Geoffrion, Arthur M., ***A Guide to Computer-Assisted Methods for Distribution Systems Planning***, Sloan Management Review, 16:2 (1975:Winter) p.17

A Guide to Computer-Assisted Methods for Distribution Systems Planning*

Arthur M. Geoffrion, University of California, Los Angeles

A great many computer-based methods have been devised for distribution planning problems, especially those involving facility location. This article develops a framework of criteria by which to judge the practical usefulness of any such method, and applies it to obtain critical evaluations of the principal available approaches. *Ed.*

Introduction

Every firm that ships a variety of products from several plants (or other supply sources) through a number of warehouses (or distribution centers) to a widely dispersed clientele has a continuing need to monitor and readjust its distribution system in response to changing conditions. Changes in plant capacity, addition and deletion of product lines, changes in the economics of warehousing and transportation, shifting markets and competitors' actions all impact the effectiveness of a physical distribution system in complex ways. Many firms have found that a comprehensive computerized model of a distribution system can be an effective tool for dealing with such complexities, and that it is able to disclose important opportunities for improvement that are likely to escape conventional analysis based on traditional tools.

The purpose here is to provide a nontechnical guide to the evaluation of available computer-assisted approaches and models for distribution system planning. Such a guide is needed because the number of available alternatives has grown to a bewildering variety, and because some fundamental misconceptions have come to prevail in both executive and staff analyst circles. This is probably due, at least in part, to lagging communication of the considerable progress which has occurred in this field over the last few years. This article attempts to organize the principal available alternatives in terms of the key factors affecting their applicability, to evaluate them according to the main criteria affecting their usefulness and, hopefully, to clear up some common misconceptions in the process.

It is appropriate to begin by clarifying the scope of "distribution plan-

---

* This work was partially supported by the National Science Foundation and the Office of Naval Research. The author gratefully acknowledges the helpful comments of J. R. Davis, M. Fieldhouse, B. M. Khumawala, A. T. Langeveld, S. G. Taylor, and C. I. Tiplitz.

ning" as the term is used here. It refers primarily to intermediate-term (one to five years) planning of the principal configuration and flow aspects of distribution system design, involving either modifications to an existing distribution system or the design of a new one. The following primary questions (and often secondary ones as well) need to be answered:

Q1. How many warehouses should there be?[1]
Q2. Where should they be located (given a list of current and plausible candidate locations)?
Q3. What size should each warehouse be (including selection among specific expansion and contraction projects under consideration)?
Q4. Which warehouse should service what customers?
Q5. How should each plant's output be allocated among warehouses and customers for each product?
Q6. What should the transportation flows be on an annual basis throughout the entire distribution system?
Q7. What is the breakdown of cost savings and customer service implications associated with the best distribution system design compared to a projection of the current system to the target period?

Notice that inventory control, order processing, packaging, materials handling, vehicle routing, and other matters of a current operational nature are not addressed directly. For planning purposes it suffices to assume that these tactical functions are performed as economically as possible consistent with a desired level of customer service, with their cost consequences carefully woven into the individual cost elements of the planning model. Space limitations preclude our delving into this important aspect of the art of modeling.

It is important to recognize that all of the above questions must be resolved *simultaneously* rather than piecemeal. Each impacts on the others.

The cost elements of a planning model are:

C1. Transportation costs between plants, warehouses and customers;
C2. All warehouse and inventory costs;
C3. Costs and savings of expanding, opening, and closing warehouses;
C4. Production costs by product at the plants.

Revenues from sales must, of course, be incorporated into the model (as negative costs) when the above design questions have a significant influence on demand. The sum of all these costs must be minimized subject to all necessary restrictions, including the following:

R1. The stipulated production capacities of each plant must not be exceeded.
R2. The size of each open warehouse must be between prescribed lower and upper limits.

---

[1] The term "warehouse" is used generically in this paper for the sake of brevity, even when some other term such as "distribution center" might be more appropriate.

