Determining Emergency Medical Service Vehicle Deployment in Austin, Texas

DAVID J. EATON
Lyndon B. Johnson School of Public Affairs
University of Texas at Austin, Austin, Texas 78712

MARK S. DASKIN
Department of Civil Engineering
Northwestern University, Evanston, Illinois 60201

DENNIS SIMMONS
The City of Austin
P. O. Box 1088, Austin, Texas 78767

BILL BULLOCH
Nash, Phillips-Copus Company
6200 La Calma Drive, Austin, Texas 78752

GLEN JANSMA
Jansma Consulting, Inc.
2101 Travis Heights Boulevard
Austin, Texas 78704

In 1978 the city of Austin, Texas, began a study of its emergency medical service (EMS) system to determine what services should be delivered, by whom, via what number and types of equipment, and sited at which locations. The resulting plan called for a system quite different from that in operation since 1976: four advanced life-support and eight basic life-support vehicles were to operate from two EMS-only stations and 10 shared-use fire stations. The plan, which the city council unanimously passed in 1980, has saved $3.4 million (1984 dollars) in construction costs and $1.2 million (1984 dollars) per year in operating costs. Average response time has been reduced despite an upsurge in calls for service, and equitable service is provided to those who need it most.

Austin, Texas, and the surrounding area have been growing rapidly over the past 15 years. Austin's population increased 49 percent between 1970 and 1980, and the city estimates that the population will increase 44 percent between 1980 and 1990. This growth has generated demands for expansion and improvement in city services.

In 1974-75, dissatisfied with the private ambulance services in Travis County, the city purchased an ambulance operating contract and started operation of a municipal emergency medical service (EMS) department. It intended to provide the most advanced level of care from freestanding

0092-2102/85/1501/0096$01.25

FACILITIES/EQUIPMENT PLANNING — LOCATION
HEALTH CARE — AMBULANCE SERVICE

INTERFACES **15**: 1 January-February 1985 (pp. 96-108)

EMS stations.

In January 1978 the EMS Department initiated a joint city/University of Texas study to review the plans for EMS implementation. The study, which was conducted as a graduate policy research class at the LBJ School of Public Affairs, was to determine the number and location of vehicles to deploy so as to provide quality service to all segments of the population at a reasonable cost.

Two specific cost-containment strategies were to be considered. In Austin, only 20-25 percent of the calls for EMS service require the advanced skills of paramedics. One way to reduce costs while maintaining quality service was to design a system in which personnel with less training would be dispatched in lower-cost basic life-support (BLS) vehicles to non-life-threatening events. Highly trained paramedics in advanced life-support (ALS) vehicles would respond to life-threatening (critical) calls.

A second strategy to be considered was to house EMS personnel and vehicles in existing fire stations when possible and, when not, to build freestanding facilities only on city-owned land. This strategy would limit the feasible locations and, therefore, was likely to reduce system coverage of demand. The study was to examine the extent of this degradation in service so that decisions could be based on explicit cost/service trade-offs.

Because neighborhood groups, medical personnel, and various departments of the city government expressed interest in the study, it was conducted as a seminar open to the public, allowing the different interests to be represented and rapid consensus to be achieved on the recommendations.

Analysis Detail and Data Collection

As in many other studies of emergency service deployment, the city had to be broken down into "analysis zones." In selecting a level of detail for the analysis, the study team considered two factors. First, a sufficiently large number of analysis zones had to be used to distinguish clearly between alternative locations. On the other hand, increasing the number of zones increases the effort required to code and analyze data as well as the computational burden involved in solving the model. Second, the zone structure had to be compatible with available sources of travel-time data, call-history information, and population data.

We opted for a network of 358 zones developed by the Austin Transportation Study Office (ATSO) for traffic modeling. Fortunately, the zone structure is most detailed in the areas of greatest EMS demands, such as the central business district.

EMS response times were approximated by the interzonal vehicular travel times already estimated by the ATSO. We compared selected estimated interzonal vehicular travel times with samples of actual EMS response times and found no statistically significant differences.

Two surrogates for EMS demand were employed: a sample of calls and population data. For the 4,331 calls during December 1976 and January, June, July, and August of 1977, we obtained from EMS files the location and time of each call, the cause of the incident, and EMS response and service times. The addresses were converted into zones to insure compatability with the travel-time data.

