Interfaces
Vol. 34, No. 3, May–June 2004, pp. 208–214
ISSN 0092-2102 | EISSN 1526-551X | 04 | 3403 | 0208

informs®
DOI 10.1287/inte.1030.0069

Schlumberger Optimizes Receiver Location for Automated Meter Reading

Srinagesh Gavirneni
Operations and Decision Technologies, Kelley School of Business, Indiana University, Bloomington, Indiana 47405, sgavirne@indiana.edu

Lloyd Clark
Austin Technology Center, Schlumberger, Austin, Texas 78726, ldclark@sugar-land.oilfield.slb.com

Gábor Pataki
Department of Operations Research, University of North Carolina, Chapel Hill, North Carolina 27599, pataki@email.unc.edu

In its efforts to provide automated meter-reading services to the utility industry, Schlumberger faces the problem of deploying receivers efficiently. In a geographic region, the problem is to install the minimum number of receivers on existing utility poles so that all wireless meters in that region can transmit their readings to at least one. Schlumberger first encountered this large-scale problem in a project it ran for Illinois Power. It found solving this problem manually very time consuming, and it had no ability to evaluate the robustness of the resulting solution. We proposed and implemented an optimization-based approach that reduced the duration and cost of implementations. In addition, we gave planners the ability to answer what-if questions.

Key words: programming: integer; industries: electric.
History: This paper was refereed.

Schlumberger provides technology services, such as oil-field services, resource management, transactions-based technology and associated systems, and semiconductor test equipment. It has offices in over 100 countries, employs more than 50,000 people, and had revenues of $8.75 billion in 1999. Schlumberger's resource-management services (RMS) division is a leading provider of electricity, gas, and water meters with more than 200 million units installed around the world. In 1999, RMS had revenues of $1.38 billion, employed over 15,000 people, and had sales offices and factories in more than 35 countries. In the late 1990s, it started offering professional business services for the utility industry. Through consulting, meter deployment and management, data collection and processing, and information analysis, RMS helps clients optimize networks, achieve operating efficiency, and increase customer loyalty.

Electricity is the most versatile, widely used form of energy available. Unlike other forms though, it cannot be stored. So providers must carefully monitor demand and meticulously control production. Production companies that generate, transmit, distribute, or trade electricity need accurate measurements of production and consumption. They must analyze data promptly to ensure that they manage this energy resource correctly and that power is always available when people want it. As the industry becomes increasingly deregulated, to remain competitive, to reduce costs, and to increase efficiency, utilities and municipalities want to monitor their customers' use of power closely. They increasingly rely on such companies as Schlumberger RMS to develop solutions to meet these operations and maintenance needs. Such systems can include a wide range of value-added services, such as automated meter reading, in-home messaging, security controls, remote disconnect, and other two-way services, that can give a utility provider a competitive advantage.

We developed a solution to the problem of optimally locating hardware for automated meter reading (AMR). The meters send out their readings as radio-frequency (RF) signals, which are captured by a receiver nearby and relayed to a central data warehouse. A receiver can capture readings from at most 540 meters within the range of a few hundred yards. This range of coverage cannot be increased without adversely affecting the life span of the battery in the meter. In most cases, the ideal place for receivers is high on the electricity poles the utility company owns. The coverage area of a pole is the area from which meters can successfully transmit to a receiver located on that pole. The coverage area depends on the properties of the pole (height, material, and so forth) and its surroundings (buildings, trees, rural or urban, and so forth). To minimize the cost of the receivers, we must solve the following problem: Select the minimum number of poles, and assign every meter to a selected pole without exceeding the capacity (that is, meters assigned to them) of the poles.

