Again, many different versions of this algorithm can be implemented. We have used a similar definition as in the ST Heuristic for the CVRP and have had success. Define $T_j = \{x_j\}$, for each $j = 1, 2, \ldots, n$ with $m = n$. Then, the values c_{ij} can be calculated exactly with little effort. We call this version the *Seed-Tours* (ST) Heuristic for which computational results are reported in Table 2.

5.1 A Lower Bound on Fixed Partition Policies

In order to assess the quality of the solutions produced by the Location Based Heuristic for the IRP, we must be able to compute a good lower bound on the best solution within the class of fixed partition policies.

For any fixed partition policy $\mathcal{P}$, let the partition be $\{X_j\}_{j=1}^m$ where the set X_j is served every $t(X_j)$ units of time, with a load of $t(X_j)D(X_j)$. Let $Z(\mathcal{P})$ be the total cost per unit time for this policy. Let $\overline{X}_i$ be the set in the partition $\{X_j\}_{j=1}^m$ that includes customer (retailer) x_i. Then we have the following:

$$\begin{aligned}
Z(\mathcal{P}) &= \sum_{j=1}^{m} \phi(X_j) \\
&= \sum_{j=1}^{m} \Big(\frac{1}{t(X_j)} [L_0(X_j) + K(X_j)] + \frac{1}{2} t(X_j) \sum_{i \in X_j} h_i D_i \Big) \\
&= \sum_{j=1}^{m} \Big(\frac{1}{t(X_j)} \frac{1}{D(X_j)} \sum_{i \in X_j} D_i [L_0(X_j) + K(X_j)] + \frac{1}{2} t(X_j) \sum_{i \in X_j} h_i D_i \Big) \\
&= \sum_{i=1}^{n} \Big(\frac{1}{t_i} \frac{D_i}{D(\overline{X}_i)} [L_0(\overline{X}_i) + K(\overline{X}_i)] + \frac{1}{2} t_i h_i D_i \Big) \qquad (\text{where } t_i \equiv t(\overline{X}_i)) \\
&\geq \sum_{i=1}^{n} \min_{t_i^* \geq 0} \Big\{ \frac{1}{t_i^*} \frac{D_i}{D(\overline{X}_i)} [L_0(\overline{X}_i) + K(\overline{X}_i)] + \frac{1}{2} t_i^* h_i D_i \;\Big|\; t_i^* D(\overline{X}_i) \leq Q \Big\}.
\end{aligned}$$

Let $g_\theta(x_i)$ be the cost of the minimum cost tour (including only transportation and fixed order costs), starting and ending at the depot, that serves a set of customers S with $x_i \in S$ and $D(S) = \theta$. If for a specific value of θ, no tour satisfies these conditions (e.g., $\theta < D_i$) then assign an infinite value to $g_\theta(x_i)$. The value θ is called the *total demand rate* of the tour. Then,

$$\begin{aligned}
Z(\mathcal{P}) &\geq \sum_{i=1}^{n} \min_{t_i^* \geq 0} \Big\{ \frac{1}{t_i^*} D_i \frac{g_{D(\overline{X}_i)}(x_i)}{D(\overline{X}_i)} + \frac{1}{2} t_i^* h_i D_i \;\Big|\; t_i^* D(\overline{X}_i) \leq Q \Big\}, \\
&\geq \sum_{i=1}^{n} \min_{t_i^* \geq 0, X_i^* \subseteq N} \Big\{ \frac{1}{t_i^*} D_i \frac{g_{D(X_i^*)}(x_i)}{D(X_i^*)} + \frac{1}{2} t_i^* h_i D_i \;\Big|\; t_i^* D(X_i^*) \leq Q \Big\}.
\end{aligned}$$

Let $D_i^* \equiv D(X_i^*)$, then the lower bound is

$$\sum_{i=1}^{n} \min_{t_i^* \geq 0, D_i^* \geq D_i} \Big\{ \frac{1}{2} h_i D_i t_i^* + \frac{D_i}{t_i^*} \frac{g_{D_i^*}(x_i)}{D_i^*} \Big| t_i^* D_i^* \leq Q \Big\}.$$

To solve this, for each customer x_i and each value of θ for which $g_\theta(x_i)$ is finite, solve the following problem:

$$f_\theta(x_i) = \min_{t \geq 0} \Big\{ \frac{1}{2} h_i D_i t + \frac{D_i}{t} \frac{g_\theta(x_i)}{\theta} \Big| t \leq \frac{Q}{\theta} \Big\} \tag{9}$$

Then, the lower bound on all fixed partition policies is:

$$\forall \mathcal{P}, \ Z(\mathcal{P}) \geq \sum_{i=1}^{n} \min_{D_i \leq \theta \leq D(N)} f_\theta(x_i). \tag{10}$$

Unfortunately determining the values $g_\theta(x_i)$ is in general NP-hard, since the Traveling Salesman Problem is a special case. Hence we use a dynamic programming procedure to find lower bounds to these values based on the following simple observation. The tours that define the values $g_\theta(x_i)$ are simple tours; no customers are visited more than once. This is the constraint which makes the computation intractable. Hence, we relax this constraint and allow customers to be visited more than once. This clearly provides a lower bound on the original values $g_\theta(x_i)$. It has however the misleading property that a tour that serves a set of customers with total demand rate θ may actually be visiting a set of customers whose total demand rate is less than θ. This will not cause any problems since the computed value will still represent a lower bound on the cost.