A special 1976 census taken in Austin provided information on the total population, the number of Black, Anglo, and Spanish surnamed (Hispanic) residents as well as the number of persons over 62 in each census block. These data were regrouped into zones.

Model Selection and Data Analysis

Three classes of models were used in the study: (1) a data analysis package developed during the study, (2) computer mapping programs, and (3) a location model [Eaton et al. 1979, 1980].

A call history analysis package (CHAP) was developed to analyze the daily and weekly patterns for each type of call. Clear patterns emerged. For example, strokes most frequently occur during working hours between 8:00 am and 8:00 pm. Violence-related calls peak in the late evening and early morning hours between 8:00 pm and 4:00 am, with Friday night being an especially busy period for such calls. These patterns were interesting but were not deemed relevant to the long-term location options. We thought that tactical vehicle-relocation policies [Kolesar and Walker 1974; Berman and Rahnama 1983; Masog 1981] would better deal with these demand patterns. In addition, we examined what potential improvements in average response time might accrue by changing staffing levels over the day [Daskin and Stern 1980].

CHAP was also used to determine the total number of calls, the number of critical and noncritical calls, and the number of calls requiring transport to the hospital in each zone. These data were used as inputs to both the location model and a computer mapping package which created two- and three-dimensional displays of the information.

We plotted patterns of critical and noncritical calls and statistically compared the relative frequency of calls in different areas of the city; no statistically significant differences existed. Thus, vehicle locations designed to serve critical calls well would also serve noncritical calls well.

Population data were analyzed similarly to identify residential patterns of ethnic groups. Clear patterns emerged: Blacks live predominantly in east Austin and Hispanics in south and east Austin. These data were used in addressing community concerns and as a proxy for EMS demand, as minorities historically generate more EMS calls per capita than the population at large.

In Austin, only 20-25 percent of the calls for EMS service require the advanced skills of paramedics.

Finally, CHAP was used to compute a variety of performance measures including the distributions of response time, time spent on-scene, travel time to the hospital, as well as the number and location of interdistrict responses for each ambulance. These data were used to gain insight into the performance of the system and to try to validate the location models considered by the project.

We employed a number of criteria for selecting a location model for EMS services. First, the model had to be related to saving lives, the goal of an emergency medical service. A system's ability to reduce mortality is related to the skill levels of the staff and to the time it takes to reach emergency scenes. Changing vehicle locations can increase the fraction of demands served in a given time and reduce the average time needed to serve all locations. Although both coverage and response time objectives are important, EMS officials preferred coverage because in 1975 they had adopted a five-minute maximum response-time goal. Their intentions reflected a desire to improve upon performance regulations related to the 1973 EMS Act, which prescribes that urban EMS systems be able to cover 95 percent of all demands within 10 minutes [ReVelle et al. 1977].

Second, the model had to be inexpensive to use, because it would be necessary to run it many times to compute the implications of different options. Third, the models had to be understandable to the city staff responsible for interpreting the results. Finally, we had to be able to validate any selected model.

The relatively high incremental effort necessary to develop and validate an appropriate simulation model [Savas 1969; Berlin and Liebman 1971] made simulation less attractive than optimization. In addition, such models are descriptive rather than prescriptive.

Stochastic queueing models [Larson 1974, 1975] embedded in location-seeking heuristics [Jarvis 1975] were considered in greater detail. One such model, CALL/CZSR [Fitzsimmons 1973] was tested using the 1977 locations but it failed the validation tests. Even with a 20 percent reduction in demand, the model overestimated the mean response time by 10 percent. Nevertheless, the model was used to assess qualitatively the effects of queueing on mean response times and the fraction of calls experiencing a specified response time.

A variety of deterministic location-optimization models was considered. These include set covering, in which the objective is to cover nodes with the minimum number of vehicles [Toregas et al. 1971; Berlin and Liebman 1971; Jarvis et al. 1975; Walker 1974; Plane and Hendrick 1977]; maximal location covering, designed to maximize the number of demands covered by a fixed number of vehicles [Church and ReVelle 1974]; and the P-median formulation, which finds the locations that minimize the average response time [Hakimi 1964; Berlin et al. 1976]. The set covering model was excluded because it fails to consider differences in demand across the network. Also, by requiring total coverage it can suggest the need for an excessive number of vehicles.

