Consistency and Robustness in Location-Routing

S. Salhi

School of Mathematics and Statistics
The University of Birmingham
Edgbaston, Birmingham, B15 2TT
e-mail: S.Salhi@Birmingham.ac.uk

G. Nagy

School of Computing and Mathematical Sciences
The University of Greenwich
Wellington Street, Woolwich, London, SE18 6PF
e-mail: G.Nagy@Greenwich.ac.uk

Abstract

Many researchers believe that it is inappropriate to rely on routing decisions when locating depots. Such a conception has unfortunately led to a shortage of published work in this area. In this study we briefly review the recent work on location-routing heuristics and show using robustness analysis that these models are as reliable as location-first routing-second methods. We employ simulation to show that combined models consistently produce solutions of higher quality than sequential ones. Selection criteria based on simulation are developed for choosing the most appropriate locational solution. Computational results based on problem sets of 400 customers and 15 to 25 depots are reported.

Keywords heuristics, location-routing, simulation, consistency, robustness

Studies in Locational Analysis Issue 13 June 1999 3-19

1 Introduction

In distribution management the location of the depots, the size of the vehicle fleet and the schedule of the delivery routes play an important part in the success of the company. These transportation-related activities are usually addressed independently, although in practice they are interrelated. The assumption usually used when solving the plant location problem is that a tour consists of a visit to a single customer. This 'strong' hypothesis may be appropriate only when customers require full truck loads.

The combined location-routing problem (CLVRP) is an agglomeration of two already difficult problems, namely depot location and vehicle routing. We note that both sub-problems and the CLVRP itself are all NP-hard. Locational decisions, which correspond to long-term commitments, are usually taken at a strategic level due to their high investment, whereas routing decisions are solved at an operational level.

In the O.R. literature applied to distribution studies, there is relatively little about the combination of routing and depot location when compared to the intensive material produced in depot location (see Mirchandani and Francis (1990), and also Drezner (1995)) and in vehicle routing (see Laporte (1992)). An interesting survey paper, which addresses both exact and heuristic approaches for the CLVRP, is given by Laporte (1988). Balakrishnan, Ward and Wong (1987) also survey location-routing and highlight some interesting research avenues in this area. The CLVRP has a large variety of applications which arise from different areas of distribution such as public and private sectors, engineering, environment, etc. A recent summary of past work on the CLVRP as well as new developments and possible applications can be found in Salhi and Fraser (1996).

The slow progress in the development of location-routing methodology is mainly due, in our view, to the fact that many location scientists believe that CLVRP models are not easily or directly usable. The question which strikes many researchers is how to rely on routing decisions, which are likely to change from time to time, when solving the problem of location which is of a more strategic nature. They may believe that it is unjustifiable to use CLVRP models because of the above difference in the nature of its sub-problems. We think that the above observation does not invalidate the CLVRP as a research topic; however it needs to be investigated how the CLVRP model works on a time horizon.

Thus, the aim of this study is to

(i) demonstrate how important it is to integrate location and routing, by showing that location-routing methods consistently produce better solution

than sequential algorithms,

- (ii) propose a selection rule for choosing the most appropriate locational solution based on simulation, and
- (iii) show, using robustness analysis, that location-routing models are stable enough to be reliable for practical use.

The paper is organised as follows. In section 2, we review the literature on location-routing with an emphasis on recently developed nested methods. In section 3, we analyse the consistency of location-routing methods and present two selection rules to choose high quality locational solutions. Robustness analysis is dealt with in section 4. Computational results are provided in section 5 and we summarise our conclusions and highlight some research issues in section 6.

2 Location-routing models

The CLVRP can be formulated as a 0-1 ILP problem and can be solved optimally only for small sized problems. Heuristics seem to be the best way forward to approach these hard combinatorial problems. Heuristics have the advantage of producing more than one single solution so providing the user with flexibility in choosing the right solution which may include other unquantifiable and urgent factors. In addition, heuristics are capable of finding good solutions in a reasonable amount of computing time. They are also easy to understand, to modify and to implement. In this section, we review three main heuristic approaches for location-routing problems, namely sequential, iterative and nested. Since the first two approaches are well known in the literature (see Laporte (1988)) we shall discuss these briefly and concentrate more on the third one which is recently developed by the authors and also it is the one which we consider for assessing the consistency and the robustness of location-routing.

