R-1023-NYC
SEPTEMBER 1972

An Algorithm for the Dynamic Relocation of Fire Companies

Peter Kolesar and Warren E. Walker

THE
NEW YORK CITY
RAND
INSTITUTE
545 MADISON AVENUE NEW YORK NEW YORK 10022 (212) 758-2244

Rand

The New York City-Rand Institute is a nonprofit research institution founded, as its Articles of Incorporation state, "... primarily to conduct programs of scientific research and study, and provide reports and recommendations, relevant to the operations, planning or administration of the City of New York." The Institute was established in 1969 by the City of New York and The Rand Corporation as a center for the continuing application of scientific and analytic techniques to problems of urban life and local government. It is governed by a Board of Trustees appointed jointly by the City and Rand.

R-1023-NYC
SEPTEMBER 1972

An Algorithm for the Dynamic Relocation of Fire Companies

Peter Kolesar and Warren E. Walker

This study was sponsored by the City of New York. Its contents, however, do not purport to represent the official views or policy of the City of New York.

THE
NEW YORK CITY
RAND
INSTITUTE
545 MADISON AVENUE NEW YORK NEW YORK 10022 (212) 758-2244

Rand

Published by The Rand Corporation

PREFACE

This work was undertaken as part of the studies of the deployment of fire resources being performed by The New York City-Rand Institute for the Fire Department of the City of New York. The purpose of the work was to develop a relocation algorithm for the Department's planned computerized Management Information and Control System. In this Report the algorithm is described, examples of its application are given, and test results are documented. The algorithm has been implemented in an interactive mode on a time-shared computer; a description of this implementation is included. Although designed to solve a problem for the New York City Fire Department, the algorithm should be applicable to other cities.

SUMMARY

When all the fire companies in a region are engaged in fighting fires, protection against a future fire is considerably reduced. It is standard practice in many urban fire departments to protect the exposed region by temporarily relocating outside fire companies in some of the vacant houses. Situations requiring such relocations arise an average of ten times per day in New York City and the Fire Department of the City of New York (FDNY) currently makes its relocations according to a system of preplanned moves. This system was designed at a time when alarm rates were low and is based on the assumption that only one fire is in progress at a time. Because of the high alarm rates currently being experienced in parts of New York City this assumption is no longer valid, and the preplanned relocation system breaks down at the times when it is needed most.

This Report describes a computer based method for determining relocations which overcomes the deficiencies of the existing method by utilizing the computer's ability to: (1) store up to date information about the status of all fires in progress and the location and activity of all fire companies, (2) quickly generate and compare many alternative relocation plans.

The method, which will become part of the FDNY's real-time Management Information and Control System (MICS), is designed to be fast and to require little computer memory.

The relocation method consists of four interrelated problems each of which is solved by the application of simple decision rules, called heuristics.

1. *When should relocations be made?* A call for relocations will be made whenever the fire protection being provided to any area of the City falls below a given minimum level.

2. *Which vacant houses should be filled?* The houses to be filled are chosen to bring fire protection in all areas above minimum levels while moving as few companies as possible.

3. *Which available units should be moved?* A function is used to compare alternative relocations which expresses the "cost" of relocation in terms of response time to future fires. The function takes into account such factors as alarm rates, relocation distance, expected response times and expected durations of the serious fires. The set of companies which produce the lowest "cost" relocation are selected to be moved.

4. *Which specific relocating units should be assigned to each of the vacant houses being filled?* The set of relocating units is assigned to the set of houses to be filled so that the total distance travelled by the relocating units is minimized.

After giving some background of the problem and the objectives of relocation, we give the problem a mathematical programming formulation and then describe the heuristic algorithm to be used for generating relocations in the MICS. The remainder of the Report is devoted to a discussion of an example, a rigorous test of the algorithm using a computer simulation model of Fire Department operations, and a description of the current use of the computer algorithm by dispatchers in an interactive time-shared environment. Results of testing indicate that the proposed algorithm is a significant improvement over existing methods, particularly in crisis situations.

ACKNOWLEDGMENTS

Many persons have contributed to this effort and we could not hope to properly acknowledge all of them here. Special mention must go to Jan M. Chaiken, who contributed several of the key ideas used in stage 2 of the algorithm, and to Jack Hausner and Carol E. Shanesy, who "made it work" by assembling the necessary data and writing the computer programs. In earlier work, Arthur J. Swersey took a different approach to the relocation problem and, although his ideas were not incorporated here, they did influence our thinking.

The continual feedback from members of the New York City Fire Department kept us from straying too far from reality. We are particularly indebted to Deputy Chiefs Homer G. Bishop and Francis J. Ronan, whose ideas and comments have been incorporated throughout.

CONTENTS

LIST OF TABLES

LIST OF FIGURES

I. INTRODUCTION

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. These methods are adequate at low alarm rates but 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 frequently and, if not solved quickly, can lead to serious situations. An average of 10 such problems occur per day in New York City and a particularly serious incident such as a 5th alarm fire in Manhattan could deplete the borough of half of its fire-fighting units. For these reasons a high priority has been placed by the Fire Department of the City of New York (FDNY) on having a system for making relocations which will be able to make rapid and satisfactory relocations under extreme conditions.

In this Report, we describe 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 Management Information and Control System of the FDNY--an on-line, real-time system. By using the computer's capability to store and update information about company status and to evaluate alternative plans rapidly, the algorithm overcomes the deficiencies of existing methods of making relocations. In addition, some of the methods described can be used without a computer to improve manual relocation methods.

Succeeding sections describe the mathematical formulation of the relocation problem, the algorithm developed for its solution, and the testing we have done with the algorithm. First we explain why we have taken this particular approach, since an understanding of our objectives is crucial to an understanding of the finished product.

Our goal was to develop a procedure for relocating fire companies 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; it would have to be determined during the study. But we decided that our definition would be operational. It would mean that we, the Fire Department, and the public would all agree that the procedure was making the right kind of relocations. The integer programming and optimization approach that we developed can be viewed as a device by which we achieved the modest goal of producing "good" relocations.

After considerable analysis of the problem and discussion with the Fire Department, we chose such a modest goal because we felt that the objectives of the Department with respect to relocation were too ambiguous to allow us to formulate "the problem" and find "the optimal solution." This difficulty is common to many situations encountered by the analyst attempting to solve problems related to the distribution of municipal services. In industrial problems it is usually less difficult to specify a reasonable objective function--some sort of minimization of costs or maximization of profits are usually meaningful, quantifiable objectives. In municipal problems the objectives are often unquantifiable or even unknown. One frequently is confronted by adversary situations in which a gain to one segment of the public comes only at loss to others.

In the case of relocation, it is not difficult to formulate several possible objectives, but they lead to policies which are unacceptable to the Fire Department. For example, if the Department's objective were to provide equal first unit response time to all areas of the City (equity), the available units should be spread out rather uniformly. If, however, the goal were to minimize the City-wide average response time to alarms (efficiency), the companies should be highly concentrated in the areas where the expected fire incidence is greatest. While the Fire Department's current policies achieve a result somewhere in between these two extremes, their actual objective seemed impossible to define. We, therefore, decided on a sequential process with several objectives, which has led to relocation policies whose results appear acceptable to the Department.

We discuss these objectives in Section II. In Section III, we present a mathematical formulation of the relocation problem, a heuristic solution to which is presented in Section IV. In Section V, a simple example is used to demonstrate the way the algorithm works. The results from the heuristic algorithm for this case are compared with the exact results of the integer program. A description of some simulated test results based on a real scenario of fires is given in Section VI. In Section VII, we describe an on-line implementation of the algorithm which is currently being used for demonstration and testing purposes.

II. DEFINING THE OBJECTIVES

PRIMARY OBJECTIVE

One way of balancing equity against efficiency is to set a minimum coverage standard for every area in the City. One such standard discussed in the literature is to guarantee that there is at least one fire-fighting company within x minutes (or y miles) of every alarm box.[7] It would not be difficult to provide such coverage if x (or y) were known and were constant over the City. In practice, however, the minimum standard varies in different areas depending on the hazards associated with the areas, and it would be difficult for the Fire Department to specify the values of x (or y) for each area.

An alternative is to let the way fire-fighting units are allocated to areas implicitly define the minimum coverage standards for those areas. Fire companies are not uniformly distributed over the City but are concentrated in some areas and spread out in others. This distribution is the result of a complexity of forces--some political, some operational, others historical. In working with this distribution, the Fire Department has implicitly decided how it wishes to balance equity against efficiency in the short run. In the long run, of course, the Fire Department may modify the distribution by building new firehouses. By assuming that the Department is satisfied with the distribution of fire companies, we can define a minimum coverage measure which will maintain approximately the same relative distribution of companies and, therefore, will maintain approximately the same variation in response times (or distances) which exist between the areas.

This is accomplished by requiring that, for every point in the City, at least one of the three closest engine houses and at least one of the two closest ladder houses contain an available company. However, since New York City has a large number of fire alarm boxes (over 16,000) and they are spread rather uniformly throughout the City (roughly one at every other street intersection), it is possible to simplify this minimum coverage definition by using the alarm boxes as our "points." The definition of minimum coverage which we are using is that at least one of the closest three engines and at least one of the closest two ladders must be available for every alarm box in the City. Relocations will be recommended whenever minimum coverage is not being provided, and the houses to be filled and the companies which relocate will be selected so as to guarantee minimum coverage.

While this standard of coverage has met with the general approval of the FDNY and is the standard we have applied in most regions of New York City it need not be applied inflexibly. The coverage standard can vary in order to provide appropriate coverage to areas which are particularly isolated or have other special characteristics. For example, some areas might require that the closest ladder company be available, while in others only one of the three closest ladders should be available. The algorithm discussed in Sections III and IV is flexible enough to easily accommodate such variations. However, in order to keep the following discussion simple, we proceed as if a uniform criterion is applied. No loss of generality is involved, for our framework allows us to easily define and implement a minimum coverage standard requiring that m_j out of the closest n_j units of a specified type be available to alarm box j.

