IMPROVING THE DEPLOYMENT OF NEW YORK CITY FIRE COMPANIES

Edward J. Ignall
Peter Kolesar
Arthur J. Swersey
Warren E. Walker
Edward H. Blum
Grace Carter
Homer E. Bishop

July 1974

P-5280

IMPROVING THE DEPLOYMENT OF NEW YORK CITY FIRE COMPANIES

Edward J. Ignall, Peter Kolesar, Arthur J. Swersey,
Warren E. Walker, Edward H. Blum, and Grace Carter
The New York City-Rand Institute
New York, New York 10022

and

Homer Bishop, Deputy Assistant Chief
The New York City Fire Department
110 Church Street
New York, New York 10007

July 26, 1974

This paper was prepared for presentation at the 45th Joint National Meeting of the Operations Research Society of America and the Institute of Management Sciences which was held in Boston, Massachusetts on April 22-24, 1974.

ABSTRACT

How many fire companies does New York City need and where should they be located? Given a fire alarm of unknown severity, how many companies should be dispatched to it? These two questions are fundamental issues in the deployment of the City's fire-fighting resources.

Since 1968, the New York City Fire Department and The New York City-Rand Institute have carried out a joint project to improve the delivery of Fire Department services in the face of skyrocketing demand. In November 1972, two historic deployment changes were implemented: (a) six of the 375 fire companies in the City were disbanded and seven other companies were permanently relocated; and (b) in high fire incidence areas of the City, an adaptive response policy was implemented. Under adaptive response, fewer companies are initially dispatched to potentially less serious alarms. This is in contrast to the traditional dispatching policy where the same number of companies are dispatched to each alarm.

The joint Fire Department-Rand Institute project and the analyses which led to these and other improvements and the wide range of mathematical models used are described. The changes have resulted in savings to the Fire Department of over $4 million per year, a reduction in the workload of fire companies and a more equitable distribution of fire companies throughout the City.

ACKNOWLEDGMENTS

The success of this project has depended on the knowledge and support of many people, both in the Fire Department of New York and The New York City-Rand Institute.

In particular, we would like to acknowledge the support of John T. O'Hagan, Fire Commissioner and Chief of Department, whose faith and confidence in the work encouraged its completion and speeded its implementation. Several Planning Officers in the Division of Planning and Operations Research were involved directly in the joint project. Particularly important contributions were made by Deputy Chiefs Francis Ronan, Robert Brown, and Joseph Hess.

Rae Archibald made substantial contributions to the work, both while he was leader of the project at the Institute and as a Deputy Commissioner in the Fire Department.

Jan M. Chaiken and John E. Rolph of The Rand Corporation and Carol Shanesy of IBM played key roles on the project while they were working at the Institute.

I. INTRODUCTION AND BACKGROUND

New York City has the largest fire department in the world with over 13,000 men in uniform and an operating budget for the fiscal year 1974-75 of 375 million dollars--larger than the entire operating budget of the City of Boston. New York City also has the most severe fire problems in the world--more incidents, more real fires, and more false alarms.

In 1968, when a joint relationship between The New York City-Rand Institute and the New York City Fire Department was begun, the Department was faced with unprecedented demands for its services. Annual fire alarms had increased from well under 100,000 in 1958 to almost 250,000 in 1968. At 8 o'clock on a summer evening in the south Bronx, one of the busiest areas of the City, typically about half the fire companies (a company consists of a pumper or ladder truck and the men assigned to it) were unavailable because they were already busy. When a fire alarm came in, only two or three of the five units assigned were available to respond.

Major fires were a nightly occurrence, often several burning at once. Communications channels and dispatching centers were clogged with alarms--it was not unusual during peak evenings for it to take five minutes or more to dispatch units to an alarm. The number of fire company responses to alarms was increasing and the men in the busiest companies were being overworked. Increasingly militant unions were demanding more fire companies and getting the City to add them, each at an operating cost of $600,000 per year. In short, skyrocketing alarms were straining the system to the breaking point.

In the previous year, 1967, Fire Commissioner Robert O. Lowery and Chief of Department John T. O'Hagan, recognizing these problems, appealed to the Mayor for some analytical assistance. In 1968 when Mayor Lindsay and his budget director Frederick O'R. Hayes asked The Rand Corporation to come to New York to work on the City's problems, the Fire Department was one of the four agencies chosen for initial work.

Since that time, a joint project has been underway--an operations research-systems analysis team from The New York City Rand Institute has been working with officers of the New York City Fire Department in an effort to improve fire protection in the City. This joint effort has covered many aspects of fire protection including prevention and new technology for extinguishing fires.

In this paper we describe one major part of this work--the effort to improve the deployment of New York City fire companies. Many changes have resulted from this effort. We focus here on two important changes which took place in November 1972: (1) the Fire Department instituted a major redeployment of its resources--disbanding six fire companies and permanently relocating seven others, and (2) in areas of high demand, the basic response policy to alarms was changed. The new response policy, called adaptive response, tailors the number of units dispatched to a prediction of the seriousness of each incoming alarm.

These changes have resulted in significant improvements in fire protection, a reduction in fire company workload and savings to the City of well over $5 million per year. Other changes resulting from the New York fire project and their impacts can be found in [2, 20].*

The rest of this paper describes the deployment problem, these two major deployment changes, and the technical work and analysis that led to them.

The Deployment Problem

The basic deployment problem is to have fire-fighters and their equipment in the right place at the right time. To do this economically in New York City is an enormous undertaking. There are 375 fire companies in New York City. The location of all but the Staten Island companies are shown in Fig. 1. The map in Fig. 2 shows that fire alarms are concentrated in areas such as the south Bronx, central Brooklyn, Harlem and the lower East Side. The dark shaded areas on the map are regions having alarm rates of over 6,000 alarms per square mile per year--almost 1 alarm per hour 24 hours a day. Since alarm rates peak in the evening and in the summer, at busy times the companies in these areas are continuously in action.

The overall deployment problem the Department faces can be broken down into three subproblems: allocation, initial dispatch, and relocation. These subproblems are interrelated but the overall problem can be better

---

*Figures in square brackets identify references located at the end of this document.

