J Geograph Syst (2001) 3:181–197

Journal of Geographical Systems

Spatial decision support system for home-delivered services

Wilpen Gorr, Michael Johnson, Stephen Roehrig

H. John Heinz III School of Public Policy and Management, Carnegie Mellon University, Pittsburgh, PA 15213-3890, USA (e-mail: roehrig@andrew.cmu.edu)

Received: 27 September 2000 / Accepted: 22 March 2001

Abstract. We present a spatial decision support system for the non-profit sector, designed to assist planning in the area of home-delivered services such as meals on wheels. Using data collected from existing programs, current and forecasted demographic data, and a set of algorithmic tools, we provide a system for evaluating current meals on wheels facilities, and for making incremental facility location decisions that satisfy coverage and equity requirements.

Key words: Geographic information system, spatial decision support system, facility location models, spatial forecasting

JEL classification: C53, C61, L31

1 The home-delivered services problem

The public, through tax dollars, grant-making foundations, and corporate and private donors, fund the non-profit sector in the United States. With its funding, this sector provides valuable human services for underprivileged and needy segments of the population – the abused, under educated, poor, homeless, addicted, elderly, and so forth. While the provision of goods and services is guided in the private sector by market mechanisms, and in government by the political process, the non-profit sector does not benefit from an underlying, organizing mechanism. Instead, it must rely directly on objective planning methods including clear statements of mission, design of corresponding service delivery systems, and objective performance measures. Such performance measures include *efficiency*, or the extent to which a fixed level of service is provided using the minimum amount of resources, *effectiveness*, or the extent to which client population needs are met, and *equity*, or the extent to which services provided (which may not meet total client needs) are nevertheless delivered in a manner perceived as fair. Society desires that all who need non-profit services have access to them.

Facility location planning for non-profit delivery systems is a critical, but difficult problem depending on both the geographic distribution of target populations and the typically small size and limited catchment areas of facilities. We find that there are three distinct types of human services in regard to facility location planning: home delivered services (e.g., meals on wheels and home nursing), site delivered services (e.g., literacy training and youth recreation), and residential programs (e.g., elderly nursing homes and drug treatment programs). The most difficult of the three, and one that this paper addresses, is home-delivered services. Site delivered services depend on travel preferences and limitations of clients. Residential programs are less sensitive to location, but should take into account travel of family members and care givers who visit residents.

Our work on models for home delivered meals was motivated by a request from a group of major grant-making foundations in Pittsburgh, Pennsylvania. Their job, in part, was to evaluate proposals to fund expansions of existing service delivery facilities or to build new ones. The foundations desired a means of evaluating the merit (objective performance) of such proposals, and believed GIS to be an enabling technology for a solution approach. Prior to the models presented in this paper, the program officers of the foundations had no means to determine the gaps and overlaps in coverage of target populations.

Several features of the non-profit, service-facility location problem make it ideal for a spatial decision support system (see Alter 1980; Densham 1991 and Sect. 2 below):

1) *Map interface:* the natural user interface for the problem is a GIS map showing kitchens, corresponding catchment areas, client stops, the street network, etc. with tools to perform spatial operations – this is no doubt a spatial problem. Our spatial DSS is based on ArcView GIS with map display and menu items for configuring and locating new kitchens. A map of existing gaps provides a needs analysis for facility planning that can be shared and is the key means for coordinating efforts of multiple stakeholders.

2) *Computational limitations:* while comprehensive, optimal solutions are desirable, they appear to be computationally infeasible for this problem (see Sect. 2). One approach to building a spatial DSS would be to start with an optimal solution displayed and then allow the user to refine it using DSS tools and expertise. We leave for future research the construction of comprehensive algorithms that provide near-optimal solutions in a reasonable amount of time. In this paper we focus instead on generating feasible but approximate solutions using an intuitive interactive heuristic.

3) *Semi-structured decision problem:* there are some aspects of the decision problem that are well structured and others best left to expert judgment – this is a key feature of DSS. We provide algorithms and functionality to design high-performance new kitchens incrementally. The user chooses a site and kitchen configuration based on judgment and/or feedback from a previous trial kitchen within the immediate environs. Our traveling salesman-based algorithm estimates a high-performing catchment area for the kitchen. Within the context of filling service gaps, there are a number of examples in which expert judgment is critical. These examples include user knowledge of: (a) a church qualified for and amenable to home-delivered meals service delivery; (b) a potential site that lacks a single piece of equipment to provide needed meal capacity; (c) a new senior apartment complex to be built, greatly increasing demand and (d) the desire of a nearby kitchen to serve clients nominally in the current kitchen's catchment area.

4) *Incrementalism decision setting:* The decision setting, like many in the public sector, is incremental rather than comprehensive (e.g., see the classic, Lindbloom 1968). For example, many organizations *independently* establish and operate meals on wheels kitchens in Allegheny County, Pennsylvania: county agencies, service arms of religions organizations and independent non-profit organizations.

5) *Decoupled solution space:* the decision problem of filling gaps in coverage of the spatial distribution of the elderly population is by nature decoupled, requiring only one or two new kitchens per gap – this makes the problem more amenable to using judgment and a DSS.

In Allegheny County, Pennsylvania, there are 63 volunteer meals on wheels facilities delivering daily hot meals to slightly over 4,000 home-bound elders. We estimate below that a little over 80% of the demand is met by current facilities, with a net benefit, in avoiding costs of residential nursing homes, to be in excess of $100 million annually. The typical facility has four to six routes driven by volunteers with a dozen stops and one or two clients per stop. Generally, meals on wheels uses existing facilities – church, school, and senior center kitchens. The major fixed costs of a new facility are thus organizational: gaining commitment from a facility, securing funding for operations, recruiting and organizing networks of drivers, and designing routes. Federal subsidies have the provision that recipients be home-bound persons aged 64 or older and that delivered meals be at least 140 °F. With insulated carriers, the temperature requirement translates into a 45-min time limit for delivery of the last meal on a driver's route.