Sequential models:

One commonly used way to approach the CLVRP is to decompose the problem into two distinct subproblems namely location and routing. In the first stage the locational problem is solved and the arising set of depots is then used to give a vehicle routing problem which is solved in the second stage. Usually, exact methods are used in the first stage and heuristic algorithms in the second. Such an approach - although it is simple and easy to implement - suffers from the drawback of suboptimising the whole problem, as noted by Rand (1976). The effect of ignoring the routing aspect when locating depots was also shown in Salhi and Rand (1989).

Iterative models:

Most combined location-routing heuristics in the literature fall into this category. These are extension of the sequential models where the information on routing is fed back into an approximate location problem which is then solved to find the new selected depots. The process is repeated until there is no change in the solution either in routing or in location. A notable exception is the "tree-tour heuristic" of Jacobsen and Madsen (1980). They also developed two iterative methods which combine existing routing heuristics with some known location methods. Perl and Daskin (1984) used an iterative approach consisting of three phases, namely location, allocation and routing. A combination of heuristics and exact methods were used to solve these three subproblems. The iterative process is repeated until no further improvement is detected. Salhi and Fraser (1995) developed an iterative process similar to the one of Perl and Daskin (1984) except vehicles of different capacities are used and the updating at the location phase is carried out using a different approximation of the route configuration. Encouraging results were obtained when compared to the sequential method.

Nested models:

The iterative approaches, though they have recently been shown to produce better quality solutions than the sequential methods, still have some drawbacks. They let the locational algorithm run until the end and then restart it using some more routing information. Thus, if the routing solution does not provide enough information at a given iteration for the location phase, the new solution found by the method may not be as informative as one would have hoped for. From the modelling point of view, there is another shortfall of iterative methods. These methods treat the two constituent components of the location-routing problem as if they were on the same footing. Observe, that a location-routing problem is essentially a location problem, with the routing factor taken into consideration. So, instead of treating the two sub-problems as equal we observe a hierarchical structure, with location as the main problem and routing as a subordinate one. We refer to this approach "nested" because the routing stage is embedded into the location phase. These nested methods for the CLVRP were successfully developed by Nagy and Salhi (1996a, b). This concept of hierarchy was also strongly emphasised by Balakrishnan, Ward and Wong (1987). These three methodologies - namely sequential, iterative and nested - are shown in Figure 1.

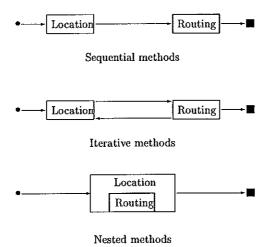


Figure 1. An illustration of the three types of solution methods for location-routing problems.

Since nested methods are not as widely known as the previous two approaches we shall provide the main points of the method. A general skeleton of a nested method for location-routing problems can be described as follows:

- 1. Obtain an initial locational solution.
- Use a suitable multi-depot routing algorithm to find the routing cost for the whole system.
- 3. Find a neighbouring locational solution for exploration.
- Define a subset of the set of all depots and customers using a suitable rule.
- Apply either a suitable multi-depot routing method or an appropriate route length approximation technique for the above set of cities to find the routing cost for this subset.
- 6. Repeat steps 3 to 5 until the neighbourhood is completely explored.
- Select the best neighbouring location (even if it does not lead to an improving solution).

8. Repeat steps 2 to 7 until a suitable stopping rule is met.

In step 2, the multi-depot routing problem is solved using the algorithm of Salhi and Sari (1997). This was chosen for its speed (as several runs have to be executed), and for its capability of producing good quality solutions.

In step 3, the neighbourhood of a current solution is defined by dropping an existing depot, adding a new depot to the existing set of open depots, and finally simultaneously closing a given depot and opening a new depot from the set of customers served by such a depot. This step originates from the local search heuristic initially introduced by Kuehn and Hamburger (1964).

In step 4, we rely on the assumption that the influence of a change in location is a subset of the set of all cities, incorporating the depot concerned and its customers, and also some neighbouring depots with their associated customers. Such an area is referred to as a "region". The reason behind using these "regions" is that they save us having to re-calculate routing costs for the entire system, when we know that routing is unlikely to change everywhere in the system. A good choice of region would avoid both excessive computational time associated with too large regions and loss of accuracy associated with too small regions. There are several ways of constructing these regions. Nagy and Salhi (1996a) developed two variants of these methods; one using the idea of "distance", the other using the idea of "proximity" inspired by computational geometry.

