Covering Problems

Lawrence V. Snyder

Lehigh University

Forthcoming in *Foundations of Location Analysis*, H. A. Eiselt and V. Marianov (eds.), Springer.

1. Introduction

The mail-order DVD-rental company Netflix chooses distribution center locations so that most of its customers receive their DVDs within one business day via first-class U.S. mail. Similarly, many municipalities aim to have fire crews reach 911 callers within a specified time, such as four minutes. Both of these are examples of the notion of *coverage*, a concept that is central to several classes of facility location models; it indicates whether a demand location is within a pre-specified radius (measured by distance, travel time, cost, or another metric) of its assigned facility. Homeowners are covered if they are within four minutes of the nearest fire station, and Netflix customers are covered if they are within one mailing day of a distribution center. Note that in the fire-station example, municipalities typically want to cover *all* residents (while minimizing the number of service stations to open), whereas Netflix wants to cover as many customers as possible (subject to a limit on the number of warehouses it may operate at any time, as specified by its capital budget). The fire-station problem is an example of the *set covering location problem* (SCLP), while Netflix's problem is an example of the *maximal covering location problem* (MCLP). This chapter discusses both problems.

The SCLP was first introduced by Hakimi (1965) and was later formulated as an integer programming problem by Toregas, et al. (1971). The MCLP was introduced by Church and ReVelle (1974). Both models, and their variants, have been applied extensively to public-sector facility location problems, such as the location of emergency medical service (EMS) vehicles (Eaton, et al. 1985), fire stations (Schilling, et al. 1980), bus stops (Gleason 1975), wildlife reserves (Church, Stoms, and Davis 1996), and emergency air services (Flynn and Ratick 1988). They have been applied to a much more limited extent in the private sector (see, e.g., Nozick and Turnquist 2001).

The SCLP and MCLP are closely related to the *p*-center problem, which aims to locate at most *p* facilities to minimize the maximum distance, among all customers, between the customer and its assigned facility. In the *p*-center problem, the coverage radius itself constitutes the objective function. See the Introduction chapter of this book for a more thorough discussion of the relationships among these classical models.

Like most location problems, the SCLP and MCLP may be defined as continuous problems (in which facilities may be located anywhere on the plane) or as discrete problems (in which they may be located only at the nodes of a network). In this chapter we consider the latter approach.

The remainder of this chapter is organized as follows. In Section 2, we discuss classical papers on the SCLP (in Section 2.1) and on the MCLP (in Section 2.2), present the results of computational experiments, and discuss more recent variations. In Section 3, we discuss the impact that these models have had and the bodies of research they have inspired, focusing generalized notions of coverage. Finally, we conclude in Section 4 and suggest some possible future research directions.

2. Historical Contributions

We present the classical models for the SCLP by Hakimi (1965) and Toregas, et al. (1971), and for the MCLP by Church and ReVelle (1974), in Sections 2.1 and 2.2, respectively.

2.1 The Set Covering Location Problem

2.1.1 Hakimi (1965)

Although the generic (non-location) set-covering problem had been formulated prior to Hakimi's (1965) seminal paper on the SCLP, Hakimi's work is important for, among other things, introducing the notion of coverage into facility location models. Hakimi's proposed solution method, which involved the use of

Boolean functions, never proved to be efficient enough to warrant its use in practice; rather, the SCLP is generally solved using integer programming (IP) techniques, first proposed by Toregas, et al. (1971). We discuss Hakimi's model and briefly outline the Boolean-function approach in this section. In Section 2.1.2, we present the IP method of Toregas, et al.

We consider a graph $G = (V, A)$ and assume that every node in V is both a customer (demand) node and a potential site for a facility. (However, one can easily extend the models below to handle the cases in which some customers may not be facilities, or some facilities are not customers, i.e., do not need to be covered. Below, we use terms like "customer i" and "facility j" as shorthand for "the customer located at node i" and "the facility located at node j.") Let $n = |V|$. The distance between nodes i and j is given by d_{ij}, and the maximum allowable distance between a customer and its nearest opened facility—the "coverage distance"— is given by s. If $(i, j) \in A$, then d_{ij} is the length of the arc (i, j), and otherwise it is the shortest distance from i to j on the graph. (We use the term "distance" throughout, but the parameters d_{ij} and s may just as well represent travel times, costs, or another measure of proximity.) Therefore, facility j covers customer i if $d_{ji} \le s$. We define

$$V_i = \{j \in V : d_{ji} \le s\},$$

that is, V_i is the set of nodes that cover customer i. Note that every V_i is nonempty, assuming that $d_{ii} = 0$ for all i.

The objective of the SCLP is to find the minimum-cost (or minimum-cardinality) set of locations such that every node in V is covered by some node in the set. The application that Hakimi cites for the SCLP is that of locating policemen along a highway network so that every intersection (vertex of the graph) is within one distance unit of a policeman. Subsequently, the problem has found a much broader range of applications, as discussed earlier.

We will assume that facilities may be located only at the nodes of the network, not along the edges. Note that it may be optimal to locate along edges, since the well known "Hakimi property"—which says that an optimal solution always exists in which facilities are located at the nodes, rather than along the edges, of the network—does *not* apply to the SCLP. (Hakimi introduced his famous property in an earlier paper (Hakimi 1964) in the context of the *p*-median problem, not of the SCLP.) A very simple counterexample consists of two nodes connected by a single edge of length 1 and a coverage distance of 0.5. If facilities are allowed on the edges, the unique optimal solution consists of one facility (located in the middle of the edge), whereas the optimal nodal solution consists of two facilities, one at each node.

On the other hand, a problem in which facilities may be located on edges may be converted to a node-only problem by inserting dummy nodes onto the edges, taking advantage of the fact that there are only a finite number of possible optimal locations along edges. See Church and Meadows (1979) for details.

In some applications, it is desirable to use a different coverage distance for each customer—for example, if customers have service agreements that specify different response times. In this case, the coverage distance is customer dependent, s_i, and the set V_i is given by $V_i = \{j \in V: d_{ji} \le s_i\}$. The analysis below changes in only minor ways.

The SCLP is closely related to the graph-theoretic *vertex cover problem*, whose objective is to find a subset of nodes in the graph such that every node in the graph is adjacent to some node in the set *and* such that no strict subset of the set has the same property. Such a set of nodes is called a *cover*. The optimization version of the vertex cover problem seeks the minimum-cardinality cover, and this problem is a special case of the SCLP in which $s = 1$ and $d_{ij} = 1$ for all $(i, j) \in A$. Indeed, this special case is the problem considered by Hakimi (1965), although he is usually credited for introducing the more general SCLP, since he presented the problem explicitly in a facility location context. In this section we will assume, following Hakimi, that $s = d_{ij} = 1$, though in subsequent sections we will allow s and d_{ij} to be arbitrary. Hakimi notes that the assumption that $d_{ij} = 1$ is not overly restrictive, since if the arc lengths are greater than 1, one could simply introduce dummy nodes along the arcs, one unit apart, assuming that the arc lengths are integers. Of course, this modeling trick comes at considerable computational expense, especially since Hakimi's method relies on an enumerative approach whose computational complexity increases exponentially with the number of nodes.

In the remainder of this section, we describe Hakimi's (1965) approach to solving the SCLP. As noted earlier, this method is not commonly used today and is discussed here primarily for its historical interest.

Recall that V_i is the set of nodes that cover node i; given the assumption of unit arc-lengths and unit coverage distance, V_i is simply the set of nodes that are adjacent to i, plus i itself. Let S be a subset of the node set V. For each node i, we define a Boolean (binary) variable x_i that equals 1 if $i \in S$ and 0 otherwise. With a slight abuse of notation, we can write

$$S = \bigcup_{i \in V} x_i i ,$$

where $x_i i$ is taken to equal the set $\{i\}$ if $x_i = 1$ and the null set otherwise. We also define X_i as the sum of the Boolean variables for the nodes in V_i; that is,

$$X_i = \sum_{j \in V_i} x_j .$$

Here, $\sum$ represents Boolean summation, analogous to the "or" operator, in which $1 + 1 = 1$. Then $X_i = 1$ if and only if S contains a node that covers node i.

