Volume 60 Number 1 Spring 1984

ILACS: A Location-Allocation Model for Retail Site Selection

MICHAEL F. GOODCHILD
Professor of Geography
The University of Western Ontario
London, Canada

Retail site selection is a major area of application for location-allocation models, so it is surprising that, despite some notable recent exceptions, almost all research has been on public sector problems. This paper identifies three major issues in the development of appropriate location-allocation models for retailing: the calibration of suitable consumer behavior models, the choice of appropriate objective functions, and the cost-effectiveness of solution procedures. The paper proposes two models which can be applied to remote cities using readily available data and solution methods which operate in continuous geographical space. The models have been implemented in an interactive graphics package. An example application is given in the final section of the paper.

INTRODUCTION

Location-allocation is the simultaneous location of central facilities, and the allocation of dispersed demand to them, so as to optimize some objective function. Defined in this way, location-allocation models should be appropriate to the location of any central facilities designed to serve a dispersed population, including retailing. Yet almost all applications of these models since their inception in the early 1960s have been to public sector services (see, for example, Hodgart 1978; Leonardi 1981). This paper presents two location-allocation models designed for application in retail site selection. The first section of the paper expands on the basic definition given above; this is followed by sections which present the conceptual framework and the retail models. The last section examines an application to the choice of sites for a restaurant chain in London, Ontario.

The demand for a retail service is assumed to be located at a set of points, each represented by a pair of coordinates. Associated with each point is a weight representing the volume of demand or potential market.

In most urban applications, each demand point is located at the centroid of a statistical zone such as a census tract, and its weight represents the demand in the zone. Ideally, access to the service would be optimized by locating a facility at each of the demand points, but, because of scale economies in the operation of facilities and their high fixed costs, it is usually more profitable to provide service from a few central locations. So, for example, a city whose population is represented by 200 census tract centroids may be served by no more than 15 appropriately located shopping centers.

The earliest of the standard location-allocation problems to be formulated was the p-median (Hakimi 1964; Cooper 1964; Teitz and Bart 1968; Maranzana 1964; and, for an early review, see Scott 1970). This requires the central facilities to be located so as to minimize the total distance traveled in satisfying the demand, either by having each consumer travel to a facility or by taking the service to each one individually. Total distance is minimized when each individual obtains service from the nearest facility and when each facility is located optimally with respect to the trade area which it serves. Of course, neither the trade areas nor the central facility locations are known in advance.

Although the total distance traveled to obtain the service is related to the total cost of transportation, the relationship is not precise, as the marginal cost of transportation is usually a decreasing function of distance. In addition, the cost of transportation rarely accrues in a simple fashion to the provider of the services. So, instead, the p-median objective is often seen as optimizing some intangible notion of accessibility. As such, it is criticized for its concern for the total accessibility in the system rather than for some minimal standard of individual access. This latter notion has led to the formulation of the p-center problem (Hakimi 1964) which locates facilities so as to minimize the maximum distance traveled in the system between any demand point and the nearest facility.

The solutions of each of these problems have certain specific characteristics. Demand is assumed to be served by the nearest facility, so service areas are simple geometric figures drawn around each one. In both the p-median and p-center problems, these are Thiessen and Alter (1911) or Dirichlet (1850) polygons whose edges are the perpendicular bisectors of lines drawn between neighboring facilities. The p-center problem tends to produce solutions in which each service area is roughly equal in areal extent. When demand density varies in a study region, the total demand served by each facility will also vary proportionately. This is illustrated in Figure 1 in the form of a graph of the relationship between demand density and the density of facilities per unit area, which is the inverse of

FIGURE 1

Relationship between Facility and Demand Densities for Various Location-Allocation Solutions

