Chapter 17

LOCATION-ALLOCATION ON A MICROCOMPUTER: THE PLACE PACKAGE

M.F. Goodchild

Location-allocation is the term given to a set of techniques designed to aid in locating fixed facilities from which service can be provided to a dispersed population. This definition covers a wide variety of application fields, from retailing, where the location analyst knows that the choice of location can make the difference between success and failure of a business, to public agencies seeking to provide some adequate and fair level of social service to everyone in a community.

Ideally one would like to be able to provide a service wherever and whenever it is needed. This is especially true of emergency services such as fire protection, which forms the basis of the example used later in this chapter. The response time to a fire is a simple function of the distance between the fire and the nearest fire station, so that given a sufficiently high density of fire stations, it would be possible to reduce the response time to any fire almost to zero. The cost of providing such a density of service would, however, be prohibitive, and in reality the benefits of reducing response time must be weighed against the costs of providing a high density of service locations. Retailing provides similar examples: the operator of a convenience food store chain is forced to provide a large number of outlets because consumers are unwilling to travel long distances, whereas the service of a department store can be offered from a few, highly attractive locations.

The problem of providing a service to a dispersed population can be overcome in some cases by making the service mobile. The economics of milk retailing make house-to-house delivery profitable in some societies, while in others the consumer is willing to collect milk from fixed food stores. Libraries are an example of a public service which is sometimes provided from mobile rather than fixed sites. Two factors appear to determine whether a service will be mobile or fixed. First, services have the option of being mobile only if the consumer, or the person demanding the service, would be satisfied with a service which is available only during certain limited hours. Second, a service will remain fixed as long as the number of consumers willing to use each site is adequate to justify the existence of the site. It may therefore 'go mobile' in order to generate adequate demand. For this reason, a service which is offered from fixed sites in high-density urban areas may often be mobile in low-density, rural areas.

Location-allocation is concerned only with the optimum location of fixed central facilities, and has been applied over the past two decades to such services as schools, ambulances, garbage disposal sites, electricity generating stations, day care centres and retail activities. A recent volume edited by Ghosh and Rushton (1987) contains a wide cross-section of applications. Other reviews of location-allocation in an applied context can be found in Handler and Mirchandani (1979), Hodgart (1978), Larson and Odoni (1981) and Massam (1980).

The essential ingredients of a location-allocation approach to optimum site selection are as follows. First, there must be some dispersed pattern of demand for a service, located at a set of points. In many cases this will simply be the population distribution, located at the centres of census reporting units, but many services are directed only at certain segments of the population, or at some other distribution. For example, the demand for day care facilities might focus on the distribution of preschool children, while the location of facilities for fighting forest fires would be determined by the pattern of risk.

Secondly, there must be some means for determining the number of central facilities required. In some cases this will be known, and some of the facilities may already exist. In other cases the model will be expected to determine the optimum number and perhaps sizes of facilities as well as their locations. Larger and more expensive facilities may be needed where the associated catchment areas contain relatively heavy demand.

Thirdly, there must be an objective to be optimised. The earliest location-allocation models sought to minimise the total distance separating facilities and demand; in other words, the total distance over which the service has to be delivered. One important variant on this simple model is to minimise cost, where cost is related to distance by a transport cost function, and this can be further generalised to the case where cost is the sum of the expense of transporting the service and the expense of building and maintaining the facilities. This Plant Location Problem, as it is known, results in a compromise between one large central facility with maximum transportation costs, and many small, dispersed facilities with very high construction and maintenance costs.

In public sector services, where costs are frequently intangible, the efficiency of provision of a service is often seen as less important than its equity of provision. In other words, the best system of service is one which minimises the range of service quality offered to the dispersed population. This is often taken to imply that one should seek locations which will minimise the furthest distance travelled by a consumer, rather than minimising the average or total distance.

The solution of a location-allocation problem has two major elements, as its name implies. The solution procedure is designed to find the optimum locations, and also, if necessary, the optimum number and size of facilities, and in addition to determine how the demand should be allocated to each location; in other words the optimum service or catchment areas of each facility. In some cases these will be outside the control of the modeller, since they will be determined by consumer choice, as in retail applications. But in other cases, particularly in the public sector, it is common for the agency providing the service to dictate the set of consumers or the catchment area served by each facility.