Thus, the catchment area of a kitchen depends on its location, meal production capacity, carrying capacity of delivery vehicles, 45 min time limitation for delivery, number of drivers/routes per day, density of clients and street network, travel or "turf" barriers, and efficiency of routing.

Section 2 of this paper presents an overview of the location and vehicle routing literature, and suggests that exact algorithms for our problem, even if implementable, are computationally impractical. Section 3 provides a GIS and traveling salesman-based algorithm for estimating the optimal catchment area of a kitchen. Section 4 reviews our GIS-based methods of forecasting point locations of delivery stops. Section 5 applies the methods of Sects. 3 and 4 to estimate service gaps and suggest new kitchen locations, and Sect. 6 concludes the paper.

2 Location-routing models and spatial decision support systems in the literature

The general planning problem this paper addresses is the home-delivered meals location-routing problem, referred to as HDM-LRP. In this problem, the goal is to simultaneously choose "depot" (kitchen) locations and delivery routes. The depots provide "products" (meals) to spatially dispersed customers via routes driven by multiple vehicles, each of which leave a depot, visit multiple customer locations, and then return to the depot when customer deliveries are finished.

This problem is a combination of two well-known planning problems: a *location-allocation problem*, in which customers are assigned to potential depots and a *multi-depot vehicle routing problem* (MDVRP), in which multiple vehicle routes are designed originating from and returning to depots. Since the MDVRP is a generalization of the single-depot vehicle routing problem, which has been shown to be NP-hard (Lenstra and Rinnooy Kan 1981), the location routing problem is NP-hard as well. Thus, optimal solutions are unlikely to be generated for problems of realistic size. Indeed, Laporte, Norbert and Talliefer (1988) have solved LRPs to optimality for at most 80 nodes. As will be shown, a typical HDM-LRP contains on the order of thousands of nodes. Thus, heuristic solutions are likely to be the only feasible method for solving practical instances of the LRP.

Bodin and Golden (1981), Assad (1988), Laporte (1988) and Min et al. (1998) have all presented comprehensive taxonomies of vehicle routing problems and in particular the LRP. Relevant characteristics of HDM-LRP include:

- *Single stage* – products originate at the depots rather than at a central plant;
- *Deterministic* – the nature of location/routing parameters such as customer demand is known and fixed with certainty;
- *Multiple facilities*;
- *Multiple vehicles*;
- *Uncapacitated vehicles* – for delivery of relatively small meals, vehicle space is not likely to be a constraining factor;
- *Capacitated facilities* – in practice, volunteer staffing at kitchens limits number of meals and available to service clients;
- *Route length limits* – limits on the length of routes driven by delivery vehicles
- *Primary facility layer* – kitchens serve as origins and destinations of delivery routes and not as transshipment points;
- *Multiple-period planning horizon* – a planning organization must account for changing levels of demand as customer demographics shift over time;
- *Unspecified time windows with no deadline* – customers generally do not require specific meal delivery times;
- *Multiple objectives* – vehicle operating costs; fixed facility location cost and equity/fairness considerations;
- *Real-world data* – the subject of Gorr et al. (2000).

We examine some of these problem characteristics in more detail. A primary motivation for this paper is that a shift occurred in locations of populations likely to use HDM services, from the central city to suburbs, and from inner-ring suburbs to outer-ring suburbs and rural areas. Many of the service gaps of meals on wheels occur in the suburbs and rural areas. While Gorr et al. (2000) examine the cause of this shift in more detail, it suffices to note that any scheme that designates certain kitchens and catchment areas for the current period is likely to be outdated, especially in areas such as Pittsburgh that have suffered "white flight" from urban to suburban communities starting in the 1960's, and long-term population losses. Because volunteers staff HDM facilities, it is probably not reasonable to solve a planning model that prescribes one configuration of kitchens in one period and another, quite different configuration in another period. Laporte and Dejax (1989) have examined dynamic LRPs. However, we will avoid the added data requirements and computational difficulty associated with explicitly incorporating dynamic considerations by solving a single-period planning model.

Another related modeling issue of interest is stochasticity of model data. We (Gorr et al. 2000) have found that it is simply not possible to identify locations of all clients currently receiving HDM services in our study area, Allegheny County, Pennsylvania. Thus, we have had to design procedures to estimate and forecast customer demand. Given known shifts over time in populations likely to need HDM services, it would be correct to regard customer demands as stochastic. While LRPs under uncertainty (Laporte et al. 1989) have been studied, the computational requirements are excessive for an initial effort to solve HDM-LRP. Thus, we treat customer demand (and other model parameters) as fixed and known with certainty.

Finally, objectives of HDM-LRP are of interest. Traditional LRPs (Perl and Daskin 1985; Laporte et al. 1988) have minimized the sum of fixed facility location costs, vehicle operating costs and vehicle routing costs, the latter usually proxied by total distance traveled. However, in the context of HDM, it is often not clear what fixed facility costs are or how they ought to be measured. This is due to the fact that while HDM kitchens are often provided free of charge by churches, schools, community centers and other organizations, there are large non-monetary costs such as those associated with recruiting volunteer drivers. Moreover, there are significant equity considerations associated with HDM service planning, mainly ensuring that all potential clients in a region have access to this valuable service. Giannikos (1998), List and Mirchandani (1991) and ReVelle et al. (1991) have addressed equity considerations in location-routing problems in the context of hazardous waste transport and disposal, however these models have not explicitly incorporated the Hamiltonian tour constraint that makes LRPs so difficult. Thus, we treat HDM-LRP as, fundamentally, a multi-objective planning problem relying on expert judgment using a decision support system. Moreover, the interactive solution method presented in this paper is motivated, in part, by the fact that fixed facility costs cannot often be well defined.