In the late 1990s, Schlumberger had projects underway to install receivers to automatically read millions of meters spread over thousands of square miles. In 1998, Schlumberger was installing AMR for a major energy utility in Illinois. The network was to cover 15,000 square miles and collect data from approximately 1.1 million gas and electricity meters. In 1999, Schlumberger signed a 15-year agreement to provide metering data and asset-management services to PECO Energy. Under this agreement, it is responsible for collecting data from more than 2 million electric and gas meters in the PECO service area. Around the same time, Schlumberger was awarded a contract to supply the Pittsburgh Water and Sewer Authority with AMR services for approximately 83,000 residential, commercial, and industrial customers. In all these projects, the firm must solve the receiver-location problem, and thus it needed a solution procedure that can give good solutions reasonably quickly.

The Problem

We are given the locations of m meters numbered $1, 2, \ldots, m$ and of p poles numbered $1, 2, \ldots, p$. We denote the set of meters by M and the set of poles by P. Each pole has a coverage area associated with it, the area from which meters can successfully transmit to a receiver located on that pole. To minimize the cost of the required receivers, we must determine the minimum number of poles and assign every meter to one (and only one) selected pole. In addition, each receiver has a predefined capacity limit, K, which restricts the number of meters assigned to that receiver. The objective is to select, without violating the capacity restriction, the minimal number of poles such that all the meters are covered.

A Simple Example

For example, consider the following case of 10 meters, eight poles, and their coverage areas (Table 1). To someone unfamiliar with how quickly the possible number of combinations explodes, this problem may seem easy. The first solution approach that comes to mind is one that selects the poles in a greedy manner. To use the greedy procedure to solve this problem, we compute the number of meters each pole covers. These numbers are listed in the column labeled Step 0. Because pole 6 covers the highest number, we pick that pole. In Step 1, for each unselected pole we compute the number of meters in its area that have not been covered by previously selected poles. We pick the pole with the highest number. We proceed in this manner, breaking ties arbitrarily, until we have covered all the meters. For this example, the greedy solution consists of poles 1, 3, 4, 5, and 6. Is that the best possible solution for this problem? To establish optimality, we have to evaluate all the possible combinations for this problem; the combinations that contain at least three poles number 219. By enumerating all the possible combinations, we establish that the solution consisting of poles 2, 3, 4, and 8 is optimal. Thus, the greedy solution is not optimal, and we cannot confidently assume that for large problems the greedy solution will be even close to optimal.

Pole	Meters covered	Step 0	Step 1	Step 2	Step 3	Step 4	Step 5
1	1, 2, 3	3	2	1*	—	—	—
2	2, 3, 9	3	2	1	0	0	0
3	5, 6, 7	3	2	1	1*	—	—
4	7, 9, 10	3	1	1	1	1*	—
5	3, 6, 8	3	3*	—	—	—	—
6	1, 4, 7, 9	4*	—	—	—	—	—
7	4, 5, 9	3	1	1	1	0	0
8	1, 4, 8	3	1	0	0	0	0

Table 1: The table shows data for an example containing 10 meters and eight poles. A greedy procedure selects, in each step, the pole (denoted by an asterisk) that has the highest number of as yet unassigned meters in its coverage area. Such greedy procedures do not guarantee optimality and can often generate inferior solutions. In this case, the greedy procedure selects five (1, 3, 4, 5, and 6) poles while the optimal solution has only four (2, 3, 4, and 8) poles.

Existing Solution Method

In early 1998, Schlumberger signed a contract with Illinois Power and was planning the deployment of these receivers in Decatur, Illinois. To evaluate the proposed solution approaches, we extracted a small instance (containing 4,208 meters and 2,393 poles) of the problem from the larger problem consisting of 116,600 meters and 20,631 poles. Throughout this paper, we use *small problem* and *large problem* to refer to these two problem instances. Data (latitude and longitude) on the geographical location of the poles and meters was available.