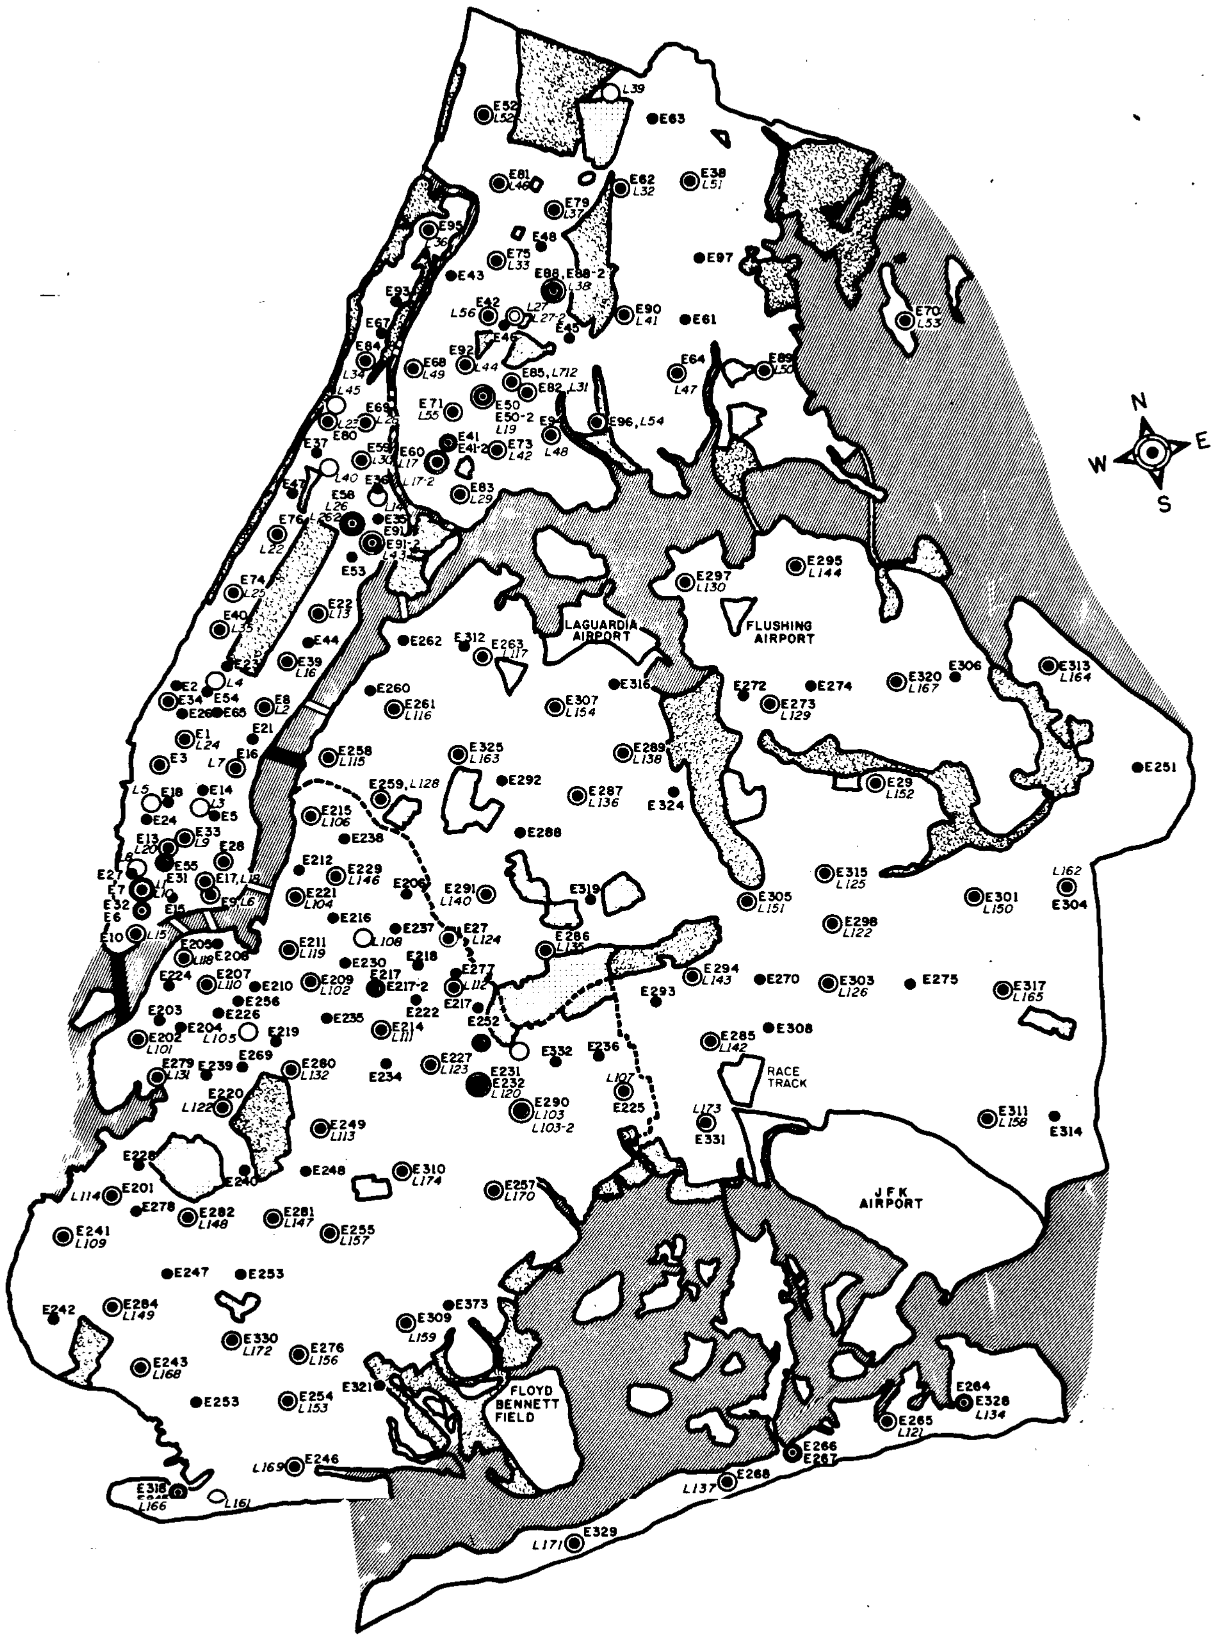

Fig. 1. New York City fire company locations

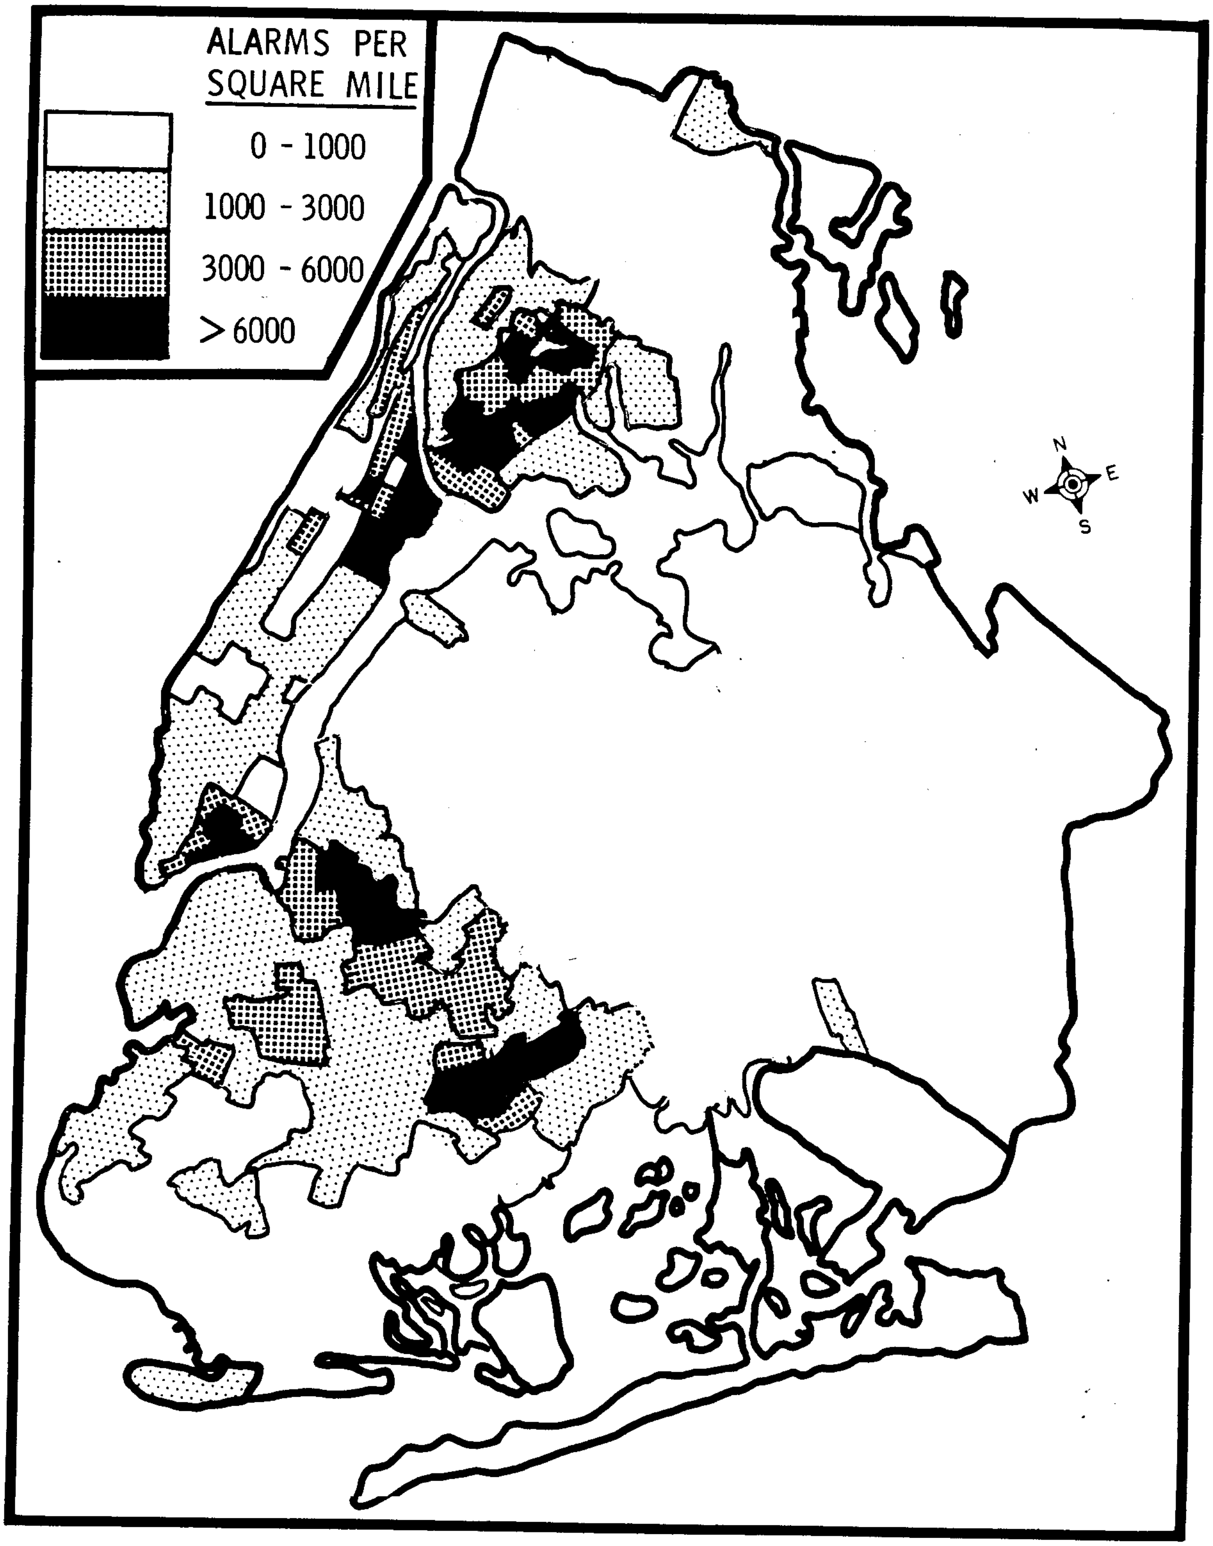

Fig. 2. Geographic distribution of fire alarms

understood and much useful analysis can be carried out by considering them separately.

- The allocation problem is: how many companies are needed and how should they be allocated to regions of the City? Clearly, more units are needed in the south Bronx, an area with many fires and poor housing, than in eastern Queens--a residential area. But how many more--and how many total units are needed in the City--370?--250?--500? This is a long-range planning problem. We are speaking of "permanent" allocations--decisions which often take a long time and large amounts of capital to implement.
- The dispatch problem is tactical and recurrent: how many and which units should be sent to each incoming alarm? When an alarm is reported by telephone, the dispatcher can usually tell from the caller's description of the fire how many units are required. But, most alarms are reported from the telegraph boxes on the street, and then the dispatcher knows only that someone has pulled the handle of the box; it could be a major fire or a false alarm.
- The relocation problem is also tactical but less frequent: how should the Department change temporarily its allocation of companies in reaction to serious fire situations? When one or more serious fires break out in an area, the companies working at these fires leave gaps in coverage. How should units be reallocated temporarily to fill these gaps in coverage? For example, a serious fire in a Times Square Hotel could tie up half of the companies in Manhattan for many hours.

The Department deploys two basic types of fire-fighting units: engine companies and ladder companies. An engine company (5 to 7 men manning a pumper) hooks up to the fire hydrant and delivers water onto the fire. A ladder company (with about the same complement of men manning a truck equipped with an aerial ladder) is responsible for life-saving. Its men carry axes and special tools which they can use, if needed, to force their way into a building and break windows to allow the hot gases to escape. They are also specially equipped and trained to perform their life-saving operations.