Having defined in some detail the HDM-LRP model, we now address solution techniques. As mentioned above, the difficulty of LRP in general precludes exact techniques for realistically sized problems. However, researchers have developed a number of heuristic methods for solving LRP. Min et al. (1998) categorize heuristic solution techniques for LRP as:

- Location-allocation-first, route-second
- Route-first, location-allocation-second
- Savings/insertion
- Improvement/exchange
- Others

Madsen (1983) provides examples of location-allocation first, route-second and route-first, location-allocation second, both integrated with a savings heuristic. Or and Pierskalla (1979) use an improvement/exchange heuristic to combine solutions to location-allocation and vehicle routing problems solved, essentially, in parallel. Perl and Daskin (1985) use a combination of three combinatorial optimization problems, the Multi-Depot Vehicle Dispatch Problem, the Warehouse Location-Allocation Problem and the Multi-Depot Routing Allocation Problem to generate solutions to the original LRP. Renaud et al. (1996) use tabu search to solve LRP, though fixed facility costs are ignored. The literature is inconclusive as to the relative efficacy of these alternative approaches; to our knoswledge, no researchers have compared worst-case performance for LRP heuristics (though Li and Simchi-Levi 1990 have done so for multi-depot vehicle routing problems). It is likely that meta-heuristic approaches such as tabu search, simulated annealing or genetic algorithms may be very competitive with more traditional approaches.

In this paper, we use an interactive version of the location-allocation first, route-second approach. This strategy is motivated by our inability to incorporate monetized fixed facility costs into an explicit optimization model, the need to allow users to apply equity considerations where necessary, and the incremental decision setting of non-profit organizations.

We briefly address a number of implementation issues that have been addressed in the vehicle routing literature. Assad (1988) and Bartholdi et al. (1983) have detailed the key role of the dispatcher in implementing recommendations of computerized (and in the case of Bartholdi et al. 1983, *non-computerized*) vehicle routing systems. Often, preferences of the dispatcher transform an optimization-based DSS into a "satisficing" DSS. Assad also emphasizes the role of benefits measurement in generating acceptance of vehicle routing systems. In our case, since no comprehensive planning method for HDM over a service area as large as a county has existed previously, such benefits may be difficult to quantify.

We know of only two papers in the literature that directly address the meals on wheels problem. Wong and Meyer (1993) applied network optimization models to the meals on wheels problem at the operational level. The value of this study is limited due to the very small sample sizes used. They compared two routes used in practice, an urban route and a rural route, with two optimal routes generated by a single depot vehicle routing procedure. The optimal routes for the urban and rural were marginally better, 1% and 9% shorter in travel distance respectively. They also solved the p-median problem to find the ideal kitchen location for a single kitchen's clients. The ideal kitchen had a 4% savings in total travel distance to clients from the kitchen. This savings is small and it is not clear that it would translate into savings in route lengths. The paper of Bartholdi et al. (1983) applied the concept of a space-filling curve to determine delivery routes. Both of these papers address only the routing aspect of the overall location-routing problem.

This paper focuses not just on solving the location-routing problem for home-delivered meals and implementing the solution in a decision support system, but as well on taking advantage of the spatial, interactive and incremental nature of our solution to HDM-LRP in order to design a spatial decision support system (SDSS). We thus turn our attention to a survey of SDSS principles and SDSS applications related to the location-routing problem. SDSS address semi-structured planning problems by integrating expert knowledge and specialized solution algorithms in a problem domain for which the input data have a spatial component or for which the solution space has a spatial dimension (Armstrong and Densham 1990; Ayeni 1997). These systems can address a number of spatial data applications: transformation, synthesis and integration, updating, forecasting, impact analysis and optimization (Birkin et al. 1996). SDSS has been a popular area for research applications. Gould and Densham (1991) have surveyed a wide range of SDSS developed, and a number public-sector, service delivery SDSS applications have been documented since then (see e.g., Wong and Meyer 1993; Lolonis 1994; Begur et al. 1997, and Xin 1999).

Strategic planning problems such as HDM-LRP represent a problem domain in which a traditional normative modeling approach, generating "ideal" solutions intended to be implemented at once, do not reflect actual planning practice. In this particular planning domain, "what-if" questions and small changes to the status quo are more likely to represent the real needs of managers (Densham and Rushton 1988).

The SDSS literature contains examples of models that are more highly structured and operationally oriented than HDM-LRP, however. For example, Begur et al. (1997) design daily visit schedules for nurses visiting homebound clients, while Wiegel and Cao (1999) design daily visit schedules for appliance repair professionals. Even in these cases, the spatial analytic strength of GIS is combined with practical needs of schedulers to incorporate sensitivity and exploratory analyses of visit schedules.

Wong and Meyer (1993) present a comprehensive, well-reasoned exposition of the role of SDSS in home delivered meals service planning and delivery. Unfortunately, the modeling and algorithmic deficiencies noted in their work above, and the very incremental nature of their results (one urban and one suburban delivery route generated and compared to the status quo) prevent their model from consideration as a comprehensive SDSS solution to HDM-LRP.

To conclude, the "generic" LRP has yielded a variety of explicit mathematical formulations that incorporate real-world considerations and solution techniques that appear quite promising. Moreover, there are a number of real-world applications of LRP and aspects of VRP implementation in general that are relevant for this study. The combination of GIS and incremental, interactive algorithms embodied by spatial decision support systems allows us to take advantage of the difficulties in modeling in closed form, and finding near-optimal solutions to HDM-LRP, to design a SDSS that is more appropriate to decision maker needs than an algorithmic, "black box" approach. Finally, the scale and scope of our proposed SDSS appears to exceed that of another related application for HDM previously documented in the research literature.

3 Algorithms for a DSS approach

Because our location problem is sufficiently large that a closed-form, or even computable, optimal solution would be difficult to achieve, we concentrate on finding accurate solutions to sub-problems, and bring these together in a way that is useful to an actual decision maker.

The problem of organizing the supply of home-delivered services typically includes constraints that are often waved away in global mathematical formulations. In particular, there is usually some pre-existing infrastructure that cannot simply be discarded; it needs to be augmented in a gradual move toward optimality. Client locations change over time, so a solution must be robust against such changes, or at least must anticipate them. Supplying services efficiently is crucially dependent on existing street networks, so a solution concept must operate at this micro-level.