The next section of this chapter clarifies this general introduction by giving a precise definition of the basic location-allocation model used in the PLACE package. This is followed by a general description of the structure of the package, and the final section describes an application to the optimal selection of fire station locations in a Canadian city.

THE LOCATION-ALLOCATION MODEL

PLACE assumes a discrete model of geographical space, in other words that the space in which facilities are to be located can be described in terms of a fixed number (n) of possible locations. These might correspond to the street intersections in a city, or to central points in the census reporting zones. Clearly this represents an approximation to reality, and the quality of that approximation has been discussed by Goodchild (1979) and others. The demand for the service is assumed to be aggregated to these discrete places, and the amount of demand at place i is denoted by w_i, with i ranging from 1 to n. It is possible for the weight at a place to be zero, and this will be quite common in some applications because the set of places must give a reasonably complete coverage of the study area.

The number of facilities to be located is assumed to be known in advance by PLACE, which contains no methods for determining the number as part of the solution. In many applications some of these facilities will already exist and must be held fixed during the solution process, while the remainder represent new facilities which are allowed to move over the set of places during the search for optimal locations. Each of the p facilities is identified by a value of the index j.

PLACE allows the user to declare certain of the n places as possible candidate locations for facilities, and to exclude the remainder from the search. This is often useful in cases where only a limited number of places may be available, perhaps because of zoning restrictions. Candidate places are identified by a variable C, C_j being set to one if place j is available, and zero otherwise. It is important that the set of places identified in the initial representation of the study area include an adequate coverage of the possible locations of facilities.

Since all location-allocation models include some form of distance in their objective criteria, PLACE includes two methods for determining the distances between all pairs of places, by calculating straight line distances using coordinates, and by finding the length of the shortest path through a network. In the first case the user must supply the locations (x_i,y_i) of all places, and in the second the user is required to input the lengths of a set of direct links connecting selected pairs of places, such as the links of a street network. It is assumed that the links provide at least one path between all possible pairs of places, in other words that the network is fully connected. The distance between place i and place j is denoted by d_{ij}.

The location problem consists of finding the place j occupied by each of the p facilities in the optimum solution. The corresponding allocation problem is to determine which central facility is to serve each dispersed place, or demand point. In PLACE this is assumed to be under the control of the modeller, rather than determined by consumer choice, and in all PLACE models the optimum solution is obtained by allocation to the nearest available facility. Allocation is denoted by the value of x_{ij} (x_{ij} is set to one if facility j serves the demand at i, and zero otherwise).

The simplest model solved by PLACE is the minimisation of total distance travelled, known as the p-median problem. This is clearly achieved by serving each dispersed place from the closest central facility, but it is not clear where those central facilities should be located. It turns out that in the typical city, with low density suburbs and high density core, the optimum locations of facilities will be relatively spread out in the periphery and closer together in the core, but not in proportion to the density of demand: each core facility will serve rather more demand, and each peripheral facility rather less.

The p-median model can be written mathematically as follows: Find x_{ij}, $i=1,n$, $j=1,n$ to:

$$\text{Minimise} \sum_{ij} w_i \, x_{ij} \, d_{ij} \qquad (1)$$

Subject to:

$$x_{ij} = \{0,1\} \text{ for all pairs i,j} \qquad (2)$$

$$x_{ij} <= x_{jj} \text{ for all pairs i,j} \qquad (3)$$

$$\sum_j x_{ij} = 1 \text{ for all i} \qquad (4)$$

$$\sum_j x_{jj} = p \qquad (5)$$

Objective 1 defines the criterion for optimality, which is that the total distance separating demand and service be minimised. Constraints 2 through 5 specify the conditions under which this criterion should be optimised, as follows. Constraint 2 requires all elements x_{ij} to be zero or one, so no partial allocations are to be allowed. Constraint 3 indicates that demand can be allocated from i to j only if the demand at j is allocated to j, in other words that a facility exists at j. Constraint 4 specifies that the demand at i must be allocated to a facility somewhere, and Constraint 5 specifies the number of facilities to be located.

