

INTRODUCTION TO A METHOD FOR THE
ALLOCATION OF FIRE COMPANIES

Kenneth Lloyd Rider

November 1974

P-5319

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

The Rand Corporation
Santa Monica, California 90406

I. INTRODUCTION

One of the fundamental problems facing the management of any municipal fire department is how to allocate its limited number of fire companies throughout the city. Even if the city's fire house locations once made sense, they should be reevaluated periodically since, as a city changes over the years, its fire experience also changes. Formerly well maintained neighborhoods may become run-down and suffer increasing fire hazards, or vacant land may be built up and create a need for fire protection where none existed before. On the other hand, urban renewal may turn a problem area into one of low risk.

A decision must be made about where to locate a new fire house whenever an outmoded fire house is scheduled to be closed. Should the new house be at the same location as the old one? Should it be put into an area of increasing fire risk? Or should it be put into a low-density area that currently has a low level of fire protection? Similar questions arise if a company is to be eliminated or if a company is to be added to the department.


A major concern in the allocation of companies is the travel time to fires throughout the city. Ideally the fire department might like to minimize its average travel time to fires throughout the city while providing equal average travel times in all areas of the city. Unfortunately it is usually impossible to do both at the same time.

Figure 1 illustrates the problem that arises if two regions of a city have widely differing rates of fire incidence. If the companies are allocated to minimize the average travel time in the city, many more companies will be assigned to the high-incidence region than to the low-incidence region and the resulting travel times within the two regions will not be equal. If this allocation is used, residents of the low-incidence region will possibly be concerned that they are getting substandard protection and are being penalized for having few fires.


If the average travel time is made equal in both regions, the travel time to many fires in the high-incidence region will be increased in order to reduce those in the low-incidence region, and the average travel time throughout the city will be larger than its feasible minimum. Under this policy, therefore, total fire losses in the city would probably

PLACING FIRE COMPANIES



HIGH DEMAND



LOW DEMAND



THIS SOLUTION MINIMIZES TOTAL RESPONSE
TIME TO ALL FIRES



THIS SOLUTION EQUALIZES COVERAGE

(But Total Response Time Is Higher)





Fig. 1. Two Allocation Strategies for Fire Companies

be greater than under most other policies. In addition, the companies in the high-incidence region would have a higher workload than the companies in the low-incidence region, which is a serious consideration in some cities.

In practice, therefore, some compromise must be made between the two objectives in order to obtain a reasonable and useful allocation of fire companies. This paper will present a systematic method of making a compromise based on estimations of fire alarms, hazards, geography, and workload.

The allocation method is derived from the consideration of managerial options under the constraint of a limited fire department budget [1]*. Its goal is not to determine specific fire house locations, but to reveal the effects that the tradeoff between minimum and equal travel time has on the allocation of fire companies to regions of a city. By varying a tradeoff parameter, allocations ranging from equal travel time to equal workload and minimum average travel time can be generated. Fire department personnel could then compare the travel times, city-wide and in different regions, that would result from the various allocations. Based on these calculations, the department could then decide which trade-off parameter would be most desirable for their city and where, in general, fire companies should be allocated.

The model does not recommend specific sites, nor does it take into account such special problems as irregularly shaped response areas. The role of the model is to provide a tool to aid the judgment of managers responsible for making decisions regarding the fire protection of their cities. We will start by defining some terms and concepts employed by the model. Then the model itself will be described and some examples of its use will be given.

* Figures in square brackets denote references listed at the end of this document.

II. MEASURING THE DEMAND FOR FIRE PROTECTION

There are several ways in which the services of a fire department are demanded by a community. One type of demand is related to the hazards in an area. For example, one part of a city may have gas storage tanks or high rise buildings, which could prove a very serious problem in case of fire. Another region of the city might have narrow streets and heavy traffic that might hinder the access of fire equipment to a fire. This type of demand is related to fire-fighting needs in case a fire occurs and might be referred to as *potential demand*.

Another type of demand reflects the fire experience actually realized in the city. Realized demand is reflected by the actual incidence of fires in a region. Even if the potential hazards in a region are small, if that region suffers a high fire-incidence rate and calls heavily upon fire department resources, its *realized demand* is large.

There are many ways to measure these demands for fire protection. We will discuss some specific measures and explain why they are useful.