Planners first tried to solve this problem using a geographical information system (GIS). They input the locations of the meters and the poles into the system and printed them out on paper. The planners then manually tried to locate the subset of poles that would cover all the meters. While they obtained a solution to the problem, they could not judge its quality. They also realized that this method was not scalable and would not work well for large problems. They sought the help of a senior research scientist in the Austin Technology Center. The scientist, an expert in digital signal processing, had limited experience with optimization. He used MATLAB to implement the greedy procedure and automatically generated a solution to the small problem. The planners could easily improve the solution by manual inspection. They used the greedy procedure followed by manual manipulation successfully to solve the small problem and found a solution containing 171 poles. It was in fact an optimal solution to the small problem. However, after all the effort they put into obtaining this solution, they realized that the approach had the following deficiencies:

(1) It was extremely time consuming and took three to four weeks to solve the small problem. The planners expected that it would take no less than six months to solve the large problem, far too long.

(2) Because the amount of data (about the locations of poles and meters) used was very large, there was a high probability that some of that data was wrong. If so, they would have to re-solve the problem with the corrected data, consuming even more time.

(3) The approach was unable to handle the capacity limitations of receivers. The total number of meters assigned to a pole (or receiver) is called the receiver's load.

In practice, poles should not have large loads because (1) at any time, a receiver can "talk" to only one meter. When a receiver has a large load, many meter-to-receiver calls will find it busy, increasing the number of calls and speeding up the depletion of the meter's battery and its need for maintenance visits. Also, (2) as the number of customers increases in the area, it is trivial to assign a new meter to an already installed receiver but cumbersome to install a new receiver. While each receiver has a theoretical capacity limit of 540, in practice the limit is close to 300. In the small problem, this capacity limit played no role because the maximum number of meters in a pole's coverage area was only 120. However, for the large problem, the planners had to incorporate a capacity restriction.

(4) Planners could not answer such what-if questions as (1) How does the solution change if we use only those poles that are taller than 25 feet? and (2) What is the effect of not using poles made of wood?

Because of these deficiencies the planners looked for other ways to solve the problem. Another scientist, who specialized in optimization and operations research at Schlumberger, became involved in the discussions in June 1998. Soon, we developed a mathematical-programming approach that proved effective in addressing most of these concerns.

An Integer-Programming Approach

The pole-selection problem belongs to a well-studied class of optimization models, called facility-location problems (Daskin 1995). When the poles are assumed to have infinite capacity (that is, they can accommodate all the meters within their coverage areas), the problem belongs to the simplest subclass, called set-covering problems (Dell'Amico et al. 1997).

Problem Formulation

We formulated the pole-selection problem as an optimization problem with a linear objective function and linear constraints, a linear program. Because of their structure, linear programs can be solved very efficiently. When a linear program contains integer variables, it is called an integer program. Formulating the problem as an integer program requires a mathematical representation of the problem.

Recall that M is the set of meters numbered $1, 2, \ldots, m$ and P is the set of poles numbered $1, 2, \ldots, p$. Each pole j has a set of meters, $C_j \subset M$, that it can cover. We want to find a set of poles, S, such that $S \subset P$, $\bigcup_{j\in S} C_j = M$, and the cardinality of S is the least possible. This is a classic set-covering problem, which is known to be NP-hard (Nemhauser and Wolsey 1988). The number of feasible solutions to this problem grows exponentially with the size of the problem, making it very difficult to solve large problems to optimality even with a fast computer. (We give an integer-programming formulation of this problem in the Appendix.)

Solving the Small Problem

To demonstrate our approach's effectiveness in reducing the time and effort needed to solve this problem, we started with the small problem consisting of 4,208 meters and 2,393 poles. To use the integer-programming formulation, we sought the feasible combination of the poles and meters relying on data in MATLAB. We took two or three days to develop the necessary MATLAB procedures. Of possible combinations of meters and poles, only 476,769 were feasible because of distance considerations. The integer-programming formulation contained 2,393 binary variables, 4,208 constraints, and 476,769 nonzero coefficients. We submitted this problem to CPLEX and obtained a solution of 171 in about five minutes of computational time, the same solution the business unit obtained manually in three to four weeks. We then assigned meters to the poles in a greedy manner. That is, we sequentially assigned to each pole selected all the meters in its cover set that were not already assigned to another pole, obtaining a feasible assignment that did not increase the number of poles.