If we consider the focus of the objective to be the worst level of service offered rather than the average level, then this suggests that one should minimise the maximum distance rather than the total distance, as follows:

$$\text{Minimise} (\max_{i,j} x_{ij} \, d_{ij}) \qquad (6)$$

The common term for this is the p-centre problem, and it results in a pattern of service which is strikingly different from the p-median solution. Facilities are evenly spaced over the city, so that those in core areas are allocated a much larger share of demand than those in the low density periphery.

Another useful model is obtained by assuming that some standard distance S exists, and by maximising the number of people within this distance of their nearest facility. It is a simple matter in PLACE to replace distances by travel times, so this model would fit an emergency service which must be located to cover as much demand as possible within some standard response time. Alternatively, one might wish to minimise total distance travelled subject to the constraint that no person be further than some specified distance S from a service.

All of these objective functions can be handled as examples of the same general form:

$$\text{Minimise} \sum_i \sum_j x_{ij} \, c_{ij} \qquad (7)$$

by appropriate definition of c_{ij}, an approach which allows one central algorithm to solve many problems. This strategy has been termed 'distance editing' because it can be seen as the systematic editing of the weighted distance matrix $w_i d_{ij}$. PLACE currently recognises six forms of distance editing, allowing it to produce solutions for a total of seven different objective functions.

PLACE IMPLEMENTATION

The PLACE project was begun in 1982 in an attempt to satisfy what was seen at that time as a need for location-allocation codes which could run in the kinds of small computers, typically 8-bit and 32k or 64k, then becoming available to planners in the third world. Up to that time the most widely available general location-allocation software consisted of a package put together by the Department of Geography at the University of Iowa (Rushton, Goodchild and Ostresh, 1973), but designed exclusively for large mainframes.

The first generation of the package was written in interpreted BASIC and required a complete overhaul of many of the algorithms to avoid dependence on large quantities of core memory. Subsequent revisions have moved to compiled BASIC and have taken advantage of the larger core memory available in recent microcomputers. Complete technical details of the design of the package can be found in Goodchild and Noronha (1983), which also serves as the user manual.

The method used to search over the space for the best facility locations is known as a vertex substitution algorithm (Teitz and Bart, 1968). Since a complete search of all possible combinations of facility locations would be impossible for all but the smallest problems, the algorithm begins with a trial solution, and then attempts to improve it by systematically moving each facility in turn to all other possible locations. This algorithm cannot be guaranteed to yield the true optimum solution, but has been shown to do so in most cases, and to give a solution close to the optimum in all other cases.

A typical application of PLACE consists of four stages, using four of the package modules. The first stage is to generate a weighted distance matrix, either by calculating straight line distances between places or by processing the nodes and links representing the network which connects the places. Depending on the objective function to be used, the user may then invoke some form of editing of the distance matrix to replace weighted distances by an appropriate set of values. The third stage finds an optimum or near-optimum solution by taking an initial solution, including any existing and fixed facilities, and searching for an optimum using the vertex substitution algorithm. Finally the user may wish to evaluate his or her solution by generating a number of statistics using a fourth module.

The PLACE package has been structured as a toolbox, to be adapted to a given situation, rather than as a black box. It requires a reasonably sophisticated level of understanding of the basic model on the part of the user. In this it reflects the notion that location-allocation provides a guide rather than a definitive solution to a problem, and must be modified by the user to take account of the unique aspects of every locational decision. There is, after all, no single, optimum solution to any planning problem. Tools such as location-allocation exist to improve the quality of debate over locational decisions, not to remove the responsibility for making them.

APPLICATION TO FIRE STATION LOCATION

This final section describes an application of the package to the planning of fire stations in a Canadian city, updated from Waters (1977). A second example, to the selection of locations for waste disposal, can be found in Goodchild and Donnan (1987).

London, Ontario, is a city of some 270,000 population located midway between Toronto and Detroit. Its fire protection problems are typical of many North American cities and illustrate the difficulties of providing a system of central services to a city which is in a constant state of growth and spatial redistribution.

The fire risk in nineteenth century North American cities was much higher than today because of poorer construction standards and the use of open fires, and the travel speeds and capacities of fire fighting equipment were much lower. The nineteenth century city therefore required a relatively high density of fire protection facilities. As the city has grown into the twentieth century it has been necessary to extend fire protection to new areas, but at the same time to reduce the number of facilities in older sections, and this has been achieved largely by relocation, so that the total number of stations in the modern city of 270,000 is roughly the same as it was in the 1890 city of 60,000. Location-allocation provides the means for continually evaluating and readjusting the pattern of fire protection services in response to the city's growth and changing patterns of demand.