Because the EMS staff preferred a coverage objective, and a working code for the maximum location covering model was available, it was adopted. Despite the simplicity of the model statement (see appendix), the formulation allowed us to analyze a variety of policy options, including changes in the number of vehicles, the response time used to define coverage, and the allowable candidate vehicle locations. Each option could be assessed in terms of its ability to serve a variety of call-based and population surrogates for demand.

Vehicle Deployment Planning

The essence of the Austin EMS vehicle deployment problem was finding a compromise among competing objectives: minimizing response times; minimizing the number of vehicles deployed; and maximizing system performance as measured by coverage of eight demand surrogates — critical, noncritical and total calls, total population, Black, Hispanic, and Anglo population, and elderly citizens. In addition, the city sought to plan for future growth and to find means to reduce costs without reducing services. Two cost-containment policies the city wanted to consider were (1) joint use of facilities by EMS and fire personnel, and (2) adoption of a two-tiered (advanced and basic life-support) response capability.

In the best of all possible worlds, the study team would have sought to develop a complete multiobjective trade-off analysis for all eight demand surrogates, a range of fleet sizes, and a variety of maximal response times. To conduct such an analysis would require roughly half a billion optimization runs, a prohibitively large number.

Our strategy was to use maximal covering to generate the best and worst coverage levels for each surrogate that would have been produced by a complete multiobjective analysis. We called the difference between these bounds the "range of controversy" because it defined the range of trade-offs among objectives. We generated this range by solving the maximal covering model for each of the eight demand surrogates for a limited number of combinations of fleet size and coverage time (Figure 1). Maximizing coverage for one demand surrogate provided the maximum possible coverage for that surrogate using the given fleet size and coverage time. Optimizing for each of the seven remaining surrogates in turn and measuring the coverage of the first surrogate by each of the seven selected sets of locations identified the range of controversy for the objective.

Real trade-offs exist. Twelve vehicles located to maximize coverage for the Black population can cover 97 percent of the Black population within five minutes. The 12 sites which best serve Anglos can reach only 60 percent of the Black population within five minutes.

Some sites were selected by several objectives simultaneously in various optimization runs. For example, in 21 maximum covering solutions designed to generate the range of controversy for selected combinations of two fleet sizes, eight demand surrogates and four response times, 146 possible sites could have been selected. The model identified only 60 sites; 26 locations were selected more than once (Figure 2). One zone in south Austin was selected in 13 runs.

A next step was to allow any interested parties to use optimization to compute how system performance would be affected by incorporating their preferences in the objective function or the constraint set. For example, one optimization run (for eight vehicles) computed that coverage of total calls would be reduced by four percent if two of the eight sites — one site requested by east Austin neighborhood groups and a second location adjacent to the University of Texas — were arbitrarily designated as EMS vehicle station sites.