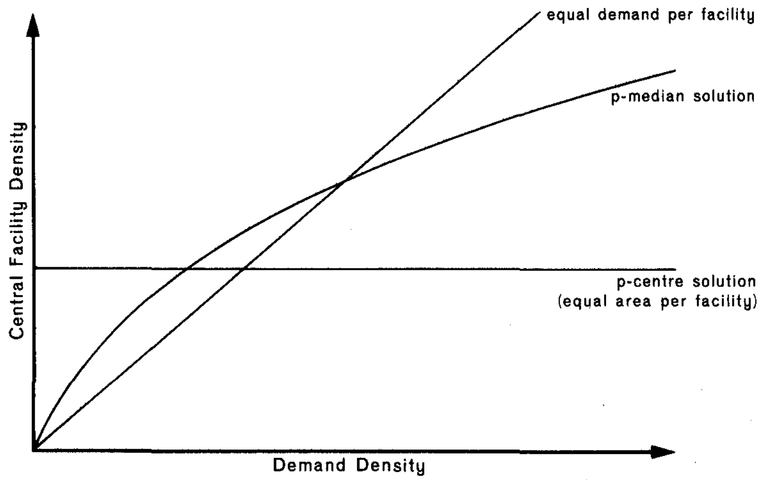

the mean area served by each facility. Since the demand served by each central site will in general vary in such a solution, the most appropriate applications are to services where capacity is not a major concern.

Solutions to p-median problems yield curvilinear relationships between demand density and facility density by providing relatively more facilities per unit area in areas of high density (Figure 1). This can be seen as a compromise between a linear relationship, which would place facilities so that each one's service area contained an equal share of demand, and the p-center solution. In the context of service provision, the linear relationship can be seen as an equalizing of service capacity per capita, whereas the p-center solution provides equality of spatial accessibility.

With this brief introduction to location-allocation problems, the next section discusses a conceptual framework for applications in the retail sector.

CONCEPTUAL FRAMEWORK

The retail sector presents three general issues which must be addressed before any appropriate model or models can be devised. First, allocation presents special problems in retailing because of its enormous complexity and the lack of reasonably accurate models. Second, it is clear that objective functions must reflect retail sector concerns for profitability and competition. Finally, any effective site selection strategy for retailing must address issues of speed, cost of analysis, and practicality. The three sections that follow address each of these sets of issues in turn.

Allocation

As formulated in the introduction, location-allocation models assume that both the locations chosen for the facilities, and the manner in which demand is allocated to them, will be so as to optimize an abstract objective function. However, in retailing, only the locations can be controlled: the allocations of demand are the outcome of consumer spatial behavior, as consumers are clearly free to choose which stores they will visit. In some sectors of the industry, behavior may approximate the simple nearest-center model, particularly in convenience shopping and in choices among branches of a single chain. But this pattern is distorted by the influences of multiple-purpose shopping trips, impulse shopping, and trips made from daytime rather than residential locations.

The location-allocation literature has termed this kind of implied allocation through consumer spatial behavior voluntary, as distinct from the controlled allocation that is possible in many public sector problems where service areas can be prescribed by the planner. In order to solve any location-allocation problem in a voluntary allocation context, it is necessary that the planner be able to predict consumer spatial behavior with respect to hypothetical sites; in other words, to determine the allocation of demand to a set of sites which may not yet exist. In practice, this has been approached by observing consumer spatial behavior with respect to a set of existing sites, and then using this finding to calibrate a model of spatial behavior which can be applied to new ones (Goodchild and Booth 1980; Beaumont 1980; Hodgson 1978, 1981; Achabal, Gorr, and Mahajan 1982; Ghosh and McLafferty 1982). Let x_{im}, $i = 1,n$, $m = 1,q$ denote consumer trip counts between the n demand points and q existing sites. A suitable class of models has the general form of

$$x_{im} = E_i A_m f(d_{im}) \tag{1}$$

where E_i is interpreted as a function of the demand in origin zone i, A_m as the attraction of the destination store m, and f as a decreasing function of the distance between them.

If total demand at each demand point is inelastic, that is, every consumer will always shop whatever the spatial distribution of stores, then the number of unknowns is reduced and the model becomes:

$$x_{im} = P_i A_m f(d_{im}) \Big/ \sum_k A_k f(d_{ik}) \tag{2}$$

where P_i is the total observed demand in zone i. The model is then said to be production-constrained. This is more likely to be appropriate when it is applied to low-order convenience goods and when the system includes all stores offering the good. On the other hand, demand is likely to be elastic for higher-order goods and when the set of stores considered is a subset of all the stores offering a given good. So, for example, the production constraint will be invalid when analyzing behavior with respect to the stores of convenience food chain A if B is also present in the market, and may also be inappropriate to the complete set of convenience food stores if substantial substitution of 24-hour supermarkets occurs.