Measures of Potential Demand

In any real city, it is likely that the rate of fire incidence and the level of fire hazards vary from region to region. In order to discuss the measurement of fire hazards without being concerned about fire incidence, therefore, we will consider a hypothetical city with a very low alarm rate. Since the actual fire incidence in this city is very small, equalizing fire-protection coverage throughout all of the city's regions would be a suitable strategy. At first glance it might seem that a distribution of companies in proportion to each region's area would equalize coverage over the city. However, this would be so only if the travel-time characteristics and the hazards were the same in all regions.

If a particular region had a large proportion of buildings of hazardous construction or contents, the fire department might feel that, for *equal coverage*, that region would require twice as many companies per unit area as the other regions of the city. Even if the building hazards are the same, if the companies in region 1 take twice as long to travel the same distance as the companies in other regions, between four and sixteen times as many companies per unit area would

be needed to insure the same average run time to fires in region 1 as to fires in other regions. The multiplier on the number of companies needed per unit area for equal coverage is called the *hazard factor*. It will be used when we are considering the strategy of equalizing coverage, and we know that an equal distribution of fire companies will not do this. In a sense, the hazard factor and the number of companies in a region provide a measure of potential risk in that region. The larger the hazard factor, the greater the risk associated with a given number of companies.

Measures of Realized Demand

We can measure average realized demand in a region by calculating the average amount of time, in each hour of the day, that units in the region spend fighting fires. For example, if on the average there are two fires per hour in a region of a city and each fire takes two companies 45 minutes to put out, then, during each hour, three hours of company time are spent fighting fires in that region.

A relatively large realized demand in one region of a city generally implies that there are more structural fires and that there is more work for fire companies to do in that region than in other regions of the city. If a fire department were concerned solely with workload, and not with the extent of fire protection coverage, it could base its fire company allocation on realized fire alarm experience alone. By finding the average time spent fighting fires in each region of the city, and by allocating companies in proportion to this time, each fire company in the city would have approximately the same workload regardless of the potential hazards in various parts of the city.

Up to this point, the size of regions has not been mentioned. If, during the busiest hour of the day, the fire companies in region 1 spend an average of three company-hours working and companies in region 2 spend an average of ten company-hours working, we do not know how serious the fire situation is in either region unless the relative size of the two regions is taken into account. Region 2 could be five times as large as region 1 and therefore have a smaller amount of time worked relative to its area. In order to compare realized fire demand in different regions, we define a *busy density* for each region by dividing the company-hours spent fighting fires in the region during the busiest hour of the day by the area, in square miles, of that region. The busy density is equivalent

to the average number of companies working in that hour per unit area. The use of this measure, instead of the number of company-hours worked per hour, gives us flexibility in defining the regions of a city. They can be large or small, but their fire-incidence intensities (alarms per unit area) will be directly comparable.

The hazard factor and the busy density are input variables in the allocation formulas and are used to help determine a reasonable distribution of fire companies.

III. MEASURING FIRE COMPANY ALLOCATIONS

We require two measures of fire company allocation to facilitate the comparison of allocations between regions. The first measure is *company density*, defined as the number of companies per square mile in the region. We use this intensive measure, since it is more meaningful, in terms of talking about levels of protection, to say that there are three companies per square mile than to say that there are nine companies in a region of unspecified area. Its use frees us from any particular definition of a region. An allocation that equalized the fire company workload over a city would be specified by saying: let the company density be proportional to the busy density over the city. This produces the same result as is obtained by dividing the city into arbitrary regions, calculating the actual time worked in each region, and making the number of companies in each region proportional to the time worked. Even though there is no saving in computation, it is not necessary to specify regions before determining the allocation. In the following discussion, however, we will use both the *intensive* "per area" measures and *extensive* absolute numbers to make the explanations clearer.

The second measure breaks up the company density to reflect the availability and the busy time of the companies. This involves finding the average number of companies that are busy at one time in a region, and subtracting this number from the total number of companies in that region. The result is the average number of companies available for service. Dividing the average number of companies available in a region by the region's area yields the *availability density*. This is the average number of companies per unit area that are ready to answer alarms. The identical results can be obtained by subtracting busy density from company density. So,

- availability density = company density - busy density, where
- availability density measures the average available service level
- busy density measures the average fire-fighting service being demanded
- company density measures the total service level being provided.

This may be rewritten as:

$$\text{company density} = \text{availability density} + \text{busy density}$$

to show the division more clearly.