It is possible to simplify the minimum coverage definition even further by noticing that, in general, several alarm boxes will have the same three closest engines or the same two closest ladders. The aggregation of all alarm boxes having the same three closest engines we call an engine response neighborhood (written engine RN for brevity). A ladder RN is defined as the set of alarm boxes having the same two closest ladders. The set of engine and ladder RN's each form non-overlapping partitions of the City. They are defined separately since the coverage standard is to be applied separately. The definition of minimum coverage can now be restated as: there must be no engine RN with its three engines unavailable and no ladder RN with both its ladders unavailable.

The use of RN's results in a considerable reduction in the calculations required to check on coverage. For example, in the Bronx, a borough of New York City, there are over 2000 alarm boxes but fewer than 50 ladder RN's.

It should be pointed out that the definition of unit availability for purposes of minimum coverage has not yet been made clear. We do not wish to relocate into the house of a company which is responding to an alarm, returning from an alarm, or is due back soon from a working fire. Therefore, we consider a company to be unavailable only if it is working at a fire which will last for a "considerable" length of time, a concept which will be defined precisely later.

SECONDARY OBJECTIVES

Once the houses to be filled have been selected, on the basis of the minimum coverage criterion, there may be many available companies which could be moved into those houses. Of course, we would never move a company if, by moving it, the minimum coverage criterion would be violated. In choosing among the possible moves, it seems reasonable to use the following secondary criteria:

- Don't move a company "too long" a distance.
- Don't move a company which is "too busy."
- Don't move a company which is protecting "too big" an area.

The trade-offs between these criteria are not immediately clear. We sought a simple function which would reflect the attractiveness of one move relative to another move and found that, by using a response time measure, all these secondary factors could be taken into account.

Since the primary objective of minimum coverage is basically equity (alarm rates are not considered and the minimum coverage measure prevents anyone's protection from being severely degraded), we sought an efficiency measure as the secondary objective. Accordingly, we select the relocation* which yields the minimum total expected response time to alarms

*The set of individual company moves made at one time is called a relocation.

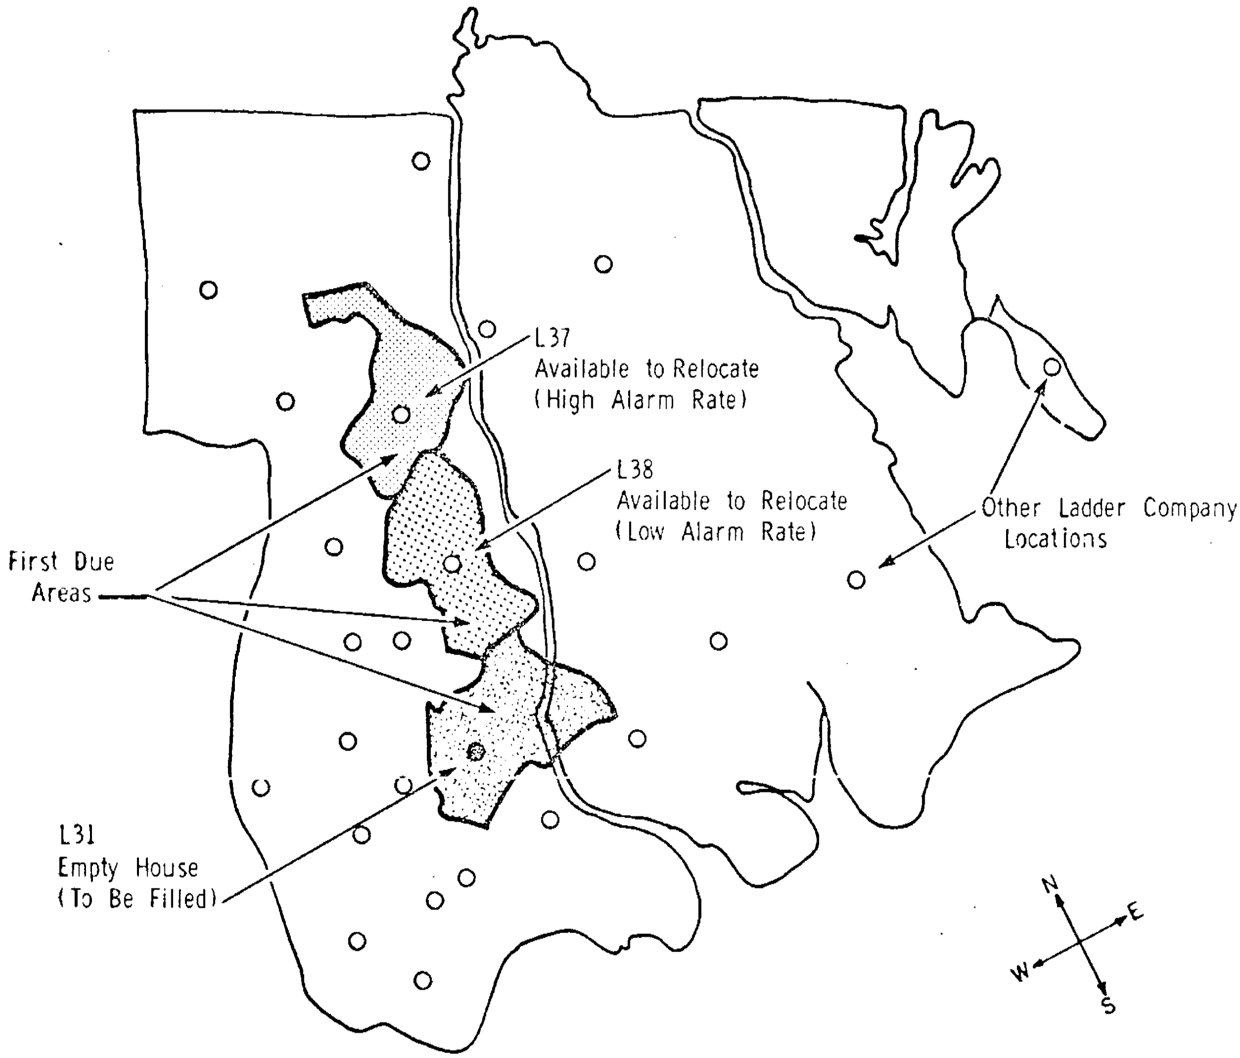

Fig. 1 - A representation of the relocation problem

which occur during the period of interest. The following simple (and simplified) scenario provides the background for the development of this secondary objective function.

Referring to Fig. 1, suppose ladder 31 has just responded to a serious fire. Its house is empty and we wish to evaluate possible relocations into it. The houses of ladder companies 37 and 38 are currently covered, and it is feasible to relocate either company into ladder 31's house without violating the minimum coverage standard. We want to evaluate which move is superior and if, indeed, any move should be made.

First, consider the data required to make the evaluation. Let R_{31} denote the region in which ladder 31 would be the closest company to all alarm boxes if it were available. Let A_{31} and λ_{31} dencte respectively the physical area and the alarm rate of this region. Let R_{37} be the region in which ladder 37 is currently the closest available company, and A_{37} and λ_{37} again be respectively the area and alarm rate in R_{37}. We have similar definitions for R_{38}, A_{38}, and λ_{38}. We assume that the alarm process is Poisson and that all the above parameters have been estimated. In addition, r_{ij}, the time required to relocate a company from location i to location j, is assumed to be known.

The expected response time of the first arriving company to an alarm in the regions served by these companies depends on which companies are available to respond to the alarm. If the closest company is not in quarters, the second closest company must play the role of the closest company with a consequent increase in response time. Exact calculations of expected response times of the first arriving fire company in any region can be made if the alarm rates and response times associated with individual alarm boxes are known. Alarm rates can be estimated from historical data, and actual response distances or times to each box can be measured directly.

Of course, the data and computation requirements for doing this are formidable, so, that for the present, we use an approximation for response times based upon abstract mathematical models and empirical evidence. The key idea used in the approximation is that expected response distance in a region is proportional to the square root of the area served per fire company. The constant of proportionality depends upon the arrangement of the companies, the nature of the street patterns, and the distribution of alarms.* Response times in regions served by a single company such as R_{31}, R_{37}, and R_{38} are approximated as follows: Constants c_1 and c_2 are determined such that $c_1\sqrt{A_i}$ is an estimate of the expected response distance of the closest responding unit in R_i, and $c_2\sqrt{A_i}$ is an estimate of the expected response distance of the second closest responding unit in R_i. So, $c_1\sqrt{A_i}$ is the expected response distance of the closest responding unit when company i is available. If it is unavailable, but its neighbors are available (as is more or less the case if minimum coverage is being guaranteed) then $c_2\sqrt{A_i}$ is the expected response distance of the closest responding unit. During an interval of length t we have $\lambda_i t$ alarms on the average in R_i. Denoting the average response velocity in R_i by v_i and ignoring some of the dynamic behavior occurring in R_i, we have $c_1\sqrt{A_i}\ \lambda_i t/v_i$ or $c_2\sqrt{A_i}\ \lambda_i t/v_i$ as the expected total response time in R_i, according to whether company i is available or unavailable during the interval t--the duration of the fire which is causing the problem. We are assuming that the other alarms which occur during the interval are not serious enough to require the service of one of the other companies for a long period of time. If any alarm is serious, another relocation problem occurs. Our deliberately myopic view of the problem ignores such second order effects by taking a snapshot of a dynamic system.

For notational simplicity, let

$$\alpha_i = \frac{\lambda_i \sqrt{A_i}}{v_i} .$$

We now return to the simple scenario where ladder companies 37 and 38 are candidates to relocate into ladder 31's house, and we evaluate the cost of relocating ladder 37 in terms of expected total response time. Let t denote the duration of the fire at which ladder 31 is working, and let $T > t$ be an arbitrary interval long enough so that any company which might be relocated would be able to return to its own quarters before T. Then, we have for the expected total response time over the interval [0, T] in regions R_{31}, R_{37} and R_{38},

$$(c_2 - c_1)[\alpha_{37}(t + r_{37,31}) + \alpha_{31}r_{37,31}] + c_1T(\alpha_{31} + \alpha_{37} + \alpha_{38}).$$

This calculation is based on the assumption that ladder 37 spends a time $r_{37,31}$ travelling to ladder 31's house, stays at that house until ladder 31 returns from the fire at time t, and then returns home. So r_{31} is covered by a second closest company for $r_{37,31}$ hours and by a closest company for $T - r_{37,31}$; R_{37} is covered by second closest company for $t + r_{37,31}$, etc.