In 1976 the city was served by nine fire stations, each equipped with a minimum of one pumper

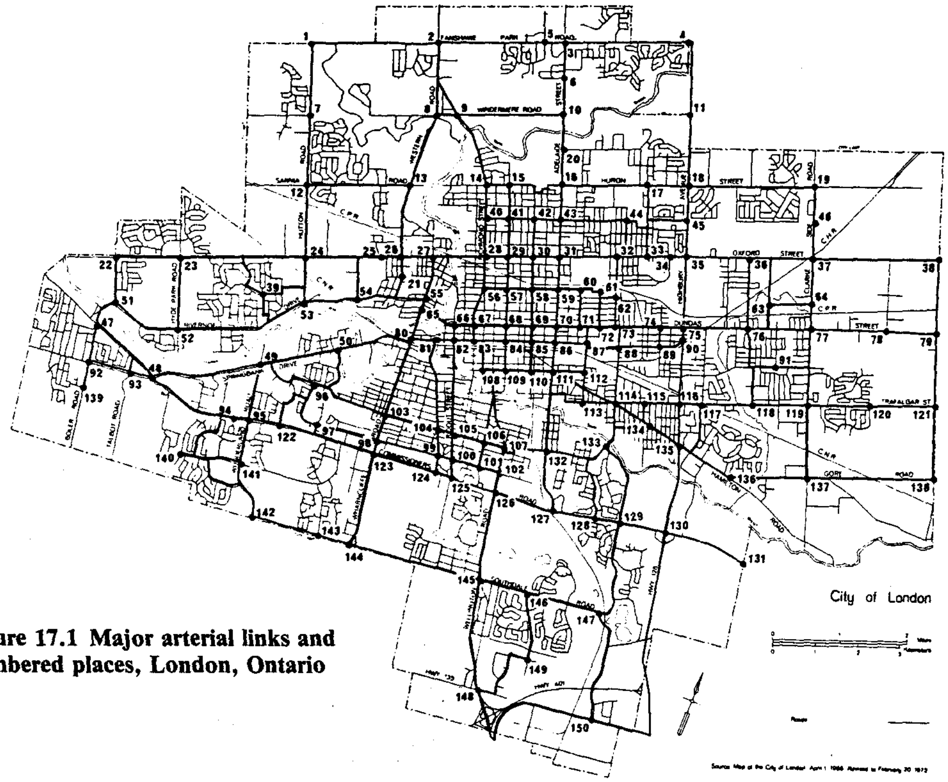

Figure 17.1 Major arterial links and numbered places, London, Ontario

truck. The demand distribution was obtained by locating all of the 2,459 alarms occurring in 1973. These were then aggregated to the nearest major intersection in the city street network, giving a set of 150 weighted places (Figure 17.1), connected by 250 major arterial links. This network is an accurate representation of the set of routes used by fire equipment in the city. The 1976 fire stations were located at nodes 8, 24, 30, 45, 84, 88, 95, 111 and 145.

Studies of the London area have shown that the speed of a responding fire truck is approximately constant over the major arterial links, so it was assumed that response time was directly and linearly related to distance from responding fire station to the nearest major intersection to the alarm. Minimisation of distance would thus serve as a surrogate for the minimisation of response time.

Once the database is established, it is possible to examine a large number of scenarios in a short time. For purposes of discussion, only three such runs will be reviewed here. The three situations considered are:

(1) relocation of all 9 existing stations to minimise total response time (the 9-median solution);
(2) holding the existing 9 fixed, and locating a tenth station so as to reduce expected response time as much as possible (the 10-median solution with 9 fixed sites);
(3) as the previous scenario, but with maximum rather than expected response time as the objective (the 10-centre solution with 9 fixed sites).

Using the existing 9 stations as the starting solution, the vertex substitution algorithm examined every possible replacement of an existing site by one of the 141 places not currently occupied, making that swap which would reduce the total response time as much as possible. Eight complete cycles of the swapping process were necessary: on the ninth cycle no swap could be found to improve the solution. The average distance between the 2,459 alarms and the 9 fire stations declined from 1.04 miles to 0.88 miles at the end of the relocation, two facilities remaining in their previous locations.