Although the literature in this area is small, several operational problems can be identified in addition to those already mentioned. First, there is some doubt that these spatially separable behavior models, which include the commonly used Huff and MCI models, are capable of generalization to hypothetical sets of sites. Work on misspecification errors and spatial bias (Fotheringham 1981) suggests that it would be essential to recalibrate the model in each market, and possibly even for each set of sites.

Second, there is little consensus on the appropriate objective function to be used. In the context of the public sector, Hodgson (1981) has argued that minimization of either total or maximum distance traveled by consumers makes little sense in applications where a substantial proportion of consumers chooses not to patronize the nearest offering of the service. Hodgson and others have suggested that sites should be selected so as to maximize the total satisfaction obtained. Noting that the constrained version has the form of a utility-based choice model, they argued that the location-allocation objective should be to maximize the total utility of all trips predicted by the choice model. Another alternative would be to use an objective function of maximizing total demand, or market share in the case of more than one chain. In a recent paper, Achabal, Gorr, and Mahajan (1982) maximized the difference between total revenue and fixed cost by assuming revenue proportional to total consumer expenditure at each store, and then used an MCI model to predict allocations. Ghosh and McLafferty extended this model by using a multicriteria framework to deal with uncertainty about a competitor's future actions.

Finally, and perhaps most critically, these arguments require the attractions to be determined a priori for hypothetical sites, which raises two problems. First, the attraction of a store is not only a composite of its physical, built characteristics, but also to some extent a function of such short-term variables as price and advertising and of factors such as crowding, which are themselves functions of the consumer spatial behavior itself (Goodchild 1978). These latter variables will be impossible to project for hypothetical store locations, particularly because of the causal feedback implied in crowding. Second, the planner in this context has the ability to modify the attractions for new stores through physical design factors. Since the attractions will affect the objective function, it is not at all clear how they should be determined and to what purpose. Hodgson (1981) was able to allocate attraction by varying the number of clerks and therefore the capacity in each of a number of license offices, under a constraint on the total number of clerks in the system.

In summary, for higher-order goods where substantial overlap between service areas is likely, it is necessary to calibrate and apply models of spatial consumer behavior in order to be able to predict market areas and total consumer inflow at each site. Because such models are expensive to calibrate and because their coefficients tend not to be generalizable, this option is unlikely to be practical until substantial further basic research has been undertaken. Until then, applications of location-allocation in the retail sector will be limited to low-order convenience goods where spatial behavior tends to be approximately nearest-place and where simple models of elasticity of demand with respect to distance can be used.

Objective Functions

The objective functions used in the public sector represent a goal for the entire set of locations. In retailing, we can distinguish two cases: (1) problems in which competitive locations are included but the objective is written for the client's sites only, and (2) problems in which competitive locations are ignored. The first case is appropriate when consumers patronize the nearest place irrespective of available brand, which is most closely approximated by convenience food store or gasoline station choice. In this case, an appropriate objective function would be to maximize either the total demand allocated to the client's sites or the market share. Note that this function is limited to that fraction of convenience shopping trips which originate from fixed points, such as residences or places of work: impulse shopping, multiple-purpose trips, and stops during the journey to or from work cannot be readily characterized as "nearest place." We assume that the locations of competitive stores are known, and that the objective function is to be optimized by selecting sites for the client only.

In the second case, the locations of competitive stores are ignored.

Demand is therefore likely to be elastic with respect to distance, since a greater distance from a client's store makes substitution by a competitor more likely. Suppose that the probability that a randomly chosen consumer will select one of the client's stores is a decreasing linear function of his or her distance from the nearest one. Assume that no consumer is so far from the nearest store that this probability is reduced to zero. Then, maximizing total demand is equivalent to minimizing total distance, or to the p-median objective given earlier. The problem thus reduces to a well-known one irrespective of the parameters of the linear demand function.

Thus both cases can be treated as variants of standard location-allocation models. The difference between them lies in the approach to competition. The first, the market-share model (MSM), is a conservative strategy of avoidance, of finding the holes in the competition's coverage of a market. The second, the competition-ignoring model (CIM), is more aggressive, ignoring competitors and concentrating on maximizing market penetration by selecting locations which are highly accessible to the consumer.