If ladder 38 relocated into ladder 31's house, the expected total response time over the interval [0, T] would be

$$(c_2 - c_1)[\alpha_{38}(t + r_{38,31}) + \alpha_{31}r_{38,31}] + c_1T(\alpha_{31} + \alpha_{37} + \alpha_{38}),$$

while, if no relocation were made, we would have

$$(c_2 - c_1)\alpha_{31}t + c_1T(\alpha_{31} + \alpha_{37} + \alpha_{38}).$$

The last term of each expression is invariant so we can compare the alternatives by comparing only the first terms, which can now be written in a general form. Letting c_{ij} denote the "cost" in expected total response time of relocating available company i into empty house j, we have

$$c_{ij} = (c_2 - c_1)[\alpha_i(t + r_{ij}) + \alpha_j r_{ij}],$$

and the "cost" of making no relocation is just $(c_2 - c_1)\alpha_j t$. Using similar reasoning, we can evaluate the cost of more complicated actions such as moving available company k to the house of available company i while the latter relocates into empty house j. The cost of this "successive moveup" is

$$(c_2 - c_1)[r_{ij}\alpha_j + r_{ki}\alpha_i + (t + r_{ij} + r_{ki})\alpha_k].$$

Before we use these functions to evaluate some actual relocation plans, we note that each of the three secondary factors--relocation travel distance (r_{ij}), the "busyness" of a company (λ_i), and the size of the region protected by a company (A_i, c_2 and c_1) are all explicitly included in in the cost functions. In addition, another element appears that perhaps had not been anticipated: the duration of the relocation. According to the cost function, it is possible that a different relocation would be made for a short incident than for a long incident. In fact, it is even possible to determine how long the incident must be before it becomes advantageous to relocate.

Figure 2 is typical of the many graphs we made of relocation costs for real and hypothetical problems. The curves are for the situation shown in Figure 1 in which ladder company 31 is working at a fire and ladder companies 37 and 38 are available to relocate. The graphs depict average total response time as a function of the duration of the incident for four alternatives:

(1) No relocation (ladder 31's house remains uncovered).

(2) Move ladder 38, which is close to ladder 31, but is a busy company.

(3) Move ladder 37, which is farther away from ladder 31, but is less busy. The Fire Department calls such move a "leapfrog" relocation.

(4) Relocate ladder 37 into ladder 38's house, and relocate ladder 38 into ladder 31--called a "successive moveup."

For any given value of t, the best policy is the one whose cost function is the smallest. Examination of the graph indicates that it does not pay to make a relocation for less than about 15 minutes. If the fire will last more than that, the best plan is to move ladder 37 into ladder 31's house. If the incident lasts over a half hour, there is a clear advantage to this relocation. Note that if ladder 37 could not be moved it would not be worthwhile to make any relocation for a fire lasting

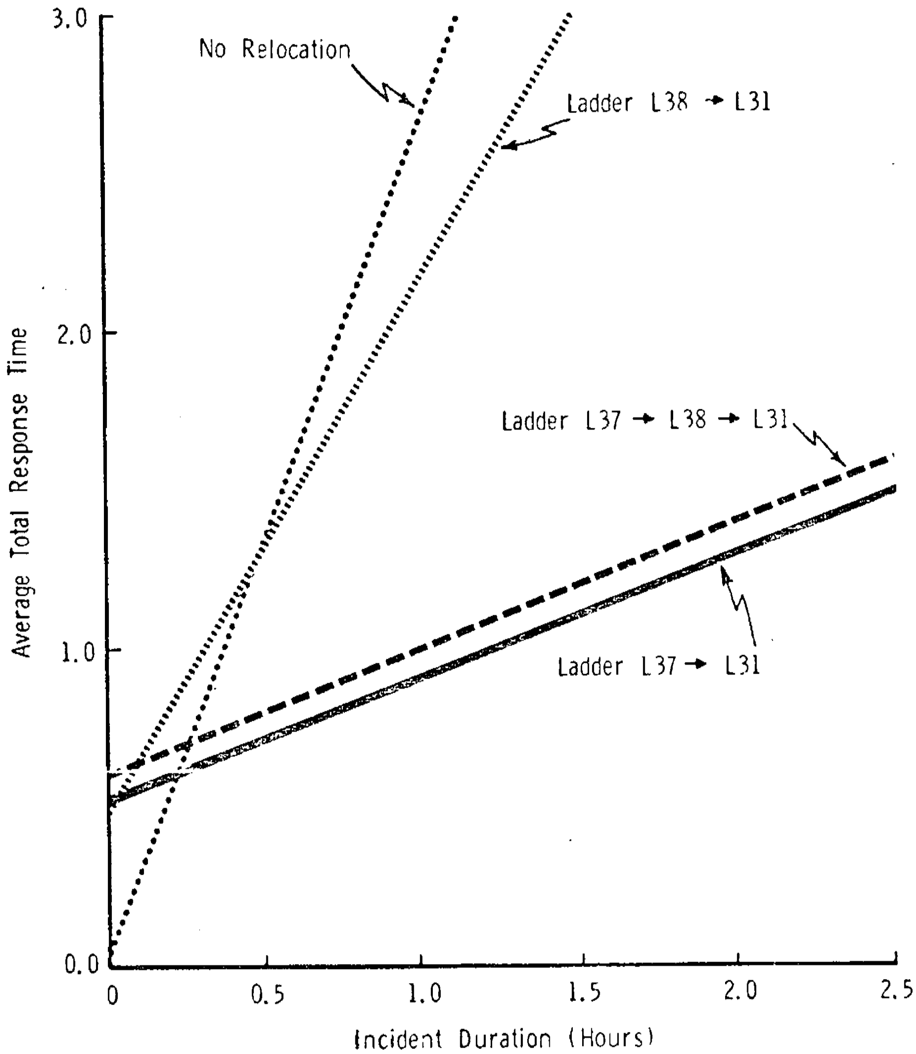

Fig. 2 - A comparison of relocations into the house of Ladder 31

one-half hour or less. We also note that the successive moveup of 37 to 38 to 31 is slightly worse than just relocating ladder 37 to ladder 31's house.

Examination of similar curves for many situations led us to some generalizations about relocations. First, we saw that "successive moveups" have about the same response time "cost" as some simple relocations. But they also move twice as many companies, thereby increasing the inconvenience to the men--firemen cook their own meals in their own house, keep dry changes of clothes there, etc.--as well as increasing communication and control problems. For these reasons we eliminated successive moveups from further consideration. Second, in most situations we looked at, there seemed to be a clear response time "cost" advantage to relocating if the relocation was to last more than one hour. In general, an incident which will last an hour or more is easily identifiable by a chief when he first arrives at the fire. So, rather than requiring the exact duration of all incidents (the value of t in the cost function), we decided to make relocations only for fires which are expected to last more than one hour.

For such "serious fires" the relative response time cost rankings do not change very much regardless of their actual duration. That is to say, if ladder 37 to ladder 31 looks better than ladder 38 to ladder 31 for a one-hour fire, it also looks better for a two-hour fire. So our response time cost calculations are based on the average duration of a serious fire--about one hour. Should the Department be able to make better predictions of incident duration in the future, those predictions could, of course, be used to make more accurate cost estimates.

III. MATHEMATICAL FORMULATION

In this section, we will present a mathematical formulation of the relocation decision in which the problem is broken into four stages which are solved sequentially:

(1) Determination of the need for a relocation by establishing if uncovered response neighborhoods exist.

(2) Determination of the empty houses to be filled in order to cover all response neighborhoods with a minimum number of moves.

(3) Determination of the available companies to be relocated to minimize expected response time.

(4) Determination of relocation assignments to minimize total travel distance.

The last three stages are integer linear programming problems which can either be solved exactly using standard algorithms, or approximately, using the heuristic procedures presented in Section IV. We show how the second and third stages can be combined into a single integer linear program which eliminates the suboptimality inherent in their separation.

STAGE 1 - DETERMINATION OF THE NEED FOR A RELOCATION

A call for relocations will be made whenever an uncovered response neighborhood exists anywhere in the City. Uncovered RN's will be detected by a program called the trigger which will be called periodically (say, every minute) or whenever the system learns that a fire requires the use of at least three engines and two ladders. A search could be made for uncovered RN's after any alarm or higher alarm, but during busy hours these will occur frequently and doing so would interfere with other functions of the MICS. The trigger will process all changes in the status of companies received by the system since the last time it was called. If a status change results in an RN becoming uncovered, the RN is added to a list of uncovered RN's. If, after all status changes are processed, this list is empty, the next stage of the relocation algorithm is not called. If the list is not empty, the second stage of the algorithm is entered with the list of uncovered RN's as input data.

Generally, the rest of the relocation algorithm will be executed twice, once for engines (if there are uncovered engine RN's) and once for ladders (if there are uncovered ladder RN's). The discussion which follows applies equally well to either case.

STAGE 2 - DETERMINATION OF THOSE EMPTY HOUSES TO BE FILLED

The second stage of the algorithm determines which of the uncovered houses should be filled so that the number of companies relocated is minimized. Formally, this problem belongs to a class of integer programming problems known as "covering problems."

Suppose that there are K uncovered RN's and L vacant houses whose busy companies appear in the labels of these RN's. Let $x_j=1$ if house j is to be filled and $x_j=0$ otherwise. Then, the problem can be stated as:

$$\min \sum_{j=1}^{L} x_j$$

$$\text{subject to} \quad \sum_{j=1}^{L} a_{ij}x_j \geq 1 \qquad i = 1, 2, \ldots, K$$

$$x_j = 0, 1 \qquad j = 1, 2, \ldots, L$$

$$\text{where } a_{ij} = \begin{cases} 1 & \text{if the jth house's busy company appears in the label of the ith RN,} \\ 0 & \text{otherwise.} \end{cases}$$

(The matrix of a_{ij}'s is known as the incidence matrix for the covering problem.)

