MATHEMATICAL PROGRAMMING APPLICATIONS IN THE ANALYSIS OF THE DEPLOYMENT AND UTILIZATION OF FIRE-FIGHTING RESOURCES

Peter Kolesar

March 1973

P-4988

ABSTRACT

This paper gives a brief review of some applications of mathematical programming in the analysis of the deployment of fire engines. We discuss a semi-Markov decision process (linear programming) formulation of the problem of deciding how many fire engines to dispatch to a new alarm, a staged integer programming formulation of the problem of relocating fire engines to rebalance city-wide protection when extremely large fires deplete part of the city, and simple non-linear integer programming formulations of strategic resource allocation problems. All analyses have resulted in policy recommendations that have influenced current operations of the New York City Fire Department.

MATHEMATICAL PROGRAMMING APPLICATIONS IN THE ANALYSIS OF THE DEPLOYMENT AND UTILIZATION OF FIRE-FIGHTING RESOURCES

Peter Kolesar
The New York City-Rand Institute
New York, New York

I. INTRODUCTION

Since 1968 The New York City-Rand Institute has been engaged in a joint research program with the New York City Fire Department. Under this research program, we have examined problems relating to fire communications, water delivery systems, fire insurance, and fire prevention in addition to looking at some traditional operations research type problems in the deployment of fire-fighting resources. A general discussion of the research program can be found in [1], and a somewhat detailed presentation of much of the OR work is given in [4]. Here, I will discuss our uses of mathematical programming in the analysis of such questions as:

- How many fire engines should be dispatched to a new alarm?
- Which fire engines should be temporarily relocated when a large fire depletes one part of the city of its fire protection?
- How many fire engines should be permanently assigned to each region of the city, and how should this number vary by time of day?

First, I should say a few words about the motivation of the New York City Fire Department in sponsoring and participating in this research. At first glance, fire would not seem to be one of New York City's most pressing problems. But, in the last ten years, while the number of firemen and fire engines has stayed essentially constant, the incidence of fires in structures has doubled, false alarms have increased five times, and the cost of running the Fire Department has more than doubled. That cost is very high--the budget of the New York City Fire Department is well over $250 million, exceeding the entire budget of the City of Boston. With such a scale of operations, even small percentage improvements can result in very significant dollar savings. Further, the finances of the City are so greatly strained that there is great interest in better utilization of the Department's resources.

Other characteristics make analysis of the deployment of fire-fighting resources a fertile research field. Many of the problems encountered involve physical processes (fire incidence, travel times, communications, etc.) which are quite amenable to scientific study. While the basic phenomena are stable enough to be studied successfully, the growth in alarm rates has been so rapid that the existing system has been hard pressed. For example, at the very high alarm rates now current, the old manual system of dispatching is strained to the breaking point, and a key element in our research program has been the design of a computer-based dispatching and information system. Two of the problems discussed below were investigated with the view toward developing computerized algorithms for assisting in making tactical deployment decisions. The traditional method of making these decisions was to reference all fire alarms to the nearest fire alarm box. For each alarm box (about 15,000 in New York City) there is an index card containing the dispatching rules. These advance plans designate which fire engines respond to the initial alarm, which additional units respond if the fire escalates beyond the capabilities of the companies initially dispatched (a second alarm), and which engines relocate on higher alarms in order to balance protection. These plans presume that all fire-fighting units are available for their assignments. Ten years ago, that presumption was largely true, today it is not and the plans break down since the designated units are often not available. In addition, the plans contained on these alarm assignment cards do not take into account other information such as how the alarm is reported, the time of day, or the season. The proposed computerized system--which is still several years from implementation--will consider the entire system status, that is, all fires in progress, and all fire engine location and availability information, as well as other factors.

In addition to the formulation and solution of mathematical programming problems, many other kinds of analysis have been necessary in the design of this system, including probabilistic modelling, structuring of information flows, and statistical data analysis. In viewing this research effort in perspective, I would say that mathematical programming has played an important, although not a dominant, role.

II. HOW MANY FIRE ENGINES TO DISPATCH

When an alarm is received, it is not known whether it signals a fire at all, let alone whether it signals a serious fire. If such knowledge were available, the dispatcher could match the number of engines sent to the needs at the scene. Acting in the absence of perfect information, we might send too few units with resulting losses of life and property at the fire in question, or we might send more units than are needed, and, if a subsequent alarm is received for a serious fire while these units are needlessly occupied, other losses could be incurred. We view the decision of how many units to dispatch to a particular alarm (in hand) as being dependent on the number and location of available units, the estimated seriousness of the alarm, and the probability that another and more serious alarm will be received in the near future. Arthur Swersey has formulated a finite state semi-Markov decision problem having, as the objective, the minimization of the long-run average of a utility function of the response time to fire alarms. His model explicitly considered the following factors:

(1) The potential seriousness of the alarm. The higher this value, the more units we would tend to dispatch. Such information is often directly available when the alarm is reported by telephone. If the alarm is reported by telegraph box, the probability that it signals a serious fire can be estimated statistically from the history of the alarm box and the immediately surrounding area.