We had successfully demonstrated the effectiveness of the integer-programming approach. Impressed, the business unit decided to use it for solving future pole-selection problems.

Solving the Large Problem

For the large problem, we used the integer-programming approach exclusively. This problem contained 116,600 meters and 20,636 poles with 1.63 million possible combinations of meters and poles. It was much too large to tackle using the traditional approach. Before submitting the problem to CPLEX, we did some preprocessing to reduce its use of memory:

(1) If the set of meters covered by a pole j_1 was a subset of the set of meters covered by pole j_2, we removed pole j_1 from the formulation because it would never be selected.

(2) If the set of poles covering meter i_1 was a subset of the set of poles covering meter i_2, we removed meter i_2 from the formulation because a pole selected for i_1 would also cover i_2.

These two procedures (Daskin 1995, p. 95) reduced the size of the formulation by 50 to 60 percent. We submitted the reduced formulation to CPLEX, and after about three hours of computation, it arrived at a solution that selected 1,431 poles. Depending on how the meters were assigned to the selected poles, we noticed that 10 to 12 percent of the receivers were assigned 360 or more meters. Before implementing the solution, we had to address this limitation.

Capacity Restriction

We had been operating under the assumption that a receiver could cover an infinite number of meters in its range. However, each receiver can cover only a finite number of meters, K, and the total number of meters assigned to it must be less than that. (We give corresponding integer program in the Appendix.)

The resulting integer program is much larger than the integer program we described previously. For example, for the small problem containing 4,208 meters and 2,393 poles, the new formulation contained 144,225 binary variables, 6,601 constraints, and 953,538 nonzero coefficients. Thus, the capacitated problem could be expected to require greater computational resources. In reality, we were unable to solve to optimality the capacitated version of the small problem. We used a heuristic approach by selecting only the closest K meters for each pole.

The theoretical upper limit for K was 540. However, to plan for additional meters in the future and to provide performance guarantees, we needed a lower capacity limit, which might increase the number of poles required. To study this trade-off between the capacity limit and the number of poles, we heuristically solved the finite capacity problem for various values ($K = 170$, 212, 270, 310, 360, and ∞) of capacity using the following approach:

(1) For each pole j, we removed all meters in C_j except for the closest K.

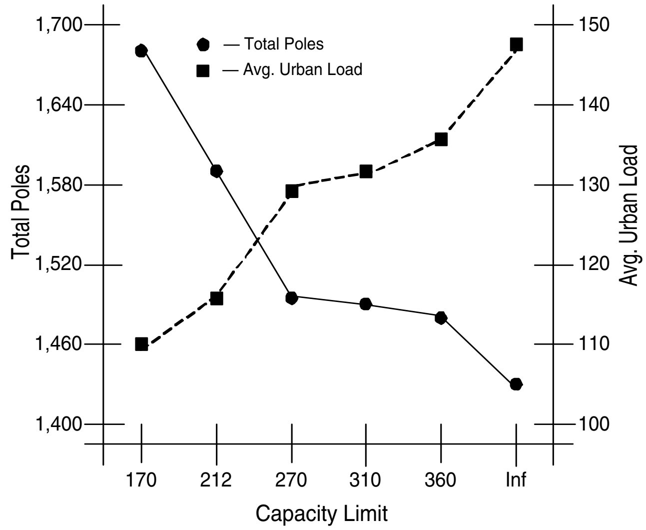

Figure 1: This figure shows the sensitivity of the optimal solution to the capacity limit. The capacity is shown on the horizontal axis while the total number of poles and the average urban load are shown on the left and the right vertical axis, respectively. For both curves, the slope changes significantly around the capacity of 270.

(2) Because every pole then had at most K meters in its coverage area, we could solve the problem without the capacity restrictions. We can use the formulation (Problem UNC-SC) from the Appendix to select the minimum number of poles needed to cover all meters.

(3) We assigned each meter to the nearest selected pole.