Finally, we define the Boolean function f, which takes as inputs the vector of Boolean variables for the nodes and returns a single Boolean value:

$$f(x_1, \ldots, x_n) = \prod_{i \in V} X_i .$$

Since node i is covered if and only if $X_i = 1$, we have the following theorem:

Theorem 1. S contains a covering of V if and only if $f(x_1,\ldots,x_n) = 1$.

The advantage of using the function f is that it allows us to use Boolean algebra to construct coverings of V. Although this approach still involves enumerating all coverings, it allows us to do so without enumerating all subsets of V to identify them. In particular, we will create a "minimum sum of products," i.e., the smallest possible sum of products of x_i variables that is logically equivalent to $f(x_1,\ldots,x_n)$. This method involves eliminating terms that are implied by others, then using Boolean algebra to simplify the resulting formula until we have an expression consisting of the sum of simple products of variables such that no product is implied by (contains) any other. The method is best explained by use of an example.

Example 1. We illustrate the method using the sample network in Figure 1.

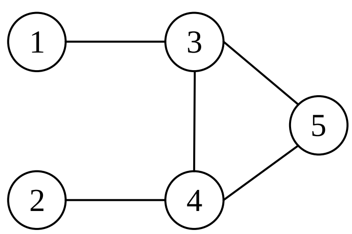

Figure 1. Sample network.

Using the adjacencies depicted in the figure, $X_1 = x_1 + x_3$, $X_2 = x_2 + x_4$, and so on. Therefore,

$$f(x_1,\ldots,x_n) = (x_1 + x_3)(x_2 + x_4)(x_3 + x_1 + x_2 + x_4 + x_5)(x_4 + x_2 + x_3 + x_5)(x_5 + x_3 + x_4).$$

By Theorem 1, to find all coverings of the graph, we need to find all possible values of $\{x_1, \ldots, x_5\}$ that make $f(x_1,\ldots,x_n) = 1$, i.e., such that all terms in the product above equal 1.

To begin, note that the first term is contained in the third. Since we need each term to equal 1, the third term equals 1 if the first does; therefore, we can eliminate the third term. Similarly, the fourth term contains the fifth, so we can eliminate the fourth term. The resulting expression is:

$$f(x_1,\ldots,x_n) = (x_1 + x_3)(x_2 + x_4)(x_3 + x_4 + x_5).$$

Boolean algebra contains two distributive laws; one says that, for any Boolean variables *x*, *y*, and *z*,

$$x + (yz) = (x + y)(x + z).$$

Applying this law to the last two terms, we get

$$f(x_1,\ldots,x_n) = (x_1 + x_3)(x_4 + x_2x_3 + x_2x_5).$$

The other Boolean distributive law says that

$$x(y + z) = xy + xz.$$

Applying this law to multiply the two terms, and repeatedly applying both Boolean identity laws (which say that $x + x = x$ and that $xx = x$), we get

$$f(x_1,\ldots,x_n) = x_1x_4 + x_1x_2x_3 + x_1x_2x_5 + x_3x_4 + x_2x_3 + x_2x_3x_5.$$

Finally, by the Boolean redundance law ($x + xy = x$), we can remove the second and last terms:

$$f(x_1,\ldots,x_n) = x_1x_4 + x_1x_2x_5 + x_3x_4 + x_2x_3.$$

Therefore, the covers for the graph in Figure 1 are:

{1, 4}, {1, 2, 5}, {3, 4}, {2, 3}.

All but {1, 2, 5} are minimum covers. □

Hakimi was optimistic that this enumerative approach would prove to be practical: "…since the subject of simplification of Boolean functions has been widely studied and there are efficient digital computer programs for such a purpose, the above formulation is feasible." Twenty-first-century readers, however, will recognize that the enumerative approach is impractical for large instances. Moreover, since the vertex cover problem is NP-complete (Garey and Johnson 1979), no polynomial-time exact algorithm for the SCLP exists. However, more efficient approaches than Hakimi's exist; we discuss a mathematical-programming-based approach in the next section

2.1.2 Toregas, et al. (1971)

Toregas, et al. (1971) formulate the SCLP as an integer programming (IP) problem and use standard mathematical programming methods to solve it. We discuss their approach next.

The IP has one set of decision variables:

$$x_j = \begin{cases} 1, & \text{if a facility is opened at node } j \\ 0, & \text{otherwise} \end{cases}$$

for $j \in V$. Note that this variable x_j has no relation to the Boolean variables x_i defined in Section 2.1.1.

The IP is formulated as follows:

$$\begin{array}{llllr} \text{(SCLP)} & \text{minimize} & z = \sum_{j \in V} x_j & & \textbf{(1)} \\ & \text{subject to} & \sum_{j \in V_i} x_j \ge 1 & \forall i \in V & \textbf{(2)} \\ & & x_j \in \{0,1\} & \forall j \in V & \textbf{(3)} \end{array}$$

The objective function (1) computes the total number of facilities opened. Constraints (2) require at least one node from the coverage set V_i to be opened for each node i. Constraints (3) are standard integrality constraints. Here, we do not assume that $s = d_{ij} = 1$ (as we did in Section 2.1.1); any values for these parameters may be used in determining the coverage sets V_i.

This formulation is virtually identical to that of the classical set-covering problem; here it is discussed in the context of location theory in particular. It is well known that the set-covering problem typically has a small integrality gap; that is, the optimal objective value of the LP relaxation (denoted z_{LP}) is close to that of the IP itself (Bramel and Simchi-Levi 1997), and often the LP relaxation even has all-integer solutions.

In fact, ReVelle (1993) argues that many facility location problems have this property and discusses "integer-friendly programming" techniques for several classical problems.

However, there do exist instances of the SCLP whose LP relaxations do not have all-integer optimal solutions (otherwise the problem would not be NP-hard). An example is as follows.

Example 2. Consider the network depicted in Figure 2. In this example, $s = 1$. An optimal solution to the LP relaxation of (SCLP) is given by $x_1 = x_2 = x_3 = 0.5$, $x_4 = 0$, with an objective value of $z_{\text{LP}} = 1.5$. □

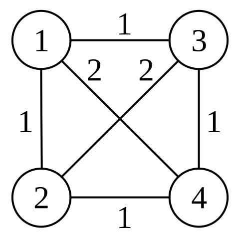

Figure 2. Network for which SCLP does not have all-integer LP relaxation.

Since the coefficient of each x_j is 1 in the objective function of (SCLP), it is clear that the objective function value is integer for any solution to the IP. Since z_{LP} is a lower bound on z^*, the optimal objective function value for the IP, and since z^* must be integer, we can say that

$$z^* \geq \lceil z_{\text{LP}} \rceil,$$

where $\lceil a \rceil$ denotes the smallest integer greater than or equal to a. Therefore, Toregas, et al. propose adding the following cut to (SCLP):

$$\sum_{j \in J} x \geq \lceil z_{\text{LP}} \rceil. \qquad \textbf{(4)}$$

We denote the resulting problem (SCLP-C). The new cut may eliminate some fractional solutions, and the LP relaxation to (SCLP-C) may have an all-integer solution as a result.

For the example in Figure 2, (SCLP-C) does indeed have an integer solution: $x_1 = x_2 = 1$, $x_3 = x_4 = 0$, for example, with $z^* = 2$. (It also has optimal fractional solutions, e.g., $x_j = 0.5$ for all j, but the simplex method would find integer solutions since these represent extreme points of the feasible region.)