The output of stage 2 is a set of M ($M \leq L$) vacant houses which are to be filled, which is the input to stage 3. In many real problems there are alternative minimal coverings, that is, several sets of houses to be covered which require the same number of moves. In these cases the desired covering is that which, together with the set of companies which relocate, gives the minimum total expected response time for all alarms which occur during the period of interest. In order to find this optimal covering either all the alternative coverings must be included in the stage 3 problem or the stage 2 and stage 3 problems must be solved simultaneously. We show later how the stages may be solved simultaneously.

In truly extraordinary circumstances there may be no feasible solution to the covering problem. In such cases the department uses emergency allocation procedures which are not discussed here.

STAGE 3 - DETERMINATION OF THE AVAILABLE COMPANIES WHICH RELOCATE

The next decision to be made is: Which available companies should relocate? We formulate this problem as if we were not only selecting the companies to relocate but also, at the same time, assigning them to the empty houses. We shall not use the specific assignments so generated but only the set of relocatees. Actual assignments are generated in the next stage. The problem is formulated as an integer linear program which contains a "transportation" subproblem, and has, in addition, a set of constraints which assures that our coverage criteria are not violated.

We let $j = 1, 2, \ldots, M$ refer to the empty houses to be filled, $j = M+1,\ldots,M+N$ refer to the available companies, and $k = 1, 2, \ldots, L$ refer to the RN's associated with the available companies.

The objective function to be minimized is the total expected response time during the relocation incident. As shown in the discussion of relocation costs given in the introduction, we have for the "cost" of relocating available company i into empty house j

$$c_{ij} = (c_2 - c_1)[\alpha_i \ (t + r_{ij}) + \alpha_j r_{ij}].$$

Now consider the total cost of M moves. We get the total cost of each relocation of available company i_j to empty house j, $j = 1, 2, \ldots, M$

$$(c_2 - c_1) \sum_{j=1}^{M} [\alpha_{i_j} \ (t + r_{i_j j}) + \alpha_j r_{i_j j}] = \sum_{j=1}^{M} c_{i_j j}.$$

Let $x_{ij} = 1$ if available company i is assigned to relocate into empty house j, and $x_{ij} = 0$ otherwise. We introduce a "dummy" empty house 0 so that when available company i is not moved, $x_{io} = 1$. The integer linear program to be solved is:

Find $\{x_{ij}\}$ to

$$\text{minimize} \sum_{j=1}^{M} \sum_{i=M+1}^{N+M} x_{ij} c_{ij}$$

subject to

$$\sum_{i=M+1}^{N+M} x_{ij} = 1 \qquad j = 1, 2, \ldots, M$$

$$\sum_{j=0}^{M} x_{ij} = 1 \qquad i = M+1, M+2, \ldots, M+N$$

$$\sum_{i=M+1}^{N+M} a_{ik} x_{i0} + \sum_{i=M+1}^{N+M} \sum_{j=1}^{M} a_{jk} x_{ij} \geq 1 \qquad k = 1, 2, \ldots, L$$

$$x_{ij} = 0, 1 \qquad \text{for all } i, j.$$

The objective function and the first two sets of constraints are just a transportation problem, and the meaning of $\sum x_{ij} = 1$ is that all available companies are assigned somewhere and all empty houses are filled. As in stage 2, $a_{ik} = 1$ if available company i serves response neighborhood k and is zero otherwise. The last set of constraints requires that no RN associated with an available company be uncovered.

STAGE 4 - SPECIFIC RELOCATION ASSIGNMENTS

The output of stage 3 is a specific set of assignments or "moves" of available companies to the empty houses being filled. We have found through testing that we can improve upon these moves through application of a secondary objective. For several reasons Fire Department management is concerned with the distance that relocating companies must move. One reason is that shorter relocation distances mean less of a burden on the relocating companies and greater availability times. Another is that by keeping the relocating distance down, one tends to keep companies in areas in which they are familiar with street patterns as well as with particular fire-fighting problems.

Thus, having a set of relocatees which has been selected in stage 3 by balancing distances, alarm rates, and response times, in stage 4 we find the assignment which minimizes total travel distance. Of course this would not make sense if it upset our response time "cost" objective, but in most cases the resulting assignment increases the relocation cost (expected response time) very little, and can significantly reduce the total distance travelled (see Section V).

We now let the indices j = 1, 2, ..., M refer to the M empty houses *selected by stage 2* and the indices i = 1, 2, ..., M refer to the M *available companies selected by stage 3*. Again, r_{ij} denotes the time required for a unit to relocate from "full" house i to "empty" house j, and $x_{ij} = 1$ if available company i is assigned to empty house j and it is zero otherwise.

Mathematically, we have a traditional assignment problem to solve:

Find $\{x_{ij}\}$ to

$$\text{minimize} \quad \sum_{j=1}^{M} \sum_{i=1}^{M} r_{ij} x_{ij}$$

subject to

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

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

$$x_{ij} = 0, 1.$$

A ONE-STAGE INTEGER L.P. FORMULATION WHICH AVOIDS SUBOPTIMALITY

Solving the relocation problem in two stages, the first finding a set of houses to be filled, the second finding a set of relocatees, may yield suboptimal solutions. A formulation which would solve both stages at once would be of interest, if only for testing how well the two-stage algorithm performs. One could, of course, generate all solutions to stage 2, try each in stage 3 and select the best solution. Still better would be to combine both stages. This can be done by solving a combined stage 2-stage 3 problem which consists of a modification of our formulation of the stage 3 problem in which we let $j = 1, 2, \ldots, L$ index all the empty houses, $j = 0$ be the dummy house and $i = L+1, L+2, \ldots, N+L$ index all the available companies. We add to the objective function a large fixed cost Q for each relocation so that the optimal solution will automatically contain a minimum number of moves. The programming problem is:

$$\text{minimize} \quad \sum_{j=1}^{L} \sum_{i=L+1}^{N+L} x_{ij}(c_{ij} + Q)$$

subject to

$$\sum_{j=0}^{L} x_{ij} = 1 \qquad i = L+1, L+2, \ldots, L+N$$

$$\sum_{j=1}^{L} \sum_{i=L+1}^{N+L} a_{jk}x_{ij} + \sum_{i=L+1}^{N+L} a_{ik}x_{i,0} \geq 1 \qquad k = 1, 2, \ldots, K$$

$$x_{ij} = 0, 1 \qquad \text{for all } i, j.$$

In this formulation, note that we deal with all empty houses, not the specific subset of size M determined by solving stage 1, and the set of RN's labelled k= 1, 2, ..., K must be the set of RN's associated with both full and empty houses. There are (L+1) · N variables and L+N+K constraints. For a difficult example involving simultaneous 2nd and 3rd alarms in the Bronx, there would be at least 200 variables and 150 constraints. It would require far more computation time to solve this problem than would be reasonable in the Management Information and Control System. We therefore derived a heuristic algorithm to solve the problem using the Stage 2-Stage 3 partitions. We must, of course, be concerned with whether the heuristic algorithm gives us optimal or near optimal solutions. We discuss such matters in the next section, and we indicate the steps taken to avoid the suboptimality inherent in partitioning the problem into stages.

IV. A HEURISTIC ALGORITHM

We now turn our attention toward a method for solving the optimization problems we have formulated. All these problems are integer programs in zero-one variables. Stage 2 has the special structure of a set covering problem, stage 3 is a transportation problem with additional constraints, and stage 4 is an assignment problem. All these problems can be solved exactly using one of several special algorithms. Unfortunately, except for the assignment problem, exact algorithms require more computation time and computer memory than can be afforded in a real-time system.[1] Relocation is only one of many on-line functions to be carried out by the proposed Management Information and Control System, so we must be careful with our use of computer time and storage. Thus, we have designed heuristic algorithms to obtain approximate solutions to the problems. In our testing we have compared the results obtained using exact algorithms[4] to those obtained using the heuristics in order to check how far from optimal our solutions are likely to be. We have not yet formulated a problem for which the heuristic obtained a non-optimal solution.

STAGE 2 - DETERMINATION OF THOSE EMPTY HOUSES TO BE FILLED

The heuristic rule for selection of a house to fill is to select that house which is associated with the largest number of uncovered RN's. The rule may be applied several times. After each application the covering problem is reduced by the elimination of the house selected to be filled and all RN's which will be covered as a consequence. This procedure continues until all RN's are covered. Furthermore, since these evaluations can be made extremely rapidly, we repeat the entire procedure several times using alternate starting points. The method can be summarized as follows:

(0) Set $j = 1$.

(1) Fill the company location associated with the jth largest number of uncovered RN's. If there is a tie, fill the location of the company with the highest alarm rate in its first due area.

(2) Fill the location which now belongs to the largest number of uncovered RN's. If there is a tie, fill the location of the company with the highest alarm rate in its first due area. Keep going until no RN's are uncovered. Count the number of houses which are filled.

(3) Repeat steps (1) and (2) for j = 2, 3, ..., P.*

(4) If there are no ties, the set of empty houses to be filled will be the one which requires the fewest number of houses to be filled. If several sets are tied with the fewest number of houses to be filled, solve stage 3 separately for each set and select that set which produces the minimum stage 3 cost.

STAGE 3 - DETERMINATION OF THE AVAILABLE COMPANIES WHICH RELOCATE

The heuristic rule for determining which of the available companies to move into a given house is to choose the available company with the lowest relocation cost. The relocation costs are the c_{ij}'s described above. To facilitate the selection of relocatees for consideration we create for each of the houses to be filled a ranked list of candidate relocatees consisting of available companies ordered by their c_{ij}.