Our initial objective is to identify, and then fill, existing gaps in service. A second step is to identify service overlaps, and suggest alternative catchment areas to eliminate them. Our strategy is therefore based on

1. Estimating client densities in a spatial environment,
2. Using data from a subset of current service providers to estimate existing catchment areas,
3. Determining existing gaps in coverage,
4. Identifying and evaluating candidate locations for new providers, and
5. Providing tools for evaluating coverage overlaps, and for assessing the impact of reducing or eliminating them.

A primary objective of non-profit service providers is to maximize coverage of the client population. Because this population is not uniformly distributed spatially, distribution facilities in more sparsely populated areas will have physically larger catchment areas, at least for a fixed distribution capacity. But when coupled with limits on distribution time, the 45-min limit on meals on wheels delivery, for example, it becomes more difficult for such a distributor to service a large client base. Countering this is the possibility that average travel speeds may well be higher, perhaps mitigating the travel distance penalty.

A second objective is to minimize facility costs. Locating facilities in densely populated areas is cost efficient, but conflicts in an obvious way with the goal of complete coverage. Population density in this application depends on the distribution network (i.e., the street network), since simple physical proximity (1/2 mile away, but across a river) doesn't always imply a short travel time.

A third objective is to minimize travel costs. To do this, delivery routes need to be intelligently designed. For a single delivery route, and a known set of customers, this is the classical traveling salesman (TSP) problem. Typically, however, a facility will have multiple "salesmen", forcing a selection of multiple routes. We deal with this non-optimally, by assigning geometrically defined sectors of the catchment area to individual delivery people.

Sampling and algorithmic details

To model the delivery of services, we assume initially that facility locations and client densities are known. We discuss the estimation of client density later, in Sect. 4. Since actual client locations change frequently, we rely on a sample of likely locations, randomly generated using the given density and projections of Census data on the elderly population at the block group level. We do this by selecting points from a two-dimensional uniform distribution. Each point can be assigned an amount of the good or service to be delivered. In the meals on wheels application, for instance, a single client location may require the delivery of multiple meals. The distribution of meals per location was estimated from historical data.

Each sample point is located within the existing street network, using standard GIS procedures. If the point is not within a certain distance of a street (typically 1/8 mile), the point is discarded. Sampling continues until the estimated number of clients has been generated. In our application, using the ArcView GIS software, this and other procedures are coded in the Avenue scripting language. There are other ways to perform this sampling. For instance, one might first choose a street segment, and then choose a point along the segment. However, the method we used is quite simple, and the resulting distribution is faithful to the street network, since we insist that sample points be close to a street.

For each facility, the number of routes r is known. Each facility has a fixed capacity c, so the next task is to attempt to supply this number of clients, subject to other constraints such as time limitations and geographical boundaries.

The algorithm is a simple one. The area surrounding the facility is broken up into r equal-angle wedges, using lines extending radially out from the facility location. Each of these is assigned to a driver. Next, each wedge is bounded by a circular segment whose center is at the facility and whose radius is, initially, rather small (typically 1/2 mile). All of the sample points lying within this pie-shaped geographical area constitute the client set for one driver. Using the street network, and a heuristic TSP algorithm available in ArcView, we find the fastest route that starts and ends at the facility and stops briefly at each client location. If the total time is acceptable, the radius of the pie-shaped sector is increased. The goal is to find the largest radius such that the client load within the sector can be serviced within the specified time limit. This process is repeated for each sector. The resulting collection of sectors (which may have different radii because of client locations and the details of the street network) is taken to be the catchment area for the facility.

For the catchment area estimates reported here, a "stop time" of one minute was used for each client; this was estimated by actually riding along on several existing routes. We used the Dynamap 2000 street maps from Geographic Data Technology. These maps include estimates of travel times based on street type and length. We also considered the possibility that the TSP heuristic in the GIS software package might give routes that were far less than optimal for this difficult combinatorial problem. We took a sample of sectors and the enclosed stop points and compared the built-in heuristic with a known optimal TSP procedure (the Held-Karp algorithm as implemented by the *Concorde* group, available at http://www.keck.caam.rice.edu/concorde.html). We could find no differences between the heuristic and optimal solutions, at least for our small sample. Another potential problem is inaccuracy in finding the shortest paths between stops, the main data used by a TSP algorithm. We could find no straightforward way to check this, and simply relied on the existing routines in ArcView.

A typical catchment area, for a facility with five delivery routes, is shown in Fig. 1. The smaller dots are clients, while the facility is shown by the larger dot. The sector to the southwest was truncated by a geographical boundary (the county line), and a portion of another catchment area is seen to the northwest. For this facility, 70 meals were available, but only 66 could be delivered within the 45-min time limit. In addition, there appears to be unmet demand in the vicinity of the facility, particularly to the northeast.

Another catchment area is shown in Fig. 2, for a facility with a capacity of 150 meals and six drivers. Our algorithm predicted that a maximum of 95 meals could be delivered. Several features of this catchment area are worth noting. First, the client population is very unevenly distributed. In fact, there are very few clients to the northeast of the facility, and many un-serviced clients south of a river near the facility. This poses a problem for the delivery algorithm, since it allocates equally spaced sectors for each delivery route, and several of these sectors enclose almost no clients. A decision maker would of course notice this immediately, and would attempt to re-allocate delivery resources to the southwest.

The GIS display suggests, however, that this may not improve coverage substantially. Note that the existing routes to the south and southwest have reached their 45-min delivery times, and any additional routes in that direction must follow the same path over the only bridge available. This suggests that a superior solution might include a new facility located to the south of the river.