Toregas, et al. therefore propose a two-step solution procedure for the SCLP:

1. Solve the LP relaxation of (SCLP). If the optimal solution is integer, STOP.
2. Otherwise, solve the LP relaxation of (SCLP-C) using the optimal objective value from step 1 in the right-hand side of (4).

Even with constraint (4), the LP relaxation may not have an integer solution. Toregas, et al. report that they found no such instance in their computational experiments, though we found several such instances in ours; see Section 2.1.3.1. In fact, Rao (1974) gives two counterexamples; in one, the addition of cut (4) results in a fractional solution, and in the other, the addition of cut (4) results in an integer but non-optimal solution. (See also the reply to Rao's note by Toregas, ReVelle, and Swain (1974)).

Toregas, et al. also discuss the relationship between the SCLP and a variant of the p-median problem in which each customer may only be served by facilities that are within a distance of s. The formulation is obtained simply by forcing the assignment variable to be 0 for facility–customer pairs that are more than s units apart, or, alternately, by indexing the assignment variables for each customer i over facilities j in V_i, as opposed to all facilities j in V. (We omit the formulation here.)

The optimal objective value of this p-median variant changes with s. For sufficiently large s, the objective function value is no different from the p-median without distance constraints; as s decreases, the objective function value increases as a step function; and for sufficiently small s, the problem is infeasible. Toregas, et al. argue that the solution to the SCLP provides some information about the feasibility of this problem. In particular, for a given value of p, the smallest value of s for which the p-median variant is feasible is equal to the smallest value of s for which the SCLP has an optimal objective value of p. On the other hand, the solution to SCLP does not provide any information about the breakpoints of the step function that relates the *p*-median objective to *s*.

2.1.3 Experiments and Variants

In Section 2.1.3.1, we discuss the results of our computational experiment related to (SCLP). In Section 2.1.3.2, we discuss a technique for reducing the problem size of the SCLP, and in Section 2.1.3.4, we discuss a variant involving fixed costs.

2.1.3.1 Computational Experiment

We performed a computational experiment to confirm the results reported by Toregas, et al.—namely, that the LP gap for (SCLP) is small, and that cut (4) produces integer solutions. For each value of $n = 50$, 100, 200, 400, 800, we generated 100 random instances of the SCLP. Parameters were generated as follows:

- x- and y-coordinates were drawn from $U[0,100]$
- Distances were calculated using the Euclidean metric
- The coverage distance *s* was drawn from $U[0,140]$ (140 ≈ maximum possible distance between two points in 100 × 100 grid)

For each instance, we solved the LP relaxation of (SCLP) using CPLEX v10.2.0 to obtain z_{LP}. If the optimal solution to the LP was not integer, we added cut (4) and solved the LP relaxation to (SCLP-C) to obtain $z_{LP\text{-}C}$. If the optimal solution was still not integer, we solved (SCLP) as an IP to obtain z^*. (If either of the LP relaxations resulted in integer solutions, their objective values give us z^*.)

The results are given in Table 1. The columns labeled "% Integer" list the percentage of instances for which the LP relaxation produced an integer optimal solution. The columns labeled "Avg LP Gap" and "Max LP Gap" list the average and maximum, respectively, of the LP gap, measured as $(z_{LP} - z^*) / z^*$ for (SCLP) and $(z_{LP\text{-}C} - z^*) / z^*$ for (SCLP-C).

Table 1. Performance of LP relaxations of (SCLP) and (SCLP-C).

	(SCLP)			(SCLP-C)		
n	% Integer	Avg LP Gap	Max LP Gap	% Integer	Avg LP Gap	Max LP Gap
50	94.0%	0.0068	0.2500	98.0%	0.0000	0.0000
100	87.0%	0.0104	0.1667	92.0%	0.0000	0.0000
200	88.0%	0.0074	0.2500	90.0%	0.0000	0.0000
400	73.0%	0.0217	0.3350	82.0%	0.0003	0.0250
800	76.0%	0.0195	0.2500	82.0%	0.0017	0.0714
Total	83.6%	0.0132	0.3350	88.8%	0.0004	0.0714

The LP gap for (SCLP) is small and tends to be smaller for smaller values of n. The worst gap we found was 33.5% for a problem with $n = 800$. The addition of cut (4) reduces the LP gap substantially (from 0.0132 to 0.0004, on average), but does not guarantee integer solutions—even with the cut, 11.2% of instances had fractional optimal solutions. Several of these instances also had integer optimal solutions, though CPLEX did not find these. In general, CPLEX solved the IP in well under one minute on a laptop computer, even for the largest problems.

2.1.3.2 Row and Column Reduction

The size of (SCLP) can often be reduced substantially by using row- and column-reduction techniques. These methods exploit the coverage structure by eliminating rows and columns that are dominated by others. In particular:

- A facility j_1 dominates another facility j_2 if it covers all of the customers that j_2 does; that is, if $j_2 \in V_i$ implies $j_1 \in V_i$ for all $i \in V$. In this case, there is no reason to open facility j_2 since j_1 covers all the same customers and possibly more. Therefore we can set $x_{j_2} = 0$, or equivalently, eliminate the column corresponding to j_2.
- A customer i_1 dominates another customer i_2 if every facility that covers i_2 also covers i_1; that is, if $V_{i_2} \subseteq V_{i_1}$. In this case, if constraint (2) holds for i_1 it also holds for i_2, and therefore we can eliminate the row corresponding to i_2.

Row and column reduction techniques are appropriate for the SCLP because of the binary nature of coverage. Most facility location problems with distance objectives cannot generally accommodate these techniques, except heuristically, since under most metrics it is impossible for a facility to dominate another, i.e., to be closer to every customer than another facility is.

These techniques were proposed by Toregas and ReVelle (1972). See also Daskin (1995) and Eiselt and Sandblom (2004) for thorough discussions and examples of row- and column-reduction techniques.

2.1.3.3 Facility Fixed Costs

If the facilities each have a different fixed cost f_j, then the problem becomes to choose facilities to cover all demands at minimum possible cost. This problem can be formulated simply by replacing the objective function (1) with

$$\text{minimize} \qquad z = \sum_{j \in V} f_j x_j$$

The SCLP as formulated above is a special case in which $f_j = 1$ for all j. The linear-programming-based solution methods described in Section 2.3 can easily accommodate this variation. So can the Boolean-function approach: at the final step, we simply choose the cover that has the smallest total fixed cost.

2.2 The Maximal Covering Location Problem

2.2.1 Church and ReVelle (1974)

2.2.1.1 Introduction and Formulation

Whereas the SCLP has the form

minimize number of facilities opened

subject to cover all demand,

the maximal covering location problem (MCLP) has the inverse form:

maximize demand covered

subject to limit on number of facilities opened.

The SCLP treats all demand nodes as equivalent since the coverage constraint applies equally to all. In the MCLP, in contrast, nodes are weighted by the demand that they generate, and the objective favors coverage of larger demands over smaller ones.

As the number of allowable facilities increases, the demand covered naturally increases as well. The modeler can plot a tradeoff curve depicting the performance of a range of solutions along these two dimensions; the decision maker can then choose a solution based on this tradeoff.

Our notation in this section is identical to that in Section 2.1, with the addition of two new parameters: a_i is the demand at node i per unit time, and p is the maximum allowable number of facilities. We also introduce a new set of decision variables:

$$y_i = \begin{cases} 1, & \text{if customer } i \text{ is covered by some facility} \\ 0, & \text{otherwise} \end{cases}$$

The MCLP is formulated by Church and ReVelle (1974) as follows:

$$
\begin{array}{llll}
\text{(MCLP)} & \text{maximize} & z = \sum_{i \in V} a_i y_i & & (5) \\
& \text{subject to} & \sum_{j \in V_i} x_j \geq y_i & \forall i \in V & (6) \\
& & \sum_{j \in V} x_j = p & & (7) \\
& & x_j \in \{0,1\} & \forall j \in V & (8) \\
& & y_i \in \{0,1\} & \forall i \in V & (9)
\end{array}
$$

The objective function (5) computes the total demand covered. Constraints (6) prohibit a customer from counting as "covered" unless some facility that covers it has been opened. Constraint (7) requires exactly p facilities to be opened. Constraints (8) and (9) are standard integrality constraints. (In fact, it suffices to relax constraints (9) to $0 \leq y_i \leq 1$ since integer values for the y_i are optimal if the x_j are integer.)