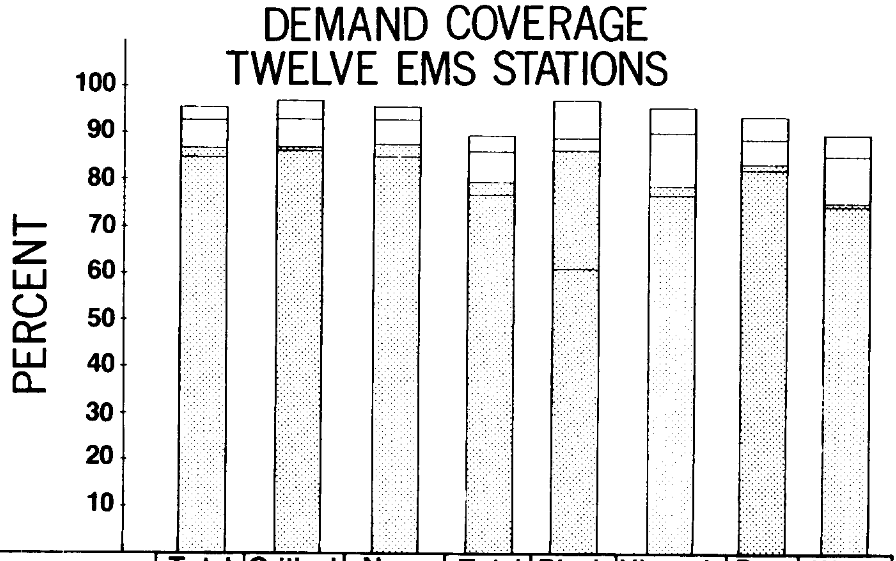

Demand		Total Calls	Critical Calls	Non-critical Calls	Total Pop.	Black Pop.	Hispanic Pop.	Pop. ≥62	White Pop.
KEY Maximum Mean City Minimum	max.	95.4	96.8	95.4	88.9	96.6	95.2	93.1	89.2
	mean	92.6	92.8	92.6	86.0	88.6	89.7	88.8	84.4
	min.	85.0	85.5	84.8	79.4	60.3	76.3	81.8	74.3
	city	86.8	86.6	86.9	76.7	85.6	78.7	83.0	74.6

Figure 1: The range of controversy for 12 EMS stations and a five-minute response time: Eight optimization runs were performed, each maximizing coverage to one of the demand surrogates. The "maximum" value indicates the largest percentage of the demand surrogate that can be covered; this value is obtained by maximizing the coverage of the indicated surrogate alone. The "minimum" value is the smallest coverage of the indicated surrogate using the locations obtained from the remaining seven optimization runs. The "mean" indicates the average coverage of the surrogate over the eight optimization runs. The "city" value is the coverage associated with the implemented city-selected EMS station sites. A considerable range of coverage exists for some demand surrogates such as Blacks. Thus, real trade-offs exist.

Computing how site preferences or aversions could degrade system performance probably contributed to developing a community consensus to support the eventual EMS plan for locating stations. The process recognized that it is often easier to achieve consensus *in action* rather than *on objectives.*

The joint-use policy was analyzed by constraining the set of candidate locations to zones either already housing a fire station or in which the city owned property. For example, constraining the run discussed above to 16 feasible sites further reduced coverage by six percent. Comparison of the constrained and unconstrained results allowed the city staff to assess the opportunity cost (as measured by coverage) of the joint-use policy.

Further city growth, anticipated in fringe areas, was captured by runs in which the coverage time was increased. These runs led to a more dispersed set of vehicle locations.

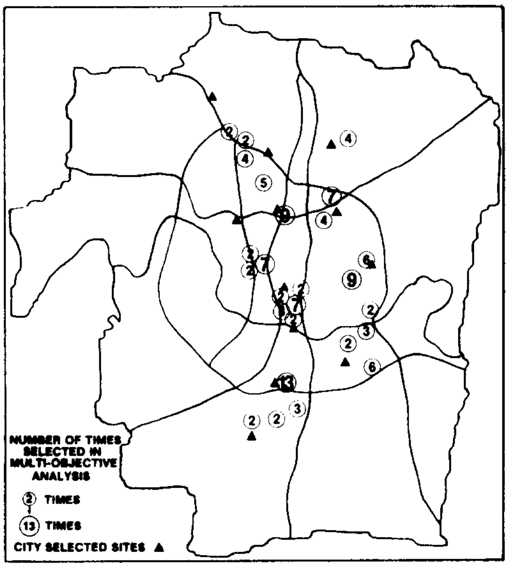

Figure 2: Locations selected more than once and the city-selected sites: In 21 runs allowing 146 possible locations, only 60 sites were selected. Twenty-six sites were selected more than once; one zone was selected in 13 runs.

After all the computer analyses, the two policy issues still remained — how many vehicles should Austin deploy and where should they be located? Furthermore, if the Austin EMS were to implement a two-tiered system and share facilities with the fire department, how many ALS and BLS vehicles should there be and how should vehicles be allocated?

To determine the total number of vehicles, we considered national EMS standards; the federal Emergency Medical Services Act of 1973 recommends that an urban EMS system reach 95 percent of all calls within 10 minutes [ReVelle et al. 1977]. A trade-off curve showing how increasing the number of vehicles would affect coverage demonstrated that 12 vehicles would cover 95 percent of all calls within five minutes and 95 percent of both critical and noncritical calls (Figure 1). In addition, 12 vehicle configurations were identified that can cover about 90 percent of each of the five population groups within five minutes.

We used maximal covering runs as one basis for selecting sites. Four additional considerations were (1) the pattern of city ownership of land, (2) the policy of joint use of facilities by city departments, (3) an EMS department goal of responding to all calls within five minutes, and (4) expectations concerning Austin's future growth. Figure 3 shows the final EMS station locations.