The role of the fire department is to protect lives and property by preventing and extinguishing fires. The impact which changes in deployment will have on fire protection is difficult to measure directly. How many lives could be saved or how much less property lost if 6 fire companies were added to the force? There is really no way to know. We have used the response times of fire companies to fires as an operational measure of effectiveness, which relates well to the Department's objective of reducing loss of life and property. We presume that, if a new policy results in shorter response times, it will also result in fewer lost lives and less property damage, even though we do not know the magnitude of the effects.

In providing fire protection, the Department is constrained by the funds budgeted for men and equipment and by restrictions, often political, on how they can be used. For example, New York City firemen's work hours, like those of most workers, are specified in their contract. In addition, their work hours are specified in the New York State Constitution. Consequently, changes require legislative action in addition to the usual labor-management agreements.

Traditional Solutions

As the alarm rate increases (see Fig. 3), the City's fire protection gets worse since more and more units are busy and those available respond from further away. In addition, fire company workload goes up--with each unit making more responses to alarms. As a result, traditional ways of deploying units no longer work well. Two examples illustrate how apparently sensible deployment tactics break down at high alarm rates.

One traditional approach to meeting increasing demands was to add fire companies. In 1967, one company was added; in 1968, nine companies were added, each at a cost of about $600,000 per year. Not only was this solution expensive, but it also did not work very well. One of the reasons

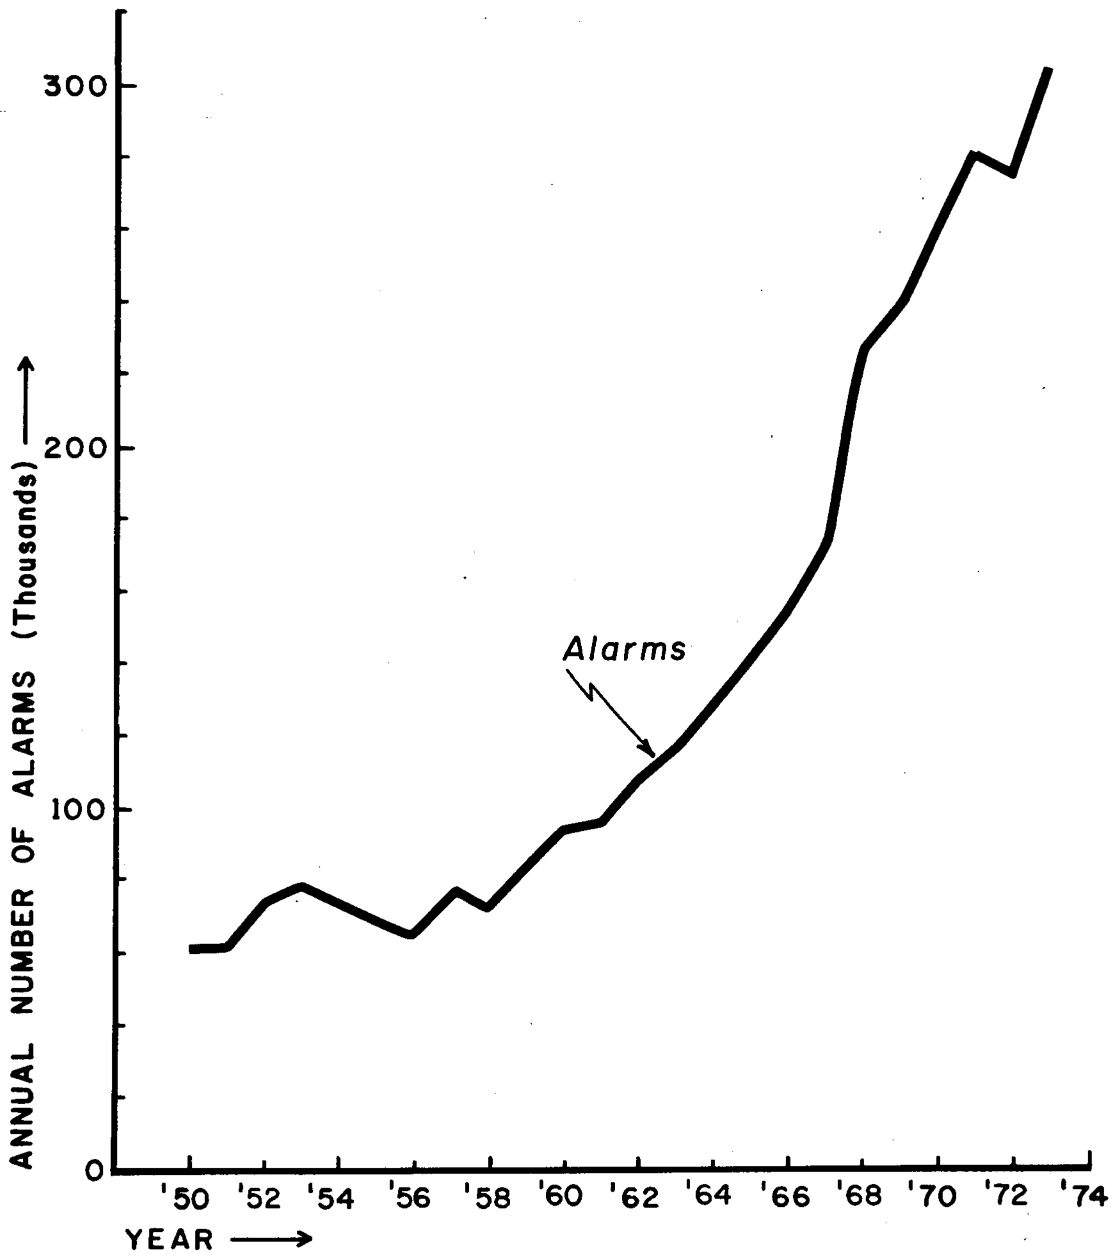

Fig. 3. Annual fire alarms, 1950-1974

for adding new companies was to relieve workload in the busiest regions. Take an example. In 1966, Engine Company 82 in the south Bronx responded to over 6,000 alarms and was the busiest engine company in the City, and, for that matter, in the world! To relieve its workload, the Fire Department, in July of 1967, created a new company, Engine 85, and put it into the same fire house as Engine 82. Its first full year of operation was in 1968. It was expected that Engine 82's workload would be cut in half. In fact, in 1968 Engine 82 was still the busiest company in the City and Engine 85 was the second busiest! (See Table 1.) Instead of helping the busy unit, there were now two busy units.