Church and ReVelle cite White and Case (1973) as formulating a similar model to (MCLP) that maximizes the number of demand nodes covered, rather than the total demand. Case and White's model is therefore a special case of the MCLP in which $a_i = 1$ for all i.

Church and ReVelle also present an alternate formulation that uses a new decision variable $\bar{y}_i$ defined as $\bar{y}_i = 1 - y_i$; that is,

$$\bar{y}_i = \begin{cases} 1, & \text{if customer } i \text{ is not covered by any facility} \\ 0, & \text{otherwise} \end{cases}$$

In the alternate formulation, constraints (6) are replaced by

$$\sum_{j \in V_i} x_j + \bar{y}_i \geq 1 \qquad \forall i \in V$$

The revised constraints say that if node i is not covered by any facility ($\sum_{j \in V_i} x_j = 0$), then $\bar{y}_i$ must equal 1. The objective function (5) can be rewritten as

$$\text{maximize} \qquad \sum_{i \in V} a_i (1 - \bar{y}_i) = \sum_{i \in V} a_i - \sum_{i \in V} a_i \bar{y}_i \, , \tag{10}$$

or equivalently,

$$\text{minimize} \qquad \sum_{i \in V} a_i \bar{y}_i \, , \tag{11}$$

since the first term in (10) is a constant. The revised objective (11) minimizes the uncovered demand. The revised formulation is then given by

$$\text{(MCLP2)} \qquad \text{minimize} \qquad z = \sum_{i \in V} a_i \bar{y}_i \tag{12}$$

$$\text{subject to} \qquad \sum_{j \in V_i} x_j + \bar{y}_i \geq 1 \qquad \forall i \in V \tag{13}$$

$$\sum_{j \in V} x_j = p \tag{14}$$

$$x_j \in \{0,1\} \qquad \forall j \in V \tag{15}$$

$$y_i \in \{0,1\} \qquad \forall i \in V \tag{16}$$

The two formulations are mathematically equivalent, as are their LP relaxations.

The MCLP is NP-hard (Megiddo, Zemel, and Hakimi 1983). In the next two sections, we describe heuristic and exact approaches to solving it, all of which are discussed by Church and ReVelle (1974).

2.2.1.2 Heuristic Solution Methods

Like many facility location problems, the MCLP lends itself nicely to greedy heuristics such as the Greedy Adding (GA) heuristic, which Church and ReVelle (1974) credit to Church's (1974) Ph.D. dissertation. The GA heuristic begins with all facilities closed, then opens p facilities in sequence, choosing at each iteration the facility that increases coverage the most. (For a discussion of greedy and other heuristics for facility location problems, see Current, Daskin and Schilling 2002.)

Solutions from the GA heuristic are nested in the sense that all of the facilities in the solution to the p-facility problem are also opened in the solution to the (p+1)-facility problem. Optimal solutions to the MCLP are not, in general, nested in this way. Therefore, Church and ReVelle also suggest an alternate heuristic, called the Greedy Adding with Substitution (GAS) heuristic, which attempts to rectify this problem by allowing an open facility to be closed and a closed facility to be opened at each iteration.

Like any heuristic, the GA and GAS are not guaranteed to find the optimal solution. The GAS, however, tends to perform well in practice, and both heuristics execute very quickly.

2.2.1.3 Linear Programming Approach

Church and ReVelle propose solving (MCLP2) directly using linear programming and branch-and-bound. Like the SCLP, the LP relaxation of the MCLP often yields all-integer solutions: Church and ReVelle report that approximately 80% of their test instances had integer solutions; we found an even higher percentage in our computational tests (Section 2.2.2.1). Branch-and-bound may be applied to resolve fractional solutions to the LP relaxation, but Church and ReVelle also suggest a method that is effective when solving the same problem for consecutive values of p.

The method takes as input a fractional solution to the p-facility problem and an integer solution to the (p – 1)-facility problem. It is effective when the (p – 1)-facility solution covers all but a few nodes. We illustrate the method using an example.

Example 3. Consider an instance of the MCLP for which the total demand across all nodes is 100 units. Suppose we have found an integer solution to the 4-facility problem and that it covers all but two nodes, for a total of 91 demand units covered. The two uncovered nodes (we'll call them 1 and 2) have demands of 4 and 6, respectively. Suppose further that the LP relaxation to the 5-facility problem is fractional and covers 98 demand units. Finally, suppose that the minimum a_i among all nodes i is 3.

Now, the optimal integer solution with $p = 5$ cannot cover all of the nodes, since the LP relaxation has an objective value of 98. In fact, the IP solution may cover at most 97 demand units, since at best it leaves the 3-demand node uncovered. We can create an integer solution to the $p = 5$ problem by adding node 2 to the $p = 4$ solution. Since the $p = 4$ solution covered 91 demands, not including node 2, this new solution covers 91 + 6 = 97 demands. This solution must be optimal for $p = 5$ since 97 is an upper bound on the objective value. An optimal solution for the problem with $p = 6$ can now be found by adding node 1 to the $p = 5$ solution; the resulting solution covers all demands. □

Church and ReVelle refer to this method as the "inspection" method. It can be summarized as follows. Let $z_{\mathrm{IP}}(p)$ be the optimal p-facility objective value of (MCLP), that is, the optimal demand covered by p facilities, and let $z_{\mathrm{LP}}(p)$ be the optimal p-facility objective value of the LP relaxation of (MCLP). We assume that we know the integer optimal solution with p – 1 facilities and that the optimal solution to the LP relaxation with p facilities is not integer. Let $a_{\min} = \min\{a_i : i \in V\}$ and $a_{\Sigma} = \sum_{i \in V} a_i$. We summarize the inspection method in the following theorem. (Church and ReVelle illustrate this method with an example, rather than stating it formally as a theorem.)

Theorem 2. Suppose the following conditions hold:

1. $z_{\text{LP}}(p) < a_{\Sigma}$
2. $z_{\text{IP}}(p-1) + a_i = a_{\Sigma} - a_{\min}$ for some node i that is not covered in the optimal solution to the (p – 1)-facility problem.

Then an optimal solution to the p-facility problem consists of the optimal solution to the (p – 1)-facility problem plus node i.

Church and ReVelle report that, of the 20% of their test instances whose LP relaxations did not have integer solutions, half could be solved using the inspection method. The other half was solved via branch-and-bound.

2.2.1.4 Mandatory Closeness Constraints

Church and ReVelle discuss a variant of the MCLP in which we require that *all* customers be covered within a secondary coverage distance t ($t \geq s$). For example, we might want to maximize the demand covered within 50 miles but require all demands to be covered within 100 miles. This model, known as the MCLP with Mandatory Closeness Constraints, can be viewed as a hybrid between the MCLP and the SCLP since it has a max-coverage objective plus a hard coverage constraint.

This problem can be formulated simply by adding the following constraint to either formulation of the MCLP:

$$\sum_{j \in U_i} x_j \geq 1 \qquad \forall i \in V,$$

where $U_i = \{j \in V : d_{ji} \leq t\}$. The resulting model can be solved using linear programming and branch-and-bound.

Suppose we solve the SCLP and find that, for a given instance, the minimum number of facilities that covers all demand nodes with a coverage distance of t is p^*. Generally there are many optimal solutions to this problem. The MCLP with mandatory closeness constraints gives us a mechanism for choosing among these, by selecting the solution that also maximizes the demands covered within some distance s. In particular, we solve the MCLP with mandatory closeness constraints using p^* as the number of facilities to open and t as the secondary coverage distance.