Meters are denser in urban areas than in rural areas. Therefore, the average load of receivers is higher in urban areas. In fact, in rural areas $|C_j|$ hardly ever exceeds 100. Moreover, once the poles have been selected, the method in Step 3 assigned meters to poles quite evenly. The average load of a pole was close to the maximum in urban areas (and low in rural areas) (Table 1, Figure 1).

Our results indicated that by using a capacity limit of 270, we stayed well below the maximal possible capacity, and the number of poles required to cover the meters increased by only five to six percent to 1,501. We decided to use a capacity limit of 270. Illinois Power is implementing the solution with 1,501 poles in Decatur, Illinois (Table 2).

Using the Tool

Project planners use the tool as follows:

(1) Collect the geographical data (longitude, latitude) of the poles and the meters using a GIS. Export this data into a text file so that it can be easily input into MATLAB.

(2) Import the data into MATLAB, compute the distances between the poles and the meters, and deter-

K	Total poles	Average urban load
170	1,678	110.4
212	1,603	116.0
270	1,501	130.9
310	1,495	132.2
360	1,488	134.0
∞	1,435	147.6

Table 2: This table illustrates the sensitivity of the solution to the capacity limit. It shows the total number of poles selected and the average urban load for five different capacity choices. For both measurements, performance changes suddenly around the capacity of 270.

mine the coverage area and the meters inside it for each pole.

(3) For each pole, retain only the K (the capacity limit) closest meters in its coverage area. This ensures that the solution generated by the integer program does not violate the capacity restriction. Output the feasible combinations of poles and meters in the format needed by GAMS.

(4) Import the pole-meter combination data into GAMS and solve the resulting integer program. As it solves the problem, review the best solution identified so far and the upper limit on its distance from the optimal solution. Using this information, decide when to stop the solver. GAMS outputs the current solution (the poles selected) to a text file.

(5) Input the text file containing the selected poles into MATLAB and assign meters to the poles in any fashion. (The assignment does not alter the number of poles selected, and Step 3 ensures staying within capacity limit.) Output data on the poles and the meters assigned to them into a text file.

(6) Input the text file containing the solution into the GIS and plot it on a map of the region. Review the solution and accept it for implementation or prepare what-if questions.

(7) Perform the what-if analysis and iterate until you reach an acceptable solution.

We had many ideas for automating this process and developing a good user interface. However, we have not pursued them because most users had technical backgrounds and did not ask for automation and a user interface. In addition, they are using the tool in a strategic manner mainly during the early stages of a project, making further investments unattractive.

Benefits to Schlumberger

Schlumberger benefited in three areas from the integer-programming-based tool for pole selection.

The tool was able to reduce the time the firm took to determine the receiver-installation sites from three to six months to less than one month, allowing it to generate revenue sooner.

Schlumberger used fewer planners per project than it would have. For the small problem, the integer-programming approach saved three person-weeks of planning time. Extrapolating to the large problem, we estimate that the savings in planning time would be around 784 person-weeks. At a cost of US$1,000 to $2,000 per person-week, that would be a savings of US$0.78 to $1.57 million.

In addition, the project managers conceded that the solution determined by this tool would be at least 10 percent better than one determined manually. We based this estimate on the difference between the integer program's solution for the small problem and the first manual solution for it. Because the problem was small, we were able to scrutinize the solution and improve it. However, we probably could not make such an improvement for large problems. Assuming that the load (receiver to meter ratio) remains the same, the firm would need 14,160 receivers to cover the 1.1 million meters. By using the integer-programming method, it might reduce that number by 1,415. The resulting savings in equipment and labor costs, at a rate of US$200 to $400 per receiver, would be about 0.28 to 0.57 million dollars. Thus, we estimate the total savings from this new approach to be between US$1.06 to $2.14 million.

This tool had an unexpected organizational benefit. While other units of Schlumberger have used operations research methods for years to solve strategic and operational problems, this project was RMS division's first to use operations research methods. As a direct result of the project's success, RMS managers realized that recruiting people with operations research backgrounds could help them to reduce inefficiencies and improve their competitiveness.