To select each site we weighed the optimization results and the institutional factors. For example, zone 258 was the site selected most frequently, but the city owned land in zone 257 and it built an ALS facility there. Zone 14, the site of a

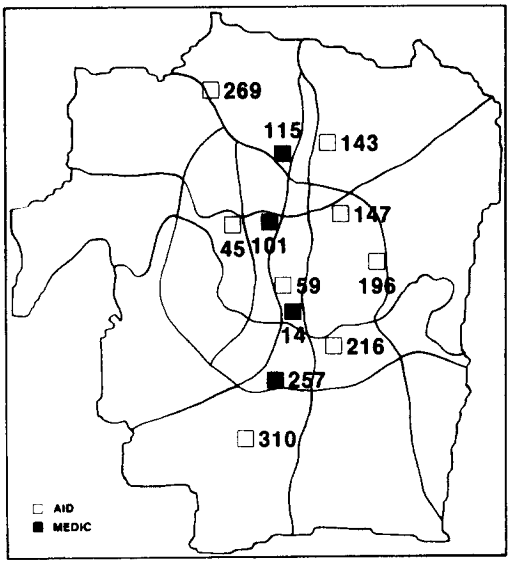

Figure 3: The sites and vehicle assignments the city selected were four advanced life-support (ALS or medic) vehicles located along a north-south line and eight basic life-support (BLS or aid) vehicles located around the city.

fire station, was chosen because it is less than one minute away from zones 7, 9, and 166, which together were selected in 12 optimization runs. Of the 12 sites selected, 10 were in zones where joint-use of fire stations was possible.

A decision on the number of advanced life-support vehicles was influenced by the call analysis, which indicated that between 20 and 25 percent of all calls involve life-threatening risks. It might be reasonable to identify three of the 12 sites for advanced life-support units. However, to reduce fatalities, the decision was made to opt for four ALS stations. Two paramedic supervisors are on duty at all times; they deliver ALS care and thus provide an additional buffer during periods of heavy ALS demand.

> . . . strokes most frequently occur during working hours . . . . violence-related calls peak in the late evening and early morning hours . . . .

ALS vehicles were sited to respond best to critical calls. BLS units were allocated to sites both to enhance the current coverage of critical and noncritical calls and to anticipate projected growth in Austin's fringe areas. The system performs well, covering 87 percent of all calls within five minutes. However, the city-selected sites perform below the maximum coverage levels attained for each of the eight demand surrogates (Figure 1). This reduction in coverage results from (1) the need to serve a variety of demand types (critical and non-critical calls as well as population groups) and not just any single demand surrogate, (2) the joint-use policy, and (3) the need to plan for future coverage. The city sites do provide more multiple coverage (more than one vehicle covers an area within five minutes) than the set of sites chosen to cover (single-cover) all calls.

The 1980 plan incorporated a "floating zone" policy to reduce the risk of response time delays further during periods of high ALS demand. When any two ALS units respond to calls, the other vehicles shift from EMS stations to prearranged "stand-by" sites; these have been selected to cover demand with a reduced number of vehicles [Masog 1981]. These policies appear to work: the EMS director cannot recall an incident when no ALS vehicle or supervisor was available for responding to an incoming ALS call.

Implementation

The city's plan for 12 vehicles deployed in a two-tiered system from jointly used facilities was unanimously approved in 1980 by a joint physician/citizen EMS advisory committee, the Austin Planning Commission, and the Austin City Council. Austin voters approved three separate bond issues to renovate and construct the facilities during a period when most municipal bond items were voted down. In each election, Austin voters approved the EMS bonds by the largest or second largest margins of any proposition on the ballot. To date, 10 of the 12 sites have been brought on line.

Benefits

The benefits of the EMS plan are

— Capital and operating cost savings,

— Reductions in average system response time,

- Improved prehospital medical care,
- Equity improvements, and
- Other hard-to-quantify benefits.

The current two-tiered, joint-use system is less expensive than an all-advanced life-support system based in 10 free-standing EMS stations. Because of joint use, the city saved $3.3 million (1984 dollars) in construction costs alone. The tiered system will save Austin $1.2 million in 1984 operating costs. Over the 23-year lifetime of the construction bonds, the projected present value of the two-tiered, joint-use system is on the order of $19.6 million less costly than a system like that in use prior to 1980.

To conduct such an analysis would require roughly half a billion optimization runs . . .