In Christofides, Mingozzi and Toth (1981), a dynamic programming procedure is implemented to find a lower bound on $g_\theta(x_i)$ for each customer x_i and for each value of θ ($D_i \leq \theta \leq D(N)$). The procedure was designed for the CVRP, but also works for this problem. Let $\psi_\theta(x_i)$ be the cost of the minimum cost route, without 2-loops (cycles of the form $\{\ldots, x_k, x_l, x_k, \ldots\}$), starting and ending at x_0, passing through x_i and with a total demand rate of θ. It is clear from the construction that $\psi_\theta(x_i) \leq g_\theta(x_i)$, $\forall i, \theta$, and hence replacing g with ψ in (9), still yields a lower bound on any fixed partition policy. The complexity of this lower bound is $O(n^2 \sum_{i=1}^{n} D_i)$.

The lower bound (10) can be further improved by using the observation that in any fixed partition policy every customer has exactly one vehicle arriving and leaving its location. Based on this, a subgradient procedure can improve the bound in much the same way as in Christofides, Mingozzi and Toth (1981) for the CVRP. The improvement comes from the fact that the set of routes obtained in the computation of the lower bound have cycles and are not customer disjoint. The idea is to assign a penalty on each customer and to recalculate the lower bound. Adjusting the penalties using the standard formula of Held, Wolfe and Crowder (1974) will result in new penalties and the lower bound is recomputed. After a series of iterations without an improvement in the lower bound we stop the procedure.

5.2 Computational Issues

As in the CVRP, we implement during Phase II an enhancement phase to better approximate the connection costs in the IRP. That is, the connection costs are accurate when exactly one customer is connected to a seed. As soon as more customers are added, the connection costs become only approximations.

Specifically, for every set of multipliers, as we search for a feasible solution to the location problem, using the GREEDY procedure described in Section 3.2, we implement the following procedure to construct a feasible inventory-routing solution.

In the procedure below, we assume the concentrators are indexed from 1 to m in increasing order of the knapsack solutions. Therefore, concentrator 1 is the "best" concentrator, while concentrator m is the "worst". In the procedure below, S_j represents the set of customers that are served with seed customer j and L_j represents the length of the nearest insertion tour for the customers in S_j.

for $v = 1, 2, \ldots, m$ **do begin**

  for $i = 1, 2, \ldots, n$ CONNECTED[i]=FALSE

  select concentrators $x_1, x_2, \ldots, x_v$

  for $j = 1, 2, \ldots, v$ **do begin** $S_j = \{x_j\}$, $L_j = 2d_j$ **end**

  while $\exists i$ **such that** CONNECTED[i]=FALSE **do begin**

    for $i = 1, 2, \ldots, n$ **do**

      if CONNECTED[i]=FALSE **then begin**

        for $j = 1, 2, \ldots, v$ **do begin**

          let $\hat{c}_{ij}$ = **cost of adding** i **to tour** S_j **using nearest insertion**

          let $\hat{s}_{ij} = \min_{t \geq 0} \left\{ \frac{K(S_j)+L_j+\hat{c}_{ij}}{t} + \frac{1}{2} t \sum_{k \in S_j \cup \{i\}} h_k D_k \middle| t \leq \frac{Q}{D(S_j)+D_i} \right\} - \phi(S_j)$

          end

        let $\hat{s}_i = \min_{1 \leq j \leq v} \{\hat{s}_{ij}\}$

        let MININDEX[i]=argmin$_{1 \leq j \leq v} \{\hat{s}_{ij}\}$

```
            let PENALTY[i]=ŝ_i − min_{j≠MININDEX[i]}{ŝ_ij}
            end
    let i* = argmax_i{PENALTY[i]}
    let j* = MININDEX[i*]
    add i* to tour j* using the nearest insertion procedure
    let S_j* ← S_j* ∪ {i*}
    let L_j* ← L_j* + ĉ_i*j*
    end
end
```

Once a feasible solution to the IRP is found using this procedure its cost is compared to the cost of the current best solution and the better one is kept.

