

**Facility Location Model for Home-Delivered Services:
Application to the Meals-on-Wheels Program**

Wilpen Gorr
Michael Johnson
Stephen Roehrig

H. John Heinz III School of Public Policy and Management
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Abstract

We present a GIS-based decision support system for the non-profit sector, designed to assist strategic and tactical decision making 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 series of algorithmic tools, we provide a system for evaluating current meals on wheels facilities, and for making facility location decisions that satisfy coverage and equity requirements.

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, undereducated, poor, homeless, addicted, elderly, and so forth.

While the provision of goods and services is performed by the private sector using market principles, and is performed by the government using the political process, the non-profit sector cannot benefit from such underlying organizing principles. 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 service 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 hardest problem of the three, and one that this paper

addresses, is home delivered services. Site delivered services depends 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 major grant-making foundation in Pittsburgh, Pennsylvania. Their job, in part, was to make grants to support existing and new facilities, but they had no information on geographic gaps and overlaps in the coverage of target populations with services. Furthermore, they had no means of estimating service areas (catchments) of new facilities. This is where GIS is valuable.

Facility location for home delivered services must take into account the point distribution of demand (stops) and efficient routing of service providers to periodically deliver services. The example used in this paper is the meals on wheels program. In Allegheny County, Pennsylvania, there are 63 volunteer meals on wheels facilities delivering daily hot meals to 4,017 home-bound elders. We estimate below that approximately 80 percent of the demand is met by current facilities (i.e., there are geographic gaps in coverage), 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. The program uses existing facilities - church, school, and senior center kitchens. The major fixed costs of a new facility are 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 minute time limit for delivery of the last meal on a driver's route.

In particular, the service area of a kitchen depends on its location, meal production capacity, carrying capacity of delivery vehicles, 45 minute 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 are 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 Sections 3 and 4 to estimate service gaps and suggest new kitchen locations, and Section 6 concludes the paper.

2. Location-Routing Models in the Literature

The general planning problem this paper addresses is the location-routing problem for home-delivered meals, referred to as HDM-LRP. In this problem, the goal is to simultaneously choose "depot" (kitchen) locations that 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, i.e. catchment areas are defined, and a *multi-depot vehicle routing problem* (MDVRP), in which multiple vehicle routes originating from and returning to depots are designed. 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 3,000 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, Jayaraman and Srivasta (1998) have all presented comprehensive taxonomies of vehicle routing problems and in particular the LRP; we refer the reader to their category definitions.

Relevant characteristics of HDM-LRP include:

- *Single stage*, i.e. products originate at the depots rather than at a central plant;
- *Deterministic*, i.e. the nature of location/routing parameters such as customer demand is known and fixed with certainty;
- *Multiple facilities*;
- *Multiple vehicles*;

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

We examine some of these problem characteristics in more detail. A primary motivation for this paper is the shift in locations of populations likely to provide *and* use HDM services, from the central city to suburbs, and from inner-ring suburbs to outer-ring suburbs and rural areas. While Gorr, Johnson and Roehrig (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. Knowledge about likely trends in kitchen sizing over time would be extremely helpful for HDM planners who must negotiate with strong-willed volunteers. 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 have found (Gorr, Johnson and Roehrig 2000) that it is simply not possible to identify with confidence exact locations of all clients currently receiving HDM services in our study area, Allegheny County, Pennsylvania. Thus, we have had to design procedures to estimate customer demand. Given known shifts over time in populations likely to need HDM services, it would be correct to regard customer demands as stochastic. Again, though, while LRPs under uncertainty (REF) 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, Norbert and Taillefer 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 associated with recruiting volunteers for kitchens or with persuading potential clients to accept service from one type of provider, e.g. the county, as opposed to another, e.g. a church. 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. 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, Jayaraman and Srivasta (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, Laporte and Boctor (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 knowledge, 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; as mentioned above, this strategy is motivated in part by our inability to

incorporate monetized fixed facility costs into an explicit optimization model and in part by the need to allow users to apply equity considerations where necessary.

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. Actual observations of HDM volunteers by these authors have verified this judgement and provided another motivation for use of an interactive LRP algorithm. 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. Finally, Assad emphasizes the role of geographic information systems in locating customers and facilities, measuring travel distances and displaying model results, considerations that are central to this paper.