Both strategies are essentially short-run, as they ignore the possibility of response by the competition, either in reaction to or in anticipation of site selection by the client. MSM assumes that the competitor's locations are known in advance and will remain fixed throughout the planned lifetime of the store. CIM, which ignores all information on competitors' locations, is possibly the most rational strategy given total uncertainty about competitive action and reaction. A recent paper by Ghosh and Craig (1984), which adds to an extensive literature on the application of game theory to competitive location, can be seen as taking the opposite extreme by assuming rational and perfectly informed behavior on both sides.

Cost-Effectiveness

One of the factors most likely to affect the cost of applying a location-allocation model in site selection is the method of spatial representation used. Two methods, discrete and continuous, have been described in the literature and each has advantages. In continuous space, all locations are assumed feasible and distances are measured according to some metric. Various papers have investigated the use of alternatives to straight-line distance, particularly the Manhattan or city-block metric (see, for example, Drezner and Wesolowsky 1983), which is argued to be a better approximation to the real distance traveled between points in areas with a gridiron system of roads or streets, although little attention has been paid to the requirement that x and y axes be aligned with the street system.

On the other hand, discrete space assumes that travel is confined to a network, and demand is located at nodes on the network. The feasible locations are confined to the network also. A theorem due to Hakimi (1964) shows that the solutions to p-median problems must lie at nodes, so the search for optimal locations can be confined to the vertices of the network and to the points where weight is located. Although the theorem is not true of the p-center problem, making the assumption can nevertheless give an efficient heuristic for this problem as well.

Discrete space has obvious advantages over continuous space. Distances can be evaluated accurately along network links. Algorithms based on a finite set of feasible locations are clearly more efficient, and there is no risk of identifying optimum retail locations in lakes or on airport runways. But there are several strong counterarguments. The accuracy of a discrete space transport network is short-term, so continuous space may be a better hedge against an uncertain future. Discrete space requires vastly more data in the form of a detailed street or road network, whereas the coordinate locations required for continuous space solutions are readily available in standard sources. For example, a convenience store location study in discrete space would require data on every street in order to evaluate accurately the short distances traveled from home to store in this market. In continuous space, the demands could be represented at Enumeration District centroids, which are readily available in machine-readable form. Finally, continuous space algorithms have much smaller requirements for core memory than comparable discrete space algorithms. This difference can be critical because it is desirable to work with a demand-point base which is as disaggregate as possible (Goodchild 1979).

The retail sector is extremely complex, so an accurate model of consumer behavior or competitive response can be expected to be complex also. There is a direct cost to complexity, however, in the effort which must go into data collection and calibration. We have already argued that with the current state of the art, the parameters of a spatial interaction model cannot be expected to survive transfer from one city to another, or even from one set of stores to another in the same city. A client must therefore pay off the time and expense of complex modeling against the increased accuracy to be gained from it. At the same time, a complex model is much more sensitive to future uncertainty, as, for example, in the response of attraction parameters to short-run variations in price or levels of crowding at a store. It is possible, then, that as well as being cheaper, a simple model is a more rational response to future uncertainty.

Ideally, the number of sites to be located should be determined through an analysis of the benefits of high market penetration against high fixed costs. Few retailers are likely to have access to the necessary market penetration forecasts and consumer behavior models (see, for example, Achabal et al. 1982) to predict dollar sales for hypothetical locations. Instead, many convenience store chains operate with some notion of threshold, or the minimum market required to support a single store. In the absence of better information, the number of sites is likely to be determined by dividing the potential market by the threshold. So, for example, a certain fast-food chain that believes its threshold to be a population of 30,000 would aim to locate 8 stores in a city of 240,000.

This approach leads to a basic paradox. In order to maximize overall market share, it is likely that each store will serve a different proportion of the total market, with some stores operating below threshold and some above. Maximizing total share does not in general lead to equal shares for each store. If this is viewed as undesirable, an appropriate constraint can be added to the model.