Adding Engine 85 did not relieve the workload of Engine 82 because of the unexpected interplay of two seemingly logical dispatching rules which together comprised the traditional dispatching policy: (1) always dispatch the units closest to an alarm, and (2) send three engine companies and two ladder companies--if they are available, but send at least one engine and one ladder. Our analysis showed the Department that, since, in busy periods, many companies are not available, as little as one engine company and one ladder company are often sent. Then, our simulation model showed that the new company was drawn into the role of "filling out" the response to alarms, primarily providing an unneeded second or third engine at trash fires and false alarms. As a result of the dispatch policy, instead of the same number of responses being spread over more companies, the total number of responses in the area was increased. Figure 4 illustrates the "filling out" effect: if Engine 85 did not exist, only two engines would be sent to the indicated alarm; with Engine 85, three engines are sent.

The traditional relocation policy for filling temporary gaps in coverage also breaks down under busy conditions. Consider the example in Fig. 5. A large fire breaks out at the indicated location and the five closest engine companies are sent to work at the incident. The traditional relocation policy consists of advance plans which, in this case, designate two companies to relocate (indicated in the figure by the dotted lines). In the example, which is quite typical of actual situations at high alarm rates, one of the companies designated is not available so that the gap which should have been filled is not filled; moving the other company, although it is available, would be a mistake since it would create a new gap in coverage.