In step 5, two approaches may be used to compute the routing cost. One applies the routing algorithm of Salhi and Sari (1997), the other is based on a new route length approximation formula recently developed by the authors, see Nagy and Salhi (1996b). Variants of the nested method using the routing algorithm are referred to as REGIONAL and the ones using route length approximation are called ESTIMATE. Empirical testing has shown that both approaches perform better than the sequential method which is used here as a benchmark. In addition, ESTIMATE was found to be significantly faster than REGIONAL.

3 Analysing the consistency of location-routing methods

Balakrishnan, Ward and Wong (1987) observed that modelling difficulties may arise when combining the long-term strategic problem of location with the short-term tactical vehicle routing problem. Customer demands may change over time, and the routing solution can be changed from time to time, to cater for such changes in demand. Depot locations, however, cannot be modified at short notice.

The issue of consistency and robustness is addressed in the seminal papers by Gupta and Rosenhead (1968) and by Rosenhead, Elton and Gupta (1972). A review of robustness and of a number of related methodologies for dealing with issues of uncertainty is given in Rosenhead (1989). In these works, robustness is considered to be the *flexibility* of a plan, i.e. its ability to accommodate changes due to changed input values.

In our investigation, we seek a solution to the location-routing problem, such that the locational solution is unchanged over some given time period, but the routing solution may change following changes in customer demand. We divide our time horizon into a number of time intervals and consider that demands may be different for these intervals. We assume that the time intervals are of equal length and that demands are constant within these intervals. Furthermore we assume that the customer set remains unchanged throughout the time horizon. Note that in the literature a growing demand is usually assumed; however, we consider here the case of fluctuating demand. Each time interval corresponds to a (simulated) location-routing problem. Before we present an overview of the different solutions, we describe the notation we use in this section.

Notation:

I is the set of all cities (depots and customers), |I| = m

J is the set of time intervals, $\mid \mathbf{J} \mid = n$,

 \hat{Q}^i is the average demand of the $i^{ ext{th}}$ city during our time horizon.

 Q^i_j denotes the demand of the $i^{ ext{th}}$ city during the $j^{ ext{th}}$ time interval, $i \in \mathbf{I}$, $i \in \mathbf{J}$.

A number of solutions are looked at. They are explained in some detail in the remainder of this section. We provide a brief introduction here. Firstly, we investigate the basic solutions provided by the location-routing and the sequential methods, say solutions A and O respectively. Then, we generate a lower bound (solution B), which is a good but infeasible solution. Finally, two selection rules are developed for creating solutions C and D, which are better than the basic solution A. Figure 2 gives an illustration of the different solutions.

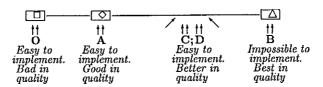


Figure 2. A representation of the analysis The basic solution (Solution A):

We first simulate the scenario when a decision on depot locations is made on the basis of average (forecasted) demand values, while in reality demands fluctuate around these average values. Now, we wish to know how much our costs are over a time horizon consisting of m simulated problems, each with a demand set $(Q_j^i, i \in \mathbf{I}, j \in \mathbf{J})$. The pseudo-code for this simulation procedure

- solve location-routing problem using average demands \hat{Q}
- take arising set of open depots D₁
- for $i = 1, \ldots, m$ do
- solve routing problem with depot set D_1 using demands Q_i
- cost of (each) solution is C_1^i final (average) cost is $C_1 = \frac{\sum_{i=1}^m C_1^i}{m}$

This simulation is called "short" simulation because it entails running the location-routing algorithm only once, with a number of additional calls to the routing algorithm. If we made a decision on depot locations using average values and chose the set of depots D_1 to be opened, we would face a cost C_1 . This simulation is based on the idea of evaluating the solution produced by our solution methods, w.r.t. the "new" information on changing demand levels, but not taking this information into account when finding the solution to the locational problem. This is practical as we do not have to change the locational solution every time. We note that the above simulation procedure using the output of the sequential method will give us solution O, with depot set D_0 and cost C_0 . We also wish to compare our solution C_1 with the solution C_0 . According to our results, which will be presented later, $C_1 < C_0$. Furthermore, in most cases we found that $C_1^j \leq C_0^j (j \in \mathbf{J})$. This shows that over the m simulated problems the solution based on our location solution using the location-routing method is consistently superior to the one found by the sequential method.