5.3 Computational Results

In Table 2 below, we present the results of the implementation of our algorithm on the fifty-customer problem from Christofides and Eilon (1969). It should be clear that the empirical performance of the heuristic depends on the relative importance of the transportation cost and the inventory cost, e.g., if the individual fixed costs (K_i) are large relative to the transportation costs, then the heuristic will perform extremely well, since the order quantities (and therefore order intervals) selected will be close to those minimizing the major part (the inventory cost) of the objective function. For that reason we choose small values for the inventory parameters which means that we evaluate the performance of the heuristic under *unfavorable conditions.* We varied the fixed order costs for individual retailers from 0 to 15, and used two different holding costs, 0.5 and 1. The vehicle capacity is exactly the one used in the CVRP, for this problem it is 160. Retailer demands were distributed uniformly between 1 and 10. The lower bound (FPP LB) is calculated as it is described in Section 5.2 using the subgradient procedure.

Table 2

Computational Results on the Location Based Heuristic

Problem No.	Size	Q	h	K	FFP LB	Alg. ST[1]	%Error	CPU Time[2]
[1]	50	160	1.0	0	471.9	502.9	6.6	455.0
[2]	50	160	1.0	5	600.9	637.2	6.0	111.0
[3]	50	160	1.0	10	705.8	742.9	5.3	587.7
[4]	50	160	1.0	15	796.4	832.1	4.5	605.0
[5]	50	160	0.5	0	346.3	373.0	7.7	103.6
[6]	50	160	0.5	5	441.4	468.0	6.0	594.2
[7]	50	160	0.5	10	518.2	543.7	4.9	718.8
[8]	50	160	0.5	15	583.5	615.9	5.5	460.3

([1]) Seed-Tours Heuristic.

([2]) In seconds on a Sun Sparc 2.

In the following table, we list certain characteristics of the solutions provided by the LBH on all eight problems. "Number of tours" represents how many sets make up the fixed partition. The "Number of retailers in each tour" specifies the size of each of these sets. The "Vehicle loads in each tour" specifies the load (as a percentage of the vehicle capacity) that is sent out every cycle to serve each set. We see that in almost two thirds of the cases the vehicle capacity is a tight constraint on the load.

Table 3

Description of Solution Provided by the LBH

	Characteristics of the Location Based Heuristic Solution		
Problem No.	Number of tours	Number of retailers in each tour	Vehicle load in each tour
[1]	4	15,20,14,1	100.0%,100.0%, 96.4%, 3.7%
[2]	4	12,16,9,13	100.0%,100.0%, 82.7%, 99.3%
[3]	5	13,14,10,12,1	100.0%,100.0%, 99.0%,100.0%, 6.7%
[4]	5	11,10,10,11,8	99.4%,100.0%,100.0%,100.0%, 80.2%
[5]	4	13,14,9,14	100.0%,100.0%, 85.2%,100.0%
[6]	5	12,8,11,9,10	100.0%,100.0%,100.0%,100.0%, 98.2%
[7]	7	9,5,10,7,8,10,1	100.0%, 76.2%,100.0%, 96.7%,100.0%,100.0%, 9.5%
[8]	8	6,7,9,6,8,9,1,4	100.0%,100.0%,100.0%, 94.2%,100.0%,100.0%, 11.0%, 65.8%

6 Applications to Other Problems

We have shown in this paper that the Location Based Heuristic works well. In fact, the methodology has much potential. Many different features of the routing problem can be incorporated into the structure of the heuristic; for example, we are currently working on a more general version of the vehicle routing problem where in addition to a capacity constraint on the total load that can be carried by a vehicle, there is a distance constraint on the length of each route traveled by a vehicle.

We also apply the technique presented here to the Capacitated Minimum Spanning Tree Problem (CMST). This problem has numerous applications in the design of local access tree networks in communication networks, see Altinkemer and Gavish (1988). The results on this problem will be summarized in an accompanying paper.

Acknowledgements

We would like to thank Prof. Michael Pinedo and Prof. Yechiam Yemini for allowing us the use of their computers. We would also like to acknowledge IBM's contribution via an equipment gift to Columbia University of the RS6000 Model 550.

References

Altinkemer K. and B. Gavish. 1988. Heuristics with Constant Error Guarantees for the Design of

Tree Networks. *Management Science* **34**, pp. 331–341.

Altinkemer K. and B. Gavish. 1991. Parallel Savings Based Heuristics for the Delivery Problem. *Operations Research* **39**, pp. 456–469.

Anily, S. and A. Federgruen. 1990. One Warehouse Multiple Retailer Systems with Vehicle Routing Costs. *Management Science* **36**, pp. 92–114.

Barcelo, J. and J. Casanovas. 1984. A Heuristic Lagrangian Algorithm for the Capacitated Plant Location Problem. *Eur. J. Opl. Res.* **15**, pp. 212–226.