Table 1

WORKLOAD OF ENGINE COMPANIES 82 AND 85

Year	Engine Company	Responses	Rank
1966	82	6234	1
1968	82	9111	1
	85	8386	2

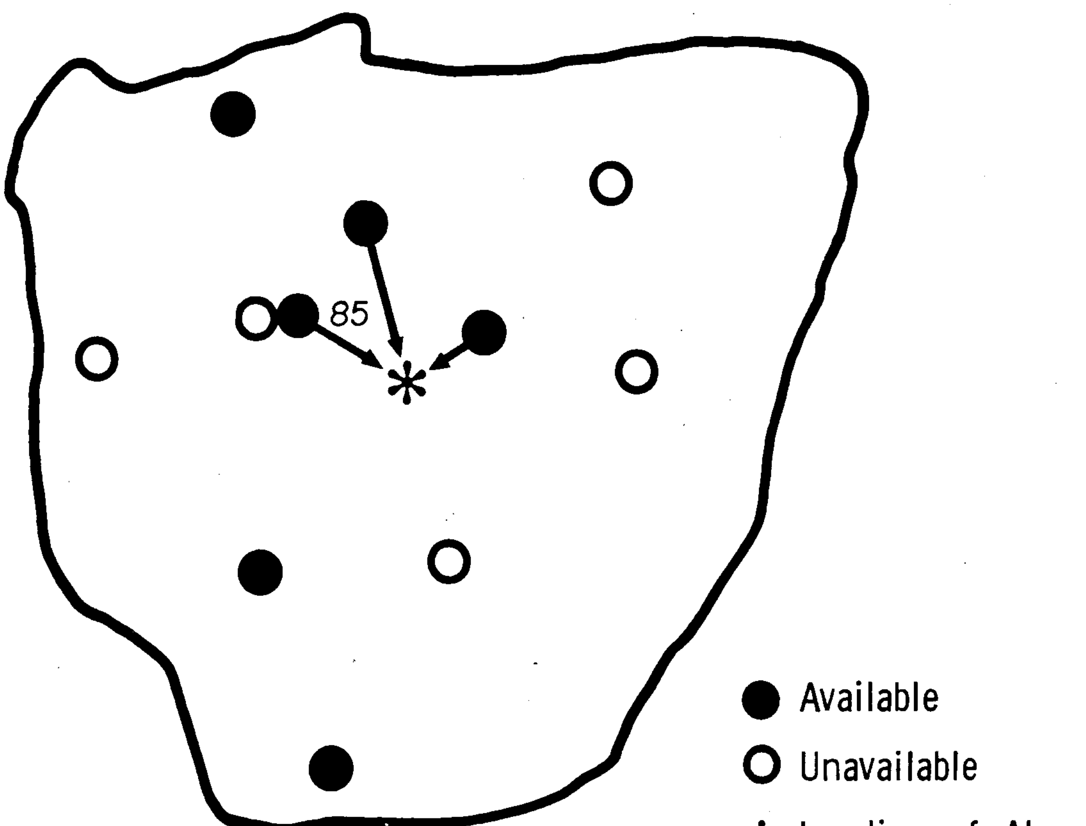

Fig. 4. The "filling out" effect

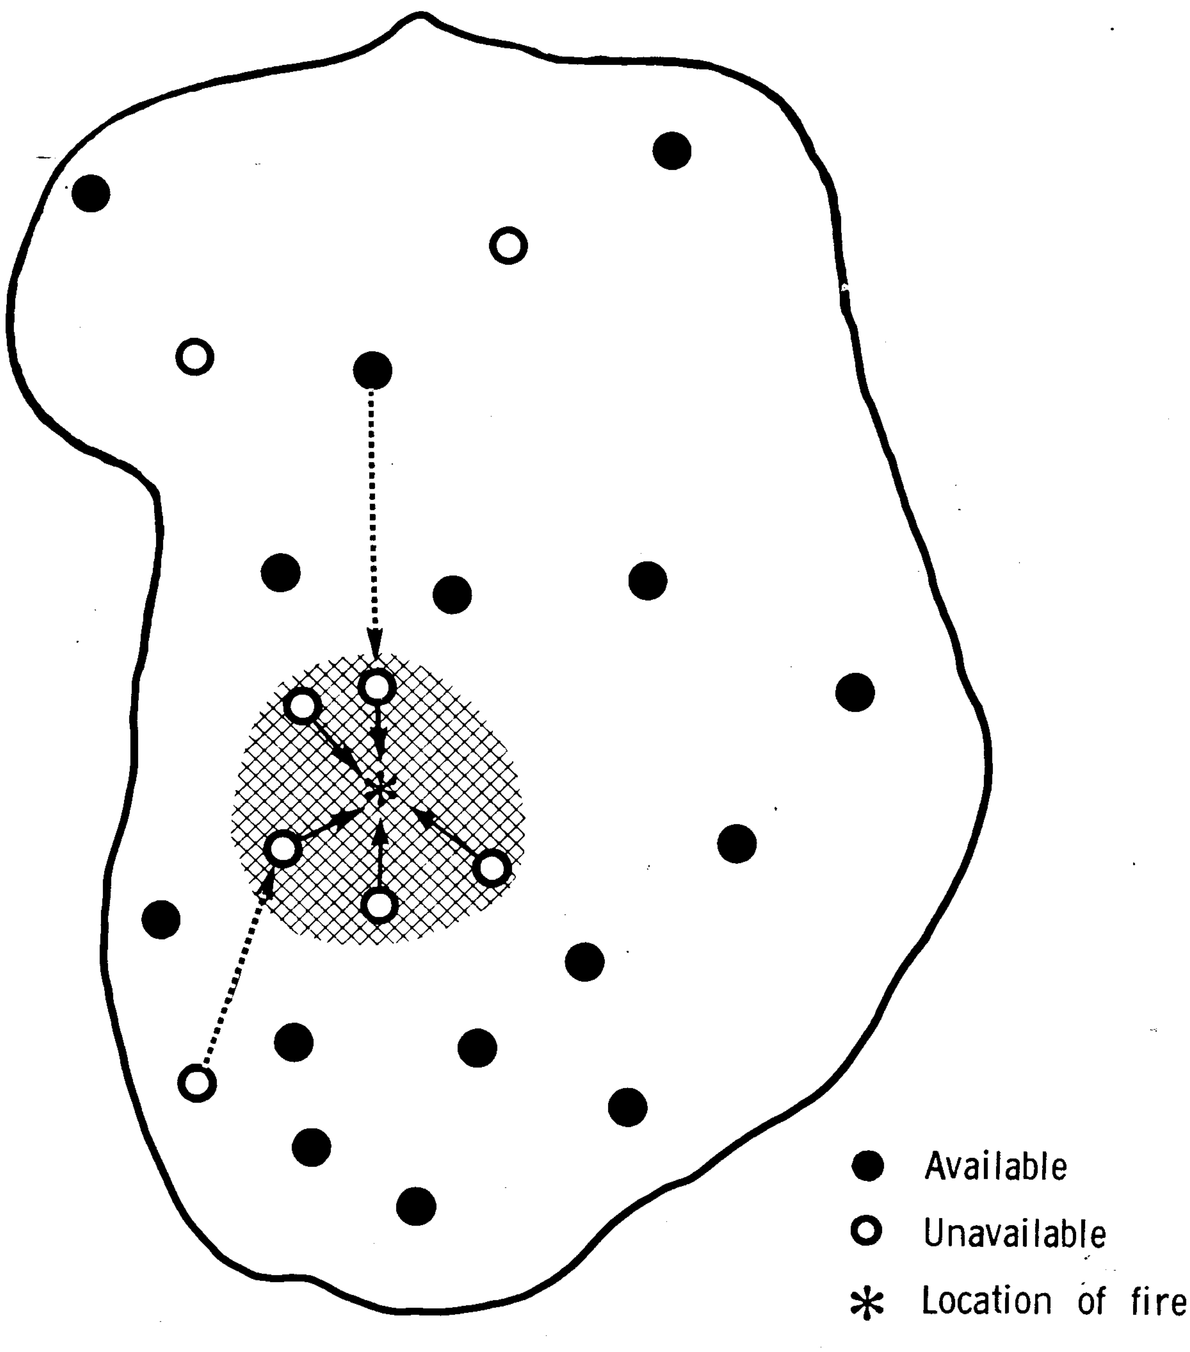

Fig. 5. Problems with the traditional relocation policy

II. FINDING BETTER SOLUTIONS

Given the inadequacy of traditional approaches, we began to look for new solutions aimed at improving fire protection and reducing workload. First, we observed first-hand what was going on in the field, by riding with fire companies and observing operations at fires.

The Fire Department had been keypunching data on every fire alarm which occurred in the City since 1962. So, when we started work in 1968, we were in the enviable position of having information on over 1 million incidents in computer-readable form--records which contained the time, place, number of companies used, and so forth. Using these data, we developed models of time of day, seasonal and geographic alarm patterns, and an understanding of the demands faced by the Department.

Then, using both the data analysis and what we had learned in the field, we built a detailed simulation model of fire-fighting operations in the Bronx [5, 7]. We tried to make the SIMSCRIPT simulation as close to reality as we could. Later, it was replaced for many applications by simpler analytic models.

At first, we used the simulation primarily as a laboratory to test and compare specific alternative deployment policies, some of which were suggested by the Department. Later, other policies were suggested by analytical models we develpped. The simulation also played an important role in the development and testing of the analytic models. One such use is given later when we discuss the square root models of response distances.

Adaptive Response

One major deployment change was the Fire Department's implementation of the adaptive response policy in November 1972. The traditional dispatching policy had been to send five units: three engine companies and two ladder companies. Does this policy make sense? A serious fire requires all five units to put it out. In fact, more units might be needed, but their arrival times (within limits) are not crucial. In fact, it is probably better not to have more than three engine companies and two ladder companies arrive initially so that the chief on the scene has a chance to deploy his forces. Thus, a reasonable maximum number of units to dispatch is five. The minimum number to dispatch is clearly one, since someone must determine the nature of the fire and call for more units if they are needed.

In deciding how many units to dispatch, we face a dilemma. Ninety-seven percent of the time two engines and one ladder are sufficient at any incident. One could argue, therefore, that you could always send two and one, or even one and one since the dispatch will be adequate such a high percentage of the time. Moreover, sending more units during busy hours may strip areas of their protection against future fires.

But, says the counter-argument, it is just those three times out of a hundred that really matter. Those are the fires at which the third engine and second ladder are really needed, when a quick and complete response by the Fire Department may save lives and property.

Thus, it would appear that a conservative policy would be to send three engines and two ladders 100 times just so that, when the third engine and second ladder are needed, they get there fast. At low alarm rates, this argument is sound but, as we mentioned earlier, at the current high alarm rate, such a dispatch would strip an area quickly of its protection against future fires. We decided that neither extreme was sound. Our idea was to try to key the number of units dispatched to the potential needs of the alarm at hand, sometimes sending 1 and 1 and sometimes sending 3 and 2.

We analyzed the question of how many units to send using a semi-Markovian decision model which considers explicitly information on the current alarm and expected future events [23]. The model assumes that alarms occur and are extinguished according to random processes. Its state space is the number of companies busy and the potential seriousness of incoming alarms. The decision variable is the number of companies to dispatch, and the objective function measures response time to serious fires. Manipulation of the model revealed the structure of

optimal dispatching rules. Solving it numerically using linear programming indicated that the optimal rule could be approximated by a simple function of the following three factors:

(1) **The probability that the incoming alarm is serious: the** greater the probability, the more units are dispatched;

(2) The expected alarm rate in the area surrounding the alarm: the greater the alarm rate, the fewer units dispatched; and

(3) The number of units available in the area surrounding the alarm: the more units available, the more units dispatched.

We tested the rule in the simulation, and found that it improves significantly the response time to serious fires.

In the example shown in Fig. 6, the traditional dispatching policy would dispatch the three closest engine companies, leaving no units available in a wide area. If the alarm rate were sufficiently high, i.e., if there were a good chance of another alarm in the next few minutes, our decision rule would dispatch only one engine company, saving the others for future alarms. Therefore, under some conditions (such as high alarm rate and low availability), it is better to send fewer units initially.

To implement this decision rule fully, a computer is needed to keep track of changing company status and to make the requisite calculations in real time. During the fall of 1974, a pilot version of the system will be implemented in the Bronx using a mini-computer--this is the first step towards implementing a computer-assisted dispatching system for the whole city. This pilot implementation will also include computer-assisted relocation to determine dynamically how to fill the gaps which develop in fire coverage [17].

However, it was possible to get much of the benefit of the new dispatching rule without using a real-time computer system. Our analysis showed that the most important factor in deciding how many companies to dispatch is the probability that the alarm is serious. So, we devised a new manual dispatching system called adaptive response,

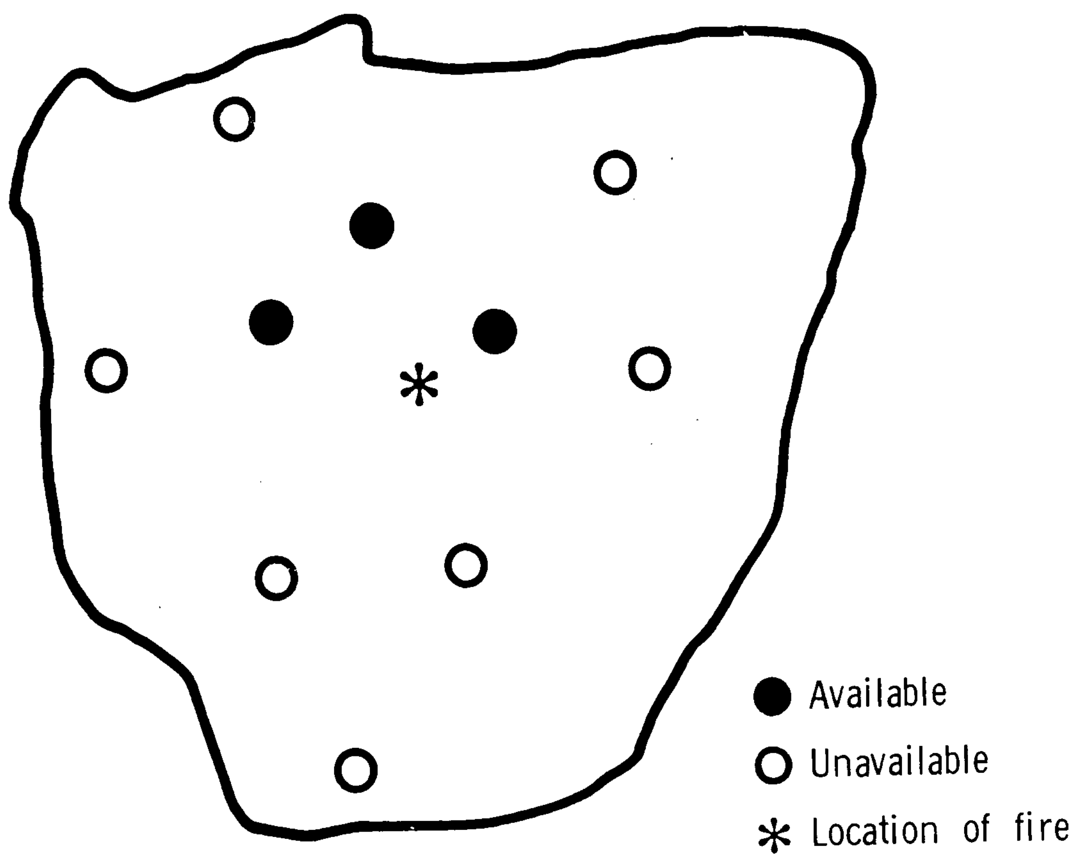

Fig. 6. A problem with the traditional dispatching policy

which uses this fact. It was this system that was instituted in November 1972. As in the old system, all alarms are referenced to the nearest alarm box. The number of companies to be dispatched under the new system depends on a prediction of the probability that an alarm from that box is serious, on the types of structures in the area, and on the time of day.