Although our solution is found to be better, we need to assess how good such a solution is and whether it is possible to find another locational solution that improves on the existing one. The remainder of this section deals with this issue.

Obtaining a lower bound (Solution B):

One way to examine the goodness of a solution is to compare against a lower bound solution. This is obtained by solving the location-routing problem m times, each time with a different demand set. The procedure provides a high quality solution which is infeasible: it violates the assumption that the depot set cannot be changed. We present a pseudo-code for this approach.

- for $i = 1, \ldots, m$ do
- solve location-routing problem using demands Q_i
- cost of (each) solution is C_2^i and the arising set of open depots is D_2^i
- final (average) cost is $C_2 = \frac{\sum_{i=1}^{m} C_2^i}{m}$

The above procedure is called "long" simulation as it entails running the location-routing algorithm m times. We observe that if we were able to change the depot set from time to time, than we should choose the sets $D_2^1, D_2^2, \ldots, D_2^m$ in succession, yielding a cost of C_2 . While this may not be possible in reality, the values of C_2 produced here are useful as lower bounds. The real cost, say C_3 , would be obviously higher since the restriction of having the same locational solution is required. It is desirable now to seek a few other locational solutions which are both feasible and have a cheaper total cost than C_1 .

A simple selection rule (Solution C):

One obvious way is to apply the short simulation model using the locational solution for each depot set $D_1^1, D_2^2, \ldots, D_2^m$ and then select the set $D_3 \in \{D_2^j; j \in J\}$ that yields the cheapest total cost over the time horizon. This is similar to solution A except that the short simulation model is used for a number of locational solutions rather than only for the "basic" solution. The pseudo-code for this approach is as follows:

- execute long simulation to find the sets D_2^j
- for each $i = 1, \ldots, m$ do
- execute short simulation for depot set D_2^i , cost is denoted by C_3^i
- final cost is $C_3=\operatorname{Min}\ \{C_3^i,i\in \mathtt{J}\}$ and choose D_3 as the configuration that

produces the cost C_3 .

The location solution found above is a solution of one of the m simulated

problems. Such a restriction has a serious effect as it limits the search and hence the obtained solution can be suboptimal and misleading. Nevertheless, such a solution can be easily implemented and we shall see later that it yields a cheaper total cost than C_1 . Our aim now is to generate a solution which relies on all the solutions to each of the m simulated problems, not just on one of them.

A more sophisticated selection rule (Solution D):

Observing the depots sets $D_2^1, D_2^2, \ldots, D_2^m$, we may be able to find a set D_4 , which yields a lower cost than D_3 . We wish to create such a set D_4 using the output of the "long" simulation. Before we describe the selection rule we decided to adopt for creating such a new solution, we present some notation. Notation:

 ν_1 and ν_2 are two threshold values; $0 \le \nu_2 < \nu_1 \le 1$ $Y_{ij} = \begin{cases} 1 \text{ if city } i \text{ occurs as a depot in the solution to the } j^{\text{th}} \text{ problem } 0 \text{ otherwise.} \end{cases}$ P_j is the probability of a city being a depot, and is calculated as:

 $P_i = \left(\sum_{i=1}^m Y_{ii}\right)/m$ Note that both Y_{ij} and P_j may be generated as by-products of the long simulation.

Our selection rule is realised by the following pseudo-code:

- apply the long simulation
- compute the probability of a city being a depot P_i for all $i \in I$
- divide the set of cities into three subsets:
- the set of cities with high P_i ; $S_A = \{i \in \mathbf{I}; P_i \geq \nu_1\}$
- the set of cities with medium $P_i; S_B = \{i \in \mathbf{I}; \nu_2 \le P_i \le \nu_1\}$
- the set of cities with low P_i ; $S_C = \{i \in I; P_i < \nu_2\}$.
- determine $D_4 \in \{S_A \cup S; \forall S \subseteq S_B\}$, i.e. the configuration which consists
 - all of the cities from S_A ,
 - some of cities from S_B, and
 - none of the cities from S_C , such that C_4 , the cost of the solution with depot set D_4 , is minimal.