The feasibility of each move must be checked against the coverage criterion (no RN's should become uncovered as a result of the move). The feasibility test is not a trivial operation. Even though, individually, a company on any relocatee list may be relocated without violating minimum coverage, if selections are made independently for each vacant house to be filled, the resulting relocation might have the same company moving into more than one house, or might leave one or more RN's uncovered by moving neighboring companies. Recall that, for ladders, the two companies comprising an RN cannot both be allowed to relocate. For engines, no relocation is feasible which leaves an engine RN without an available company.

---

*The value of P is arbitrary, but we have been using P = max(6, number of empty houses). Thus, we determine a maximum of six covering sets.

A feasible relocation is generated by successive applications of the heuristic subject to feasibility tests, starting with one house to be filled and sequencing through the others one by one. If the lowest cost move for each house produces a feasible relocation, it is used since it is optimal. Otherwise, since the algorithm is fast, several feasible relocations will be generated and the best one selected. Different feasible relocations are produced by changing the order in which houses being filled are considered, and by changing the heuristic for the first house being considered to "choose the available company with the kth lowest relocation cost." The least cost relocation generated is taken as the stage 3 solution.

The method may be summarized as:

A. Fast check for an optimal solution

1. Find the minimum value of c_{ij} over all the relocatee lists. Imagine that the move it represents is made.
2. Search the lists of the remaining houses to be filled for the smallest c_{ij} which represents a feasible move. Add this move to the relocation being generated.
3. Repeat step 2 until a relocatee has been selected for each house to be filled.
4. If each of the companies in this relocation is associated with the lowest c_{ij} element in one of the relocatee lists of houses to be filled, this relocation is optimal. Exit.

B. Generate a series of feasible relocations

0. Assign the numbers 1, 2, ..., M to the M vacant houses to be filled. Set the house indicator, i, to 1.
1. Imagine that the lowest cost move into house i is made.
2. Sequence in order through the remaining houses to be filled ($j = 1, 2, \ldots, M; j \neq i$). For each house,

j, find the feasible move associated with the smallest cost element on house j's relocatee list (it must be feasible with respect to the other moves already included in the relocation now being generated). This set of moves produces a feasible relocation.

3. Repeat step 2, first filling house i with the company associated with the second* lowest cost element on house i's relocatee list. This produces a second feasible relocation starting with house i.

4. Steps 0-3 produce two feasible relocations. They are generated by first finding the best and second best relocatees for house 1. Repeat steps 1-3 for each of the other houses (i = 2, 3, ..., M).

Part A of this method generates one relocation and Part B generates 2M relocations.* Thus, a total of 2M + 1 candidate relocations are generated from which the one with the lowest cost is selected. The relocation so selected provides the set of M available companies which relocate.

STAGE 4 - SPECIFIC RELOCATION ASSIGNMENTS

Having selected the set of relocatees as above, we now must solve an assignment problem to minimize total relocation distance. For small problems, say up to M = 5, solutions can be quickly obtained by complete enumeration of all M! permutations. For larger problems an exact and efficient algorithm may be used. We have employed the Balinski-Gomory[2] method which appears to be very effective.

---

*One need not stop after the second lowest cost. In general, step 2 may be repeated for each of the first Q lowest cost relocatees on house i's relocatee list. Our tests have indicated that going below the second rarely leads to a better solution.

V. AN EXAMPLE

In this section we present a hypothetical situation to illustrate how the algorithm works.

We consider a case in which two serious fires, a 2nd and a 3rd alarm fire, are announced simultaneously in the Bronx, a borough of New York City. To keep the discussion simple, but without any loss of generality, we analyze the relocation problem for ladders only and consider only Bronx ladder companies as possible relocatees. Figure 3 shows the locations of the two fires, the locations of the seven ladder companies working at the fires, and the region left uncovered as a result. All but three of the other ladder companies are assumed available to relocate. These three companies are identified as reserved because Fire Department policy dictates that, due to geographical isolation, they are not candidates for relocation. In the algorithm the ranked lists of available companies would reflect this policy.

In stage 1 of the algorithm we find that there are 9 uncovered ladder RN's. We then solve stage 2 and find that there are two solutions, each consisting of four of the seven empty houses, which provide a minimum covering; that is, they leave no RN's uncovered and move a minimum number of companies. The two sets differ by only one house. Stage 3, which selects the companies to relocate, was run using both of these solutions to stage 2. The same set of relocatees was found by stage 3 regardless of the stage 2 solution used. This result is not atypical of other examples run and indicates a robustness of the model. The least cost assignment resulting from stage 3 is indicated by dotted arrows in Fig. 4. The least travel distance solution produced by stage 4 is, of course, a permutation on the least cost solution and is shown in Fig. 4 by solid arrows. The solution produced by stage 4 results in a reduction of 22 percent in travel distance and only a 9 percent increase in the cost function.

We also solved this problem using integer programming.[4] The result obtained was identical to the minimum cost solution found by the heuristic algorithm. However, the heuristic was four times faster and required only one-half the amount of computer core storage.

