decision-making problem that is confronted by managers of a wide range of activities. Facility location is both a private and public sector concern, although the two application areas may present very different types of problems. A typical private sector problem is to locate a factory, warehouse, distribution facility (e.g., a truck terminal), or service center so as to minimize the firm's cost of supplying a given market. Mathematical programming models such as the plant location formulation have been presented to deal with this problem.

Our concern is with public sector location problems, in which a frequently encountered difficulty is the absence of a monetary measure, such as cost, for the assessment of a facility's performance. We can think of three types of public sector facility: ordinary, emergency (or extraordinary), and noxious facilities. Ordinary public facilities, such as libraries or health clinics, are established so as to provide a service to as many people as possible. Facilities that provide an emergency service, such as fire stations, must be located so as to provide as much coverage as possible; i.e., as many emergencies as possible must be covered. The location of noxious facilities, such as solid-waste disposal facilities, should be sensitive to their impact on neighbors.

Optimization models for public sector facilities of the three types mentioned above have attracted increasing attention from operations researchers and regional scientists in the past 15–20 years. ReVelle and Church (1978) provide a complete review of the important theoretical developments of this period.

While most of the optimization models developed to date have used a single objective, it is clear that the location of each of the three facility types is an inherently multiobjective problem. In each case—ordinary, emergency, and noxious type—the facility's performance, which is measured by a non-monetary criterion, must be compared to cost since all agencies—even critical ones such as fire and police departments—operate with fixed finite budgets. Of course, if benefits are associated with the services provided by the facility, then performance and cost are commensurable since they would share a common monetary unit. There are, however, critical problems with estimating the benefits (or disbenefits) of some facilities. In many cases there are *multiple* performance measures so commensuration among the measures as well as with cost is required. Even if a single criterion is relevant, the benefit estimation procedure may require value judgments that an analyst will not care to make: In the case of fire station location, what is the value of a life?

The two problems discussed in this chapter demonstrate the role that multiobjective analysis has to play in public facility location decisions. Emergency facilities (fire stations) and noxious facilities (power plants) are considered, but the basic constructs of the multiobjective programming models that are presented and the manner in which they are used to generate noninferior location alternatives should be applicable to the wide range of location problems that exists.

11

10.2 FIRE STATION LOCATION IN BALTIMORE, MARYLAND

Analysts at The Johns Hopkins University were contracted by the City of Baltimore to study the location of new fire stations and the relocation of existing fire companies. The study lasted for 15 months, terminating in July 1976 with a set of alternative locations and relocations that were presented to the Fire Chief and the Mayor of Baltimore. In this section the methodology used in the new location phase of the project and selected results are discussed.

The problem is discussed in the next subsection, followed by a presentation of the location model. The objectives of the location analysis are then developed, followed by a discussion of computational considerations and results.

10.2.1 The Problem

Baltimore is in the rather comfortable position of having a fire department that is one of the most highly rated among large urban areas in the United States and is recognized for its innovative approaches to management and prefire planning. The fire protection system exists, however, like most systems that serve people, in a changing environment where decisions made in the past may now be no longer optimal or even acceptable. The changing character of Baltimore's residential and commercial patterns has created a situation in which there is room to improve the fire protection system.

At present (1977) fire-fighting companies and the stations that house them are concentrated in the central part of Baltimore, providing excellent (and in some locations excessive) coverage to the center of the city, but leaving some parts of the periphery uncovered or inadequately covered. There are good historical reasons for this configuration of facilities. Many of the fire stations in Baltimore were built prior to 1920 at a time when the city's boundaries fell well within its current boundaries. As land was annexed from the surrounding county, the city grew out from its center so that previously located stations are now distant from some portions of the annexed areas. Another historical motivation for the concentration of stations in the central city was the great conflagration of 1904, which virtually leveled the portion of Baltimore that is today the high-value district (central business district). It was this catastrophe that promoted the intensive period of fire station construction in the central city prior to 1920.

1

There is yet another incentive for concentrating fire-fighting capability in the central city that is still important today and is not unique to Baltimore. The rating of fire departments and the determination of fire insurance premiums is controlled by the Insurance Services Office (ISO). This organization exists primarily to assess a fire department's ability to protect *property*. Since property value tends to be concentrated in the center of the city (hence the name "high-value district"), the fire department is under some pressure to focus its effort on that area. The nature of ISO's influence on location decisions of the fire department is formal and specific since it is communicated by the ISO standards for the coverage of structures. The standards played an important role in our analysis and will be discussed further in subsequent subsections.