2.2.2 Experiments and Variants

2.2.2.1 Computational Experiment

We performed a computational experiment to verify Church and ReVelle's claim that the MCLP often results in all-integer solutions. We set n = 50, 100, 200, 400, 800. For each value of n, we generated 100 random instances and tested three different values of p. The random instances were generated as described in Section 2.1.3.1, with one additional parameter:

- Demands a_i were drawn from $U[0,100]$

We solved the LP relaxation of (SCLP2) using CPLEX v10.2.0 and, if the solution was not all integer, we solved the IP. The results are displayed in Table 2. The columns labeled "*p*" gives the value of p in (SCLP2). The column labeled "Avg LP Gap >0" gives the average integrality gap among only those instances with a positive integrality gap, or "—" if there were no such instances. All other columns are interpreted as in Table 1.

Table 2. Performance of LP relaxation of (MCLP2).

n	*p*	% Integer	Avg LP Gap	Avg LP Gap >0	Max LP Gap
50	2	95.0%	0.0011	0.0542	0.0646
	5	96.0%	0.0019	0.0635	0.1109
	8	99.0%	0.0000	—	0.0000
100	2	100.0%	0.0000	—	0.0000
	5	98.0%	0.0002	0.0232	0.0232

	8	98.0%	0.0000	—	0.0000
200	4	96.0%	0.0016	0.0540	0.1293
	10	93.0%	0.0092	0.1308	0.3957
	16	92.0%	0.0028	0.0699	0.1296
400	4	98.0%	0.0000	—	0.0000
	10	92.0%	0.0006	0.0190	0.0280
	16	92.0%	0.0158	0.5254	0.9632
800	4	100.0%	0.0000	—	0.0000
	10	91.0%	0.0002	0.0089	0.0089
	16	89.0%	0.0195	0.4865	0.9704
Total		95.3%	0.0035		0.9704

The LP relaxation MCLP seems to generate integer solutions even more frequently than the SCLP does (at least for our test instances), an average of 95.3% of the time. When it fails to do so, the integrality gap can be quite large, though this is partly a function of the minimization objective, which may have optimal values near zero and hence any suboptimal solution may have a large error on a percentage basis.

Note that some instances had fractional LP solutions but an integrality gap of 0, as evidenced by the fact that some rows have "% Integer" <100% but an average LP gap of 0. For these instances, an optimal integer solution exists for the LP relaxation but CPLEX returned a fractional optimal solution instead.

2.2.2.2 Tradeoff Curve

Figure 3 displays the optimal objective function value of (MCLP2)—the number of demand units uncovered—as p varies for a particular random instance with n = 100 and s = 15. As expected, the uncovered demand decreases as p increases. For $p \geq 18$, all demands are covered. The convex shape is typical of tradeoff curves for the MCLP, meaning that additional facilities provide decreasing marginal returns in terms of additional coverage.

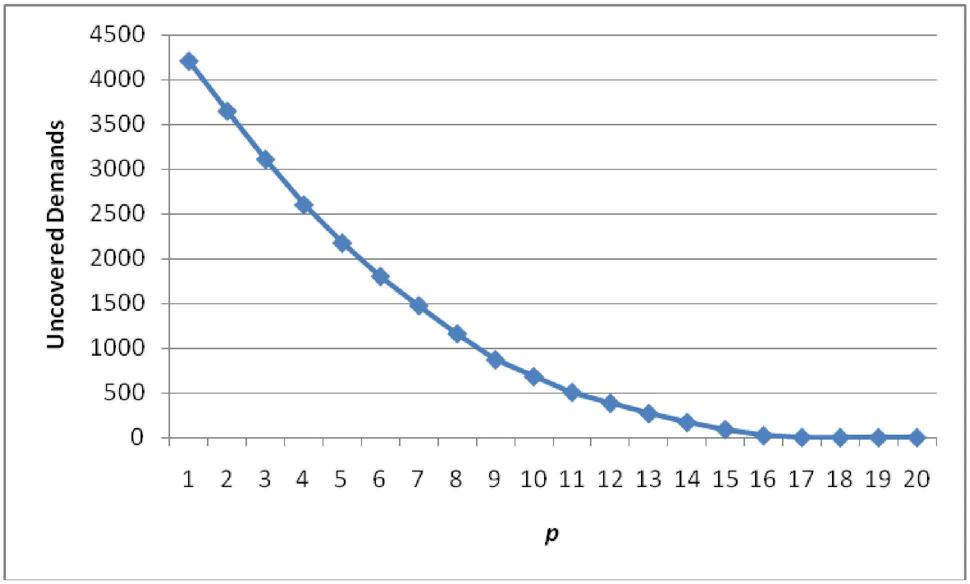

Figure 3. Tradeoff curve: demands uncovered vs. *p*.

2.2.2.3 Lagrangian Relaxation Approach

The MCLP can also be solved using Lagrangian relaxation. The key idea is to remove a set of constraints and add a penalty to the objective function for violating the constraints. The resulting problem is easier to solve but may produce solutions that are infeasible for the MCLP. By adjusting the objective-function penalties iteratively, the solutions found approach the optimal solution for the MCLP. The use of Lagrangian relaxation for the MCLP was detailed by Galvão and ReVelle (1996), although Daskin, et al. (1989) also report computational results from a similar method without providing details. See Fisher (1981, 1985) for an excellent overview of Lagrangian relaxation.

We illustrate the Lagrangian relaxation method using formulation (MCLP), though it can also be applied to (MCLP2). We relax constraints (6) using Lagrange multipliers λ_i to obtain the following Lagrangian subproblem:

$$\text{(MCLP-LR)} \quad \text{maximize} \quad z = \sum_{i \in V} a_i y_i + \sum_{i \in V} \lambda_i \left(\sum_{j \in V_i} x_j - y_i \right) \tag{17}$$

$$= \sum_{i \in V} (a_i - \lambda_i) y_i + \sum_{j \in V} \left(\sum_{i \in V : j \in V_i} \lambda_i \right) x_j$$

$$\text{subject to} \quad \sum_{j \in V} x_j = p \tag{18}$$

$$x_j \in \{0,1\} \qquad \forall j \in V \tag{19}$$

$$y_i \in \{0,1\} \qquad \forall i \in V \tag{20}$$

This problem decouples by *x* and *y* since there are no constraints involving both sets of variables. As a result, it can be solved easily. The optimal *y*-values are given by

$$y_i = \begin{cases} 1, & \text{if } a_i - \lambda_i > 0, \\ 0, & \text{otherwise} \end{cases}$$

To find the optimal *x*-values, we set $x_j = 1$ for the *p* facilities with the largest values of $\sum_{i \in V : j \in V_i} \lambda_i$. The optimal objective value of (MCLP-LR) provides an upper bound on that of (MCLP). Feasible (lower bound) solutions can be found by opening the *p* facilities that are opened in the upper-bound solution and setting $y_i = 1$ for each customer *i* that is covered by some opened facility. Lagrange multipliers can be updated using subgradient optimization, and branch-and-bound can be used if the Lagrangian procedure fails to yield a suitably small optimality gap; see Daskin (1995) for more details. Daskin, et al. (1989) report that the procedure works quite well, especially when the lower-bound heuristic is supplemented by a substitution heuristic.