In comparing the cash outflow of the two systems over 1980-1987 (additional facilities may be added in 1987), the expected cumulative undiscounted cash savings are $10.8 million or about 24.7 percent of an all-ALS system. These savings underestimate the cost reductions from two tiers of service and joint use, as land purchase expenses are not considered. As the system expands after 1987 and if more two-tiered, joint-use facilites are implemented, the cost savings will increase.

In summary, the cost savings of the EMS system are

- Construction cost savings of $3.3 million (1984 dollars),
- Annual operating cost savings of $1.2 million (1984 dollars),
- Discounted construction and operating cost savings over the 23-year lifetime of the construction bonds of $19.6 million (1980 dollars), and
- Undiscounted cash-flow savings during the first seven years before additions of $10.8 million.

These cost savings can be compared to the expense of the research to the city of Austin of approximately $30,000.

The current system responds more rapidly to calls than did the all-ALS system operated prior to 1980. The average number of EMS calls received per day has increased by 52 percent between 1980 and 1983. Despite this increase the average system response time declined by seven percent during that period.

The two-tiered EMS service greatly improves prehospital medical care. Paramedics are assigned predominantly to calls requiring their advanced training, thereby reinforcing their skills and improving morale. Standard examinations were used to assess the medical skill of paramedics prior to and after the implementation of the tiered system. Paramedic performance was found to be significantly better after 1980 than prior to adoption of the tiered system even though current performance standards are more rigorous.

A two-tiered EMS system requires that dispatchers be trained in emergency medicine. In an all-ALS system, dispatchers send an ALS vehicle to each call. If a choice must be made between sending an ALS or BLS vehicle, each dispatcher should have the training and experience to assess the medical condition of the patient from clues provided over the telephone. Because of its two-tiered system, Austin has upgraded its dispatchers; now each person who answers incoming calls is trained to at least the level of an emergency medical technician. There have been a number of incidents in which the calm advice of a dispatcher has, in fact, saved lives.

The current EMS deployment plan provides equitable service, in that the system covers 87 percent of critical and non-critical calls within five minutes. Preferential service is provided to those segments of the population most likely to require emergency medical care. For example, within a five-minute response time the current EMS system can reach a larger fraction of the elderly than of the general population. Preferential treatment is also provided to non-Anglo minorities (Blacks and Hispanics), who have historically generated relatively large per capita EMS demands.

Several other benefits resulted from the project. First, because the study was open to all interested groups, conflicts over site selection were resolved more easily. Second, the study trained 37 graduate students in methods of analyzing public service policy options. Third, the public, open-ended process fostered confidence in management science techniques among city officials. Recently joint city/university studies have been initated on lake water quality and on fire and EMS demand modeling techniques for rapidly growing outlying areas of the city.

Conclusions

Determining sites for public service facilities is always a process of choice. By accommodating the political process — the open and unambiguous clash of objectives — systems analysts can enhance the role of operations research in resolving potential conflicts over public services. In this project, systems analysis enhanced traditional methods for making decisions. Optimization techniques were used to calculate benefits and costs and to assess opportunity costs of a number of policies, but not to select an "optimal" answer. Due to the deployment plan reported here and other management innovations, the National Association of Emergency Medical Technicians on May 12, 1984, designated the Austin EMS system as the best in the United States.

. . . it is often easier to achieve consensus *in action* rather than *on objectives.*

Acknowledgments

We acknowledge the contribution of the following people, members of the Austin city staff and students and faculty of the University of Texas at Austin: R. Banks, F. Battle, M. Berman, N. Bucek, N. Bunch, S. Carter, L. Clegg, W. Coll, T. Cowan, A. Eells, F. Fitzgerald, J. Fitzsimmons, B. Flores, N. Garrison, P. Greenberg, R. Herrington, B. Johnson, P. Kalazny, S. Kelly, L. MacFadden, R. Machemehl, M. Malnory, T. Masog, M. Morris, S. Nelson, L. Nguyen, L. Nungesser, K. Rocha, J. Rock, G. Rohlich, R. Rubio, F. Sheppard, B. Srikar, E. Stern, J. Storbeck, J. Stypka, C. Treece, A. Turnquist, T. Waters, R. Wiggans, and K. Zamrazil.