Another commonly available statistic in convenience retailing concerns the range of the good. So, for example, a chain may believe that the consumer will not travel more than a certain maximum distance for the product. This is, in effect, a statement about the elasticity of demand with respect to distance, and it is at best a crude approximation: intuitively, we expect the demand curve to be smooth, rather than discontinuous at a specific range. On the other hand, the cost of calibrating an accurate demand curve may not be justified. So, in the absence of better information, the assumption of a constant elasticity of demand with distance is particularly attractive because it leads to a tractable model without requiring any further calibration.

The assessment of potential demand in retailing is equally problematic as it requires retailers to have an accurate knowledge of the socioeconomic characteristics of their potential consumers. In the short term, these are the groups that actually use the stores, or would if they were close enough to them, but in the long term, they may be groups to which advertising is currently directed. A typical market might be 15- to 44-year-olds with household incomes $25,000 and over. This again is oversimplified, implying, as it does, 100 percent penetration of this market and zero percent of all others. The relationship between the two categories is ambiguous: Is the market some weighted combination of the numbers in the two categories, or a count of those with both characteristics? Again, the cost of more sophisticated data may not be justified, and the results may be longitudinally unstable.

All these arguments lead to the conclusion that simple models like MSM and CIM, which can be applied cheaply and quickly in continuous space using readily available data, can provide adequate answers to site selection problems in the retail sector. Until more complex and accurate models have been calibrated, they represent the best possible use of available information. In the long run, they may even be preferable to complex models because of their comparative insensitivity to future uncertainty. When available, models of consumer behavior and total sales can certainly be incorporated into the location-allocation framework, but we have already seen that this inclusion can lead to unanswerable questions of retail strategy, as in the case of allocating attraction among stores.

The examples of the threshold and range given above lead to a difficult dilemma for the location-allocation model. It would be quite impossible to locate stores so that every one had exactly 30,000 potential customers within its prescribed range. To what extent, then, should a model be an accurate representation of reality, and to what extent an embodiment of the beliefs of the chain's management, if the two are clearly contradictory? Location-allocation is simply a tool, not an empirically verifiable model of reality, although some of its components, such as the assumptions made about consumer behavior, may be verifiable.

MODELS

The concepts discussed in this paper have been incorporated into ILACS (Interactive Location-Allocation in Continuous Space). (Documentation and FORTRAN code are available from the author. The data sets referred to are available from the U.S. Bureau of the Census, Statistics Canada, and numerous vendors.) The system is intended to provide fast, cheap solutions to problems using readily accessible data for any city in North America, and one of its major advantages is in the information it can provide on remote cities as an alternative to, or in advance of, a site visit.

Two data files are used. The first, referred to as the image file, contains a network of major streets or natural boundaries which can be displayed on an interactive graphics terminal as a map of a study area. Suitable image data can be obtained from the census tract-boundary files available for any North American city. The second file contains one or more measures of demand for each of a series of statistical zones, together with the zone centroid locations. Census tracts are of a suitable level of spatial resolution for use in locating stores with thresholds of 50,000 or more, but for smaller thresholds, data are readily obtainable by Enumeration District.

Three classes of stores are recognized by ILACS. Fixed stores are existing facilities owned by the client, unfixed stores are new sites whose locations are to be optimized, and competitive stores are existing facilities owned by other chains. Distances are evaluated along straight lines in continuous space, but can be modified by defining barriers to represent unbridged stretches of river or major gaps in the city street network. Trips are assumed to run around the end of any barrier lying between a demand point and the nearest site. Freeways offer increased speed, and consumers are assumed to follow the path of minimum travel time if a freeway crosses a straight line drawn from demand point to nearest site, reducing the effective distance by an appropriate amount.

Let the coordinates of the n demand points be denoted by x_i, y_i, $i = 1,n$, with weights w_i defined from the variables in the demand data file. The facilities are divided into three sets: the existing or fixed sites owned by the client and denoted by F, the optimized or unfixed sites denoted by S, and the competitive set C. The complete set of sites will be represented by T, with the coordinates of each individual site j being denoted by u_j, v_j. Finally, a_{ij} will be 1 if demand point i is allocated to site j, otherwise O, for all $i = 1,n$ and for all members j of T. The optimization problem is to find locations for the set S, and allocations for all demand points, in order to maximize or minimize some objective. Distances from all demand points to all sites will be evaluated as straight lines

$$d_{ij} = [(x_i - u_j)^2 + (y_i - v_j)^2]^{1/2} \tag{3}$$

except where modified by barriers or freeways, or unless some other metric is to be used.