2.2.2.4 Budget Constraints

We can incorporate fixed costs into the model in a similar manner as we did for the SCLP in Section 2.1.3.3. Here, the fixed cost appears in the constraints rather than the objective function. In particular, we replace constraint (7) or (14) with

$$\sum_{j \in V} f_j x_j \le B,$$

where B is a budget imposed exogenously on the total fixed costs. This constraint can be easily handled by the linear programming approach in Section 2.2.1.3, but it somewhat complicates the Lagrangian approach in Section 2.2.2.3 since determining the optimal x values now requires us to solve the following knapsack problem:

$$\begin{array}{lll} \text{maximize} & \sum_{j \in V} \left(\sum_{i \in V: j \in V_i} \lambda_i \right) x_j & \\ \text{subject to} & \sum_{j \in V} f_j x_j \le B & \\ & x_j \in \{0,1\} & \forall j \in V \end{array}$$

Although this problem can be solved quite quickly using modern codes, it is still NP-hard, and it may slow the Lagrangian procedure significantly.

2.2.2.5 Relationship to *p*-Median Problem

The MCLP can be formulated as a special case of the *p*-median problem (PMP) through a simple transformation of the distance matrix. In particular, we set

$$d_{ji} = \begin{cases} 0, & \text{if } j \in N_i \\ 1, & \text{otherwise} \end{cases}$$

That is, we redefine the distance metric so that the distance from node j to node i is 0 if j covers i and 1 otherwise. The PMP is then formulated as usual (see, e.g., Daskin 1995). The optimal solution will cover as many demand units as possible using p facilities. Any algorithm for the PMP can then be applied to solve the MCLP.

3. Extensions

The literature contains many enhancements to the SCLP and MCLP. In this section, we focus in particular on generalizations of the notion of coverage. One common criticism of the SCLP and MCLP is that they assume that all customers within a facility's coverage radius can be served by the facility, and served equally. In practice, facilities are not always available when needed, especially in the public-sector arena where facilities may represent ambulances, fire crews, etc. One approach to this issue is *backup coverage*, in which customers are required or encouraged to be covered by more than one open facility. Another approach is *expected coverage*, which accounts for probabilistic information. Moreover, in many cases the coverage benefit changes as the distance between a customer and its assigned facility changes. This dependency is captured by the notion of *gradual coverage*. We briefly discuss models for backup, expected, and gradual coverage in the next three subsections. For thorough reviews of backup and expected coverage models, see Daskin, Hogan, and ReVelle (1988) or Berman and Krass (2002).

3.1 Backup Coverage Models

Both the SCLP and the MCLP have been extended to consider solutions in which customers are covered by more than one facility. One may *require* backup coverage in order for a customer to count as "covered," or one may simply *reward* solutions for backup coverage.

3.1.1 Required Backup Coverage

It is simple to formulate a required-backup version of either the SCLP or the MCLP. In the SCLP, we simply modify constraints (2) to read

$$\sum_{j \in V_i} x_j \geq m \qquad \forall i \in V ,$$

where m is the desired number of times that each customer is to be covered. In the MCLP, we can replace constraints (6) with

$$\sum_{j \in V_i} x_j \geq m y_i \qquad \forall i \in V .$$

Then y_i must equal 0 unless at least m facilities that cover customer i are open. This constraint is likely to weaken the LP relaxation of (MCLP), however, as is typical of such "big-M" constraints.

3.1.2 Rewards for Backup Coverage

We focus on models in which $m = 2$. Extensions to these models to consider $m > 2$ are straightforward but often yield weaker LP relaxations, as discussed above. Let

$$w_i = \begin{cases} 1, & \text{if customer } i \text{ is covered by two or more facilities} \\ 0, & \text{otherwise} \end{cases}$$

The models formulated below contain a reward in the objective function for each customer that is covered twice. However, the backup-coverage reward is strictly a secondary objective; in no case should a solution with more facilities have a better objective than one with fewer facilities, even if it has better backup coverage.

Daskin and Stern (1981) propose the following model for the SCLP with backup coverage:

$$\text{(SCLP-BC)} \quad \text{minimize} \quad z = (|V|+1)\sum_{j\in V} x_j - \sum_{i\in V} w_i \quad \textbf{(21)}$$

$$\text{subject to} \quad \sum_{j\in V_i} x_j - w_i \geq 1 \qquad \forall i \in V \quad \textbf{(22)}$$

$$x_j \in \{0,1\} \qquad \forall j \in V \quad \textbf{(23)}$$

$$w_i \in \{0,1\} \qquad \forall i \in V \quad \textbf{(24)}$$

The objective function (21) enforces the hierarchical nature of the primary objective (minimizing the number of facilities) and the secondary one (maximizing twice-covered customers). It does so by multiplying the primary objective by a constant large enough that even if the primary objective is as small as possible (equal to 1), the secondary objective can never exceed it. Therefore, the solution will never open more facilities than necessary solely to improve the secondary objective. Constraints (22) require each customer to be covered at least once and prohibit w_i from equaling 1 unless customer i is covered at least twice.

Another advantage of this formulation is that its solutions avoid facilities that are dominated by others in the sense described in Section 2.1.3.2. As a result, the LP relaxation to (SCLP-BC) is more likely to have all-integer solutions than that of (SCLP) is. (See Daskin and Stern 1981 for justifications of both of these claims.)

A similar hierarchical version of the MCLP was introduced by Storbeck (1982) and reformulated by Daskin, Hogan, and ReVelle (1988). We modify their formulation somewhat in what follows.

$$\text{(MCLP-BC)} \quad \text{maximize} \quad z = (|V|+1)\sum_{i\in V} a_i y_i + \sum_{i\in V} w_i \quad \textbf{(25)}$$

$$\text{subject to} \quad \sum_{j\in V_i} x_j - y_i - w_i \ge 0 \qquad \forall i \in V \quad \textbf{(26)}$$

$$\sum_{j\in V} x_j = p \quad \textbf{(27)}$$

$$x_j \in \{0,1\} \qquad \forall j \in V \quad \textbf{(28)}$$

$$y_i \in \{0,1\} \qquad \forall i \in V \quad \textbf{(29)}$$

$$w_i \in \{0,1\} \qquad \forall i \in V \quad \textbf{(30)}$$

The objective function (25) maximizes a sum of the primary coverage (first term) and backup coverage (second term); the weight on the first term ensures that primary coverage will never be sacrificed in order to achieve backup coverage. Note that the secondary coverage objective considers *nodes* covered, rather than *demand units* covered. This is required in order for the weighting to achieve the desired hierarchy. Constraints (26) stipulate that customer i may be considered covered ($y_i = 1$) only if at least one facility in V_i is open, and may be considered twice covered ($w_i = 1$) only if two such facilities are open. Since the objective function coefficient for y_i is greater than that for w_i, the model will always set $y_i = 1$ before it sets $w_i = 1$, thus ensuring the desired coverage hierarchy.

3.2 Expected Coverage Models

The class of expected coverage models is descended primarily from the Maximum Expected Covering Location Problem (MEXCLP) introduced by Daskin (1982). Daskin's primary application is in the siting of emergency medical service (EMS) vehicles. The MEXCLP maximizes the *expected* coverage of each node, defined using probabilistic information about facility availability, subject to a constraint on the number of facilities.