It was perceived by the Fire Chief of the Baltimore Fire Department that two related problems existed in the current configuration of fire stations. First, some peripheral areas of the city were inadequately covered; second, there were several old fire stations with outmoded facilities in locations that were less than optimal. It was the charge of the analysts to identify alternative locations for new stations and determine which old stations could be closed. The problems are interrelated to the extent that companies in closed stations can be relocated to a new station. For the purposes of analysis, however, the two problems were treated separately. The station-closing problem, although it presents an interesting analytical challenge, is not discussed here. In the next section a model for determining new fire station locations is presented.

10.2.2 A Fire Station Location Model

A variety of problem-specific models for the analysis of facility locations exists. ReVelle and Church (1978) present a complete review and mathematical development of many of these models. Within the usual categorization of location models, we are interested in a *prescriptive* model for the analysis of *public* facilities that provide *extraordinary* (emergency) services over an area described as a *network*. Each of the italicized terms in the previous sentence represents a characteristic of the fire station problem that sets it apart from other facility location problem areas.

First, we are interested in a prescriptive or normative model that will identify (prescribe) optimal (actually, noninferior in the multiobjective case) location alternatives. We are not interested in the descriptive type of model that regional scientists and geographers use to explain demographic or industrial location patterns. Our interest in a prescriptive model implies the use of optimization (vector optimization in the multiobjective case).

Second, fire stations are public facilities in that they are owned and operated by a governmental agency. Third, the service provided is of an extraordinary nature in that users of the facility do not employ its services on a regular basis. Both of these characteristics affect the nature of the model's objective function. In the private facility case, the objective is usually straightforward: Minimize cost while providing a given service, or maximize profit. The public facility problem is more complex in this respect in that the notion of the payoff or benefits of a facility is dependent on the welfare gains from the facility. In the case of ordinary public facilities, e.g., libraries, a useful surrogate for welfare is to minimize average distance of people from the facility, thereby maximizing accessibility to the services provided. A more useful notion for extraordinary facilities is the maximum distance or time of a potential user from a facility such as a fire station. An objective for a fire station location problem is to maximize coverage, for which a point is covered if it is within the maximum allowable distance or time of a facility.

Fourth, location models are specific to the manner in which the area of analysis is represented. If location is allowed anywhere in a continuous land area, then a planar model is appropriate. In the Baltimore case study, the city was represented by a network with approximately 600 points, each of which was located at the centroid of an area roughly $\frac{1}{9}$ square mile in size. All adjacent points were connected by arcs that measured the corresponding interpoint distances. See Fig. 10-1 for an example of this network representation. (Note that time of travel could be used in place of distance. Distance was used, however, since time of travel was found to relate almost linearly with distance and this relationship did not vary from arc to arc.) This grid of points and arcs form a network so a network model, which restricts the analysis of potential facility locations to the grid points, was used.

Church and ReVelle (1974) formulated a network optimization model called the "maximal covering location problem" (MCLP) for public facilities that provide an extraordinary service. Simply stated, the MCLP maximizes coverage for a given number of facilities to be located. Coverage is defined in the following way: point i is covered if there is a facility located within S miles of i, where S is the "maximum service distance," which for the case of fire stations has been specified by ISO standards.

In the general statement of the MCLP there are two sets of points: I, the set of all "demand" points, and J, the set of all potential facility locations. The two sets may overlap partially or completely; i.e., points in the network may be both demand and potential facility points. We define x_j as a 0–1 integer variable that equals 1 if a facility is placed at point j and 0 otherwise. One constraint in the model is that exactly α facilities must be sited, where α is a prespecified number,