Extensions

The integer-programming approach could be used to solve extensions to this problem.

Different Types of Poles

We assumed in our initial work that all the poles were identical, which in reality is not true. Even if two poles have exactly the same coverage area, a pole that is easily accessible (say, on a road) is preferred to one that is not (say, in a pond). Similarly, a pole with a street light is preferred because the transformer used by the light can be used by the receiver as well, thus saving money. We can implement this feature in the problem by assigning different weights (preferred poles get lower weights) to the poles. Let the weight assigned to pole j be W_j. Then, we need only to modify the objective function, making it to minimize

$$\sum_{j=1}^{p} W_j Y_j.$$

This change does not significantly alter the computational complexity of the problem and the time required to solve it.

Different Types of Receivers
Various types of receivers are available to cover the meters, and these receivers differ in shape, distance of coverage area, and costs. Assume that N types of receivers are available and that we can place only one type on any pole. Let the costs of these receivers be $R_1, R_2, \ldots, R_N$. Let Y_j^n, $n \in \{1, 2, \ldots, N\}$, be a 0–1 variable, taking the value of one if a receiver of type n is placed on pole j. Let C_j^n be the subset of meters that the pole j covers if it has a receiver of type n. Then, the uncapacitated optimization problem can be formulated (Appendix, Problem MRT-SC) as an integer program.

Appendix

Integer-Programming (IP) Formulation Without Capacity Restrictions
Assuming that receivers have infinite capacity, the pole-selection problem can be formulated as an integer program as follows:

Problem UNC-SC

$$\begin{aligned} \text{Minimize} \quad & \sum_{j=1}^{p} Y_j \\ \text{subject to} \quad & \sum_{\{j \mid i \in C_j\}} Y_j \geq 1 \quad \forall i, \\ & Y_j \in \{0, 1\} \quad \forall j. \end{aligned}$$

In this formulation, Y_j is a 0–1 variable and is set equal to one if pole j is selected and set to zero otherwise. The objective is to minimize the total number of poles selected. The set of constraints indicates that for every meter, among all the poles that can communicate with this meter, at least one pole is selected. Let $\{Y_j^*, j = 1, 2, \ldots, p\}$ be the optimal solution to the integer program. Then, $S = \{j \mid Y_j^* = 1\}$.

IP Formulation for the 10-Meter, Eight-Pole Example

$$\begin{aligned} \text{Minimize} \quad & Y_1 + Y_2 + Y_3 + Y_4 + Y_5 + Y_6 + Y_7 + Y_8 \\ \text{subject to} \quad & Y_1 + Y_6 + Y_8 \geq 1 && (1) \\ & Y_1 + Y_2 \geq 1 && (2) \\ & Y_1 + Y_2 + Y_5 \geq 1 && (3) \\ & \vdots \\ & Y_4 \geq 1 \\ & Y_1, Y_2, Y_3, Y_4, Y_5, Y_6, Y_7, Y_8 \in \{0, 1\}. \end{aligned}$$

The objective is to minimize the sum of variables Y_1 through Y_6, which represent the poles selected. The constraints require that every meter is covered by at least one selected pole. For example, meter 1 is in the coverage areas of poles 1, 6, and 8. Thus, constraint (1) indicates that at least one of these poles should be selected.

IP Formulation with Capacity Restrictions
To formulate this problem, we have to introduce a new set of binary variables X_{ij}. X_{ij} is set equal to 1 if meter i is assigned to pole j, and set to zero otherwise. Thus, the integer-programming formulation for this problem is

Problem CAP-SC

$$\begin{aligned} \text{Minimize} \quad & \sum_{j=1}^{p} Y_j \\ \text{subject to} \quad & \sum_{\{i \mid i \in C_j\}} X_{ij} \leq K Y_j \quad \forall j, \\ & \sum_{\{j \mid i \in C_j\}} X_{ij} \geq 1 \quad \forall i, \\ & Y_j \in \{0, 1\}, \quad X_{ij} \in \{0, 1\} \quad \forall i, j. \end{aligned}$$