The project was funded from a number of sources including the city of Austin, the Ford Foundation, the Henry J. Kaiser Family Foundation, the Texas Department of Human Resources, the Lyndon Baines Johnson Foundation, and the Policy Research Institute of the University of Texas at Austin. Project members are grateful to J. Fitzsimmons for providing a version of his CALL/CZSR program and R. Church for an early version of his MCLP code.

APPENDIX: Use of Maximal Covering For Policy Analysis

The maximal covering location model (MCL) selects sites to maximize the total demand that can be covered by a fixed number of facilities, P, within a user-specified critical response time, S. The traditional MCL model, originally proposed by Church and ReVelle [1974], is reformulated here as a multi-objective optimization problem as follows:

$$\text{maximize} \quad \sum_k W_k \sum_i a_{ik} Y_i \qquad (1)$$

$$\text{subject to} \quad \sum_j X_j = P, \qquad (2)$$

$$\sum_{j \in N_i} X_j - Y_i \geqslant 0 \qquad \forall i, \qquad (3)$$

$$X_j, Y_i = 0, 1 \qquad \forall i, j, \qquad (4)$$

where

a_{ik} = value of the k-th demand surrogate in zone i (for example, the number of critical or noncritical calls in zone i),

W_k = weight associated with coverage of the k-th demand surrogate (weights sum to 1),

$X_j = \begin{cases} 1 \text{ if a facility is located in zone } j \\ 0 \text{ if not,} \end{cases}$

$Y_i = \begin{cases} 1 \text{ if zone } i \text{ is covered} \\ 0 \text{ if not,} \end{cases}$

N_i = the set of zones j that are within S time units of zone i (that is the zones j which can cover zone i).

The objective function (1) is to maximize the weighted sum of the covered demand surrogates. Constraint (2) states that P facilities or vehicles are to be located. Constraint (3) stipulates that demands in zone i are covered only if a vehicle is located in one of the zones that can cover zone i within S time units (that is, a vehicle is located in the set N_i). Constraints (4) are the integrality constraints.

The MCL model has modest data requirements including (1) surrogate(s) for demand, (2) an initial number of facilities to locate, (3) a zone-to-zone travel-time matrix, and (4) an initial upper limit on response time.

The formulation can explore many issues. For example, the trade-off between coverage and the number of facility sites can be explored by varying P, the number of sites to locate. Changing the coverage time S changes the set N_i of zones that can cover node i, thereby enabling the user to analyze the effects of changes in the coverage time on both total coverage and facility locations. Coverage of specific demand surrogates (such as critical calls) can be maximized by setting all weights W_k to 0.0 except the weight on the surrogate to be maximized which is set to 1. The effect of preferences for (or against) specific sites can be analyzed by setting the corresponding X_j values to 1 (or 0). The resulting demand coverage can then be compared with that which results from the unconstrained case to assess the opportunity cost of the site preferences. Finally, the joint-use policy can be analyzed by only allowing joint-use sites to enter the sets N_i of candidate facility locations.

The linear programming (LP) relaxation of the MCL model was solved using the IBM Mathematical Programming System (MPSX) package on an IBM 360/65 computer at the University of Tennessee at Knoxville. The average solution time for six problems locating one through six facilities with a coverage time of five minutes was 47 seconds. Solution times of this order of magnitude would have precluded solving hundreds of problems as was needed to explore the impacts of a large number of policy options. Instead of solving the problems optimally, we used the "greedy adding and substitution" algorithm developed by Church. A typical heuristic solution required two seconds compared to the average of 47 seconds found for the LP relaxation runs. In the six problems discussed above, the heuristic objective function averaged 98.7 percent of the LP value, with the worst case heuristic solution being 3.3 percent below the LP value.

References

Berlin, G. N.; ReVelle, C.; and Elzinga, D. J. 1976, "Determining ambulance-hospital locations for on-scene and hospital services" *Environment and Planning A*, Vol. 8, No. 5, pp. 553-561.

Berlin, G. N. and Liebman, J. C. 1971, "Mathematical analysis of emergency ambulance location," *Socio-Economic Planning Sciences*, Vol. 8, No. 6, pp. 323-328.

Berman, O. and Rahnama, M. R. 1983, "Optimal path of a single service unit on a network to a 'nonemergency' destination," *Transportation Science*, Vol. 17, No. 2, pp. 218-232.