A small-scale version of adaptive response that used these notions in a rough way had been operating since 1969. Its success suggested that there were predictable box-to-box variations in the chance that an incoming alarm would be serious, and motivated a statistical analysis of alarm patterns that confirmed the stability of box-to-box differences over time. Using 1967-1969 data, we predicted alarm rates and the expected proportion of serious fires for each alarm box in the Bronx. We then compared the predictions to actual 1970 data. For example, as shown in Table 2, we predicted for box 2277 that less than 0.5 percent of all alarms would be structural fires, while for box 2209 we predicted almost 32 percent. In 1970, both of these boxes had about the same number of alarms. In both cases, the predictions were quite close to the actual results. Box 2277 had no structural fires and therefore no serious fires (major incidents requiring at least two ladder companies) while box 2209 had 25 structural fires and 12 of these were serious. A very striking aspect of this example is that both alarm boxes are on the same street--Brook Avenue in the south Bronx--and are about three blocks apart. Under the traditional dispatching policy, if an alarm had been received from box 2277, all of the closest companies would have been dispatched to that alarm. If, in the next few minutes, another alarm had been received from box 2209, none of the closest units would have been available to respond to this alarm, and companies from further away would have had to be sent. The essence of adaptive response is to send fewer units to boxes like 2277, which have a small chance of signaling a serious fire and more units to boxes like 2209, which have a high chance.

The implementation of an adaptive response policy requires the specification for each alarm box of whether it gets a full or a reduced response. We illustrate how this specification is made. Consider, for example, the question of which alarm boxes get a dispatch of two ladders and which get only one ladder. For each alarm box, two predictions are generated: one an estimate

Table 2

STRUCTURAL FIRE PREDICTIONS FOR TWO ALARM BOXES

Bronx box number	Predicted percent structural ('67-'69 data)	Actual 1970 Data	
		Alarms	Structural fires
2277	0.4	96	0
2209	31.8	94	25

of λ, the total annual alarms at the box, and the other an estimate of P_s, the proportion of those alarms which will be serious [8, 9]. (The predictions of the proportion of serious alarms are made using empirical Bayes' methods which take into account the alarm history of the box in question as well as the alarm history at neighboring boxes.) The alarm boxes are then ordered according to increasing values of P_s. An initial specification of the response for each box is to specify a cutoff value, α, of P_s and let alarm boxes whose value of P_s is less than α receive a one-ladder dispatch while boxes with P_s greater than α receive a two-ladder dispatch.

Given the cutoff point and the ordered list of alarm boxes, we can estimate the total number of responses that will be made, the number of serious fires that will receive only a one-ladder dispatch, etc. One can also start with a specification of the total number of responses to be made and find the required α. In this way, policies can be generated which reduce workload and keep the number of one-ladder responses to serious fires fixed at previous levels, or vice versa. The Department eventually settled for a policy which would both reduce the total number of responses made (and hence reduce workload) and increase the proportion of serious fires which receive a two-ladder response.

The initial specification of the adaptive response policy generated purely on a statistical basis as described above was then modified by the Department, first by considering the type of structures near each box and then by considering special hazards. Overall, the new policy sends fewer units than the traditional 3-2 policy, thereby reducing workload, yet response time to serious fires is still reduced (see [19]).

Changing the Number and Location of Fire Companies

We turn now to the most important deployment issue: how many companies are needed and where should they be located?

Before our analysis, the Department had no useful quantitative measures of the fire protection it was providing in each of the regions of the City. Comparing protection from one region to another or deciding on overall needs was principally a matter of subjective judgment. In 1970, we developed a simple "square root" model which relates the average response distance in a region to the density of units in the region [16]. Square root models for predicting average distances have been used for some time in the physical sciences, but their theoretical justification involves assumptions which are not true of the environment we were modeling. We validated a square root model for our applications by testing it with our detailed simulation of New York City fire-fighting operations, which incorporates explicitly many of the deviations from the usual assumptions that had been used in deriving square root distance models. The graph in Fig. 7 shows the relationship we derived and also shows how we validated it by comparing its predictions to simulation data. Then, after analyzing data from field experiments on fire engine travel speeds, we were able to model response times. We learned how they depend on distance, time of day, and location in the City [18]. Our results were remarkably simple and robust. Time of day and geographical variations existed but were small, and a simple spline-like function (Fig. 8) served to predict average times given distance throughout the City. Since this work on response distances and times was completed in New York, we have participated in or been informed of response time experiments in other cities including Denver, Jersey City, and Yonkers. The results have shown the general validity of the models derived for New York.

By combining the distance and time models it was possible to predict average response times in regions throughout the City knowing only the regional alarm rate, average service time, area, and, of course, the number of companies in the region [15]. For the first time the Department had a quantitative measure of the service it was providing. Planning officers in the Fire Department then divided the City into 21 planning regions and identified the fire hazards in each. We used the mathematical models to compare response times among regions having the same fire hazards and found substantial imbalances. We showed the figures to the Chief of the Fire Department and he was startled. He knew there were differences, but was surprised at the extent and magnitude of them. Some regions having the same hazards differed in average response time by over one minute. The Department then began the long process of planning reallocations to bring the protection across regions into better balance.