R3. (Often) Each customer must be served by a single warehouse for certain products.
R4. A warehouse is eligible to serve a customer only if it is sufficiently close that the transit times under economical delivery modes are in accord with the desired level of customer service.
R5. All forecast customer demands must be satisfied (any undesirable demands would be eliminated from the forecast).
R6. Any other desired constraints on configuration, such as: lower and upper limits on the number of open warehouses, subsets of warehouses among which at least one or at most one should be open, more complex constraints on warehouse capacity and desired levels of customer service.

The task of jointly answering the given questions by minimizing the given costs over all allowable alternatives is a very difficult one for all but the smallest firms. Conventional methods of manual analysis, even when supported by good data made available through computerized accounting and management information systems, are usually not completely satisfactory by comparison with what is now possible using a more formal modeling approach. Some of the reasons why are as follows.

1. Conventional studies are often addressed to relatively specific symptomatic difficulties recognized as "problems" by management: inadequate warehouse space at some location, deterioration of customer service, excessive inventories, and so on. There is an understandable tendency to carry the analysis just far enough to make the problem "go away" without taking the extra time necessary to achieve a truly comprehensive longer-term solution.
2. Conventional studies are often carried out along the functional lines most familiar to the principal analyst: e.g., finance, transportation, manufacturing, or marketing. It is well known that the interests of these functional specialties tend to be in conflict with one another, and therefore that a deliberate effort must be made to balance these conflicting interests if the study recommendations are to be valid in a corporate-wide sense.
3. Conventional studies tend to be regional rather than national (or international) in scope. The dangers here are evident.
4. There is a tendency to oversimplify the complex "chain reactions" of cause and effect by which the influence of each decision propagates throughout the system and impinges on many other decisions. Humans usually do not have the patience and capacity needed to cope with such intricacies.
5. Finally, this limited time and capacity is responsible for the inability to consider seriously more than a few plausible decision alternatives among the usually innumerable possibilities.

These limitations of conventional manual analysis can be overcome to a large extent by adopting a computer-based distribution system model as a framework. There will still be plenty of manual work, perhaps even more than before, but the improved scope and quality of the conclusions that can be reached by this vastly more powerful approach should justify the added effort many times over.

A conceptual framework is necessary before we can usefully discuss the available distribution system planning models and their associated computational techniques. We must first establish (a) the key factors which commonly determine a model's applicability to real situations, and (b) the key criteria which determine a computational technique's ability to yield the desired insights and conclusions from a model. These two tasks are discussed in the next two sections. This framework is then applied in the form of a critical analysis of available planning models and techniques. The final section presents a summary of the main conclusions.

Important Problem Features to Be Modeled

A distribution planning model is a comprehensive collection of precise statements and assumptions about the system which describes it in sufficient detail that this formalized representation can be useful for answering the questions posed at the outset of this article. The art of designing an effective computer model is the art of compromising between degree of detail on the one hand and economy of use on the other. Essential features of the real system must be reflected in sufficient detail to assure that the model will yield valid conclusions and earn the active acceptance of management. But the level of detail should not be so great that the model's data requirements are unreasonably large, or that no available computational technique can solve it at reasonable expense.

Over the years there has been a gradual shift toward a greater permissible level of detail. This is due to the increased availability of computer-resident data sources and steady progress in computational techniques. In most cases it is no longer necessary to settle for highly simplified models that omit some of the most salient features of a firm's distribution system. Such oversimplification has been responsible for a high proportion of the past failures of distribution planning models in terms of their actual decision-making impact. Management will not implement conclusions based on an oversimplified model.

The following problem features are common to many if not most real applications. Often they have been ignored or dealt with only approximately in an ad hoc fashion in the earlier models and even in many models in use today. Experience and common sense teaches, however, that these features should be incorporated directly in the model of any distribution system to which they apply.