ILACS recognizes two objective functions that respectively define the Market Share Model (MSM) and Competition-Ignoring Model (CIM):

MSM: $$\text{Max} \sum_i \sum_{j \in S \cup F} a_{ij} w_i \tag{4}$$

where $a_{ij} = \begin{cases} 1 \text{ if } d_{ij} < d_{ik} \text{ for all } k \in \text{T (assuming no ties)}, k \neq j \\ 0 \text{ otherwise} \end{cases}$

In other words, MSM maximizes the share of the market allocated to stores in the two sets S and F.

CIM $$\text{Min} \sum_i \sum_{j \in S \cup F} a_{ij} w_i d_{ij} \tag{5}$$

$$\text{where } a_{ij} = \begin{cases} 1 \text{ if } d_{ij} < d_{ik} \text{ for all } k \in \mathrm{S} \cup \mathrm{F} \text{ (assuming no ties), } k \neq j \\ 0 \text{ otherwise} \end{cases}$$

As noted earlier, CIM will maximize total demand at the client's sites under certain minimal assumptions about the demand elasticity function, as follows. Suppose that the probability that a randomly chosen consumer will choose one of the client's stores is a decreasing linear function of the distance from the nearest one,

$$p_i = c - mz_i, \; z_i = \min_{j \in \mathrm{S} \cup \mathrm{F}} d_{ij} \tag{6}$$

and that no consumer is further than c/m from the nearest store. Then CIM maximizes market share for the client.

Both optimization problems can be solved in ILACS using either the Tornqvist et al. (1971) or Alternating (Cooper 1964) heuristic algorithms. In general terms, the former is a more reliable heuristic based on patterned search, but it is also more expensive in computing time. The latter algorithm alternates between allocation of demand and relocation of each site to the optimal location within the demand allocated to it, until convergence is obtained. It is relatively fast but more likely to converge to a suboptimal solution. ILACS provides several methods of generating starting positions for optimized sites but, in most cases, intuitive guesses provide the best approach. They can be input interactively on a display of the city using a graphics cursor or light pen.

No optimization problem can ever incorporate all the criteria used to make a locational decision. The objective functions in ILACS ensure that the sites selected are optimal with respect to demographic distributions and existing and competitive sites, but they ignore site availability, zoning, and traffic conditions. The advantage of the interactive graphic approach used in ILACS is that these factors can be considered and evaluated as distorting influences on the optimal solution. For example, it may be necessary to move some distance from an optimum to find a feasible site. At this point, the analyst can return to the terminal in order to evaluate the impact on other proposed sites, or the loss of market which will result from the move. In some cases, this process may continue for a long time as sites become available. One advantage of continuous space is that its algorithms have relatively small memory requirements, so ILACS is comparatively mobile and it may be used in the field during real estate evaluation.