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 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 percent

and 9 percent 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 percent 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

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, in particular the use of GIS and interactive algorithms. However, the special nature of HDM-LRP prevents us from taking advantage of the literature associated with algorithmic, “black box” solutions to LRP, at least initially.

3. Algorithms for a DSS Approach

Since 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. Such a system is often called a *decision support system*, since the goal is to augment a decision maker's synthesizing abilities with computer tools providing tactical support.

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 have a component that operates 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 service 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 service 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-minute 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.

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 Section 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. This is done by selecting points in the vicinity of the facility 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 product, this and other procedures are coded in the Avenue scripting language.

For each facility, we assume that the number of routes r is known. For the meals on wheels application, the number of routes corresponds to the number of volunteer drivers available. 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 semi-infinite 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.

A typical catchment area, for a facility with five delivery routes, is shown in Figure 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-minute time limit. In addition, there appears to be unmet demand in the vicinity of the facility, particularly to the northeast.




Figure 1: A Typical Catchment Area

Another catchment area is shown in Figure 2, for a facility with a capacity of 150 meals and six drivers. Our delivery 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-minute 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. These estimates of catchment areas are, on the one hand, conservative, since a near-optimal TSP heuristic is used to select the delivery routes. On the other hand, the random configuration of client locations may produce unusual or time-consuming delivery patterns.

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 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, boundaries are used.




Figure 2: Another Catchment Area




Figure 3: Overlapping Catchment Areas

Figure 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 Figure 5.

This process allows the decision maker to assign territories to facilities, 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.




Figure 4: Boundaries Around a Service Location




Figure 5: Resulting Catchment Area

4. Estimation Via GIS: Allegheny County HDM

Our purpose is to provide a tool that can be used to produce a five-year plan to plug gaps in the Allegheny County meals on wheels kitchens. Our algorithm of Section 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 Figure 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, Johnson, and Roehrig (2000) provide an estimate of 2 percent 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 following distribution in Table 1. The expected number of clients per stop is 1.3.

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

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 Figure 6.

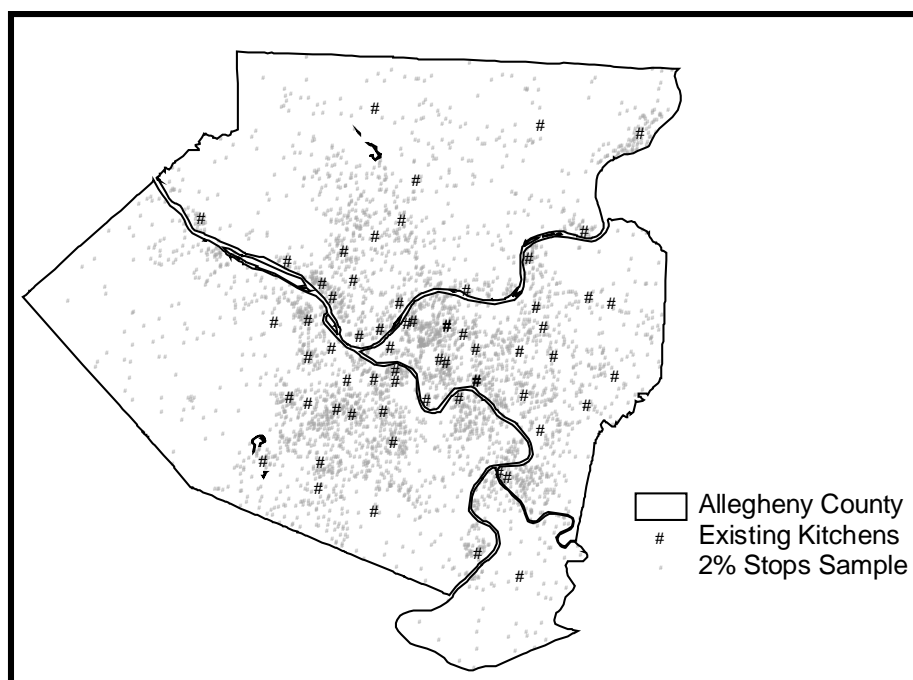


Figure 6: 2% Sample of Stops and Existing Kitchens

5. Overlaps and Analysis of 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 Section 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 often, 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 un-coverable 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.