(2) The alarm rate in a region served by several fire engines and surrounding the location of the alarm. The greater the alarm rate, the greater the chance that units dispatched now will be needed in the near future, and so the fewer units one would tend to dispatch.

(3) The availability of fire engines in the surrounding area. The more units available, the more we would tend to dispatch.

(4) The relative utility or value of the response time of the first responding unit to that of the second responding unit. The higher this ratio, the fewer units one would tend to dispatch.

The mathematical programming formulation allows the simultaneous interplay of all the above factors. Empirical data were gathered for one region of New York City, the problem solved using the linear programming representation of the decision process, parametric and sensitivity analysis carried out, and finally, the suggested policies were tested using a detailed simulation of fire-fighting operations.* For the basic reference, see Swersey [8]. Out of this work came a simple rule suggesting control limits on each of the key parameters. This rule and elaborations, which select the particular companies to dispatch, are to be implemented in the computerized control system, but, in advance of that, in November 1972 the Fire Department implemented a city-wide "adaptive response" policy which was based in large part on our analysis.

---

*For a discussion of the simulation, see the work of Carter, Ignall and Walker [2], [3].

III. TEMPORARY RELOCATION OF FIRE ENGINES

When one large fire, or several small fires, is being fought in a single area of a city, the fire houses of the working fire units are left empty, resulting in a sharp degradation in the fire protection afforded the surrounding area. It is common practice in many cities to spread out the available companies by relocating some companies into selected empty houses. Existing manual methods to perform relocations use preplanned assignments which are adequate at low alarm rates but which break down at high alarm rates when the companies preassigned to relocate are not available, or when more than one serious fire is in progress at a time.

In New York City relocation problems occur on the average about 10 times a day and, if not solved quickly, can lead to serious situations. For example, a 5th alarm fire in Manhattan could deplete the borough of half of its fire-fighting units.

We have developed a dynamic algorithm which determines when relocations should be made, which empty houses should be filled, and which available companies should be moved. The algorithm has been specifically designed to be implemented in the proposed computerized control system. By using the computer's capability to store and update information about company status and to evaluate alternative plans using mathematical programming formulations--all in real time--the algorithm overcomes the deficiencies of the existing method.

Our aim was to develop a procedure for relocation which would overcome the problems of the existing system, was implementable within the computer time and space constraints we faced, and which produced "good" relocations. It was by no means clear at the outset what "good" meant, and we were concerned with "optimality" only as it was a means of achieving our more modest goals. We discuss at length the motivation behind our formulation in Kolesar and Walker [5]. It should suffice to say here that the Fire Department's objectives, although clear in principle, were never unambiguous enough to lead to a simple objective function, and we wound up with a problem with multiple criteria, which we formulated as the following series of integer optimization problems:

Problem 1. Determination of Empty Houses to Fill

The city was completely partitioned into small regions called response neighborhoods and fire engine responsibilities for covering these neighborhoods defined. Generally, there is one fire engine per fire house and, when some fire houses are left empty (for an extended period of time), the house and perhaps some of its response neighborhoods are left uncovered. For example, if the Fire Department adopts a criterion of having one of the two closest engines available, each response neighborhood would be the region "covered" by two engines. We adopted as our criterion for the determination of empty houses to fill: have every response neighborhood covered but move as few engines as possible. This translates to the integer program known as the set covering problem.

Problem 2. Determination of the Available Companies Which Relocate

Each of the houses designated to be filled in the solution to problem 1 must borrow an engine. We have been able to define a cost c_{ij} of relocating available engine i in empty house j. These "costs" are the increments in expected response time to future fires if the i to j move is made. Thus, the determination of who moves where becomes an assignment problem with some additional constraints like those of the set covering problem which are due to coverage requirements set by the Fire Department.

Problem 3. Determining Specific Relocation Assignments

The solutions to problem 2 repeatedly yield good selections of relocatees, but we found that the specific assignments could be improved upon. By permuting the assignments, the total distance traveled by the relocating engines could often be greatly reduced with small increases in the expected response time increments. This was a symptom of our multiple criteria problem. Consequently, we use an assignment problem to permute the relocations to achieve minimum total travel distance.