Other approaches to allocating clients to drivers are suggested in the literature. For example, a p-median approach might be taken, grouping clients into p groups, where p is the number of drivers at a kitchen. However, our problem requires that clients simultaneously be assigned to both kitchens *and* drivers, that is, the collection of clients assigned to a kitchen is not known in advance. This forces an extra layer of allocation onto the problem (as discussed in Sect. 2), resulting in a considerably more difficult optimization problem. Nonetheless, improvements can certainly be made to the static wedge procedure given here. For instance, user-defined wedge angles and

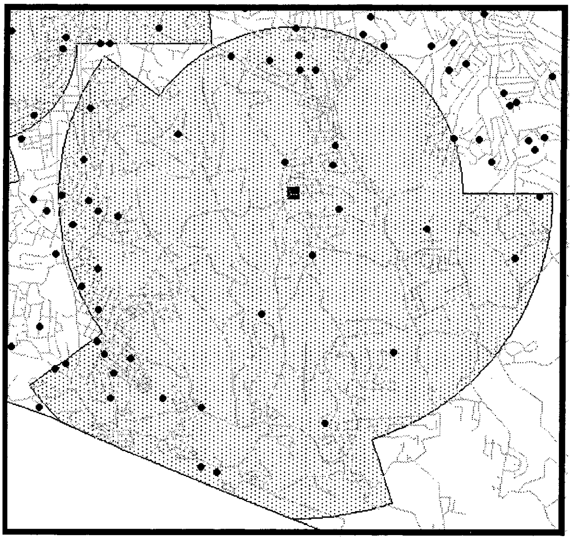

Fig. 1. A typical catchment area

orientations would allow wedges to be tailored to accommodate the clustering of clients in the vicinity of a kitchen. We dealt with some of these problems by allowing users to define barriers, described next.

During the process of estimating catchment areas for an existing set of facility locations and capacities, it was quickly noticed that maximal catchment areas often overlap. Figure 3 shows an example of this. This is probably fairly common for many organizations maintaining multiple facilities. Obviously, though, this allocation of

Fig. 2. Another catchment area

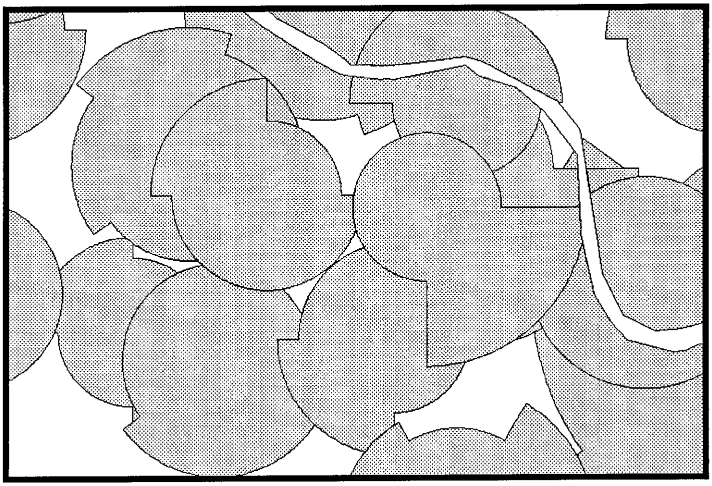

Fig. 3. Overlapping catchment areas

clients to facilities involves double counting, so should be eliminated in order to better understand real shortages or over-capacity. To estimate catchment areas without overlap, Fig. 4 shows a facility (the large central dot) for which a catchment area is to be estimated. Since it is fairly close to a number of other facilities (the large square dots), two "off-limits barrier" areas are defined, one to the northeast, and a larger one to the south and west. Defining these areas effectively removes the clients within them from consideration by the catchment algorithm. The central, roughly triangular region is the area "assignable" to the facility. The resulting catchment is shown in Fig. 5. This process allows the decision maker to assign territories to facilities,

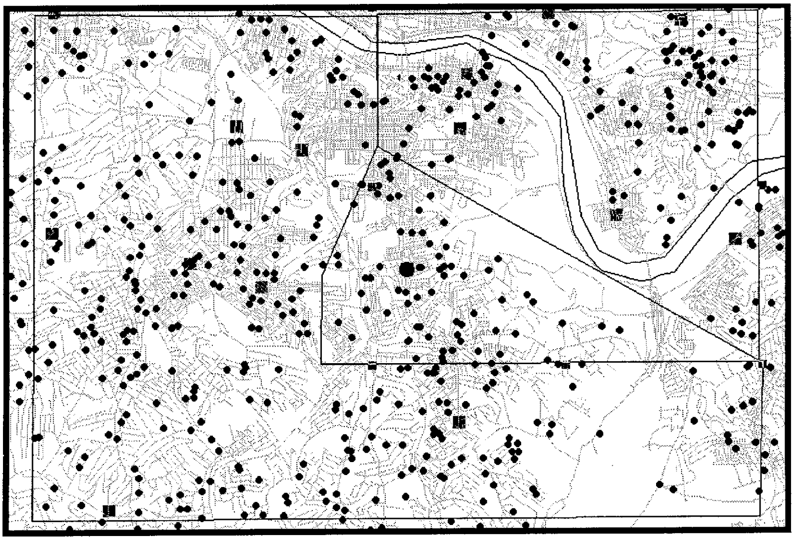

Fig. 4. Boundaries around a catchment location

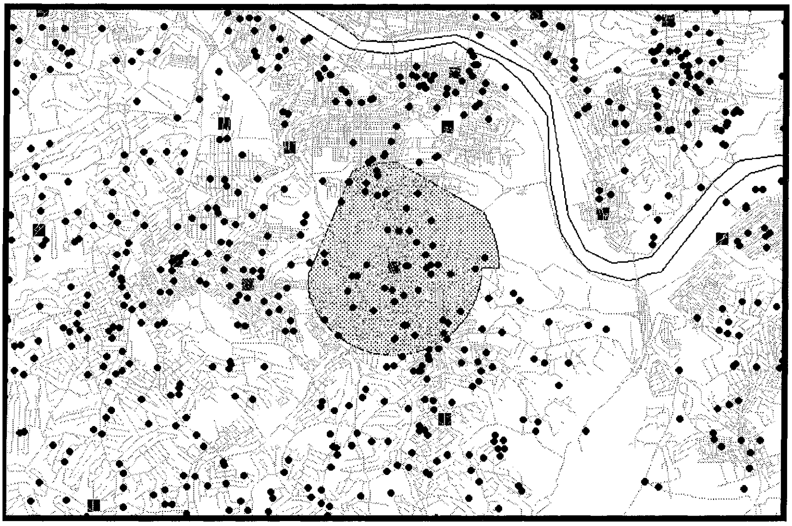

Fig. 5. Resulting catchment area

prevents multiple counting of clients, and thus makes clear where capacity is needed. Of course, if imposing these boundaries prevents a facility from using all of its capacity, this will be discovered as well. The decision maker can use a trial and error process to determine an initial set of boundaries, and then modify them to accommodate excess capacity where possible barrier boundaries are used.