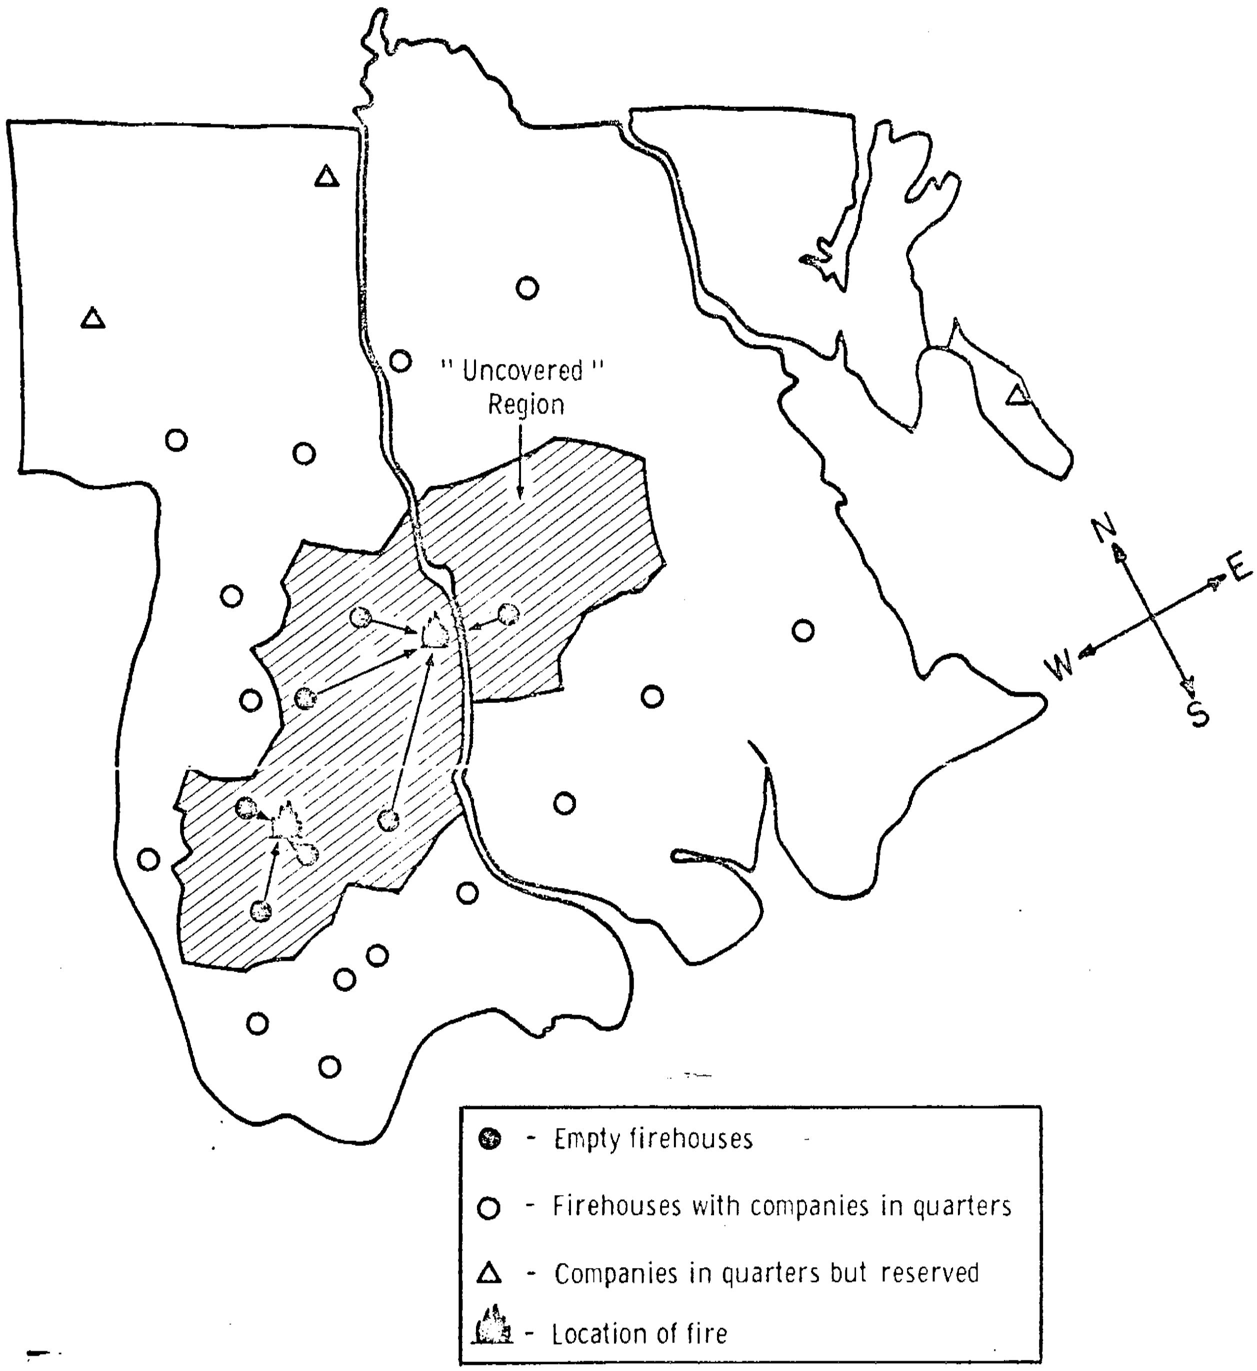

Fig. 3 - A sample relocation problem

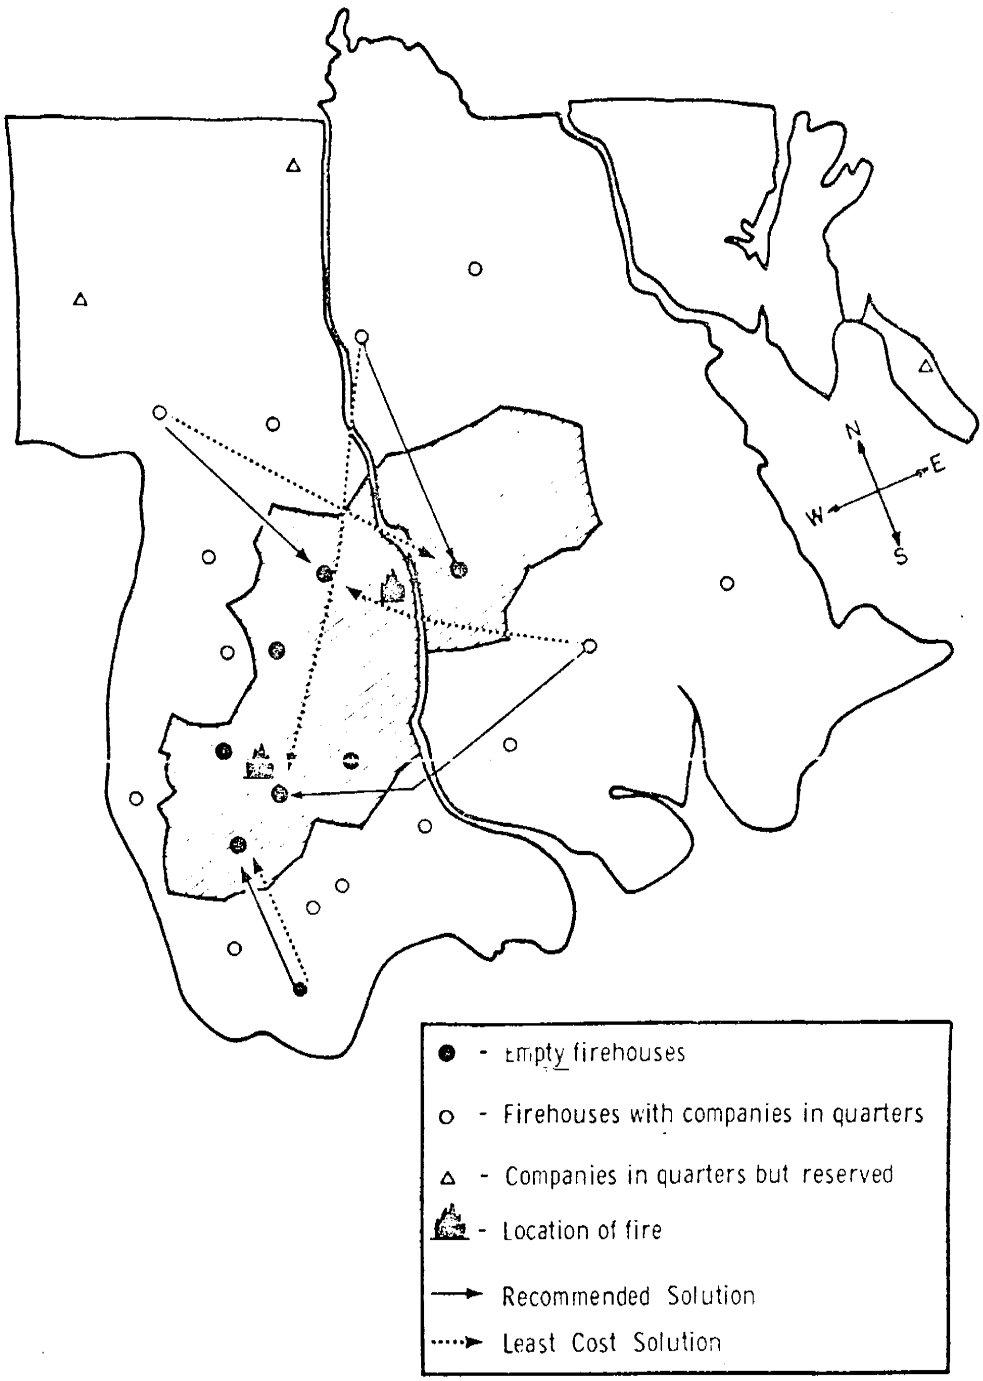

Fig. 4 - Solutions to the sample relocation problem

VI. TESTING THE ALGORITHM: JULY 4, 1969

The algorithm has been tested extensively, first with problems we manufactured--designed to present difficult or interesting situations--second in a simulation model in which over 3600 alarms were generated at random according to historical patterns, and finally, to provide a strenuous realistic test, we chose one of the worst evenings ever experienced in the Bronx, reconstructed the sequence of incidents, and simulated what would have occurred if the relocation algorithm had been operating.

The evening chosen was July 4, 1969. Typically, July 4th is a very busy day for the Fire Department; a large number of false alarms are received and many fires are fought which have been purposely set. This particular July 4th was one of the worst nights ever experienced in the Bronx. There were 288 alarms turned in on that day in the borough, almost twice as many alarms as during a normal day. To make matters even worse, these alarms did not occur uniformly during the day; over 40 percent occurred in the four-hour period from 8:00 p.m. to midnight.

The alarms which were received during the period 8:00 p.m. July 4 to 3:00 a.m. July 5 are summarized in Table 1. In this seven-hour period an average of 24 alarms per hour were received. There were 33 structural fires, including 2 three-alarm fires. In addition to the multiple alarm fires, there were 5 other very serious fires which occurred during this period. A summary of these 7 serious fires is given in Table 2. Five of the fires broke out within an hour of each other (11:38 p.m. to 12:38 a.m.). The number of fire companies needed to put out the serious fires was almost equal to the total number of companies stationed in the Bronx.

Not only did the fires occur closely in time, but they were grouped geographically in the south (see Fig. 5). Thus, it is clear that the relocation of a significant number of units into the South Bronx was required in order to maintain adequate coverage.

Table 1

ALARMS RECEIVED IN THE BRONX
8 P.M. JULY 4, 1969-3 A.M. JULY 5, 1969

Time	Number of Alarms			
	Total	Structural	Serious	False
8-9 p.m.	35	4		8
9-10	31	2		14
10-11	38	6	1	16
11-12	30	6	3	11
12-1 a.m.	7	4	2	1
1-2	14	8	1	1
2-3	12	3	—	5
Total	167	33	7	56

Table 2

SERIOUS FIRES IN THE BRONX
JULY 4-5, 1969

Time	Box No.	Number of Working Units	
		Engines	Ladders
10:45 p.m.	2791	3	2
11:38	4789	8	3
11:56	2732	10	5
11:59	3131	2	2
12:27 a.m.	2240	2	2
12:35	2550	7	4
1:12	2916	3	2
2.5 hours		35	20

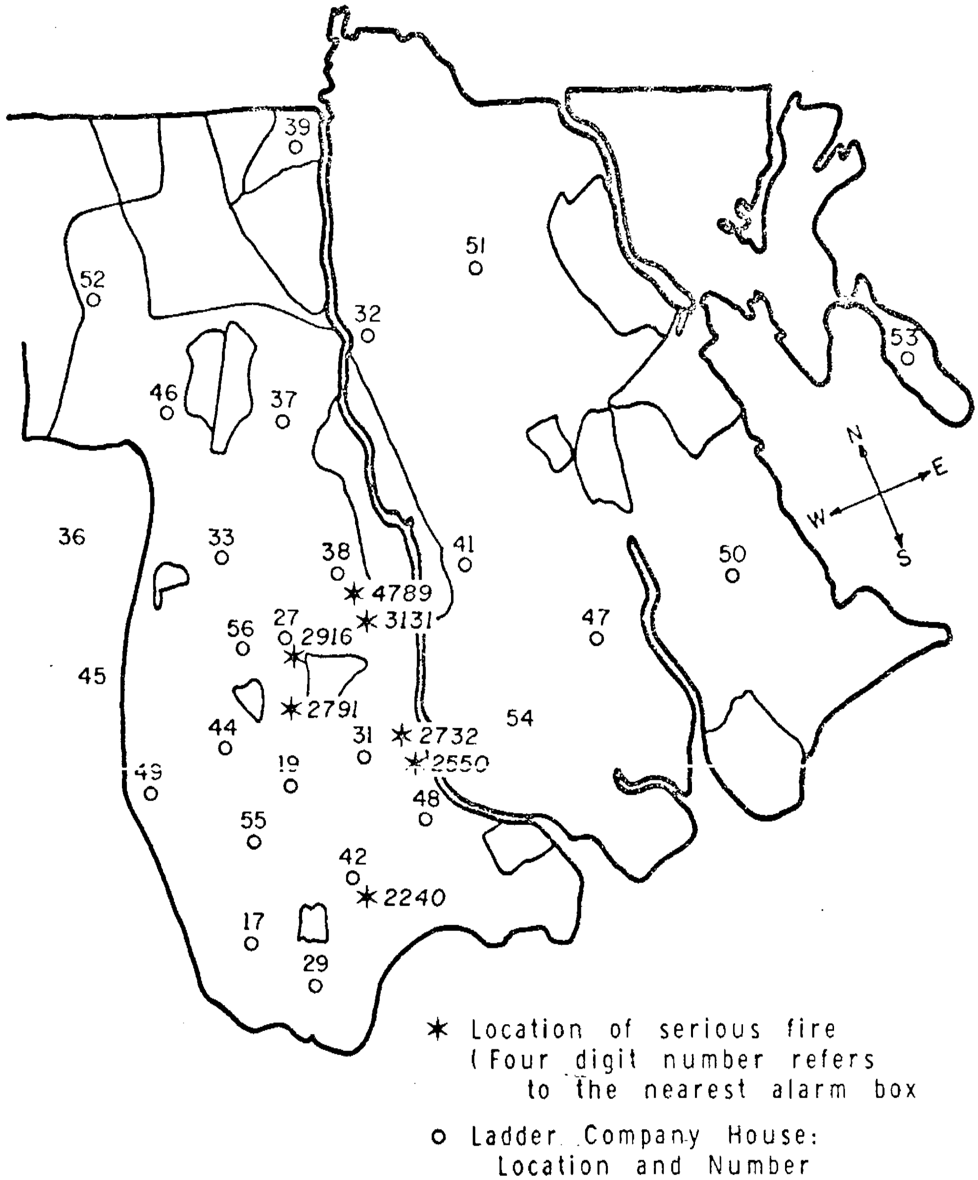

Fig. 5 - Location of serious fires in the Bronx, July 4-5, 1969

The manual alarm assignment card relocation system in current use was unable to effectively handle this July 4th situation. Under this system every alarm box has a card associated with it which, among other things, specifies the relocations to be made if there is a serious fire at that box. Such a system cannot take into account other alarms in progress which involve the companies specified to be relocated by the card. In addition, the relocations for two or more different alarms, when taken together, could cause the neighborhoods shared by these relocating units to become uncovered.