Clearly, the separation of problem 1 and problem 2 is artificial and leads to suboptimal solutions. But, the c_{ij} used in problem 2 depend on the geographic configuration resulting from the solution to problem 1 so that, in practice, we had no recourse save trying to obtain many solutions to problem 1 and using each as an input to problem 2. We have, however, been able to develop an approximate formulation linking both problems and are still testing it.

Heuristics which rapidly solve these problems have been developed and implemented [5], [7]. The algorithm has been extensively tested in simulations, and in the field on a small time-sharing system. Its actual implementation awaits the installation of the Fire Department's "Management

Information and Control" computer system. Meanwhile, the algorithm has been used recently in Denver in a way we never anticipated. It is employed to create for Denver the very alarm assignment cards we set out to replace in New York City, for, at Denver's low alarm rates, the static preplanned deployments still work!

We have also used a variant of the set covering problem in New York City to aid in making decisions about the permanent location of tower ladders and aerial ladders--different but similar types of equipment that the Fire Department wants distributed in a manner which avoids having two units of the same type respond to the same alarm.

IV. ALLOCATING FIRE ENGINES TO REGIONS OF THE CITY

Now we consider a more strategic problem, determination of the number of fire engines to station (permanently) in the various regions of the city. Suppose that the city has been partitioned into regions labeled $i = 1, 2, \ldots, m$, and that we know the utility function, $f_i(n_i)$, of the number of fire engines stationed in the region. Then, a variety of important allocation problems could be formulated. We have had some success along these lines. First, we have shown that, for any region, i, the response distance of the first arriving engine is given approximately by

$$RD_i = C_i \sqrt{\frac{A_i}{n_i - \lambda_i S_i}}$$

where

A_i = physical area in square miles

λ_i = alarm rate (average number of alarms per hour)

S_i = average total time required to service (extinguish, etc.) an alarm

C_i = an empirically determined constant depending on the geometry of the region.

Second, we have determined empirically the relation between response time and response distance. We then formulated and solved a variety of very simple integer optimization problems of the form:

Find integers $n_1, n_2, \ldots, n_m$ to

$$\text{minimize} \sum_{i=1}^{m} a_i f_i(n_i)$$

$$\text{subject to} \sum_{i=1}^{m} n_i = N$$

where the $f_i(n_i)$ are generally expected response times and the a_i are weights taking into account the hazards in the region.

A reference to this work is Kolesar and Blum [6]. Such analysis done for and with the Fire Department was partly instrumental in the decisions implemented by the Department in November 1972 to change the permanent locations of some fire-fighting units and to disband others and reassign their men to units in high hazard areas.

REFERENCES

1. Blum, E. H., "Deployment Research of the New York City Fire Project," The New York City-Rand Institute, R-968, May 1972; also appears as Chapter 7 in A. Drake, et al. (ed.), Analysis of Public Systems, M.I.T. Press, Cambridge, 1972.

2. Carter, G., and E. Ignall, "A Simulation Model of Fire Department Operations: Design and Preliminary Results," The New York City-Rand Institute, R-632, December 1970; also appears in IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-6, No. 4, 1970, pp. 282-293.

3. Carter, G., E. Ignall, and W. Walker, "A Simulation Model of the New York City Fire Department: Its Use as a Deployment Tool," Proceedings of the Winter 1973 Simulation Conference, 1973, pp. 353-370.

4. Chaiken, J., and R. Larson, "Methods for Allocating Urban Emergency Units," The New York City-Rand Institute, R-680, May 1971; also appears in Management Science, Vol. 19, No. 4, December 1972, pp. (P)-110-130.

5. Kolesar, P., and W. Walker, "An Algorithm for the Dynamic Relocation of Fire Companies," The New York City-Rand Institute, R-1023, September 1972.

6. Kolesar, P., and E. Blum, "Square Root Laws for Fire Engine Response Distances," to appear in Management Science.

7. Shanesy, C., "An On-Line Fire Company Relocation Program: Description and User's Manual," The New York City-Rand Institute, R-1152, forthcoming.

8. Swersey, A., "Models for Reducing Fire Engine Response Times," unpublished doctoral dissertation, School of Engineering and Applied Science, Columbia University, 1972.

Kolesar

MATHEMATICAL PROGRAMMING APPLICATIONS IN THE ANALYSIS OF THE
DEPLOYMENT AND UTILIZATION OF FIRE-FIGHTING RESOURCES

P-4988