Models For Fire Station Location: A Review and Improved Distance Estimation Method Tested For Winnipeg

by

Lenore Sigurdson Kersey

A thesis
presented to the University of Manitoba
in fulfillment of the
thesis requirement for the degree of
M.Sc.
in
Actuarial and Management Sciences

Winnipeg, Manitoba

ISBN 0-315-48113-7

MODELS FOR FIRE STATION LOCATION:
A REVIEW AND IMPROVED DISTANCE ESTIMATION
METHOD TESTED FOR WINNIPEG

BY

LENORE SIGURDSON KERSEY

A thesis submitted to the Faculty of Graduate Studies of the University of Manitoba in partial fulfillment of the requirements of the degree of

MASTER OF SCIENCE

I hereby declare that I am the sole author of this thesis.

I authorize the University of Manitoba to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Lenore Sigurdson Kersey

I further authorize the University of Manitoba to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Lenore Sigurdson Kersey

The University of Manitoba requires the signatures of all persons using or photocopying this thesis. Please sign below, and give address and date.

ABSTRACT

Over the last two decades, researchers have developed several models which can be of help in determining how to locate and allocate fire fighting units in such a way as to best meet emergency service objectives. This thesis begins by reviewing these models.

The fire station location models measure performance in terms of response times to emergency incidents. The largest and most variable component of response time is the travel time from source to destination points. Various methods have been suggested and used for estimating travel times, but they tend to either be insufficiently accurate or require very large amounts of data. In this study, some simple methods for estimating travel time which have been used in models are tested for Winnipeg. A new algorithm is developed and tested which provides more accurate travel times for a city which has some major barriers to travel, while at the same time having small data requirements.

The review of available models revealed that most of the models developed for fire station location assumed that units are always available when an incident arises. In order to test the validity of this assumption for Winnipeg, incident data was studied to determine utilization rates for fire service units and the distribution of inter-arrival and service times for fire incidents. Finally, a model is suggested which would be helpful for fire service planning in Winnipeg, and the procedure for implementation is outlined.

ACKNOWLEDGEMENTS

This thesis was made possible by the generous contributions of several individuals. My two advisors provided invaluable guidance and support. Dr. Fraser Dunford gave me the opportunity to work on this project. He knew how to produce clarity from confusion and taught me a great deal about research and writing a thesis. His influence was felt throughout the project, although he left the university part of the way through.

I am grateful to Dr. Attahiru Alfa, who took over as my advisor, for giving so freely of his time and wisdom. He made me believe that I could finish, and saw that I did.

Ed Burgener of the Winnipeg Fire Department kept me on track throughout, as well as providing data, reference material and inspiration.

Financial assistance for this project was provided by the Transport Institute of the University of Manitoba.

Finally, special recognition must be given to Robin for his support and patience, and to Sean and Kristin for cheering me up when the programs wouldn't work.

CONTENTS

LIST OF TABLES

LIST OF ILLUSTRATIONS

Chapter 1

INTRODUCTION

Providers of emergency service for a city have the responsibility of locating and deploying their emergency units in such a way as to provide the best possible service. In the case of a fire department, achieving optimum service means minimizing human injury and property loss by efficient allocation of limited resources.

Considerable effort has been made by researchers to develop models that can help in deciding how best to locate units. Since a way has not been found to directly predict life and property loss resulting from a certain allocation of fire units, response time to incidents has been generally accepted as a measure of performance. Researchers have attempted to find models which will suggest how to allocate available units to best meet response time objectives. Chapter 2 reviews the research that has been carried out with the objective of improving emergency service deployment.

A review of the literature on location modelling for emergency services results in the conclusion that, before a methodology can be chosen to help in locating emergency units in a particular city, the following questions must be answered:

1) What travel time estimation method can be used to provide accurate estimates of travel time from supply points to demand points?
2) Is it necessary to take into account the possibility of a unit being busy when needed?
3) What is the criterion to be used for measuring service provided? This criterion is usually stated in terms of desired response times, but can vary from minimizing average response time to providing equal response times to all areas of the city.