These problems and others were actually experienced on July 4, 1969. Had it been operating then, the computerized relocation algorithm, which determines relocations based on the status of the whole system at any given time, would have significantly out-performed the alarm assignment card system. We demonstrated this by recreating the events of July 4, 1969 using a computer simulation model of the Bronx. The simulation had been previously developed by The New York City-Rand Institute for the study of alternative deployment strategies for New York City.(3)

The period simulated lasted from 9:52 p.m. on July 4th until 4:00 a.m. on July 5th. During this period, which begins well before any serious alarms were received and lasts until well after they were ended, a total of 107 alarms were received. The simulation results were compared to the actual performance of the system during that period as reconstructed from Fire Department records, including Battalion Chiefs' reports, Communications Office logs and a memorandum to the Chief of the Fire Department from a Deputy Assistant Chief reviewing the activities of that evening in the Bronx. Some of the more interesting comparisons between the simulated and actual performance are described below. Before discussing the results we must remark, in fairness to the dispatchers who worked that night, that they were using a system which was not designed to handle the volume or complexity of alarms actually received. Our comments are reflections on the manual system rather than on the men who used it.

NUMBER OF RELOCATIONS MADE

The algorithm would have made a total of 35 relocations; in reality, 25 relocations were made. Five relocations listed on the alarm assignment cards could not be made because the units specified to be relocated were not available. Of the 25 relocations made, 10 were specified on the alarm assignment cards. The other 15 were made by the dispatchers because they felt that coverage in the South Bronx was bad and additional companies were needed to provide protection. The alarm assignment cards afforded them no help in determining how many extra units were needed, when the relocations should be made, or which companies should be moved.

TIMING OF THE RELOCATIONS

The algorithm generated its relocations gradually and continually over time while the relocations made by the dispatchers were generally made in spurts. Thus, although by 1:08 a.m. each method had made a total of 23 relocations, the algorithm had specified relocations to be made at 16 separate times, while the dispatchers had made their relocations at only 5 different points in time. The dispatchers made 7 relocations at 11:45 p.m. and 10 relocations (only one of which was specified on the alarm assignment cards) at 12:49 a.m.

COVERAGE

Under the definition of minimum coverage used as the primary objective in the relocation algorithm, every alarm box will always have at least one of its closest three engines and at least one of its closest two ladders "available." "Available" in this case includes responding to an alarm, returning from an alarm, due back from a working fire soon or actually in quarters. Much stronger measures of the coverage being provided to an area are the percentage of alarm boxes in the area for which at least one of the closest three engine companies is available in quarters and the percentage of alarm boxes with at least one of the two closest ladders available in quarters. We will use these measures in comparing the coverage actually provided on July 4, 1969 to the coverage which the algorithm would have provided.

During the first two hours of the scenario the simulation and the actual system performed comparably. The simulation made 10 relocations compared to 9 in the real system and the coverage provided at 11:52 p.m. is given in Table 3.

Table 3

COVERAGE AT 11:52 P.M. - JULY 4, 1969

	11:52 p.m.	
	Engine Coverage	Ladder Coverage
Simulation	98%	100%
Actual	97.5%	88%

After midnight, the actual situation began to deteriorate. Because of the large number of alarms in progress, the alarm assignment cards became less and less useful. It became increasingly difficult for the dispatchers to manage both their dispatching and relocation functions, so relocations began to suffer. The algorithm, however, was able to keep up with the situation and to maintain a high level of coverage. Table 4 below shows the coverage levels (where coverage is defined as above) at 12:54 a.m. and 1:36 a.m.

Table 4

COVERAGE AT 12:54 A.M. AND 1:36 A.M. - JULY 5, 1969

	12:54 a.m.		1:36 a.m.	
	Engines	Ladders	Engines	Ladders
Simulation	88%	100%	94%	90%
Actual	78%	88%	82%	79%

The algorithm never leaves a response neighborhood without minimum coverage. However, on that night, a total of 16 RN's were actually left uncovered for periods ranging from 30 minutes to 1.6 hours. Figure 6 shows the status of the ladder houses in the Bronx as it actually was at 12:54 a.m. on July 5, 1969. The shaded areas represent the uncovered response neighborhoods. There were 4 uncovered RN's at that time. These can be labelled (44, 49), (32,41), (32, 37), and (33, 37). The RN labelled (44, 49) was left uncovered after both ladder companies were dispatched to the third alarm fire at box 2732 (neither company appeared on the alarm assignment card) and no relocations were made to cover it. This RN was left uncovered for a total of 1 hour and 38 minutes.

Figure 7 shows the status of the ladder houses in the Bronx at 12:54 a.m. in the simulation. There are no uncovered response neighborhoods. It should be noted that the simulation has two more ladder companies in the Bronx at this time than were there in the actual case. These are Manhattan ladders which the simulation had relocated into the Bronx. Eventually, the dispatchers in the Bronx moved the same number of Manhattan ladders into their borough as the simulation did, but they were moved significantly later.

RESPONSE TIMES

Response times to actual alarms and area coverage are closely related, but not equivalent, measures of effectiveness. Coverage is a geographical measure while response times are associated with actual responses to incidents. While we have no data on the actual response times to the incidents of July 4, 1969, since response times are not recorded by the Fire Department, we can look at the units which were actually dispatched to the incidents and see how close they were to the alarm box to which they were sent. We compare these to the closeness of the units dispatched in the simulation. Our measure of closeness is the position on the alarm assignment card of the house from which the unit responded, since the houses are generally listed on the card in order of closeness to the box.

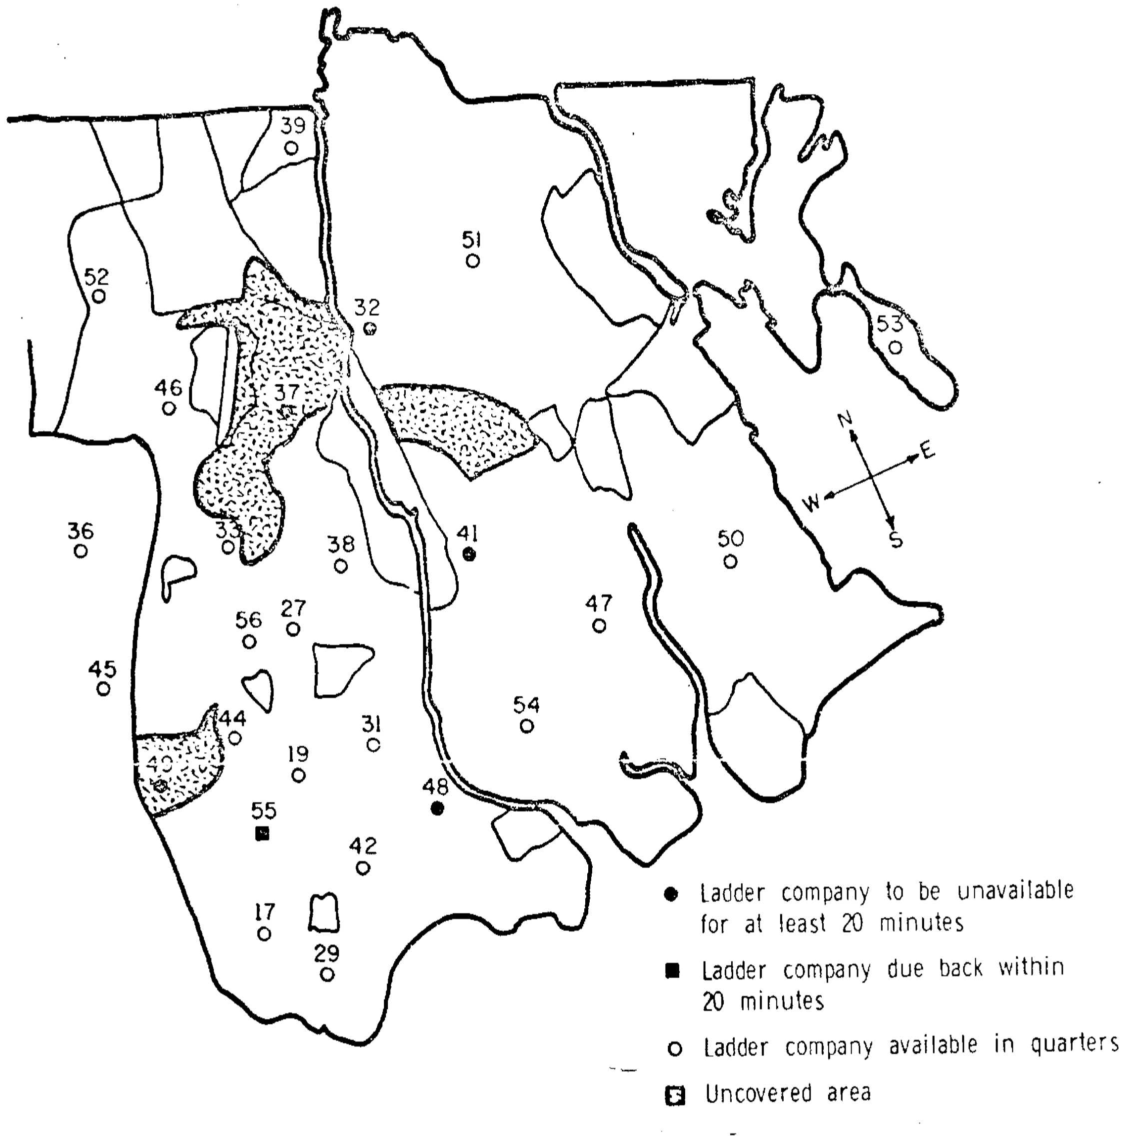

Fig. 6 - Actual Bronx ladder company status, July 5, 1969, 12:54 a.m.