Church, R. and ReVelle, C. 1974, "The maximal covering location problem," *Papers of the Regional Science Association*, Vol. 32, pp. 101-118.

Daskin, M. S. and Stern, E. H. 1980, "A framework for EMS crew scheduling," in *Modeling and Simulation*, proceedings of the Eleventh Annual Pittsburgh Conference, Vol. 11, Part 4, pp. 1443-1448.

Eaton, D. J. et al. 1979, "Location techniques for emergency medical service vehicles," Vol. I-IV, Policy Research Report 34, Lyndon B. Johnson School of Public Affairs, University of Texas at Austin, Austin, Texas.

Eaton, D. J. et al. 1980, "Analysis of emergency medical services in Austin, Texas," Vol. I, II, Policy Research Report 41, Lyndon B. Johnson School of Public Affairs, University of Texas at Austin, Austin, Texas.

Fitzsimmons, J. A. 1973, "A methodology for emergency ambulance deployment," *Management Science*, Vol. 19, No. 6, pp. 627-636.

Hakimi, S. L. 1964, "Optimum locations of switching centers and the absolute centers and medians of a graph," *Operations Research*, Vol. 12, No. 3, pp. 450-459.

Jarvis, J. P.; Stevenson, K. A.; and Willemain, T. R. 1975, "A simple procedure for the allocation of ambulances in semi-rural areas," Technical Report 13-75, Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Jarvis, J. P. 1975, "Optimization in stochastic service systems with distinguishable servers," Technical Report 19-75, Operations Research Center, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Kolesar, P. and Walker, W. E. 1974, "An algorithm for the dynamic relocation of fire companies," *Operations Research*, Vol. 22, No. 2, pp. 249-274.

Larson, R. C. 1975, "Approximating the performance of urban emergency service systems," *Operations Research*, Vol. 23, No. 5, pp. 845-868.

Larson, R. C. 1974, "A hypercube queueing model for facility location and redistricting in urban emergency services," *Computers and Operations Research*, Vol. 1, No. 1, pp. 67-95.

Masog, T. G. 1981, "Modeling techniques for a flexible emergency aid location system," MSc Thesis, LBJ School of Public Affairs, University of Texas at Austin, Austin, Texas.

Plane, D. R. and Hendrick, T. E. 1977, "Mathematical programming and the location of fire companies for the Denver Fire Department," *Operations Research*, Vol. 25, No. 4, pp. 563-578.

ReVelle, C. S.; Bigman, D.; Schilling, D.; Cohon, J.; and Church, R. 1977, "Facility location: A review of context-free and EMS models," *Health Services Research*, Vol. 12, pp. 129-146.

Savas, E. S. 1969, "Simulation and cost-effectiveness analysis of New York's emergency ambulance service," *Management Science*, Vol. 15, No. 12, pp. 608-627.

Toregas, C.; Swain, R.; ReVelle, C.; and Bergmann, L. 1971, "The location of emergency service facilities," *Operations Research*, Vol. 19, No. 4, pp. 1363-1373.

Walker, W. E. 1974 "Using the set covering problem to assign fire companies to fire stations," *Operations Research*, Vol. 22, No. 2, pp. 275-277.

A letter from Ron Mullen, Mayor of the city of Austin, states: "I . . . recall in 1978 how the City Council split two-to-five between members who doubted the study's necessity or desirability and those who were willing to "gamble" on the use of exotic operations research techniques for dealing with Austin's planning problems.

Today, Austin's City Council sees the value of operations research as an adjunct to common sense. Austin is committed to making increasing use of management science techniques in coping with our unprecedented growth and development. For example, in April I presided when a unanimous Council approved two City/University studies using systems techniques — one to improve the quality of our Highland Lakes and a second to plan emergency medical and fire services in suburban areas into which Austin is expanding.

The City/LBJ study has much to do with this shift in perspective. The study's consequences — whether measured in terms of reduced costs, improved medical care, or enhanced public accessibility to services — are real and recurring benefits to the quality of life in Austin."

Roving Reporter: A two-page color ad in *Forbes* (November 5, 1984, pp. 134-135) describes briefly the Norfolk Southern computer dispatching system that is described more fully in *Interfaces* (Vol. 13, No. 6, pp. 24-36) in the CPMS/TIMS finalist paper, "Computer aided train dispatching: Decision support through optimization" by Richard L. Sauder and William M. Westerman.