The third question can only be answered by policy-makers, who must weigh the importance of various, often conflicting, criteria, such as providing an efficient service as opposed to providing equal coverage to all. A great deal of insight into the other two questions can, however, be obtained for a particular city by studying the available information about demand for service and travel characteristics. This paper attempts to answer the first two questions for the City of Winnipeg. Chapters 3 to 5 discuss several methods that have been used in models to estimate travel times, and review data for Winnipeg to find an appropriate method for that city. A method is suggested that improves over those previously used. Chapter 6 looks at incident data in order to provide an answer to Question 2 for Winnipeg. In Chapter 7, findings are summarized, other implementation issues are discussed, and an appropriate model for Winnipeg is suggested.

Chapter 2

REVIEW OF LITERATURE

Researchers have advanced many models to try to shed light on the problem of how to locate emergency services in such a way as to minimize response times. Mirchandani and Reilly (1987) classified these models into two types: static models, which assume that all units are always available for dispatch to an emergency incident, and dynamic models, which allow for the possibility that some units will be unavailable if they are busy at other fires. These two types of models will be discussed separately below.

Models can also be classified (Mirchandani and Reilly) as evaluation and optimization models. Evaluation models give performance results for a configuration of company locations which the user suggests. Optimization models produce a configuration which is optimal, based on certain specified criteria. These terms will be used in the discussion of the models.

In general, the following procedure is followed in modelling an emergency service system involving mobile server units:

1. The area being studied is divided into small homogeneous sub-areas or zones. A central point

is chosen from which traffic will be assumed to originate or end. This point may be the geographical centre of the zone or the centre of gravity in terms of the number of demands for service.

2. A method of finding travel times between zones is determined. The travel times may either be calculated by the model or fed into the model.
3. The model then either evaluates a configuration of emergency service locations that is fed in, in the case of an evaluation model, or determines the optimum locations for units, in the case of an optimization model. In both cases the yardstick that is used is response time to emergency incidents.

STATIC MODELS

One of the first studies which specifically addressed the problem of locating fire services was carried out by Hogg (1968). She used a p-median type of model. This model locates "p" fire-fighting units in such a way as to minimize the total travel time to all fires. The study assumed that the rate of fire incidence was known or could be estimated from population densities. Travel times were estimated from run data and from the results of an experimental set of journeys, and were fed into the model. The total travel time was calculated as the sum for all zones of the travel time to the nearest station times the number of incidents in the time period being studied. The p-median model is formulated as:

$$\text{Min } Z = \sum_{j=1}^{n} \sum_{i=1}^{m} f_i t_{ij} x_{ij}$$

subject to

$$\sum_{j=1}^{n} y_j = p$$

$$\sum_{j=1}^{n} x_{ij} = 1 \qquad i=1,2,\ldots,m$$

$$y_j \geq x_{ij} \qquad i=1,2,\ldots,m; j=1,2,\ldots,n$$

$$y_j, x_{ij} \in \{1,0\} \quad i=1,2,\ldots,m; j=1,2,\ldots,n$$

where

y_j = 1 if a facility is located at site j, 0 otherwise,

x_{ij} = 1 if a facility at site j serves zone i, 0 otherwise,

f_i is the number of incidents in zone i for the time period,

t_{ij} is the travel time from zone i to site j.

The choice of minimizing total travel time is good in that it avoids locating too many stations in areas with very low demand. It can, however, result in poor coverage for the low-demand areas. The method used assumes a linear relationship between response time and fire damage, but Hogg suggested that a better knowledge of the actual relationship between these factors would improve the usefulness of the model.

In the late 1960's, a major study was undertaken by the Rand Institute for the U.S. Department of Housing and Urban Development and the New York City Fire Department. This project (see Ignall et al. 1975), which took about eight years to complete, was geared toward improving the delivery of New York City Fire Department services in the face of skyrocketing demand. Results were intended to be applicable to other large metropolitan areas as well.

At the time of the study, New York City had 375 fire companies and the most fires of any city in the world. The study looked at trying to improve effectiveness through changes in three areas: allocation of companies, dispatch policy and relocation. Relocation refers to temporarily moving companies to fill gaps caused by a busy period in a certain area.