FIGURE 2

Optimization of a New Site for a Restaurant Chain with Two Existing Sites*

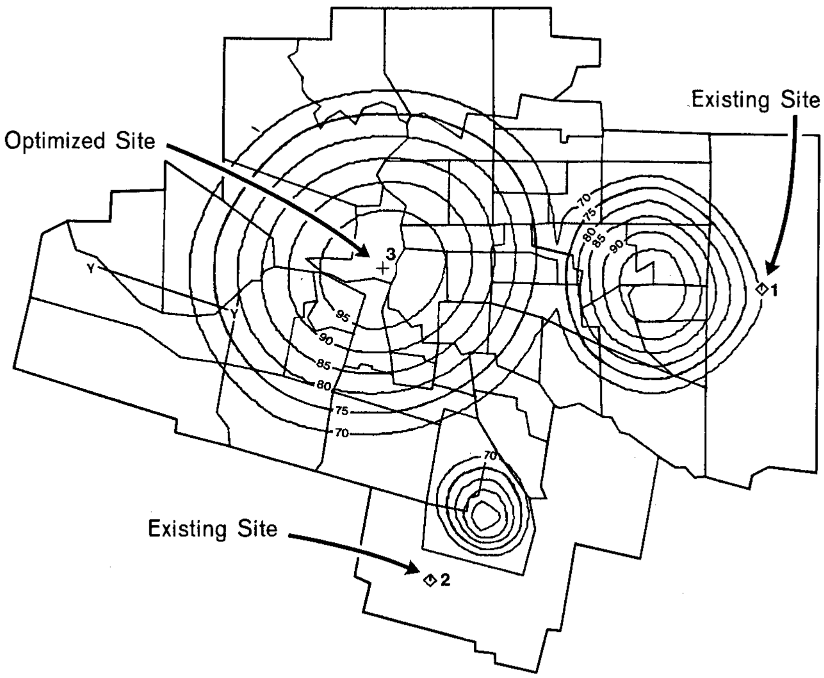

* Optimization of a new site in London, Ontario, using CIM, for a restaurant chain with two existing sites. Weights were populations aged 15 to 44 by census tract, located at geographical centroids.

Example Application

A restaurant owner wished to locate three sites in London, Ontario, and had already taken options on two. The analysis was based on census tract data, which were adequately detailed given the threshold market of 75,000 people. The client identified his market segment as aged 15 to 44 with household incomes of $25,000 and over. Daytime population patterns were of concern because of the potential lunchtime restaurant market. However, the client was unable to provide any guidance as to how these three criteria should be combined, so separate solutions were developed using each in turn as the demand distribution. All three could then be used in evaluating alternative sites.

Because the product was unique, the client opted for the aggressive CIM strategy. A barrier was located to allow for a long, unbridged stretch of the Thames River. The result is shown in Figure 2, based on the 1971 census tract outlines for the city. The contours show the effect of moving a facility away from its optimum or fixed location, and can be used to assess the disbenefits associated with suboptimal location. The contoured value at any point is the contribution to the objective function from the demand served by a given facility, if it were located at that point, and if all other facilities were held constant in their fixed or optimum location. The value contoured is expressed as a percentage of the optimum value. For example, a site on a 95 percent contour would obtain 95 percent of the share possible at the optimum. Note that the existing sites are each some distance from their respective 100 percent peaks. In fact, the optimum position for site 3 proved to be approximately the same for all three market definitions and the client entered into negotiations for a site in the immediate vicinity.

CONCLUSION

Attempts to extend location-allocation modeling from the traditional public service sector context to retailing have led to a number of difficulties in the modeling of consumer spatial behavior and in developing appropriate objective functions. This paper has argued that much of this process may be unnecessary because, although the retail problem is vastly more complex and less planned than many public sector examples, nevertheless there are strong rational arguments for the use of simple location models. We propose the use of an interactive system based on readily available census data applied in continuous space, at an initial cost not much more than that of a simple trade area analysis. If detailed information on spatial information, demand elasticity, or street networks is available, or if its collection is shown to be cost-effective, it can be readily incorporated in place of the simple assumptions made here.

Location-allocation models inevitably emphasize locational, and therefore demographic, factors at the expense of other marketing variables. Again, interactive solution allows other factors to be considered, at least subjectively, and the computer mapping of objective function contours, as shown in Figure 2, may allow more objective evaluation.

Perhaps the greatest need at the present time in this area is for a spatial data base of retail activity to complement the census-based material by providing information on existing and competitive locations. A latitude-longitude or UTM data base containing simple data on owner, floor area, and so on, would be immensely useful in allowing rapid search for viable locations in remote cities.

EXECUTIVE SUMMARY

Location-allocation is the simultaneous location of central facilities and allocation of dispersed demand to them, in order to optimize some criterion. In the retail sector, the central facilities are stores, and the allocation results from the choices made by consumers in deciding which stores to visit. The decision to open a new store or close an existing one will affect the other stores in the chain and also any competing stores through the resulting redistribution of consumers. So, in order to make optimal decisions about store locations, it is necessary to be able to predict consumer response through a spatial choice model.

The paper discusses three major issues in retail location-allocation modeling: the suitability of current spatial behavior models, including MCI, the choice of appropriate objective functions, and the cost-effectiveness of solution procedures. The site selection process moves through a series of stages of successive narrowing of options, from an initial decision which might require choosing one from a number of metropolitan areas to the development of a general strategy for the city as a whole, and finally to the consideration of specific parcels of land. The models discussed in this paper are more appropriate to the first and second stages, where the cost of modeling is a major concern. They allow a retailer to develop location strategies quickly and often without the need for a site visit. Furthermore, the paper argues that the costs associated with more elaborate modeling may not be justified even at the third stage because of numerous technical problems and uncertainties about the future.