The MEXCLP assumes that the average system-wide probability that a given facility (vehicle) is busy is given by q. If a customer is covered by k facilities, then the probability that all those facilities are busy at a given point in time is given by q^k, and the probability that at least one facility is available is $1 – q^k$. The MEXCLP defines new variables to keep track of the number of covering facilities for each customer: let

$$y_{im} = \begin{cases} 1, & \text{if customer } i \text{ is covered by at least } m \text{ facilities} \\ 0, & \text{otherwise} \end{cases}$$

for all $i \in V$ and $m = 1, \ldots, p$. Note that if customer i is covered by *exactly* k facilities, then $y_{im} = 1$ for $m = 1, \ldots, k$ and $y_{im} = 0$ for $m = k + 1, \ldots, p$. Then

$$\sum_{m=1}^{p}(1-q)q^{m-1}y_{im} = \sum_{m=0}^{k-1}(1-q)q^{m} = 1-q^{k}$$

using a standard formula for geometric sums. In other words, the first summation in the equation above expresses the probability that customer i is covered by an available facility in terms of the decision variables y_{im}. Using this approach, Daskin formulates the MEXCLP as follows:

$$\text{(MEXCLP)} \quad \text{maximize} \quad z = \sum_{i\in V}\sum_{m=1}^{p}(1-q)q^{m-1}a_i y_{im} \tag{31}$$

$$\text{subject to} \quad \sum_{m=1}^{p} y_{im} - \sum_{j\in V_i} x_j \le 0 \qquad \forall i \in V \tag{32}$$

$$\sum_{j\in V} x_j = p \tag{33}$$

$$x_j \in \{0,1\} \qquad \forall j \in V \tag{34}$$

$$y_i \in \{0,1\} \qquad \forall i \in V \tag{35}$$

The objective function (31) calculates the expected coverage. Constraints (32) allow the total number of y_{im} variables, for fixed i, to be no more than the total number of opened facilities that cover i. At first it may seem that the model needs a constraint of the form

$$y_{im} \le y_{i,m+1} \qquad \forall i \in V, m = 1, \ldots, p-1$$

in order to ensure that y_{im} is set to 1 for the correct values of m; that is, for the k smallest values of m, where k is the number of opened facilities that cover i. However, such a constraint is not necessary since the objective function coefficient is larger for smaller values of m; the model will automatically set $y_{im} = 1$ for the k smallest values of m.

Daskin (1983) proposes a heuristic for the MEXCLP based on node exchanges, and several metaheuristics have been proposed subsequently (e.g., Aytug and Saydam 2002, Rajagopalan *et al*. 2007).

The primary criticism that has been leveled at the MEXCLP concerns the assumption of a uniform system-wide availability probability, since availability might vary based on geographic area, or on the demand assigned to each facility. ReVelle and Hogan (1989) address this concern in the Maximum Availability Location Problem (MALP), a chance-constrained version of the MCLP. They formulate two versions of the model, one in which the availability probability is assumed to be the same throughout the system; the main difference between this model and the MEXCLP is that the MALP maximizes the number of demand units that are covered with at least a certain probability, whereas the MEXCLP includes the expected coverage in the objective. ReVelle and Hogan's second MALP model estimates the busy probability separately for each customer by assuming that facilities within the coverage radius of a given customer are available only to that customer. Obviously this assumption is not true but provides an easy, and fairly accurate, estimate of the availability probability. The two models are nearly identical once the availability probabilities are calculated. Galvão, Chiyoshi, and Morabito (2005) present a framework that attempts to unify the MEXCLP and MALP.

Batta, Dolan, and Krishnamurthy (1989) embed Larson's (1974, 1975) hypercube queuing model into MEXCLP to compute the availability probabilities endogenously. They find that their model disagrees substantially with MEXCLP in terms of expected coverage predicted but nevertheless results in similar sets of facilities chosen. Marianov and ReVelle (1996) formulate a version of the MEXCLP that endogenously calculates the availability property using a queuing model at each facility. The region around each customer node is treated as an $M/M/s/s$ queue, where s is the number of servers located within the coverage radius. Their model implicitly assumes that that the call rate in the neighborhood is not too different from that in adjacent neighborhoods. The resulting model is structurally similar to the MALP but uses different (but pre-computable) values for the coverage radius.

3.3 Gradual Covering Models

The models discussed in this chapter so far all assume that coverage is a binary concept: either a customer is covered or it isn't, and the distance from the customer to the covering facility is irrelevant. In practice, though, customers that are located very close to a facility (e.g., a fire station) may be served better than those located farther away, even if both customers are within the nominal coverage radius. In this case, the benefit from coverage is decreasing with the customer–facility distance, as illustrated in Figure 4(a). Moreover, some facilities (e.g., garbage dumps) are most beneficial when they are close (to reduce transportation costs) but not too close (to reduce odors and truck traffic), as illustrated in Figure 4(b).

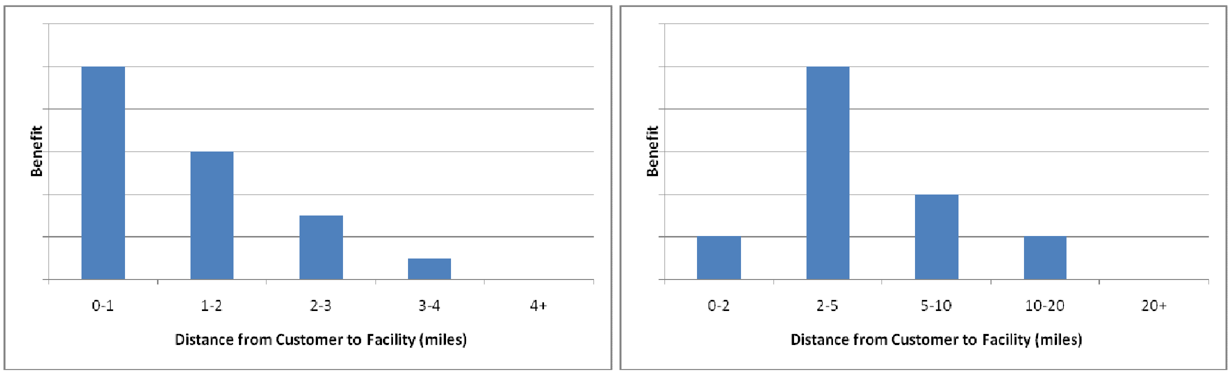

(a) (b)

Figure 4. Benefit of coverage vs. distance: (a) strictly decreasing, (b) non-monotonic.

Church and Roberts (1983) introduce the Weighted Benefit Maximal Coverage (WBMC) Model, which extends the MCLP to accommodate non-binary coverage benefits. The objective is to maximize the sum of all customers' coverage benefits (defined as the benefit per unit of demand times the demand at that customer) subject to a constraint on the number of facilities located. The formulation is a relatively straightforward modification of (MCLP) and includes a coverage variable (y) and a constraint for each customer–distance pair. (Each "distance" is really a range of distances, as in Figure 4.) The number of variables and constraints therefore grows linearly with the number of distance ranges. If the benefits are not monotonically decreasing with the distance, as in Figure 4(b), then an additional set of constraints is required to ensure that customers are assigned to their nearest opened facilities, a property that is automatic if benefits are monotonically decreasing. The resulting formulations are more complex than (MCLP), but Church and Roberts find that they still retain their “integer-friendliness”: the LP relaxation is generally very tight, and often all-integer.

4. Conclusions and Future Research Directions

In this chapter we have discussed two classical models for locating facilities to ensure coverage of customer nodes. One model, the SCLP, requires *every* customer to be covered and does so with the minimum number of facilities, while the other, the MCLP, maximizes the demand covered subject to a limit on the number of facilities. Both models have garnered considerable attention in the location theory literature, and both models (and their extensions) have been widely applied in practice, especially in public-sector applications such as EMS location.

The SCLP and MCLP are both reasonably easy to solve, in the sense that modern general-purpose IP solvers such as CPLEX can solve problems with hundreds or thousands of nodes to optimality in a few minutes on a desktop computer. This stems, in part, from the fact that the LP relaxations of both problems tend to be tight, and even yield integer optimal solutions for a large percentage of instances. Therefore, although these problems are NP-hard, they are among the easiest problems in that class.

On the other hand, many of the extensions of these models are much more computationally challenging. Daskin’s (1982) MEXCLP model, for example, or the queuing-based congestion models discussed by Berman and Krass (2002), have more complex structures than the SCLP or MCLP and therefore cannot be solved using off-the-shelf solvers, except for small instances. One important direction for future research, therefore, is the development of effective, accurate algorithms and heuristics for extensions of the SCLP and MCLP.