Bienstock, D. , J. Bramel and D. Simchi-Levi. 1991. A Probabilistic Analysis of Tour Partitioning Heuristics for the Capacitated Vehicle Routing Problem with Unsplit Demands. To appear in *Mathematics of Operations Research.*

Bramel, J. , Coffman, E. G. Jr. , Shor, P. and D. Simchi-Levi. 1991. Probabilistic Analysis of Algorithms for the Capacitated Vehicle Routing Problem with Unsplit Demands. *Operations Research* **40**, No. 6, pp. 1095–1106.

Christofides, N. 1985. Vehicle Routing. Lawler, E. L., J. K. Lenstra, A. H. G. Rinnooy Kan and D. B. Shmoys (eds.), *The Traveling Salesman Problem,* John Wiley & Sons Ltd., pp. 431–448.

Christofides, N. and S. Eilon. 1969. An Algorithm for the Vehicle Dispatching Problem. *Opns. Res. Quart.* **20**, pp. 309–318.

Christofides, N., A. Mingozzi and P. Toth. 1979. The Vehicle Routing Problem. In *Combinatorial Optimization.* John Wiley & Sons, New York.

Christofides, N., A. Mingozzi and P. Toth. 1981. Exact Algorithms for the Vehicle Routing Problem Based on Spanning Tree and Shortest Path Relaxations. *Math. Prog.* **20**, pp. 255–282.

Clarke, G. and J. W. Wright. 1964. Scheduling of Vehicles from a Central Depot to a Number of Delivery Points. *Operations Research* **12**, pp. 568–581.

Desrochers, M. and T. W. Verhoog. 1989. A Matching Based Savings Algorithm for the Vehicle Routing Problem. Cahier du GERAD G-89-04, École des Hautes Études Commerciales de Montréal.

Fisher M. L. 1991. Optimal Solution of Vehicle Routing Problems using Minimum K-Trees. Working Paper, The Wharton School, University of Pennsylvania.

Fisher M. L. and R. Jaikumar. 1981. A Generalized Assignment Heuristic for Vehicle Routing. *Networks* **11**, pp. 109–124.

Gallego, G. and D. Simchi-Levi. 1990. On the Effectiveness of Direct Shipping Strategy for the One Warehouse Multi-Retailer R-Systems. *Management Science* **36**, pp. 240–243.

Gillett, B. and L. Miller. 1974. A Heuristic Algorithm for the Vehicle Dispatching Problem. *Operations Research* **22**, pp. 340–349.

Held, M. , P. Wolfe and H. P. Crowder. 1974. Validation of Subgradient Optimization. *Mathematical Programming* **6**, pp. 62–88.

Herer, Y. and R. Roundy. 1990. Heuristics for a One Warehouse Multi-Retailer Distribution Problem with Performance Bounds. Technical Report No.916, Cornell University, Ithaca, NY.

Johnson, D. B., A. Demers, J. D. Ullman, M. R. Garey and R. L. Graham 1974. Worst-Case Performance Bounds for Simple One-Dimensional Packing Algorithms. *SIAM J. Comp.* **3**, pp. 299–325.

Klincewicz, J. G. and H. Luss. 1986. A Lagrangian Relaxation Heuristic for Capacitated Facility Location with Single-Source Constraints. *J. Opl. Res. Soc.* **37**, pp. 495–500.

Mole, R. H. and S. R. Jameson. 1976. A Sequential Route-Building Algorithm Employing a Generalised Savings Criterion. *Operational Research Quarterly* **27**, pp. 503–511.

Nauss, R. M. 1976. An Efficient Algorithm for the 0-1 Knapsack Problem. *Management Science* **23**, pp. 27–31.

Neebe, A. W. and M. R. Rao. 1983. An Algorithm for the Fixed-Charged Assigning Users to Sources Problem. *J. Opl. Res. Soc.* **34**, pp. 1107–1113.

Pirkul, H. 1987. Efficient Algorithms for the Capacitated Concentrator Location Problem. *Computers in Operations Research,* vol. 14, No. 3, pp. 197–208.

Pureza, V. M. and P. M. Franca. 1991. Vehicle Routing Problems via Tabu Search Metaheuristics. Publication CRT-747. Centre de Recherche sur les Transports, Montréal.

Rosenkrantz, D. J., R. E. Sterns and P. M. Lewis II. 1977. An Analysis of Several Heuristics for the Traveling Salesman Problem. *SIAM J. on Computing* **6**, pp. 563–581.

Simchi-Levi, D. and J. Bramel. 1990a. On the Optimal Solution Value of the Capacitated Vehicle Routing Problem with Unsplit Demands. Working Paper, Department of IE&OR, Columbia University, New York.

Simchi-Levi, D. and J. Bramel. 1990b. Probabilistic Analysis of Heuristics for the Capacitated Vehicle Routing Problem with Unsplit Demands. Working Paper, Department of IE&OR, Columbia University, New York.