In the preceding discussion we have defined the following: busy density, hazard factor, company density, and availability density. In the next section we will show how these are used to generate fire company allocations.

IV. ALLOCATING FIRE COMPANIES

An allocation of fire companies is a determination of the number of fire companies to locate in each region of a city given the total number of companies to be distributed. In allocating a limited number of fire companies in a city, fire department management is faced with the problem of reconciling several objectives within a limited budget. We have discussed two idealized objectives in allocating companies: equalizing workload based solely on realized demand, and equalizing coverage based solely on potential demand. A third strategy, minimizing average risk, is actually a compromise between these two and is dependent mostly on realized demand although it takes potential demand into account.

Fire companies perform two kinds of service to the citizens of a community: fighting fires and providing insurance. Over the course of a year it might be found that a city had, on the average, two fire companies busy extinguishing fires at any one time. This number of companies would be inadequate, however, to provide fire protection for the city, because the actual number of companies needed fluctuates throughout the day above and below the average. If an average of two companies were busy, perhaps twenty companies would be necessary to make sure that 99.9 percent of the time the city was "adequately covered." The companies that are not *fighting fires* are performing the important service of *insurance*.

Insurance is generally thought of as the distribution of risk so that a very unlikely but potentially catastrophic risk, such as loss due to hurricane, is reduced to an expected but small loss, such as an insurance premium. In the case of fire risks, the chance of a fire destroying a business or killing a family is reduced by having fire companies in excess of the average needed. The additional companies provide protection against the contingency that a greater than average number of fires, or a fire of larger than average proportions might occur.

Another reason for having extra companies is to provide adequate geographic distribution. If two companies are needed to put out a fire occurring in the southwest part of a city, it is of little comfort to know that there are even three or four companies available in the northwest part of the city, fifteen, twenty, or even fifty minutes away. Although the fire might be controllable, and the ensuing loss acceptable

to the city itself, the loss to a family or business might amount to a local disaster. A fire department may attempt to insure against this kind of contingency by making sure that there are companies, above and beyond those needed at one time to fight fires, available for service around the city in case a new fire should break out.

In our allocation method, we first take care of the average number of companies fighting fires at any one time by initially setting the company density in each region equal to the region's busy density. Then we distribute the remaining companies throughout the city in proportion to a formula involving hazard factors and busy densities. This process is stated compactly by the following formula: units are allocated to each region so that, for each region

(availability density) is proportional to (busy density)^v x (hazard factor)^{1-v}

where v is a number (with possible values between 0 and 1) that represents the degree to which potential and realized demand is taken into account. To illustrate how this formula works we look at the two extreme values for v:

- If v=1, then the formula reduces to:

(availability density) is proportional to (busy density).

Since the allocation of unavailable companies is already equal to the busy density, this means

(company density) is proportional to (busy density)

and the allocation of companies distributes the workload evenly.

- If v=0, then the formula reduces to:

(availability density) is proportional to (hazard factor).

If the hazard factors throughout the city are equal, then this formula just allocates available companies in proportion to the area of any region. This allocation will equalize travel times in all areas of the city. If the hazard factor in a given region of the city were twice as high

as in any other region, this region would be allocated twice as many available companies.

Calculations made for New York City, Yonkers, Wilmington, and Jersey City show that if v is between $2/3$ and $4/5$, the average travel time, weighted by the busy densities and hazard factors, is minimized throughout the city.

The tradeoffs between the consideration of potential and realized demand are apparent from the formula. The formula, however, was not derived intuitively, but is the end result of the solution of a mathematical programming formulation of the allocation problem [1].

V. REALLOCATING ENGINES IN YONKERS

We will take as a specific example of the use of the allocation model the problem of reallocating fire companies in the City of Yonkers, New York.^[2] Several of Yonkers' fire houses are obsolete and must be replaced. The question, therefore, is should the new fire houses be built on the sites of old ones or are there more suitable locations that can be found? An on-line computer implementation of the allocation model was used for the following calculations [3].

Figure 1 is a map of Yonkers showing the present location of engine houses and the boundaries of the four regions into which Yonkers was divided for the purpose of making the allocations. The regions were chosen to be relatively homogeneous in their fire risks and to be geographically compact. In using the allocation model, it is important that the regions used be internally homogeneous, but it doesn't matter if homogeneous regions are subdivided, since the resulting allocations will be the same.