Note that the cardinality of S_B depends on the parameters ν_1 and ν_2 , but is usually quite small since our location-routing method is stable enough as will be shown in section 4. The probability values are either too high (near

1) or too small (near 0) but there are not many lying in the middle range. Thus, it is easy to calculate C_4 .

We shall see in section 5, that $C_4 < C_3$, thus the solution we recommend for implementation is D_4 .

4 Robustness analysis

Having now shown that it is possible to choose a good location which yields good quality solutions when applying routing, it is also interesting as a secondary study to demonstrate how robust our CLVRP solution algorithms are. In other words, we would like to have a method which provides a set of depots which does not change dramatically when some customers change their demands. For instance, the sequential method is very robust since the location will not be altered if customers demand change. However, we should not be obsessed by having a very robust method which in reality yields very inferior results. In this section we will show that location-routing methods are robust enough while producing good quality solutions.

We define robustness as the ability of a method to produce similar solutions under slight perturbations to its input. This arises from the definition of the robustness of a plan or decision. A method will be considered robust in our paper if the decision it produces is robust as described in Rosenhead et al. In our case, the "slight perturbation" is the changing demand, while a "solution" is a set of depots D. We note that while consistency was based on cost (an inherently continuous measure), robustness is based on a discrete measure, the set of depots produced by the location-routing methods. Thus, "similarity" of solutions is more difficult to define. Two solutions are considered "similar", if the symmetric difference of their depot sets D_1 and D_2 has low cardinality. This means that there are only a few cities which belong to either D_1 or to D_2 , but not to both. This can be readily generalised for m depot sets D_1, D_2, \ldots, D_m . We introduce a penalty for cities belonging to some depot sets but not to all using the probability of a city being chosen for a depot site. This measure was defined in earlier and can be found using "long" simulation. If it is one or zero, then there is no penalty; otherwise the penalty is proportional to the difference of the probability from either one or zero (whichever is the smaller). Putting together these penalty values we get the robustness measure of a method:

$$R = (2 \sum_{i=1}^{n} |P_i - 0.5| /n)$$

We note that $0 \le R \le 1$. For instance, if R = 1, then for all sets of customer demands, the same locational solution is produced. An example of this is

the sequential method, as its output is not influenced by changes in customer demand. However, had we generated completely random solutions for each of the time intervals with approximately the same number of depots, then for 400 cities and about 20 depots we would still have $R \approx 0.9$. This is due to the fact that the number of cities is much larger than the number of depots and thus most cities tend to be not used as depots, making the solution more robust. We also note that even if the number of depots is allowed to vary between 0 and 400, we would still get $R \approx 0.25$ for a completely random solution. We can illustrate the robustness measure with an example. Suppose that in 9 out of 10 of our simulation runs the same solution is produced (say 20 depots out of 400 cities), however, in the tenth run we have an extra depot in the solution. The robustness measure of the method providing these solutions would then be 0.999. We also wish to note that robustness may sometimes be a misleading measure. Both Gupta and Rosenhead (1968) and Rosenhead, Elton and Gupta (1972) mention an example where a number of nearby and approximately equally good depot locations cause the robustness value to go down, as slight perturbations can change the solution set to include a different depot from this cluster. The concept of robustness can be extended to eliminate this distortion.

5 Computational results

The heuristics proposed here were written in VAX Fortran and executed on a VAX 4000-500 computer at the University of Birmingham. They were evaluated using empirical testing. We first present our data generation. Our results are analysed under three headings, namely comparison within methods, comparison between methods, and robustness. To demonstrate the consistency of the CLVRP models we based our analysis on the nested methods. A similar investigation could also be carried out for the iterative methods.

Data generation and plan for comparison:

Two sets of coordinates and demands, each consisting of 400 customers, are tested with three levels of fixed depot costs to provide solutions around 25, 20 and 15 depots respectively: this gives 15 problems in total. These data sets are generated using some of the problems given by Christofides, Mingozzi and Toth (1979). We used problems 3 and 11 to produce larger problems with 400 customers. The problem sets are numbered as 11, 12, 13, 21, 22 and 23: problem 11 is the first problem set with low depot costs, etc. We also note that customers are clustered for problems 21 to 23. Furthermore, maximum capacity and maximum distance constraints were chosen so that they are

tight for all problems. Ten consecutive time intervals were considered, with ten different demand values. These values were generated randomly, using standard distribution with the demands given by Christofides, Mingozzi and Toth (1979) as the average demands and a standard deviation of 0.01. In this study we have set the parameters ν_1 and ν_2 to 0.1 and 0.9 respectively.