The pattern of relocation revealed two significant trends. There was a tendency for facilities in the core to move outwards, indicating that the older part of the city is still overserved. On the other hand two of the most peripheral facilities moved inwards, indicating that previous relocations had perhaps attempted to anticipate future growth around the city's edge, or had gone too far in attempting to adjust service to changing demand.

Both of the remaining scenarios gave similar solutions for a tenth facility location, holding the existing nine fixed. Minimisation of total distance resulted in the selection of place 119, which is in the poorly serviced southeastern corner of the city, reducing average distance from 1.04 miles to 0.93 miles for the city as a whole. Minimisation of maximum response time led to place 118, which reduced the worst response distance in the city from 4.44 miles to 3.22 miles. Shortly after this study was completed the city located a tenth fire station between places 118 and 119. Two other facilities have since been relocated, and in each case the locational decision has followed the location-allocation solution closely.

CONCLUDING REMARKS

Location-allocation problems are computationally intensive, and so would appear not well suited to microcomputers. For problems of realistic size, approximate algorithms which cannot be guaranteed to yield the true optimum solution are essential even on large mainframes, because of the need to examine huge numbers of possible solutions. The experience with PLACE has shown that it is possible to solve real location-allocation problems on small computers, particularly given recent improvements in compilers, processor speed and central memory. Nevertheless, response can be slow. A recent test using the Ryan-McFarland Fortran compiler and an 80286 processor with 80287 math coprocessor required some 10 minutes to solve a p-median problem with 450 places and 10 facilities, and occupied all 600k of central memory.

Against this one must weigh the well-known advantages of interactive microprocessing, including an enhanced user environment, access to graphics and easy transfer of software between workstations. Location-allocation packages can be interfaced with digitising tablets to automate the collection of the basic spatial data. Perhaps the most serious factor inhibiting use of microcomputers for location-allocation at this time, apart from relatively slow response, is the frequent need to access major demographic and geographic data banks for information on street networks and population distributions. For this reason the functionality of PLACE remains well below that of its nearest mainframe equivalent (Goodchild, 1984).

REFERENCES

Ghosh, A. and Rushton, G. (eds) (1987) *Spatial Analysis and Location-Allocation Models*, Van Nostrand-Reinhold, New York.

Goodchild, M.F. (1979) 'The Aggregation Problem in Location-Allocation', *Geographical Analysis*, 11, 240-255.

Goodchild, M.F. (1984) 'ILACS: A Location-Allocation Model for Retail Site Selection', *Journal of Retailing*, 60, 84-100.

Goodchild, M.F. and Donnan, J.A. (1987) 'Optimum Location of Liquid Waste Disposal Facilities: Formation Fluid in the Petrolia, Ontario, Oilfield', in M. Chatterji (ed.) *Hazardous Waste Materials Disposal: Siting and Management*, Gower, Aldershot, England.

Goodchild, M.F. and Noronha, V.T. (1983) *Location-Allocation for Small Computers*, Monograph No. 8, Department of Geography, University of Iowa, Iowa City.

Handler, G.Y. and Mirchandani, P.B. (1979) *Location on Networks: Theory and Algorithms*, MIT Press, Cambridge, Massachusetts.

Hodgart, R.L. (1978) 'Optimising Access to Public Facilities: A Review of Problems, Models and Methods of Locating Central Facilities', *Progress in Human Geography*, 2, 17.

Larson, R.C. and Odoni, A.R. (1981) *Urban Operations Research*, Prentice-Hall, Englewood Cliffs, New Jersey.

Massam, B.H. (1980) *Spatial Search: Applications to Planning Problems in the Public Sector*, Pergamon Press, Oxford.

Rushton, G., Goodchild, M.F. and Ostresh, L.M. Jr. (1973) *Computer Programs for Location-Allocation Problems*, Monograph No. 6, Department of Geography, University of Iowa, Iowa City.

Teitz, M.B. and Bart, P. (1968) 'Heuristic Methods for Estimating the Generalized Vertex Median of a Weighted Graph', *Operations Research*, 16, 955-961.

Waters, N.M. (1977) 'Methodology for Servicing the Geography of Urban Fire', Unpublished Ph.D. Dissertation, Department of Geography, University of Western Ontario, London, Ontario.