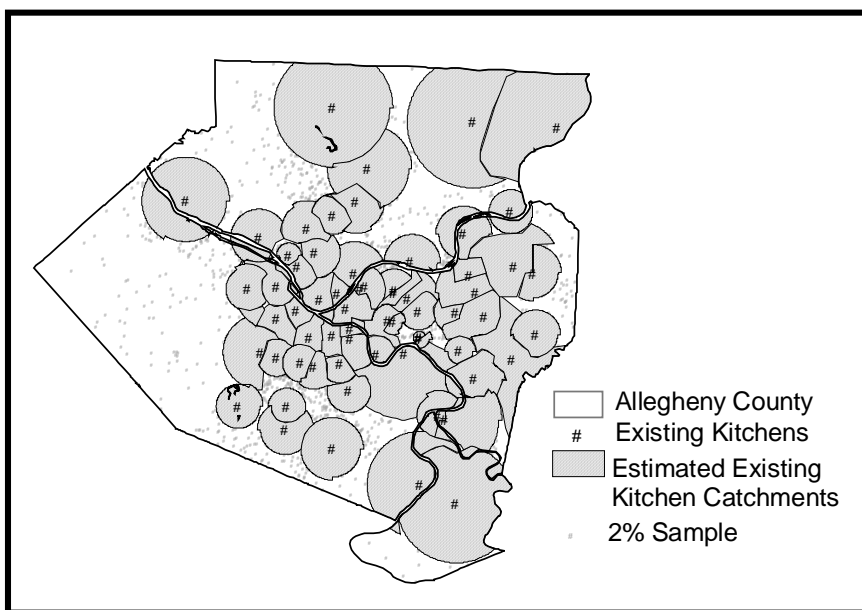


Figure 7: Non-Overlapping Catchment Areas for All Existing Kitchens

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 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.

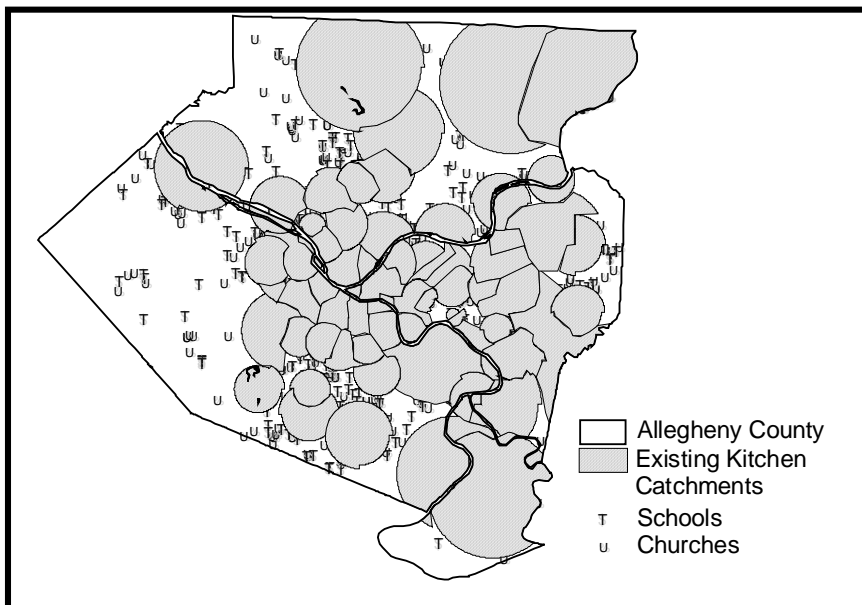


Figure 8: Schools and Churches in Catchment Gaps

Our overall approach for identifying new kitchen sites consisted of conceptually breaking the coverage gaps into smaller regions, and then counting the number of unserved clients in each. Then we located a site (from Figure 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-minute 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-minute 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.