$$\sum_{j \in J} x_j = \alpha \qquad (10\text{-}1)$$

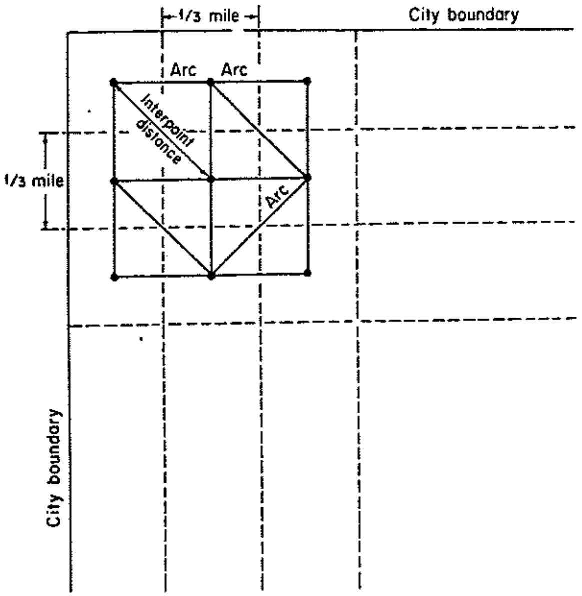

Fig. 10-1. Network representation of an area.

The notion of coverage can be incorporated by defining y_i, a 0–1 variable that equals 1 if there is a facility within S miles of point i and 0 otherwise, and N_i, the set of all potential facility sites within S miles of point i. The coverage constraints are

$$y_i \leq \sum_{j \in N_i} x_j \qquad \forall i \in I \tag{10-2}$$

or

$$y_i - \sum_{j \in N_i} x_j \leq 0 \qquad \forall i \in I \tag{10-3}$$

The constraint requires at least one $x_j, j \in N_i$, to equal 1 if y_i is to equal 1; i.e., point i can be considered covered only if at least one facility is located within S miles of it.

The objective function of the MCLP is to maximize coverage. In general, there may be some attribute of a demand point that should serve as a weight on the importance of covering it. For example, property value is an important attribute in a fire station problem and we should be more inclined to provide coverage to a point with high property value than to a low-value point. Calling a_i the amount of the attribute, e.g., property value, at point i, the objective function is

$$\text{maximize} \quad Z = \sum_{i \in I} a_i y_i \tag{10-4}$$

The entire MCLP is the maximization of (10-4) subject to (10-1), (10-3), and integrality restrictions on $y_i \; \forall i \in I$ and $x_j \; \forall j \in J$. It is of importance for computational reasons that the MCLP frequently terminates all-integer [upward of 80% of the time—see Church and ReVelle (1974)] when the integrality requirements are ignored, i.e., when the MCLP is solved as a linear program. Integrality is, of course, important since a half of a fire station is of no practical significance: We want to build a whole facility at a site ($x_j = 1$) or we do not ($x_j = 0$), but we are not interested in fractional solutions (e.g., $x_j = 0.5$). The techniques that exist for resolving fractional solutions, e.g., the branch and bound algorithm, can be expensive for the large problems encountered in practice. The MCLP is appealing, then, since relatively inexpensive linear programming routines may be used in lieu of expensive integer programming routines.

The MCLP incorporates a single attribute into its objective function. As we shall see in the next subsection, however, fire station location problems are multiobjective in nature. The Baltimore study gave rise to the multiobjective facility location (MOFLO) problem presented in Schilling (1976). MOFLO employs the same constraint set as the MCLP, but it allows for more than one coverage objective. Defining a_{ik} as the amount of the kth attribute at point i, the objective function of MOFLO is

$$\text{maximize} \quad Z = \left[\left(\sum_{i \in I} a_{i1} y_i\right), \left(\sum_{i \in I} a_{i2} y_i\right), \ldots, \left(\sum_{i \in I} a_{ip} y_i\right)\right] \tag{10-5}$$

where, in general, there are p objectives.

The MOFLO problem was used in the analysis of new fire station location alternatives in Baltimore. The objectives used in the analysis are discussed in the next subsection.

10.2.3 Planning Objectives

The point of departure for the analysis of fire station location alternatives in Baltimore was an assumption that location policy was influenced to a large extent by the property coverage standards of ISO. While ISO does exert influence over departmental decisions, as it must, several meetings with the Fire Chief and Deputy Chiefs indicated that many objectives were important. By the end of the study, six coverage objectives were identified and used in the analysis.