As in the previous case, the objective here is to select the minimum number of poles that cover all the meters. The first constraint imposes the restriction that if a pole is not selected the number of meters assigned to that pole must be zero. On the other hand, if a pole is selected, the number of meters assigned to it is restricted to be less than K, the capacity limit. Constraint (2) ensures that every meter is assigned to exactly one pole. Let us now revisit the 10-meter, eight-pole example and impose a capacity limit of two meters per pole. The resulting formulation is

$$\begin{aligned} \text{minimize} \quad & Y_1 + Y_2 + Y_3 + Y_4 + Y_5 + Y_6 + Y_7 + Y_8 \\ \text{subject to} \quad & X_{(1)(1)} + X_{(2)(1)} + X_{(3)(1)} \leq 2Y_1 && (4) \\ & X_{(2)(2)} + X_{(3)(2)} + X_{(9)(2)} \leq 2Y_2 && (5) \\ & \vdots \\ & X_{(4)(7)} + X_{(5)(7)} + X_{(9)(7)} \leq 2Y_7 && (6) \\ & X_{(1)(8)} + X_{(4)(8)} + X_{(8)(8)} \leq 2Y_8 && (7) \\ & X_{(1)(1)} + X_{(1)(6)} + X_{(1)(8)} = 1 && (8) \\ & X_{(2)(1)} + X_{(2)(2)} = 1 && (9) \\ & \vdots \\ & X_{(9)(2)} + X_{(9)(4)} + X_{(9)(6)} + X_{(9)(7)} = 1 && (10) \\ & X_{(10)(4)} = 1 && (11) \\ & Y_4 \geq 1 \\ & Y_1, Y_2, Y_3, Y_4, Y_5, Y_6, Y_7, Y_8 \in \{0, 1\}. \end{aligned}$$

Selected pole	Meters assigned
1	2, 3
3	5, 7
4	9, 10
5	6, 8
6	1, 4

Table 3: This table presents an optimal solution for the 10-meter, eight-pole example with capacity limit two. The optimal solution consists of five poles (1, 3, 4, 5, and 6) with two meters assigned to each pole.

The solution with poles 1, 3, 4, 5, and 6 with the meter assignments $2, 3$; $5, 7$; $9, 10$; $6, 8$; and $1, 4$; respectively, is an optimal solution to this problem (Table 3).

IP Formulation with Different Types of Receivers

Problem MRT-SC

$$\text{Minimize} \quad \sum_{n=1}^{N} R_n \left[\sum_{j=1}^{p} Y_j^n \right]$$

$$\text{subject to} \quad \sum_{\{(j,n) \mid i \in C_j^n\}} Y_j^n \geq 1 \quad \forall i,$$

$$\sum_{n=1}^{N} Y_j^n \leq 1 \quad \forall j,$$

$$Y_j^n \in \{0, 1\} \quad \forall j, n.$$

The objective here is to minimize the total cost of the selected receivers. Constraint (1) ensures that for every meter, among the pole, meter-type combinations in whose coverage area it falls, at least one combination is selected in the optimal solution. Constraint (2) ensures that on every pole, only one type of receiver is installed.

References

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

Dell'Amico, M., F. Maffioli, S. Martello. 1997. *Annotated Bibliographies in Combinatorial Optimization.* John Wiley and Sons, New York.

Nemhauser, G. L., L. A. Wolsey. 1988. *Integer and Combinatorial Optimization.* John Wiley and Sons, New York.

Jesse M. Pantalion, General Manager, Illinois Power Project, 2475 Federal Drive, Decatur, Illinois 62526, writes: "This is to certify that the methods and results reported in the paper 'Schlumberger optimizes receiver location for automated meter reading' were developed for and applied to Schlumberger's mapping for deployment of a wireless metering system for Illinois Power Company. In addition, these methods will be available for application to the planning of future wireless metering systems for other Schlumberger clients."