The number of fires in New York was increasing constantly and it was found that the traditional solution of adding more companies was not keeping up with providing the desired service. The dispatch policy was to send the closest units to a fire and to send three pumpers and two ladders, if available, but at least one pumper and one ladder. When the situation was analyzed, it was found that since companies were so busy (at the busiest times it was not uncommon for half the companies in the City to be responding to incidents), in most cases three pumpers and two ladders were not available. Adding extra units merely filled out the number of units being dispatched to incidents without making more units available to wait for new incidents.

A model was developed (see Swersey 1982, Ignall 1982) to aid dispatchers in implementing an adaptive response to incidents based on the following three factors:

1. The probability of an alarm being serious;
2. The expected alarm rate in the area surrounding the alarm; and
3. The number of units available in the area surrounding the alarm.

By sending fewer pumpers and ladders out to some of the incidents, it was possible to increase the number of units available for future incidents.

In order to assess how many companies were needed and where they should be located, Rand developed a "square root model" (Kolesar and Blum 1973) which could predict average response times in the regions of a city by knowing the regional alarm rate, average service time, area and the number of companies in the region. The expected average travel time, t, is calculated simply as

$$t = c(A/(n-b))^{a}$$

where A is the area of the region, n is the number of companies allocated to the region, b is the average number of companies busy in the region, and c and a are empirical constants that depend on the street configuration. This model assisted in determining whether any regions were not adequately covered and whether any could have companies removed without reducing coverage unacceptably. It assumed that units did not travel outside of their regions, which might result in inaccuracies for some cities.

A discussion follows of two other models which were produced by the Rand study: the Parametric Model for the Allocation of Fire Companies and the Firehouse Site Evaluation Model. Two other models developed by Rand at that time, the Hypercube Queuing Model and the Simulation Model of Fire Department Operations, will be discussed in the section on dynamic models.

The Parametric Model for the Allocation of Fire Companies (Rider 1975) was developed as an aid in determining the number of fire companies needed in different parts of a city. It recognized that there may be two conflicting objectives: reducing average travel time, which would suggest locating companies in the areas of greatest alarm activity, and providing equal service to all parts of a city, which would imply spreading fire companies out evenly throughout the city.

The Parametric Allocation Model provides an explicit tradeoff parameter. It uses travel time as a measure of system performance and generates allocations satisfying criteria ranging from the minimization of city-wide travel times to the equalization of average regional travel times as the tradeoff parameter varies.

The measure of travel time used is average travel time in a region (as calculated by the "square root model" mentioned above), where a region is an area of the city which is relatively homogeneous with respect to fire hazards, potential firefighting problems and alarm incidence. This model is intended to help with the problem of how many fire companies to locate in each region of a city, but does not address the question of where to locate companies within the region. It assumes that fires in a region will be serviced by a company located in that region.

The Firehouse Site Evaluation Model (Dormont, Hausner and Walker 1975) was developed to assist in deciding how many fire companies should be on duty in a city and where they should be located. It provides a way to estimate fire protection levels, measured by response time, that would result from any given arrangement of fire companies. By comparing the fire protection levels resulting from various arrangements, a fire department can make rational decisions about the location of its fire companies. It is more detailed than the Parametric Allocation Model in that it evaluates exact locations of stations and finds the response time for each demand point in a city.

The model does not by itself generate alternative firehouse configurations. However, the information provided about the travel-time and workload characteristics of proposed configurations will suggest ways of changing the arrangement to improve performance.

Travel distance can be estimated in the Firehouse Site Evaluation Model as either the right-angle distance or the straight-line distance (or some fixed multiple of straight-line distance.) No consideration is given to barriers to travel. Travel distance is converted to travel time using algorithms developed by Kolesar and Walker (1973) from their study of the travel characteristics of New York City fire companies (see Chapter 5).