Three methods were examined, namely the sequential method denoted by CLASSIC, and the recently developed nested methods REGIONAL and ESTIMATE. The results of the different simulation runs are given in Table 1.

		11	12	13	21	22	23	Average
CLASSIC	C_0	5035	5225	6226	6491	6789	7558	6221
	C_1	4877	5101	5967	6132	6395	7009	5914
REGIONAL	C_2	4811	5011	5832	6103	6233	6953	5824
	C_3	4863	5101	5943	6132	6388	7009	5900
	C_4	4854	5101	5909	6132	6357	6999	5892
	C_1	4818	5101	5903	5959	6400	6989	5862
ESTIMATE	C_2	4793	5032	5785	5903	6357	6939	5802
	C_3	4818	5101	5877	5959	6388	6977	5841
	C_4	4811	5077	5832	5933	6388	6953	5832

Table 1. Simulation results

Comparison within methods:

It can be shown that the theoretically predicted relationship between C_1, C_2, C_3 and C_4 holds for both REGIONAL and ESTIMATE and for all the six problem sets. (For CLASSIC, all the simulated solutions would still give C_0 and are thus not tabulated.) Quantifying the above relationship, we calculate percentage improvements in order to see how much improvement C_2, C_3 and C_4 represent w.r.t. C_1 . On average, C_2 is about 1.52% better than C_1 in the case of REGIONAL; for ESTIMATE, this figure is 1.02%. C_3 produces a0.23% and a0.36% improvement for REGIONAL and ESTIMATE respectively. C_4 is on average 0.37% better than C_1 if we use REGIONAL; it gives a0.51% improvement for ESTIMATE. While the gap between C_1 and C_2 is already quite small, the solution producing C_4 may cut this gap up to half its original size. Thus we may conclude that while using no simulation and relying on average values for future levels of demands gives already satisfactory solutions, the selection rule developed earlier produces even more desirable solutions, which are very near to the solution values of the lower bound C_2 . We remember that these values are likely to be unattainable by using a constant set of solutions and therefore surmise that our selection rule provides the best way of finding a single set of depots which behave well under changing conditions.

Comparison between methods:

According to Table 1, the costs C_0 produced by CLASSIC are always higher than the highest costs C_1 given by the nested methods. This serves as yet another justification for the use of nested methods. Comparing RE-GIONAL and ESTIMATE the picture is not so clear any more. While pairwise comparison of the respective values in almost all cases shows ES-TIMATE being marginally better than REGIONAL, the range of values (C_1, C_2, C_3, C_4) for REGIONAL overlaps with the range for ESTIMATE.

	CLASSIC	REGIONAL	ESTIMATE
11	1.00	0.987	0.995
12	1.00	0.981	0.987
13	1.00	0.978	0.980
21	1.00	0.996	0.991
22	1.00	0.975	0.993
23	1.00	0.992	0.993
Average	1.00	0.985	0.990

Table 2. Robustness scores

Robustness results:

The robustness values of CLASSIC, REGIONAL and ESTIMATE are tabulated in Table 2. We have already noted that the robustness value of CLASSIC is always 1. The average robustness values of REGIONAL and ESTIMATE are 0.985 and 0.990, respectively. This shows that both methods are robust. ESTIMATE is slightly more robust, perhaps due to the fact that it is more similar to CLASSIC than REGIONAL. In other words, small changes in customer demand do not influence the output of the location-routing methods used here significantly.

6 Conclusions and future research directions

A comparison which entailed simulating the behaviour of a distribution system over a period of time was carried out for the case of location-routing.

Selection criteria for choosing the most appropriate location within the framework of location-routing were proposed. These criteria were based on simulation. We found that the new results have enhanced the benefit of location-routing methods over sequential methods in all the test problems we looked at. A similar exercise could also be performed to see how consistent the iterative methods are, but in our view it is most likely that the same conclusion would be obtained.