In Yonkers, region 1 is an old, built-up district with a high fire experience relative to the rest of the city. Regions 2, 3, and 4 are semi-suburban in character. This difference is reflected in the alarm densities of structural fires and emergencies during the period 4 p.m. to midnight. The densities are .07 alarms per hour per square mile in region 1 and .01 alarms per hour per square mile in regions 2, 3, and 4.

A study of the current travel times of first- and second-due engines in each of the regions in Yonkers reveals some large regional differences, as shown in Table I. The travel times are the average times to structural fires and emergencies estimated by the computer program.

Table I
CURRENT TRAVEL TIMES OF ENGINE COMPANIES IN YONKERS

Region	No. of Companies	Area (sq. mile)	Average No. Busy	Travel Times (minutes)	
				First Due	Second Due
1	5	2.90	.38	1.19	1.79
2	2	3.80	.08	1.92	2.89
3	2	5.36	.10	2.22	3.35
4	4	5.96	.14	1.73	2.61
Average				1.52	2.29




Fig. 2. Yonkers Demand Regions and Engine Company Locations

The computer program assumes that the fire companies and fire alarms in a region are uniformly distributed throughout the region. If this is not the case, the actual travel times would differ from the calculated times. A glance at the map shows that the distribution of companies in regions 3 and 4 is not uniform. We may expect, therefore, that if the alarms are uniformly distributed throughout each region, the differences in travel times between regions will be even greater than indicated in Table I.

The first step in the allocation analysis is to try several allocations for differing objectives. In this step, we will assume that the potential demands, and therefore the hazards factors, are the same in all the regions of Yonkers. The allocations are calculated by using the formula presented in Section IV with several values of v .

Table II presents the allocations appropriate for objectives ranging from equal workload to equal coverage. Fractional allocations mean that the first-due area of a company is shared by two regions.

Table II

ALTERNATIVE ALLOCATIONS OF ENGINE COMPANIES TO REGIONS

Region	Current Allocation	Allocations For					Equal Coverage
		Equal Workload	Minimum Avg. Travel Time	Compromise			
				1	2	3	
1	5	7.2	5.5	4.7	3.8	3.0	2.3
2	2	1.4	1.9	2.1	2.3	2.5	2.7
3	2	1.9	2.5	2.8	3.1	3.5	3.8
4	4	2.5	3.1	3.4	3.8	4.0	4.2

It can be seen that a very wide range of allocations is generated by varying v in the allocation formula. In order to evaluate these allocations, we will consider the travel-time data presented in Table III.

Table III

AVERAGE TRAVEL TIMES OF FIRST-DUE ENGINES
UNDER ALTERNATIVE ALLOCATIONS

Region	Current Travel Time	Equal Workload	Travel Times For				Equal Coverage
			Minimum Travel Time	Avg. 1	Compromise 2	3	
1	1.19	1.01	1.14	1.23	1.35	1.50	1.69
2	1.92	2.22	1.98	1.89	1.81	1.74	1.69
3	2.22	2.29	2.03	1.92	1.83	1.76	1.69
4	1.73	2.13	1.93	1.85	1.78	1.73	1.69
Average	1.52	1.54	1.51	1.52	1.55	1.61	1.69

The equal coverage allocation may seem attractive overall, but there are nearly three times as many alarms in region 1 as in the next busiest region. The result of using this allocation is a city-wide average travel time .18 minutes greater than that given by the allocation that minimizes the average travel time. We can decrease this difference to .10 minutes by using compromise 3, and the increase in travel time in region 3 will be only .07 minutes. A further increase of .07 minutes in the travel time of region 3 will decrease the overall travel time by .06 minutes with compromise 2. It can be seen that we are trading off equal coverage in all regions for better balanced coverage. The best allocation for Yonkers is a matter of judgment. It would appear from the travel-time calculations, however, that with the assumption of similar hazards in all regions, one of the compromise strategies, 1, 2, or 3, would be suitable with the current number of engines used. We will select the compromise allocation 2 as a case for further discussion. This allocation differs from the current allocation of engines by having 1.2 less in region 1, .3 more in region 2, 1.1 more in region 3, and .2 less in region 4. The implementation of compromise 2 would involve, roughly, moving an engine company from region 1 to region 3. The model, however, cannot say which company should be moved or exactly where it should go (a siting model would be needed for this purpose).