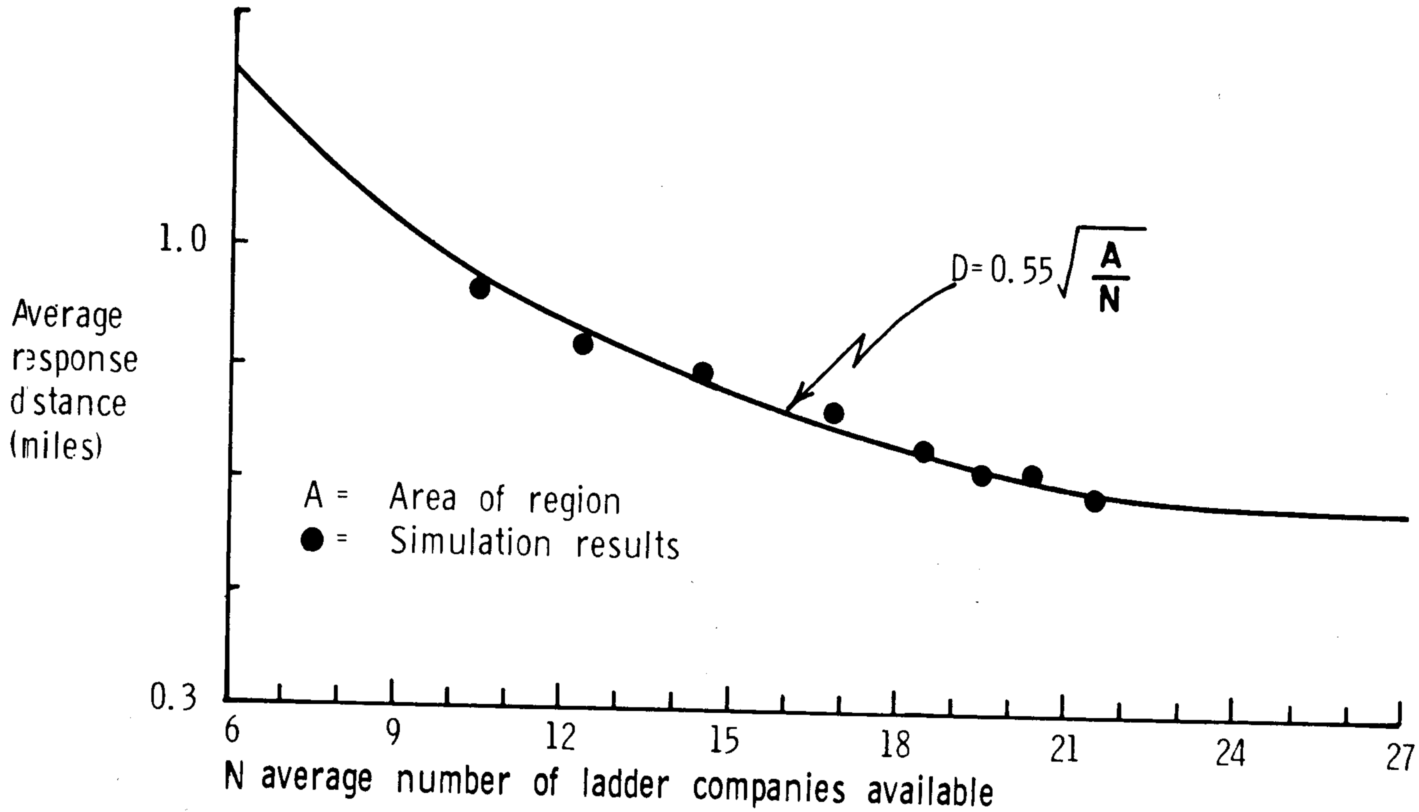

Fig. 7. The square root model for predicting average response distances

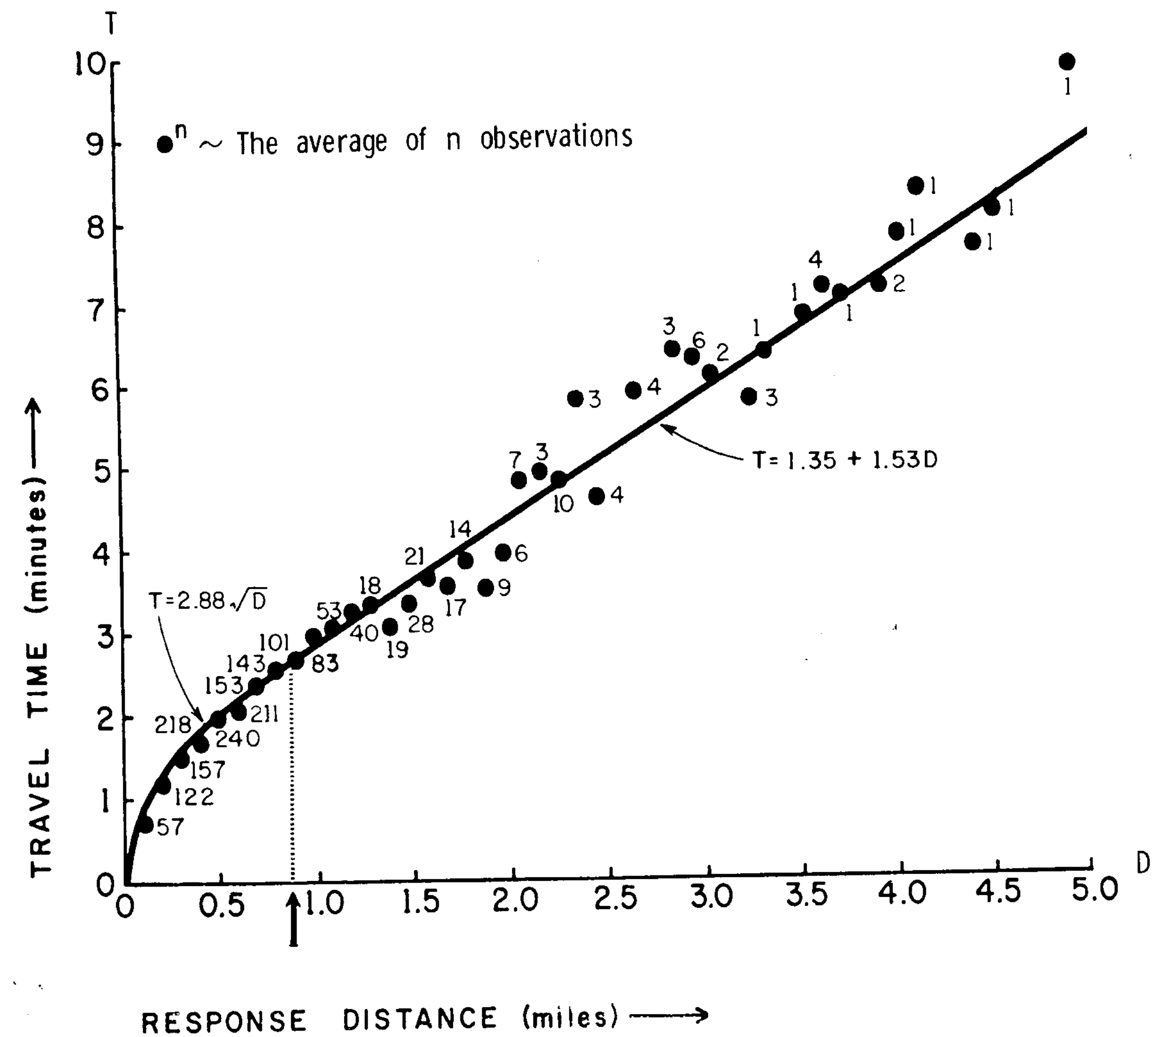

Fig. 8. An empirically determined function for relating travel times to response distances.

Late in 1971, the City faced a severe budget crisis and Mayor Lindsay asked the Department to present plans for cutting its budget by varying amounts. Our joint team presented the Department with various options which we estimated could have resulted in savings of from $1 million to $10 million per year. The response time models and the simulation played a key role in this analysis. We looked for regions where cuts could be made without significant degradation of fire protection.

As a result of these analyses and recommendations, the Department instituted sweeping changes in the deployment of its resources: six companies were disbanded and seven companies were permanently relocated to fill new fire houses or to provide protection in areas where they were needed more. Since the cost of operating a single fire company in 1972 had risen to over $700,000 per year, disbanding six companies saved the City more than $4.2 million per year. These changes took place and have proven effective even though fire alarms have continued to rise (see [19]).

A simple example illustrates the approach that was used. We found regions in which response time was very favorable. In those regions, a company could be removed while response time remained as good as or better than other regions of the same hazard. South Brooklyn was one of these regions. Before Engine 208 was removed from this region it had 10 companies and an average response time of 1.7 minutes. Removing Engine 208 resulted in an increase in average response time of less than .1 minutes--only 6 seconds. In this case, Engine 208 was moved to another region where it was needed more--thus improving balance. In other cases, the company in question was simply disbanded.

The square root model estimates only overall response time changes for a large region. Before any companies were actually moved, the Department made detailed calculations of response times to the specific alarm boxes affected, and, in addition, evaluated any special conditions, considerations, and hazards in each area affected by changes.

In total, 13 fire companies were removed. These changes were not received quietly. Firemen and their supporters, 5,000 strong, marched on City Hall demanding that the companies be put back, and, to this end, the fire unions sued the City in Federal Court. As expected, there was also reaction from the communities which lost companies. When even one fire company is moved, the local councilman, state senator, and congressman will react strongly. Imagine the reaction when 13 companies are moved. But, the Department was prepared with hard facts and detailed analyses to back up the moves. Detailed briefings were given to local groups and, in one case, additional field experiments were run to demonstrate concretely to the community that the changes were not injurious. The uproar died down and eventually the unions dropped their suit.

III. CONCLUSION

So far we have concentrated on two important results of our deployment work: adaptive response and the redeployment of fire companies. Let us now mention briefly some other results.

- Using insights from a queueing model, we found that, under certain conditions, it pays not to dispatch the closest unit [4]. The analysis, which involved optimization of a simple queueing system, showed that sometimes it was possible both to reduce average response times to fires and even out the workload distribution among fire companies by dispatching a unit further away from the incident. Figure 9 illustrates such a situation. The "city" consists of two regions of equal area--one a high alarm rate region and the other low. Each has a centrally located fire company. Suppose that two fires break out shortly after one another at locations 1 and 2. Under the usual "closest company dispatch," we would send company a to fire 1 since it is a little closer, and then would be forced to send company b all the way to fire 2. Shifting the response strategy by making company b responsible for part of what was the company a region, would balance workload between them and reduce the overall average response time to fires. Analogous situations exist in many high alarm rate parts of New York City, and the computer-assisted dispatching system being tested soon in the Bronx will incorporate an algorithm which will evaluate the desirability of sending the closest unit to an incoming alarm. For the first time, response boundaries are being changed to distribute workload more evenly among companies.
- A dynamic relocation method has been developed to replace the traditional method [17]. As we mentioned earlier, the traditional relocation policy consists of a set of advance plans for temporary relocations to make in the event of major fires