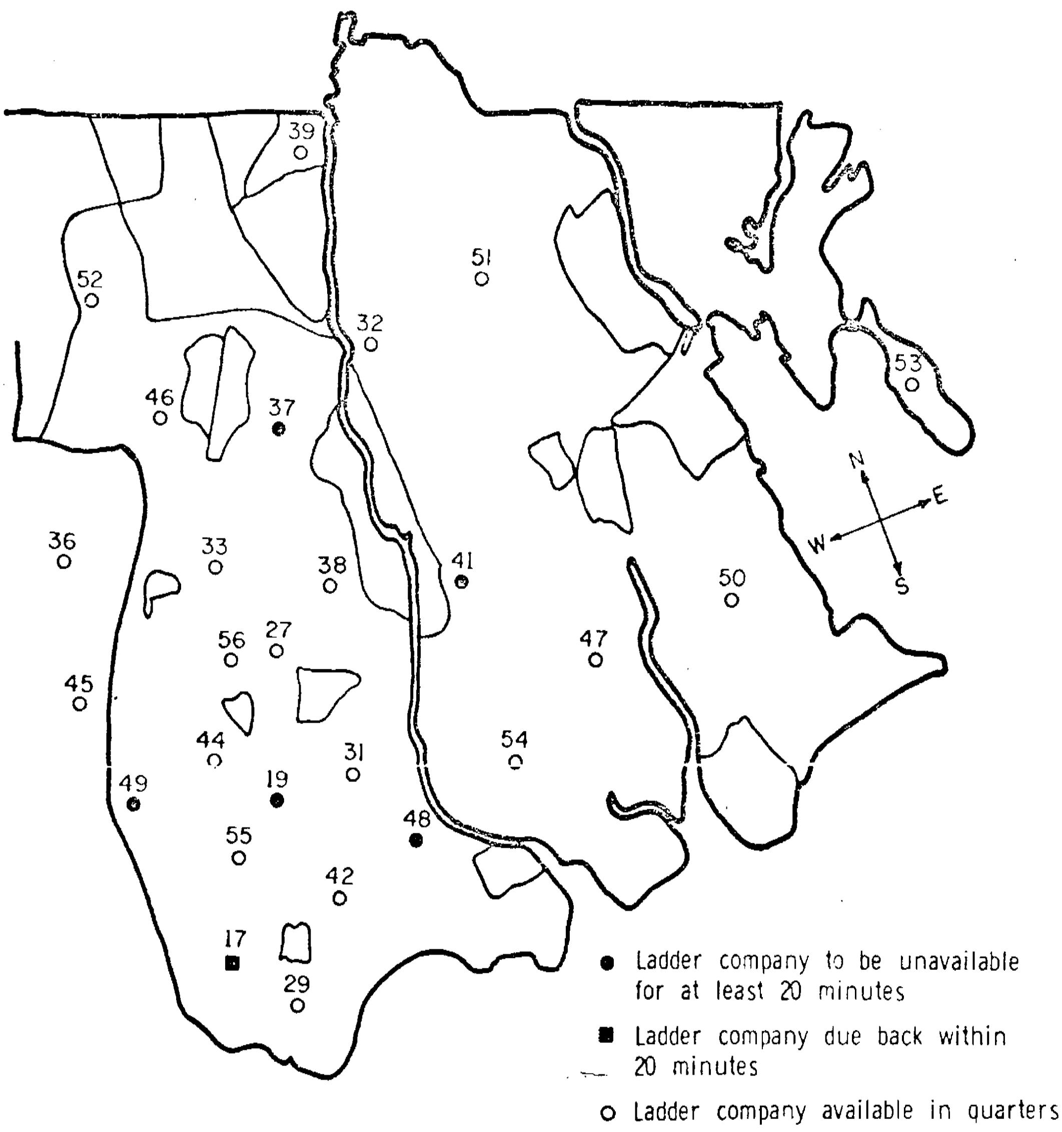

Fig. 7 - Simulated Bronx ladder company status, July 5, 1969, 12:54 a.m.

In making these comparisons we found that using the algorithm we almost never had to dispatch farther units than did the existing system. Often we were able to dispatch closer units. In one case the effect was dramatic. For the third alarm fire at box 2732 the existing system dispatched the fifth closest ladder while using the algorithm we were able to dispatch the closest ladder and hence we could have reduced the response time to this fire by about two minutes.

VII. CURRENT IMPLEMENTATION OF THE ALGORITHM

The relocation algorithm was designed to be part of a large computer-based Management Information and Control System to be implemented by the FDNY. In that system, relocation will be one of several deployment functions which will be handled together with keeping system status, updating statistical records, etc. We have found that, although the relocation algorithm was conceived as part of a large total systems package, it can function well independently and can be implemented prior to the rest of the system. We have developed a version of the algorithm which operates in an interactive mode on a small time-shared computer and is currently in the process of being implemented on a test basis in New York City--running in parallel with the existing manual system.* The program is small and rapid. In order to give an idea of how the algorithm can be used and its flexibility, we describe some of the ways in which the dispatcher is able to interact with it.

The dispatcher, who is informed of all changes company status via radio or telephone, types in each change as it occurs. The status of a company is indicated by a one-letter code, e.g., Q for "in quarters," R for "responding," S for "working at a serious fire," etc. The user acts as the trigger, asking the program to determine whether a relocation problem exists. The program responds with the number of empty houses and a list of uncovered response neighborhoods if there are any. It will generate a relocation recommendation if requested.

At this point, the dispatcher has several options available. He can implement the suggested relocation by typing in the recommended set of moves. He can ignore the suggestion and make no relocations. The algorithm will simply remind him again if any response neighborhoods are uncovered after the next status changes. There are also several other possibilities:

(a) The dispatcher can modify the suggested relocation in whole or in part based on considerations that he may be aware of but that are not included in the model. For example, he may know that a company which the computer suggested

*This program was developed by Carol Shanesy. For a description of the algorithm and an instruction manual on its use see [6].

as a relocatee has itself just returned from a very arduous fire. If he modifies part of the suggestion, he can then ask the computer to re-solve the problem, taking this change into account.

(b) The dispatcher may temporarily exclude one or more of the recommended relocatees and ask the computer to re-solve the problem without them.

(c) To help him make modifications in the recommended location, the dispatcher can ask that the computer provide him with a ranked list of desirable and available candidates for relocation into any given house.

(d) The dispatcher can delay making the relocation. He might do this if he thinks or has information that an incident will soon escalate. There are several reasons for wanting to delay a relocation decision, including the fact that some companies which are selected to be relocatees might soon be needed at the incident, and the fact that fewer relocations will generally be needed if relocations are held up until all higher alarms are in for an incident. If the dispatcher thinks that the incident might escalate, he can ask the computer to "look ahead" and generate the relocation which would be suggested if the incident did escalate, and then implement those moves which are needed at present.

The algorithm does not make any "unrelocation" suggestions. Under the current New York City Fire Department system, a company remains in its relocated quarters until the company which belongs there returns from fighting its fire. Thus, no dispatching decision need be made to unrelocate a company. However, if a neighboring company were to return to quarters first, the relocated company may be able to be unrelocated without uncovering an RN. In fact, his presence in another area might be more important for coverage.

The algorithm permits any unrelocation policy to be followed by simply informing the dispatcher whenever an RN is "over-covered" and this over-coverage involves at least one relocated unit. (An RN is considered over-covered whenever the movement of a unit out of a response neighborhood will not uncover it.) The deployment decision is then left to the dispatcher's discretion.

We associate with each working unit a status based on the type of incident at which he is working; either S for a serious incident or W for a non-serious incident. Non-serious incidents are those which generally last 60 minutes or less. These constitute about 95 percent of all alarms. Incidents expected to last over 60 minutes are classified as serious and are usually easy for the fire chief at the fire to identify. We define an RN as being uncovered if all of the companies comprising the RN are currently unavailable and are working at serious fires. If the algorithm suggests filling a house whose unit has been at a serious fire for a while but is due back shortly, the dispatcher will generally detect this and delete the relocation.

REFERENCES

1. Balinski, M., "Integer Programming: Methods, Uses, Computation," Management Science, Vol. 12, No. 3, 1965, pp. 253-313.

2. Balinski, M., and R. Gomory, "A Primal Method for the Assignment and Transportation Problems," Management Science, Vol. 10, No. 3, 1964, pp. 578-593.

3. 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, in the IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-6, No. 4, pp. 282-293, October 1970.

4. Geoffrion, A., "Implicit Enumeration Using an Imbedded Linear Program," The Rand Corporation, RM-5406-PR, September 1967.

5. Kolesar, P., and E. Blum, "Square Root Laws for Emergency Service Response Distances: Applications to Fire Service Deployment," The New York City-Rand Institute, R-895, Forthcoming.

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

7. Toregas, C., R. Swain, C. ReVelle, and I. Bergman, "The Location of Emergency Service Facilities," Operations Research, Vol. 19, No. 6, pp. 1363-1373, October 1971.