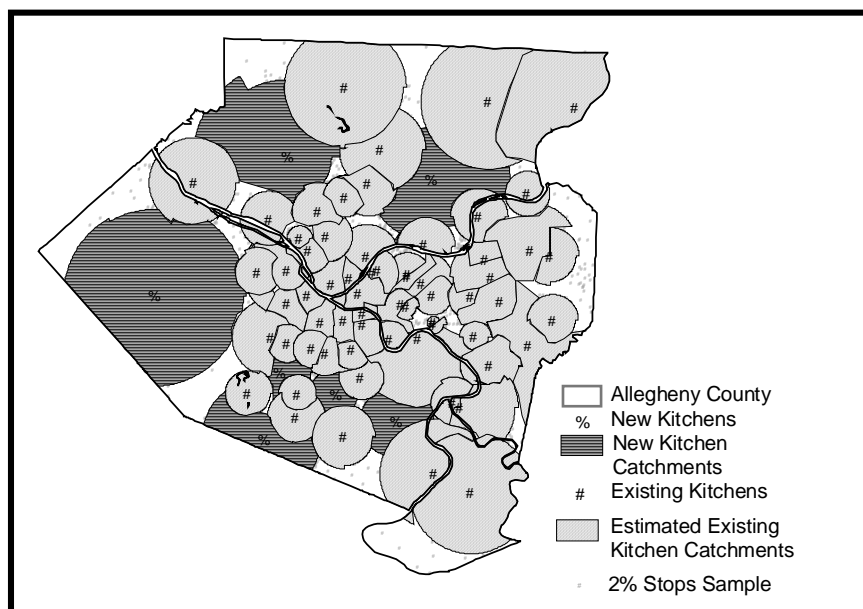


Figure 9: The Effect of Adding Seven New Kitchens

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

- Assad, A.A. 1988. "Modeling and Implementation Issues in Vehicle Routing," in B.L. Golden and A.A. Assad (Eds.), Vehicle Routing: Methods and Studies. Amsterdam: North-Holland, pp. 7 – 45.
- Bartholdi, J.J., III, Platzman, L.K., Collins, R.L. and W.H. Warden, III. 1983. A Minimal Technology Routing System for Meals on Wheels. *Interfaces* **13**(3): pp. 1 - 8.
- Bodin, L. and Golden, B.L. 1981. Classification in Vehicle Routing and Scheduling. *Networks* **11**, pp. 97 - 108.
- Giannikos, I. 1998. A Multiobjective Programming Model for Locating Treatment Sites and Routing Hazardous Wastes. *European Journal of Operational Research* **104**: 333 – 342.
- Gorr, W.L., Johnson, M.P. and S. F. Roehrig. 2000. [UNTITLED DRAFT FOR FORECASTING JOURNAL].
- Laporte, G. 1988. "Location-Routing Problems," in B.L. Golden and A.A. Assad (Eds.), Vehicle Routing: Methods and Studies. Amsterdam: North-Holland, pp. 163 – 197.
- Laporte, G. and P.J. Dejax. 1989. Dynamic Location-Routing Problems. *Journal of the Operational Research Society* **40**(5): 471 – 482.

- Laporte, G., Norbert, Y. and S. Taillefer. 1988. Solving a Family of Multi-Depot Vehicle Routing and Location-Routing Problems. *Transportation Science* **22**(3): pp. 161 - 172.
- Lenstra, J.L. and A.H.G. Rinnooy Kan. 1981. Complexity of Vehicle Routing and Scheduling Problems. *Networks* **11**: 221 – 227.
- Li, C.-L. and D. Simchi-Levi. 1990. Worst-Case Analysis of Heuristics for Multidepot Capacitated Vehicle Routing Problems. *ORSA Journal of Computing* **2**: 64 – 73.
- List, G.F. and P.B. Mirchandani. 1991. An Integrated Network/Planar Multiobjective Model for Routing and Siting for Hazardous Materials and Wastes. *Transportation Science* **25**(2): 146 – 156.
- Madsen, O.B.G. 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. and R. Srivastava. 1998. Combined Location-Routing Problems: A Synthesis and Future Research Directions. *European Journal of Operational Research* **108**: 1 – 15.

- Or, I. And W.P. Pierskalla. 1979. A Transportation Location-Allocation Model for Regional Blood Banking. *AIIE Transactions* **11**: 86 – 95.
- Perl, J. and M. S. Daskin. 1985. A Warehouse Location-Routing Problem. *Transportation Research* **19B**: 381 – 396.
- Renaud, J., Laporte, G. and F.F. Boctor. 1996. A Tabu Search Heuristic for the Multi-Depot Vehicle Routing Problem. *Computers and Operations Research* **23**(3): pp. 229 - 235.
- ReVelle, C., Cohon, J. and D. Shobrys. 1991. Simultaneous Siting and Routing in the Disposal of Hazardous Wastes. *Transportation Science* **25**(2): 138 – 145.
- Wong, D.W.S. and J.W. Meyer. 1993. A Spatial Decision Support System Approach to Evaluate the Efficiency of a Meals-on-Wheels Program. *Professional Geographer* **45**(3): 332 – 341.