Toregas et al. (1971) viewed the location of emergency services as a set covering problem. The maximum time or distance that separates a user from his closest service is viewed as the crucial parameter. An upper limit is placed on the response time or distance to any user, and consideration is then given to determining the minimum-cost spatial arrangement of service facilities that adequately serves the entire user region. If costs are identical for all possible facility locations, then an equivalent problem is to minimize the total number of service facilities required to meet the response time or distance standards for each of the users. The solution to this problem will indicate both the number and location of the facilities that provide the desired service. It is assumed that the user demands can be represented as occurring at a finite set of points and that the potential locations for service facilities are also a finite set of points. It is also assumed that the minimum distance or minimum response time between any user-node and service-facility pair is known. The formulation of this problem is stated succinctly in Merchandani and Reilly (1987) as:

$$\text{Min } Z = \sum_{j=1}^{n} Y_i$$

Subject to

$$\sum_{j=1}^{n} a_{ij} y_i \geq 1 \qquad i=1,2,\ldots,m$$

$$y_j \in \{0,1\} \qquad j=1,2,\ldots,n$$

where

m is the number of zones,

n is the number of available sites,

a_{ij} = 1 if zone i can be served by a unit at zone j without violating constraints, 0 otherwise.

A drawback of this approach is that it ignores differences in the demand levels at various points, and since it also does not take into account the possibility of units being busy when needed, it will tend to locate too few units in high demand areas. This problem can be alleviated to some extent by requiring shorter response times for high-demand areas. This study was a precursor to the "Fire Station Location Package" (Public Technology Inc. 1974), a model developed under contract with the U.S. Department of Housing and Urban Development to evaluate firehouse locations.

The Fire Station Location Package estimates travel times using a street network. The city is broken down into "fire demand zones." A street map of the region being studied is converted into a computer-readable network description in which street intersections are represented as nodes, and streets are the connecting arcs between the nodes. An estimate is made of the average speed at which a fire company would travel along each arc. The speed of travel and the length of the arc determine the average time for a fire company to traverse it. The travel time from any firehouse to any fire demand zone is then estimated by finding the set of arcs that form the shortest time path. The model therefore explicitly accounts for barriers.

This model can be used either in "descriptive" or "optimization" mode. In its descriptive mode, reports are produced providing such information as the workload for each fire company and the covered and uncovered zones in the city. A covered zone is defined as one which can be responded to by at least one station within the specified desired response time. In its "prescriptive" or optimization mode, the Model will choose from a large set of potential firehouse locations the smallest subset that will provide certain required travel times to a set of points in the city. Public Technology Inc. reported in 1977 that 52 cities and counties had used the Fire Station Location Package (Chaiken 1978).

Hendrick et al. (1974) borrowed from Public Technology Inc. locator model concepts and data base requirements in their study of fire department operations in Denver, Colorado. One aspect of this study was the employment of various location methodologies to determine whether the level of service could be maintained while reducing the number of companies. Both evaluation and optimization methods were used. The optimizing method was formulated as a set-covering model. The objective was to determine the minimum number of stations out of 120 possible locations which would satisfy required response times. Demand areas were coded with various degrees of hazard, with associated required response times. One of the features of the Denver project was the development of new concepts to generate cost-effective station configurations which place special emphasis upon the use of existing fire stations.

The question of barriers was examined as part of a response time experiment which was carried out in order to verify the use of a right-angle distance calculation combined with a formula for travel time, none of which took into account barriers. About 1600 actual fire vehicle responses were timed and these were matched with the calculated response times. The team investigated several unusually long response times which had been identified to determine whether it had been necessary for the vehicle to cross a barrier. In only two cases from the nearly 1600 observations was it obvious that a barrier had substantially lengthened the response distance, and hence the response time. They concluded that in Denver a rectilinear calculation would be adequate without adjusting for barriers.

In order to overcome the problem of the set-covering model not taking into account differing demand rates between regions, Church and ReVelle (1974) formulated the maximum covering location model to maximize the total number of covered demands. They located a fixed number of emergency vehicles in such a way that the maximum number of demands for service were covered.