4 GIS-based estimation: Allegheny County HDM

Our purpose is to provide a tool that can be used to plug gaps in the Allegheny County meals on wheels kitchens. Our algorithm of Sect. 3 needs stops locations and the number of clients at each stop in order to estimate a kitchen's catchment area.

A naïve forecast (tomorrow is the same as today) is often difficult to improve upon in forecast competitions. A map of existing meals on wheels kitchens in Allegheny County reveals, however, that there are large gaps in the coverage of elderly populations (see Fig. 7 below). Thus, even if we were able to collect the entire population of stops and clients per stop for existing meals on wheels routes, we would not have all demand points for this service. Moreover, the client population is dynamic, with clients being added and removed as time passes, and with spatial trends decreasing elderly populations in urban areas and increasing them in suburban areas. Clearly, then, we need to forecast stops using a model.

Our approach to making such a forecast has several steps. At the basis are three components: 1) five-year-ahead forecasts of elderly population at the block group level made by Claritas, Inc., 2) our estimate of the usage rate for meals on wheels (i.e., the percentage of elderly who use meals on wheels services), and 3) a frequency distribution of the number of clients per stop. Gorr et al. (2000) provide an estimate of 2% for usage rate based on a sample of stops for 25 kitchens and a regression model for removing an underestimate bias. A tabulation of clients per stop from existing kitchens yields the distribution in Table 1. The expected number of clients per stop is 1.3.

Finally, to make a stop points forecast, we used the following steps for a each block group: 1) select a random sample of points within the block group of size equal to 2 percent of the elderly population forecasted by Claritas divided by the expected number of clients per stop and 2) reverse geocode each sample point to a street address. One realization is in Fig. 6.

5 Application of the spatial decision support system

We implemented our algorithms in ArcView GIS, with a few additional menu items for drawing barriers, locating a new kitchen, setting the kitchen's design parameters,

Table 1. Distribution of clients per stop

Number of clients per stop	Frequency
1	0.80
2	0.15
3	0.02
4	0.01
5+	0.02

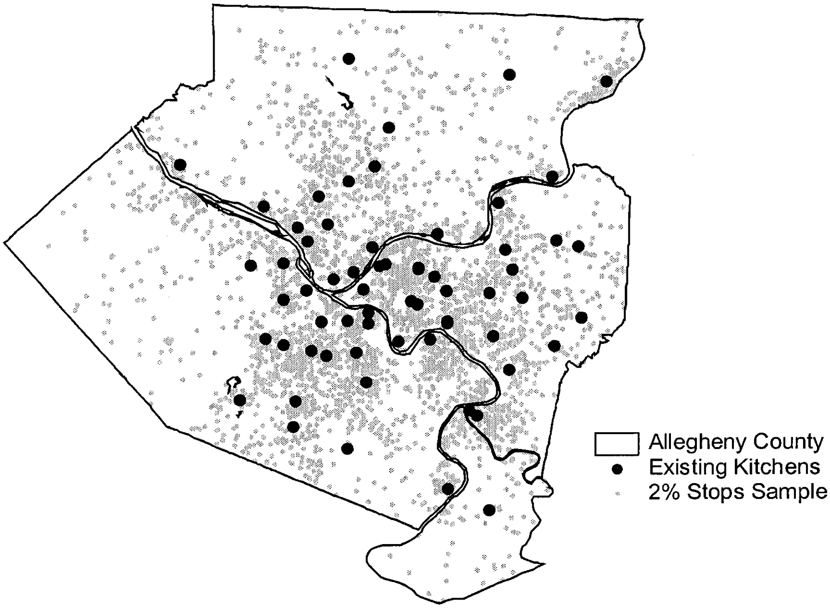

Fig. 6. 2% sample of stops and existing kitchens

and estimating the resulting catchment area. The application presented here is a gaps analysis of the current meals on wheels kitchens, and then an example set of new kitchens for plugging identified gaps.

To gauge the efficiency of the current kitchen configuration, we used a conservative demand estimate of 2% of the elderly population. This resulted in a simulated client sample of 4379 stops and 5069 total clients. Then we applied the algorithms of Sect. 3 to estimate catchment areas for all 63 kitchens. In applying these algorithms, we used the boundary technique to prevent overlaps of catchment areas. Current catchment areas have considerable overlaps and therefore are either inefficient or do not maximize coverage of clients. The result is an estimate of the maximum capacity of existing kitchens to cover client demand.

Figure 7 is the resulting estimate of the performance of the existing kitchens, projected five years into the future. Gaps in coverage of the elderly population are represented by light gray points left uncovered by the filled gray catchments. It is readily apparent that there are significant shortcomings in coverage, since about 17% of the sample stops cannot be serviced by the existing kitchens. The gaps are largely in suburban areas, and sometimes, as it turns out, in well-to-do areas (the latter determined from Census block group data). While this general trend is known to, and is a concern of, funding agencies, the DSS analysis quantifies this somewhat vague understanding, and helps focus attention on providing assistance where it is most needed. Furthermore, several of the areas estimated to be uncoverable by the existing kitchen network are not obvious at all. These tend to be densely populated bands between existing catchments, some of which could be filled by expanding existing capacities and others that require new kitchens.

An important use of the HDM DSS is to explore the utility of adding additional kitchens. The GIS displays are an excellent way to present coverage gaps to decision makers, and experiment with new kitchen locations. We illustrate this by presenting one straightforward approach to identifying new kitchen sites and quantifying their effects.