Table IV compares the first- and second-due travel times of compromise allocation 2 with the current engine allocation.

Table IV
TRAVEL TIMES FOR THE CURRENT ALLOCATION VS. COMPROMISE 2

Region	First-Due Travel Times		Second-Due Travel Times	
	Current	Compromise 2	Current	Compromise 2
1	1.19	1.35	1.79	2.03
2	1.92	1.81	2.89	2.72
3	2.22	1.83	3.35	2.76
4	1.73	1.78	2.61	2.68
City-wide	1.52	1.55	2.29	2.33

It can be seen that the city-wide first- and second-due travel times are very similar between the two allocations, but that the first-due travel time of region 3 is decreased by .39 minutes and the second-due travel time by .59 minutes. Region 1 and region 4 are the two regions where travel times are increased under Compromise 2. Region 4 changes very little, but the first-due travel time in region 1 is increased by .26 minutes. We may note, however, that the travel time in this region is still far lower than that in any other region. Furthermore, it is known from available data that on an average evening only .38 engine companies in the region are busy. The situation in region 1 is still very satisfactory while regions 2 and 3 have been improved markedly.

The allocation model can be used to answer another question. What is the effect of adding or deleting an engine (or ladder) company from the city? Compromise allocation 2 may be calculated for 14 and 12 engine companies to find new allocations, and to show their effect on travel times.

Table V
EFFECT ON FIRST DUE ENGINE TRAVEL TIMES OF
ADDING AND REMOVING A COMPANY

Region	Number of Companies		
	14	13 (Current)	12
1	1.30	1.35	1.39
2	1.75	1.81	1.87
3	1.77	1.83	1.90
4	1.72	1.78	1.84
City-wide	1.50	1.55	1.60

It can be seen from Table V that adding a company or removing a company would have little effect on city-wide average first-due travel times. If first-due travel times were the only criterion, this result would indicate that the addition of an engine company is not necessary and that the removal of a company could be easily sustained. However, an analysis of the resulting travel-time changes in the region immediately surrounding the new company or the company removed would be necessary before drawing such a conclusion.

Up to this point, we have been assuming that the hazard factors for all regions of Yonkers are the same. This assumption, however, is not true. Region 1 has the highest alarm density, the oldest housing, and the densest population of any region in Yonkers, and it is possible that one minute of travel time in this region involves the same risk as much longer response times in other regions. In fact the current allocation of 5 engines to region 1 implies that a 2-minute average travel time in region 1 provides the same coverage as a 2-minute and 40-second average travel time in the other regions.

We determine this result by using the same value of the tradeoff parameter that yielded compromise allocation 2, but we also adjust the hazard factor for region 1 so that 5 engines are allocated to that region. Table VI shows the resulting allocation. If this allocation were to be implemented, it would involve moving an engine company

Table VI
ALLOCATION D WITH INCREASED HAZARD FACTOR FOR REGION 1

Region	Engines	First-Due Response Time	Second-Due Response Time
1	5.0	1.19	2.79
2	2.0	1.91	2.87
3	2.7	1.95	2.94
4	3.3	1.88	2.84
City-wide	13	1.52	2.29

from region 4 to region 3. This result depends on the assumption that the current allocation of engines to region 1 is appropriate. A comparison of region 1 with similar regions in other cities should be used to verify that the hazard factor for that region has a reasonable value.

VI. CONCLUSION

The allocation method presented in this paper provides fire department management with the capability of generating and quickly evaluating fire company allocation strategies. A mathematical model cannot substitute for sound judgement based on experience, but it can serve as a tool to facilitate the exercise of that judgement. It is the intention of this paper to provide such a tool.

References

1. K.L. Rider, *A Parametric Model for the Allocation of Fire Companies: Formulation, Solution and Interpretation*, The New York City-Rand Institute (R-1615) Forthcoming.
2. J.Hausner, W.Walker, A.Swersey, *An Analysis of the Deployment of Fire-Fighting Resources in Yonkers, New York*, The New York City-Rand Institute (P-1566/2) October 1974.
3. K.L. Rider, *A Parametric Model for the Allocation of Fire Companies: User's Handbook*, The New York City-Rand Institute (R-1646) Forthcoming.

Rider

INTRODUCTION TO A METHOD FOR THE ALLOCATION OF FIRE COMPANIES

P-5319