Schilling et al. (1980) also used a set-covering approach in a study conducted with the Baltimore Fire

Department. The model which was developed located a certain number of facilities in such a way that the largest number of people would have a facility within the maximum allowable service distance (or time). This model could be adjusted to take into account other criteria for allocating stations such as property value, number of fires and land area. Other models were also developed as part of this study. The "capital improvement" study assumed that K stations were to be relocated and then evaluated which K stations to move and to what locations. The "reallocation model" determined how the existing companies should be allocated to the new configuration of stations to ensure that each demand point has a pumper and a ladder within its response distance standard.

Mirchandani and Reilly (1987) suggested that using travel time as a proxy measure of the cost of a certain configuration of station locations may produce less than optimum results because this method assumes that there is a linear relationship between cost and travel time. They suggested a model which can take into account a nonlinear relationship between cost and travel time by incorporating utility functions for various response times of both first and second-due units. The model cannot be solved exactly, but a solution can be obtained through heuristic methods.

The model was applied to the Albany Fire Department, incorporating separate utility functions for low-risk, property-risk and high-risk fires and for first- and second-due pumpers (engines) and ladders. Travel distance was determined by dividing the city, with a population of 102,000, into thirty-eight zones and determining travel distance between each pair of zones as the rectilinear distance measured on a map. Travel time was obtained from distance using Kolesar and Walker's (1973) travel-time/distance model.

Several researchers have extended models to incorporate probabilistic travel times (see, for example, Mirchandani and Odoni 1979; Chelst and Jarvis 1979; Daskin 1987). They show that different optimum locations may be found if the distribution of travel times is considered, rather than using the mean travel time.

DYNAMIC MODELS

Models for fire station location that take into account the possibility of units being busy when needed are not as prevalent as those that do not. The reason for this is that fire companies in most cities are busy only a small percentage of the time (Dormont, Hausner and Walker 1975). Some of the models discussed in this section were developed for ambulance location; however aspects of them may be found useful for fire station location if the utilization rate is found to be high enough that the units cannot be assumed to be always available.

In some cases simulation models have been used to verify the findings from static models by introducing the possibility of the closest unit being unavailable when a demand occurs. Uyeno et al. (1981) created an ambulance location system for the British Columbia Provincial Ambulance Service. They found rough locations with a p-median model, which minimized average response time from p ambulance bases, and fine-tuned the results with a simulation model. This model has been applied to ambulance location in Vancouver, Victoria and Edmonton.

To determine travel time, the region is broken down into sub-regions with calls assumed to take place at the central point. A most-likely route is plotted between each pair of adjacent sub-nodes and the distances converted to travel time using average ambulance speeds over various classes of roads. The travel time between any pair of non-adjacent nodes is found by applying a shortest-route algorithm, using the travel times between adjacent nodes.

A simulation model was also used by Rand (Carter 1974) as part of their study of the New York City Fire Department. It was designed to examine the effect of modifications in any of the number of companies on duty, the location of fire stations or the rules used to dispatch and deploy the available companies. Travel distance is calculated based on a rectilinear or Euclidean distance obtained from the coordinates of the fire station and incident locations.

Travel time is calculated from distance using the algorithm developed by Kolesar and Walker (1973). Chaiken (1978) reports that although this model was tested and found useful in New York City and Denver, it had not at the time of his report been used by other cities. He suggests that the reason for this is that earlier applications validated the results of the Parametric Model and the Firehouse Site Evaluation Model which are less costly to implement.

Another model which resulted from the Rand study was the "Hypercube Queueing Model" (Larson 1974). It was intended for use by police and ambulance agencies for design and evaluation of fixed sites for their units and/or response areas for the units. It assumes that only one unit is dispatched to each incident. Larson suggests that it is suitable for use by agencies which often have ten percent or more of their units busy at one time. The standard estimate of travel distance is the right-angle distance. This assumption can be overridden and a matrix of travel times substituted. Alternatively, a few selected travel times can be put in to override the distance calculation for certain source-destination combinations.

The Hypercube model is an evaluative model which gives values of certain performance measures (such as workloads of units and travel times to emergency incidents) for various arrangements of patrol areas. Larson (1975) also developed an approximate procedure for computing selected performance