Figure 8 shows some candidate sites in the uncovered areas, constructed very simply from a geocoded table of churches and K-12 schools in Allegheny County. Since many current kitchens are in similar facilities, this provides a starting point for identifying likely candidates for new kitchens. In practice, of course, a decision maker

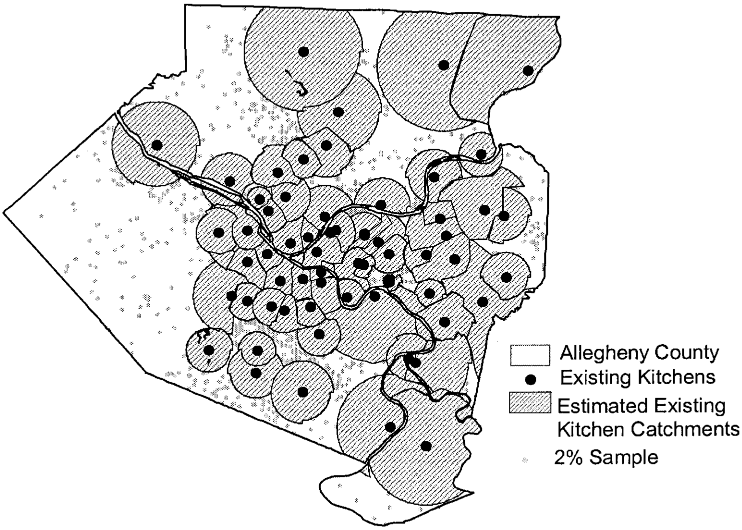

Fig. 7. Non-overlapping catchment areas for all existing kitchens

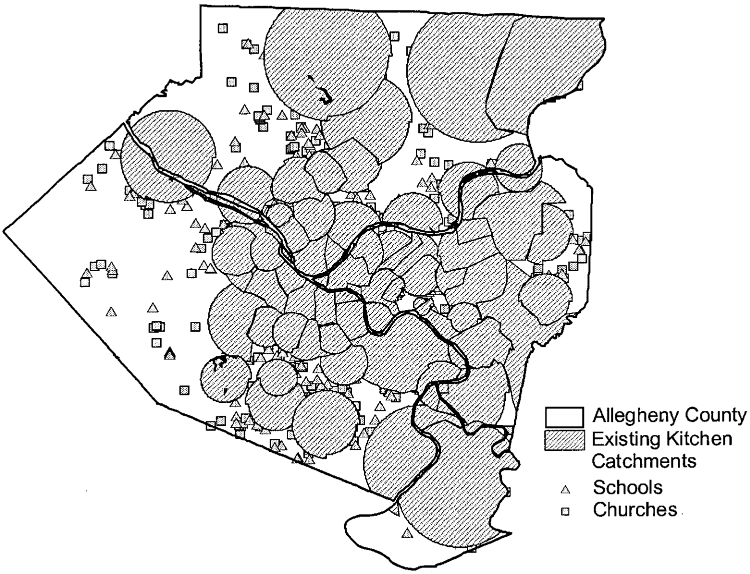

Fig. 8. Schools and churches in catchment gaps

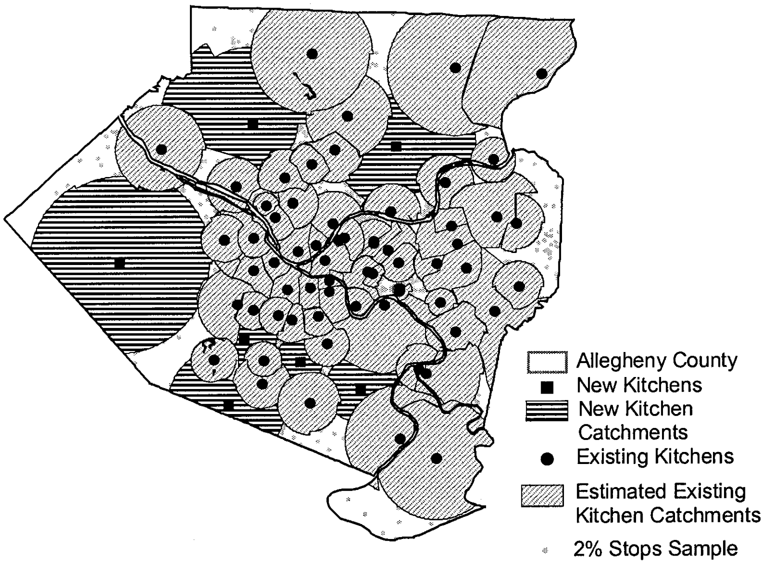

Fig. 9. The effect of adding seven new kitchens

would not be completely free to choose from among this collection, because of a lack of kitchen facilities, the inability to recruit sufficient volunteers, or many other reasons.

Nonetheless, a GIS-based analysis can show the relative value of any geographic location, and help the decision maker develop a ranking of potential sites. Our overall approach for identifying new kitchen sites consisted of conceptually breaking the coverage gaps into smaller regions, and then counting the number of un-served clients in each. Then we located a site (from Fig. 8) that was near the center of each region. For each such site, we interactively experimented with numbers of routes, policies such as "no return routes", and even extending the 45-min delivery deadline to one hour. All assumptions made about kitchen capacities and numbers of routes were consistent with the capabilities of existing kitchens.

With the help of the DSS, it quickly became evident that even a moderate number of new kitchens, strategically positioned, could have a significant impact on total coverage. Figure 9 shows that adding seven new kitchens could supply an additional 10% of the client base, yielding a total coverage of 93%. Two kitchens were assumed to supply 100 meals (which is within the scope of current kitchen capacities), while the remaining five supplied between 40 and 60 meals apiece. Each kitchen used between four and six routes.

Many of the remaining coverage gaps are in rural areas, far removed from the city of Pittsburgh (which lies at the confluence of the three rivers shown in all the county maps). The catchment areas of rural kitchens are typically considerably larger than those closer to the city, suggesting that even if they had more capacity to provide meals, they would find it difficult to deliver those meals within the 45-min limit. This poses a difficult problem for planners, for whom equity considerations loom large. Catchment gaps in more densely populated areas might best be served by expanding, where possible, the capacity of existing kitchens.