Of particular interest are stochastic and robust variants of coverage models. Although the literature on stochastic facility location models is extensive (see Snyder 2006 for a review), most such models consider cost-based objectives rather than coverage-based ones. (Notable exceptions are the expected-coverage models described in Section 3.2, and their variants.) An important topic for future study is therefore the incorporation of stochastic elements—such as demands, travel times, server availabilities, and supply disruptions—into coverage models. The resulting models are likely to be significantly more complex than their deterministic counterparts, but the stochastic programming and robust optimization literatures are vast, and many of their more sophisticated tools have yet to be tapped by the location science community.

The distinction between cost- and coverage-based models made in the previous paragraph is an important one since it is often equivalent to the distinction between private- and public-sector applications—the former is primarily concerned with cost minimization while the latter is often encouraged or mandated to provide adequate coverage to all demand locations (ReVelle, Marks, and Liebman 1970). Public-sector and humanitarian applications have gained increased attention in the OR community in recent years—for example, the 2008 INFORMS Annual Meeting featured "Doing Good with OR" as a central theme, as did the February, 2008 issue of *OR/MS Today*. The application of coverage models to EMS and other services has been a public-OR success story for decades, and the renewed interest provides an opportunity for existing and new coverage models to be applied for the public good.

References

H. Aytug and C. Saydam. Solving large-scale maximum expected covering location problems by genetic algorithms: A comparative study. *European Journal of Operational Research*, 141(3):480–494, 2002.

R. Batta, J. M. Dolan, and N. N. Krishnamurthy. The maximal expected covering location problem: Revisited. *Transportation Science*, 23(4):277–287, 1989.

O. Berman and D. Krass. Facility location problems with stochastic demands and congestion. In Z. Drezner and H. W. Hamacher, editors, *Facility Location: Applications and Theory*, chapter 11. Springer-Verlag, New York, 2002.

J. Bramel and D. Simchi-Levi. On the effectiveness of set covering formulations for the vehicle routing problem with time windows. *Operations Research*, 45(2):295–301, 1997.

R. Church. Synthesis of a Class of Public Facilities Location Models. PhD thesis, The Johns Hopkins University, Baltimore, MD, 1974.

R. Church and C. ReVelle. The maximal covering location problem. *Papers of the Regional Science Association*, 32:101–118, 1974.

R. L. Church and B. Meadows. Location modelling using maximum service distance criteria. *Geographical Analysis*, 11:358–373, 1979.

R. L. Church and K. L. Roberts. Generalized coverage models and public facility location. *Papers in Regional Science*, 53(1):117-135, 1983.

R. L. Church, D. M. Stoms, and F. W. Davis. Reserve selection as a maximal covering location problem. *Biological Conservation*, 76(2):105–112, 1996.

J. Current, M. S. Daskin, and D. Schilling. Discrete network location models. In Z. Drezner and H. W. Hamacher, editors, *Facility Location: Applications and Theory*, chapter 11. Springer-Verlag, New York, 2002.

M. S. Daskin. Application of an expected covering model to emergency medical service system design. *Decision Sciences*, 13:416–439, 1982.

M. S. Daskin. A maximum expected covering location model: Formulation, properties and heuristic solution. *Transportation Science*, 17(1):48–70, 1983.

M. S. Daskin. *Network and Discrete Location: Models, Algorithms, and Applications*. Wiley, New York, 1995.

M. S. Daskin, A. E. Haghani, M. Khanal, and C. Malandraki. Aggregation effects in maximum covering models. *Annals of Operations Research*, 18:115–140, 1989.

M. S. Daskin, K. Hogan, and C. ReVelle. Integration of multiple, excess, backup, and expected covering models. *Environment and Planning B*, 15(1):15–35, 1988.

M. S. Daskin and E. H. Stern. A hierarchical objective set covering model for emergency medical service vehicle deployment. *Transportation Science*, 15:137–152, 1981.

D. J. Eaton, M. S. Daskin, D. Simmons, B. Bulloch, and G. Jansma. Determining emergency medical service vehicle deployment in Austin, Texas. *Interfaces*, 15(1):96–108, 1985.

H. A. Eiselt and C.-L. Sandblom. *Decision Analysis, Location Models, and Scheduling Problems*. Springer, New York, 2004.

M. L. Fisher. The Lagrangian relaxation method for solving integer programming problems. *Management Science*, 27(1):1–18, 1981.

M. L. Fisher. An applications oriented guide to Lagrangian relaxation. *Interfaces*, 15(2):10–21, 1985.

J. Flynn and S. Ratick. A multiobjective hierarchical covering model for the essential air services program. *Transportation Science*, 22(2):139–147, 1988.

R. D. Galvão, F. Y. Chiyoshi, and R. Morabito. Towards unified formulations and extensions of two classical probabilistic location models. *Computers & Operations Research*, 32:15–33, 2005.

R. D. Galvão and C. ReVelle. A Lagrangean heuristic for the maximal covering location problem. *European Journal of Operational Research*, 88:114–123, 1996.

M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.

J. M. Gleason. A set covering approach to bus stop location. *Omega*, 3(5):605–608, 1975.

S. L. Hakimi. Optimum locations of switching centers and the absolute centers and medians of a graph. *Operations Research*, 12(3):450–459, 1964.

S. L. Hakimi. Optimum distribution of switching centers in a communication network and some related graph theoretic problems. *Operations Research*, 13(3):462–475, 1965.

R. C. Larson. A hypercube queuing model for facility location and redistricting in urban emergency services. *Computers and Operations Research*, 1:67–95, 1974.

R. C. Larson. Approximating the performance of urban emergency service systems. *Operations Research*, 23(5):845–868, 1975.

V. Marianov and C. ReVelle. The queueing maximal availability location problem: A model for the siting of emergency vehicles. *European Journal of Operational Research*, 93:110–120, 1996.

N. Megiddo, E. Zemel, and S. L. Hakimi. The maximum coverage location problem. *SIAM Journal of Algebraic and Discrete Methods*, 4(2):253-261, 1983.

L. K. Nozick and M. A. Turnquist. Inventory, transportation, service quality and the location of distribution centers. *European Journal of Operational Research*, 129(2):362–371, 2001.

H. K. Rajagopalan, F. E. Vergara, C. Saydam, and J. Xiao. Developing effective meta-heuristics for a probabilistic location model via experimental design. *European Journal of Operational Research*, 177:83–101, 2007.

A. Rao. Counterexamples for the location of emergency service facilities. *Operations Research*, 22(6):1259–1261, 1974.

C. ReVelle. Facility siting and integer-friendly programming. *European Journal of Operational Research*, 65(2):147-158, 1993.

C. ReVelle and K. Hogan. The maximum availability location problem. *Transportation Science*, 23(3):192–200, 1989.

C. ReVelle, D. Marks, and J. C. Liebman. An analysis of private and public sector location models. *Management Science*, 16(11):692–707, 1970.

D. A. Schilling, C. ReVelle, J. Cohon, and D. J. Elzinga. Some models for fire protection locational decisions. *European Journal of Operational Research*, 5(1):1–7, 1980.

L. V. Snyder. Facility location under uncertainty: A review. *IIE Transactions*, 38(7):537-554, 2006.

J. E. Storbeck. Slack, natural slack and location covering. *Socioeconomic Planning Sciences*, 16(3):99–105, 1982.

C. Toregas and C. ReVelle. Optimal location under time or distance constraints. *Papers of the Regional Science Association*, 28:133–143, 1972.

C. Toregas, C. ReVelle, and R. Swain. Reply to Rao's note on the location of emergency service facilities. *Operations Research*, 22(6):1262–1267, 1974.

C. Toregas, R. Swain, C. ReVelle, and L. Bergman. The location of emergency service facilities. *Operations Research*, 19(6):1363–1373, 1971.

J. White and K. Case. On covering problems and the central facilities location problem. Unpublished paper, Virginia Polytechnic Institute and State University, Blacksburg, VA, 1973.