The primary ISO activity is the establishment of standards for fire fighting that will encourage at least adequate protection for property. As a mathematical function for MOFLO, the ISO objective is to maximize property value covered by new fire stations

$$Z_1 = \sum_{i \in I} a_{i1} y_i \tag{10-6}$$

where a_{i1} is the property value of the $\frac{1}{9}$-square-mile area represented by point i. Values of the a_{i1} were obtained from Baltimore's Real Property File, which includes property values for all assessed structures in the city. Note that in our case there are some areas of high property value that are currently covered that we do not wish to cover again. To prevent further redundant coverage, the set I is defined to include only currently uncovered points.

Siting a limited number of facilities on the basis of the ISO objective alone may lead to a substantial lack of coverage in densely populated low-property-value areas. In fact, most American cities have such low-value dense residential areas so the exclusive dependence on ISO standards for location decisions may be generally ill advised. Of course, the ISO cares about people, but it is their role to insure protection of structures, which may not necessarily include the best protection for populations. When there is a sufficient budget to allow all points to be protected, the problem is moot: All people and all property are covered. When budgets are constrained, as they are, protection directed at property only may not be the best policy.

The Fire Chief articulated a population coverage objective at an early stage in the analysis. The form of the objective to maximize population coverage is identical to (10-6), only it will be called Z_2, and the population at point i is a_{i2}. Values for the a_{i2} were obtained from census data disaggregated from tract levels.

The MOFLO formulation has a form that promotes the rapid quantification of objectives. All coverage objectives are of the same form, as in (10-6), so that one need only develop attributes of interest at the grid points to identify new objectives. In this spirit, a third objective of maximizing area covered was formulated to interject a sensitivity to future changes in housing and industrial location patterns. Some areas with few people or structures would receive very low weights in the first two objectives. The area coverage objective assigned equal weight to all points, since all points represented equal areas, so that currently undeveloped areas were treated as having the same importance as developed areas to reflect their potential future development. The area coverage objective is Z_3 with $a_{i3} = 1 \ \forall_i$.

The first three objectives did not consider the importance of relative fire frequency for fire station location: Stations should be placed where fires are most likely to occur. Three additional objectives were developed with the Fire Chief and Deputy Chiefs in order to incorporate this consideration. A fourth objective was to maximize expected fires covered. Calling this objective Z_4, the a_{i4} are the expected number of fires at point i. Values for the a_{i4} were estimated from a sample of historical fire data maintained by the Fire Department.

The fire coverage objective Z_4 did not capture the variation of population and property value from point to point, and the property and population coverage objectives Z_1 and Z_2 did not include the variation of fire frequency from point to point. Measures that approached overall indices were developed by literally combining fire frequency with property value and population. The fifth objective Z_5 was to maximize *property hazard* covered, where a_{i5} is the product of fire frequency with property value at point i, and the sixth objective Z_6 was the maximization of *population hazard* covered, where a_{i6} is the product of fire frequency with population at point i. The hazard objectives are weighted versions (by fire frequency) of objectives Z_1 and Z_2.

Some of the six objectives may be collinear (see Section 4.5.2) in the general situation, especially objectives Z_2, Z_4, and Z_6. One would expect this if densely populated areas have high fire frequency, which they often do. If this is the case, then two of the population, fire, and population hazard objectives may be redundant. Even with the possibility of redundancy, it was decided to use the full set of objectives since there was interest in all of them and because the extent of the collinearity was unknown *a priori*.

10.2.4 Computational Considerations

The MOFLO problem is a multiobjective 0–1-integer programming problem that has a special constraint structure that accelerates its solution time. These two characteristics of the problem—the integrality requirements and the special structure of the constraints—have important implications for the multiobjective solution technique that may be employed. In fact, the implications result in a dilemma—the analyst's own multiobjective problem.

Of the many solution techniques available (see Chapters 6–8) our emphasis was on generating methods (see Chapter 6). Recall that generating techniques find an approximation of the noninferior set; they focus on the range of choice available to the decision maker. This emphasis on alternatives, without a recommendation of an optimal alternative from the analysts, appealed to the Fire Chief. Indeed, it is fair to claim that it was our proposed use of multiobjective generating methods that sparked the Fire Chief's initial interest in our study. He liked not being told what was "best."

As discussed in Chapter 6, however, there are several generating methods and it is here that the dilemma arises. Methods that are based on weights on the objectives—either the weighting method or the noninferior set estimation