The selection rules we used seem to us both logical and flexible. Although they produced good solutions, other suitable measures can also be worth considering. For instance, debility, which is the opposite of robustness, could be introduced into our selection rule. Debility (see Caplin and Kornbluth (1975)), is a measure to guide the decision maker to exclude the bad solutions from the analysis. In other words, a decision with a high value of debility is a bad decision and should not be pursued. Other measures may combine both robustness and debility as recently used by Giannikous and Rizakou (1995) for locating obnoxious facilities.

In this study, the same importance (weight) is given to each of the simulated problems (time intervals). A further investigation may consider different weights based on historical data and other factors to assess further the impact of variability of customer demand.

Our robustness score was based on individual sites without considering the effect of interference arising between neighbouring sites. This is important as the sites close to each other may compete between themselves for the allocation of customers and hence share the benefit. Such cities may therefore produce a less attractive robustness score although if they were considered as a single location, their contribution to the robustness score would increase.

Acknowledgments We would like to acknowledge the financial support of the Committee of Vice-Chancellors and Principals of the United Kingdom (ORS Awards) and of the School of Mathematics and Statistics, The University of Birmingham. We are also grateful to the referee whose comments improved the presentation of the paper.

References

A. BALAKRISHNAN, J. E. WARD and R. T. WONG (1987), Integrated facility location and vehicle routing models: recent work and future prospects, American Journal of Mathematics and Management Science, 7, 35-61

- [2] D. A. CAPLIN and J. S. H. KORNBLUTH (1975), Multiple investment planning under uncertainty, Omega, 3, 423-441
- [3] N. CHRISTOFIDES, A. MINGOZZI and P. TOTH (1979), The vehicle routing problem, in N. Christofides, A. Mingozzi, P. Toth and C. Sandi (eds) Combinatorial Optimization, Wiley, Chichester, 315-338.
- [4] Z.DREZNER (Ed.), (1995), Facility location: A survey of applications and methods, Springer, Berlin
- [5] I. GIANNIKOS and E. RIZAKOU (1995), The use of a flexible planning methodology in making location decisions under uncertainty, paper presented at EWGLA8, Lambrecht, Germany.
- [6] SH. K. GUPTA and J. ROSENHEAD (1968), Robustness in sequential investment decisions, *Management Science* 15, B18-B29.
- [7] S. K. JACOBSEN and O. B. G. MADSEN (1980), A comparative study of heuristics for a two-level routing-location problem, European Journal of Operational Research 5, 378-387.
- [8] A. A. Kuehn and M. J. Hamburger (1963), A heuristic program for locating warehouses, *Management Science* 9, 643-666.
- [9] G. LAPORTE (1988), Location-routing problems, in Vehicle Routing: Methods and Studies, B. L. Golden and A. A. Assad (eds), Elsevier, Amsterdam, 163-197
- [10] G. LAPORTE (1992), The vehicle routing problem: an overview of exact and approximate algorithms, European Journal of Operational Research 59, 345-358
- [11] P. B. MIRCHANDANI and R. L. FRANCIS (eds) (1990), Discrete Location Theory, Wiley, New York.
- [12] G. NAGY and S. SALHI (1996a), Nested heuristic methods for the location-routing problem, Journal of the Operational Research Society 47, 1166-1174.
- [13] G. NAGY and S. SALHI (1996b), A nested location-routing heuristic using route length estimation, Studies in Locational Analysis 10, 109-127.
- [14] J. Perl. and M. S. Daskin (1984), A unified warehouse locationrouting methodology, *Journal of Business Logistics* 5, 92-111.

- [15] G. K. RAND (1976) Methodological choices in depot location studies, Operational Research Quarterly 27, 242-249.
- [16] J. ROSENHEAD, (1989) Rational Analysis for a Problematic World: Problem structuring methods for complexity, uncertainty and conflict, Wiley, Chichester.
- [17] J. ROSENHEAD, M. ELTON and SH. K. GUPTA (1972), Robustness and optimality as criteria for strategic decisions, *Operational Research Quarterly* 23, 413-431.
- [18] S. Salhi and M. Fraser (1996), An integrated heuristic approach for the combined location-vehicle fleet mix problem, Studies in Locational Analysis 8, 3-22.
- [19] S. SALHI and G. K. RAND (1989) The effect of ignoring routes when locating depots, European Journal of Operational Research 39, 150-156.
- [20] S. Salhi and M. Sari (1997) A multi-level composite heuristic for the multi-depot vehicle fleet mix problem, European Journal of Operational Research 103, 78-95.