We recognize that the approach taken above is far from globally optimal, as would be the case if one of the more elaborate location-routing models of Sect. 2 were applied. Indeed, our approach is essentially a greedy one, reacting only to local opportunities. However, as mentioned earlier, for the meals-on-wheels problem there is no central authority that could implement a global solution. Additionally, the equity considerations commonly applied to this kind of problem are difficult to quantify. This suggests that a DSS approach, allowing decision makers to examine alternate strategies, is a useful one.

6 Conclusions

Every day, non-profit agencies and volunteer groups provide countless services to disadvantaged citizens of every stripe. This work is often very loosely organized, with little central planning. Charitable foundations and other funding agencies, seeking to maximize the benefit of their efforts, have traditionally found it difficult to evaluate community needs in any systematic way. This paper suggests that the use of geographical information systems, coupled with tools from operations research and statistics, present a real and valuable means to support these agencies.

While this work is focused on home delivered meals, HDM is but a single example among many for which the spatial insights provided by GIS are valuable to non-profits. As the cost of computer hardware and software decrease, and the use of the Internet increases, we believe it will be possible for even loosely-confederated non-profits to share maps, street networks, and other data, resulting in improved day-to-day performance and better long-term planning. The example in this paper considers both of these aspects, combining demographic forecasts with current operational data.

References

Alter SL (1980) *Decision support systems: Current practice and continuing challenges.* Addison-Wesley, Reading, MA

Armstrong MP, Densham PJ (1990) Database organization strategies for spatial decision support systems. *International Journal of Geographical Information Systems* 4: 3–20

Assad AA (1988) 'Modeling and implementation issues in vehicle routing.' In: Golden BL, Assad AA (eds) *Vehicle routing: Methods and studies.* North-Holland, Amsterdam, pp 7–45

Ayeni B (1997) 'The design of spatial decision support systems in urban and regional planning.' In: Timmermans H (ed) *Decision support systems in urban planning.* E & FN SPON, London

Bartholdi JJ, III, Platzman LK, Collins RL, Warden WH, III (1983) A minimal technology routing system for meals on wheels. *Interfaces* 13(3): 1–8

Begur SV, Miller DM, Weaver JR (1997) An integrated spatial DSS for scheduling and routing home-health-care nurses. *Interfaces* 27(4): 35–48

Birkin M, Clarke G, Clarke M, Wilson A (1996) *Intelligent GIS: location decisions and strategic planning.* GeoInformation International, Cambridge

Bodin L, Golden BL (1981) Classification in vehicle routing and scheduling. *Networks* 11: 97–108

Densham PJ (1991) 'Spatial decision support systems.' In: Maguire DJ, Goodchild MS, Rhind DW (eds) *Geographical Information Systems: Principles and applications.* Longman, London, pp 403–412

Densham PJ, Rushton G (1988) 'Decision support systems for locational planning.' In: Golledge RG, Timmermans H (eds) *Behavioral modelling in geography and planning.* Croom Helm, New York, pp 56–90

Giannikos I (1998) A multiobjective programming model for locating treatment sites and routing hazardous wastes. *European Journal of Operational Research* 104: 333–342

Gorr WL, Johnson MP, Roehrig SF (2000) 'Point demand forecasting.' Unpublished draft. Carnegie Mellon University, H. John Heinz School of Public Policy and Management, Pittsburgh

Gould MD, Densham PJ (1991) 'Spatial decision support systems: A bibliography.' Technical paper 91-9. National Center for Geographic Information and Analysis. Santa Barbara, Calif

Laporte G (1988) 'Location-routing problems.' In: Golden BL, Assad AA (eds) *Vehicle routing: Methods and studies*. North-Holland, Amsterdam, pp 163–197

Laporte G, Dejax PJ (1989) Dynamic location-routing problems. *Journal of the Operational Research Society* 40(5): 471–482

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

Laporte G, Norbert Y, Taillefer S (1988) Solving a family of multi-depot vehicle routing and location-routing problems. *Transportation Science* 22(3): 161–172

Lenstra JL, Rinooy Kan AHG (1981) Complexity of vehicle routing and scheduling problems. *Networks* 11: 221–227

Li C-L, Simchi-Levi D (1990) Worst-case analysis of heuristics for multidepot capacitated vehicle routing problems. *ORSA Journal of Computing* 2: 64–73

Lindbloom CE (1968) *The policy making process*. Prentice Hall, New York

List GF, Mirchandani PB (1991) An integrated network/planar multiobjective model for routing and siting for Hazardous materials and wastes. *Transportation Science* 25(2): 146–156

Lolonis P (1994) Knowledge-based spatial decision support systems for solving service provision problems. Unpublished Dissertation. University of Iowa, Iowa City, IA

Madsen OBG (1983) Methods for solving combined two level location-routing problems of realistic dimensions. *European Journal of Operational Research* 12: 295–301

Min H, Jayaraman V, Srivastava R (1998) Combined location-routing problems: A synthesis and future research directions. *European Journal of Operational Research* 108: 1–15

Or I, Pierskalla WP (1979) A transportation location-allocation model for regional blood banking. *AIIE Transactions* 11: 86–95

Perl J, Daskin MS (1985) A warehouse location-routing problem. *Transportation Research* 19B: 381–396

Renaud J, Laporte G, Boctor FF (1996) A tabu search heuristic for the multi-depot vehicle routing problem. *Computers and Operations Research* 23(3): 229–235

ReVelle C, Cohon J, Shobrys D (1991) Simultaneous siting and routing in the disposal of hazardous wastes. *Transportation Science* 25(2): 138–145

Wiegel D, Cao B (1999) Applying GIS and OR techniques to solve sears technician-dispatching and home-delivery problems. *Interfaces* 29(1): 112–130

Wong DWS, Meyer JW (1993) A spatial decision support system approach to evaluate the efficiency of a meals-on-wheels program. *Professional Geographer* 45(3): 332–341

Xin X (1999) A spatial decision support system for designing solid waste collection routes in rural counties. Unpublished Dissertation. University of Tennessee, Knoxville, Tennessee