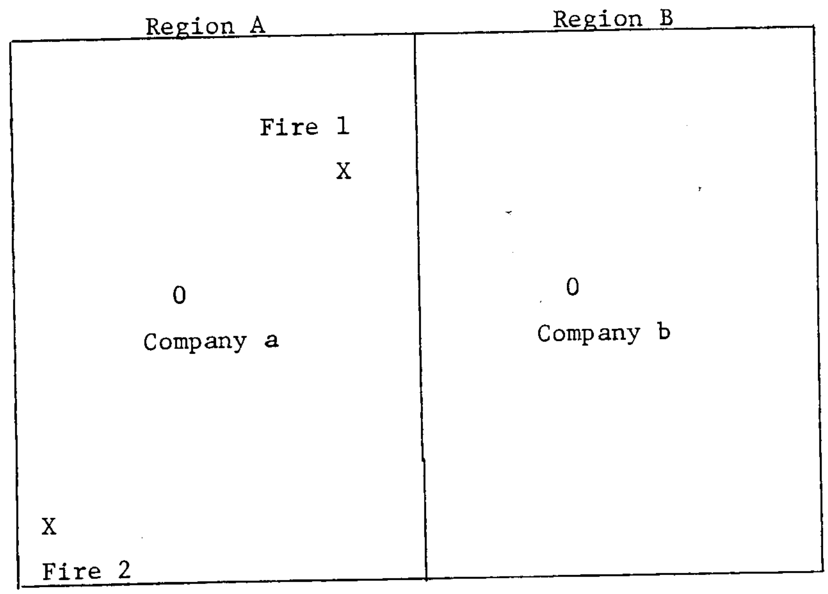

Region A has a high alarm rate. If company a is available, it responds first to fires there. Region B has a low alarm rate. If company b is available, it responds first. In the example, a fire breaks out at location 1 first, and shortly thereafter a fire breaks out at location 2. Total response time would be reduced if we send company b rather than company a to fire 1, even though b is farther away.

Fig. 9. Sending the closest company is not always the best policy

at any location in the City. At high alarm rates, these plans break down in three ways: either the companies designated to move are unavailable so that the plan cannot be implemented, or the companies are available but moving them creates new and possibly worse gaps in coverage due to the unavailability of their neighboring companies, or several small fires at nearby locations create a big gap in coverage which is not recognized by the advance plans which perforce consider only one fire at a time. We have developed an algorithm for solving relocation problems as they occur. The algorithm depends on the computer's ability to keep up-to-date status information on all companies and fires in progress. We recognize the multiple objectives inherent in relocation problems--both equity of protection for all parts of the City and speed of response are at issue and at times conflict with one another--and so formulate and solve the problem in real time as a staged series of integer linear programs. The first of these linear programs is a covering problem that assures minimum levels of fire protection to all regions. The second is an assignment problem--with additional constraints so that minimum coverage standards are still met--whose objective function balances distance moved in relocating and the size and the alarm intensity of regions being filled or vacated.

- In one of the early uses of the simulation model, we demonstrated the effectiveness of part-time fire companies [7]. All New York City fire companies had been manned 24 hours a day, but 60 percent of all alarms occur between 3 p.m. and midnight. Having some units on duty only during these hours provides a better match of companies on duty to demand. We used the Bronx simulation model to measure numerically the benefits of these units. Their use results in improvements in response time and a balancing of workload. But, as mentioned earlier, the fireman's work schedule is specified in the State Constitution, which does not allow this special tour. So, the Department negotiated with the unions to man these companies with volunteers.

In 1969, six part-time units were created and placed in high demand areas of the City. After two-and-a-half years of successful operation, the part-time units became pawns in labor-management dealings. Under union pressure, the men stopped volunteering and the program was discontinued. Although a successful program ended, there were long-term beneficial effects. The loss of these units increased the pressure on the Fire Department to search for other means of improving deployment, eventually leading to the reallocation of fire companies in November 1972.

To sum up, in 1968 when our joint project began, the main approach to meeting increasing demand was to add companies. Our analyses showed that a new response policy combined with a redeployment of units is a more effective way to meet the problem. As a result of the new deployment policies, fire protection has been improved, workload has been balanced and reduced, and the City has saved millions of dollars.

It has been over six years since we began our work, a long time for an OR/MS project. Both the Institute researchers and Fire Department management have had to be patient, to understand the complex fire-fighting system in new ways, and to work for these changes to be implemented and maintained. We have seen important changes and have developed with the Fire Department a new set of tools that broaden their options and give the City a better chance to make effective use of its fire-fighting resources.

The methods and models we have discussed have a generality which make them useful and applicable to cities besides New York. Their success in New York encouraged the U.S. Department of Housing and Urban Development to sponsor field tests in a wide range of cities of different characteristics. Based on these tests, the useful models are now being documented and made available to cities throughout the country. The end result will be a set of tools for deployment analysis which hopefully can be applied in any city.

REFERENCES AND BIBLIOGRAPHY

1. Blum, E.H., "Fire Service: Challenge To Modern Management," The New York City-Rand Institute, P-4512, November 1970. Also published in Public Management, November 1970.

2. Blum, E.H., "Urban Fire Protection: Studies of the Operations of the New York City Fire Department," The New York City-Rand Institute, R-681, January 1971. Also published in Fire Research Abstracts and Reviews, 13, 1:23-27 (1971).

3. Blum, E.H., "Deployment Research of the New York City Fire Project," The New York City-Rand Institute, R-968, May 1972. Also published in Drake, Keeney and Morse (ed.), Analysis of Public Systems, MIT Press, 1972.

4. Carter, G., J.M. Chaiken, and E. Ignall, "Response Areas for Two Emergency Units," The New York City-Rand Institute, R-532, March 1971. Also published in Operations Research, Vol. 20, No. 3, May-June, 1972.

5. 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 published in IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-6, No. 4, October 1970.

6. Carter, G., and E. Ignall, "Virtual Measures for Computer Simulation Experiments," The New York City-Rand Institute, P-4817, April 1972. Also to be published in Management Science.

7. Carter, G., E. Ignall, and W. Walker, "A Simulation Model of the New York City Fire Department: Its Use in Deployment Analysis," The New York City-Rand Institute, P-5110, October 1973. Also published in Proceedings of 1973 Winter Simulation Conference.

8. Carter, G., and J. Rolph, "New York City Fire Alarm Prediction Models: I. Box-Reported Serious Fires," The New York City-Rand Institute, R-1214, May 1973.

9. Carter, G., and J. Rolph, "Empirical Bayes Methods Applied to Spatial Analysis Problems," The New York City-Rand Institute, P-5076, August 1973.

10. Carter, G., and J. Rolph, "New York City Fire Alarm Prediction Models: II. Alarm Rates," The Rand Corporation, R-1215; forthcoming.

11. Chaiken, J.M., "The Number of Emergency Units Busy at Alarms Which Require Multiple Servers," The New York City-Rand Institute, R-531, March 1971.

12. Chaiken, J.M., and E. Ignall, "An Extension of Erlang's Formulas Which Distinguishes Individual Servers," The New York City-Rand Institute, R-567, March 1971. Also published in Journal of Applied Probability, 9, 192-197, March 1972.

13. Chaiken, J.M., and J.E. Rolph, "Predicting the Demand for Fire Protection," The New York City-Rand Institute, P-4625, May 1971.

14. Kolesar, P., "Mathematical Programming Applications in the Analysis of the Deployment and Utilization of Fire-Fighting Resources," The New York City-Rand Institute, P-4988, March 1973.

15. Kolesar, P., and W. Walker, "Measuring the Travel Characteristics of New York City's Fire Companies," The New York City-Rand Institute, R-1449, April 1974.

16. Kolesar, P., and E.H. Blum, "Square Root Laws for Fire Engine Response Distances," Management Science, Vol. 19, No. 12, August 1973.

17. Kolesar, P., and W. Walker, "An Algorithm for the Dynamic Relocation of Fire Companies," The New York City-Rand Institute, R-1023, September 1972. Also published in Operations Research, Vol. 22, No. 2, March-April 1974.

18. Kolesar, P., "Mathematical Models for Predicting Average Fire Engine Response Times," Operations Research; forthcoming.

19. O'Hagan, J.T., "Improving the Deployment of Fire-Fighting Resources," Fire Journal, July 1973.

20. O'Hagan, J.T., and E.H. Blum, "Technology Aids Fire Service," The New York City-Rand Institute, P-4872, June 1972. Also in Nations's Cities, June 1972.

21. Swersey, A.J., "Applications of Systems Analysis To Urban Fire Protection," The New York City-Rand Institute, P-4741, November 1971. Also in Proceedings of the Third Annual Meeting of the American Institute for Decision Sciences, October 27-29, 1971.

22. Swersey, A.J., "Reducing Fire Engine Dispatching Delays," The New York City-Rand Institute, R-1458, December 1973.

23. Swersey, A.J., "A Markovian Decision Model for Deciding How Many Units to Dispatch," The New York City-Rand Institute, R-1261; forthcoming.

Ignall, Kolesar, et al.

IMPROVING THE DEPLOYMENT OF NEW YORK CITY FIRE COMPANIES

P-5280