The paper proposes two locational criteria which can be used to select sites for new stores, or to choose among candidates for closure. The two criteria represent opposing strategies with respect to competitors: a conservative strategy of avoidance and an aggressive strategy which ignores the competition. Both are implemented in an interactive system which allows the user to manipulate locations at a graphic display. The data sets required are readily available for any metropolitan area in North America. The model is most suitable for application to low-order functions such as convenience stores, gasoline outlets, banks, and restaurants.

Although the paper argues that the need for cost-effectiveness and speed often requires a simple model, the location-allocation framework can readily accommodate any available data on consumer spatial behavior, demand elasticity, and the fixed and variable costs of the retail operation. The model is flexible enough to be used repeatedly as other factors force a retailer to modify an initial optimal plan, and it is portable enough to be used in the field.

REFERENCES

Achabal, D. D., W. L. Gorr, and V. Mahajan (1982), "MULTILOC: A Multiple Store Location Decision Model," *Journal of Retailing,* **58** (2), 5–25.

Beaumont, J. (1980), "Spatial Interaction Models and the Location-Allocation Problem," *Journal of Regional Science,* **20,** 37–50.

Cooper, L. (1964), "Heuristic Methods for Location-Allocation Problems," *SIAM Review,* **6,** 37–52.

Dirichlet, G. L. (1850), *"Uber die reduktion der positiven quadratischen formen mit drei unbestimmten ganzen Zahlen," Journal fur die reine und angewandte Mathematik,* **40,** 216.

Drezner, F., and G. O. Wesolowsky (1983), "The Location of an Obnoxious Facility with Rectangular Distances," *Journal of Regional Science,* **23,** 241–248.

Fotheringham, A. S. (1981), "Spatial Structure and Distance-Decay Parameters," *Annals of the Association of American Geographers,* **71,** 425–436.

Ghosh, A., and C. S. Craig (1984), "A Location-Allocation Model for Facility Planning in a Competitive Environment," *Geographical Analysis,* **16** (January), 39–56.

———, and S. L. McLafferty (1982), "Locating Stores in Uncertain Environments: A Scenario Planning Approach," *Journal of Retailing,* **58** (4), 5–22.

Goodchild, M. F. (1978), "Spatial Choice in Location-Allocation Problems: The Role of Endogenous Attraction," *Geographical Analysis,* **10,** 65–72.

——— (1979), "The Aggregation Problem in Location-Allocation," *Geographical Analysis,* **11,** 240–255.

———, and P. J. Booth (1980), "Location and Allocation of Recreation Facilities: Public Swimming Pools in London, Ontario," *Ontario Geography,* **15,** 35–51.

Hakimi, S. L. (1964), "Optimum Location of Switching Centres and the Absolute Centres and Medians of a Graph," *Operations Research,* **12,** 450–459.

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.

Hodgson, M. J. (1978), "Towards More Realistic Allocation in Location-Allocation Models: An Interaction Approach," *Environment and Planning A,* **10,** 1273–1286.

——— (1981), "A Location-Allocation Model Maximising Consumers' Welfare," *Regional Studies,* **15,** 493.

Leonardi, G. (1981), "A Unifying Framework for Public Facility Location Problems—Part I: A Critical Overview and Some Unsolved Problems," *Environment and Planning A,* **13,** 1001–1028.

Maranzana, F. E. (1964), "On the Location of Supply Points to Minimise Transportation Costs," *Operational Research Quarterly,* **15,** 261–270.

Scott, A. J. (1970), "Location-Allocation Systems: A Review," *Geographical Analysis,* **2,** 95–119.

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

Thiessen, A. H., and J. C. Alter (1911), "Precipitation Averages for Large Areas," *Monthly Weather Review,* **39,** 1082–1084.

Tornqvist, G., S. Nordbeck, B. Rystedt, and P. Gould (1971), "Multiple Location Analysis," in *Lund Studies in Geography,* Ser. C, General, Mathematical and Regional Geography, No. 12, University of Lund, Sweden.