1. *Multiple Products*

Most firms produce more than a single product (different sizes and packages are often best viewed as different "products"). When these products share the use of production or distribution facilities to an appreciable extent, it is usually necessary to treat multiple products explicitly in the model. Ignoring all but a single product clearly invites gross errors in such a situation, as does solving the problem independently for each major product and then attempting to agglomerate the results back into a single distribution system. Defining a single composite product composed of a standard mixture of individual products obviously runs afoul if there are products that cannot be made at every plant, or if demands for individual products occur in differing ratios in different geographic regions. The proper course is to define and incorporate into the model as many composite products or product groups as are necessary to reflect realistically the sharing of production and distribution facilities. The product spectrum of most firms can usually be treated realistically with from three to thirty suitably defined "products."

2. *Two Stages of Distribution: Plants → Warehouses → Customers*

Large firms commonly utilize warehouses (or distribution centers) to couple their plants (or other sources of supply) with their customers. This means that two stages of distribution should be distinguished: plant-to-warehouse and warehouse-to-customer.

Early distribution system models treated only a single stage of distribution. If the plant-to-warehouse stage is selected, then customers must be collapsed into warehouses by allocating all demands to the warehouses a priori. This is usually done to study plant location or expansion questions, but unfortunately the influence of such decisions on warehouse locations and service area boundaries is then lost. If, on the other hand, the warehouse-to-customer stage is selected for study, then plants must be collapsed into warehouses either by guessing the optimal allocation of production to the warehouses or by treating plant capacities as infinite. Either way, differential production costs among plants and plant-to-warehouse transportation economics are treated only approximately. A third possibility is to distinguish plants and customers but not warehouses, with the transportation costs along each fictitious plant-to-customer route devised so as to include an approximation to warehouse costs. The shortcoming of this approach, obviously, is that warehouse location, capacity, and economy of scale aspects are dealt with to an uncontrollable degree of approximation.

3. *Capacities for Plants and Size Limits for Warehouses*

Every distribution system model must take some cognizance of realities relating to the production or supply side. Incorporating capacity limitations on each supply location by product group is usually necessary when a product group can be made at more than one location. One may also incorporate production costs; a common trick for doing this is to include them as a component of the plant-to-warehouse transportation costs.

The model should be able to control the size of a warehouse, which can be expressed in terms of its annual throughput volume (cases, hundredweight, etc.). Minimum and maximum size limits can be used to keep the throughput of an existing facility within the range over which no major change is needed in floor space or material handling methods. Such major changes would ordinarily be treated as explicit project alternatives to be accepted or rejected, each with its own effect on permissible throughput. For a potential warehouse not currently in existence, a minimum size limit would be specified to assure a sufficient degree of operating efficiency to justify existence at all, while a maximum size would be specified to preclude gross *dis*economies of scale that would distort the assumed cost structure.

Another important use for warehouse size limits is to provide a technical mechanism for achieving greater accuracy in both warehousing and transportation cost structure. If several alternative size ranges for a warehouse can be specified, then warehouse and inbound and outbound freight costs can all be made specific to each alternative size range. A model that fails to provide for constraints on plant capacity and warehouse size is likely to suffer from an oversimplified cost structure and lack of control over the feasibility of the solution.

4. *Warehouse Economies of Scale and Fixed Charges*

Opening and closing warehouses are major decisions whose cost consequences require the introduction of indivisible fixed charges (or savings). Computations for fixed charges turn out to be quite difficult to handle. Similarly, warehouse economies of scale (unit costs that decrease with increasing volume) are often significant enough to necessitate their explicit inclusion in the cost structure of the model.

A model that can handle fixed charges exactly can usually handle warehouse economies of scale adequately. There is a standard trick for reducing the latter complexity to the former which involves allowing alternative facilities at a given location with different cost characteristics for each. This trick works satisfactorily as long as the computational method yields an exact optimal solution, and as long as there are no jumps (either up or down) in the warehousing cost curve as a function of annual throughput (such as might be associated with possible expansion or contraction projects at an existing warehouse). If either of these conditions is violated, it is generally necessary to impose size limits on the alternative facilities and also additional constraints which stipulate that at most one of the facilities may be used at a given location. This requires the use of problem features 3 and 7.

5. *Each Customer Serviced by a Single Warehouse*

Most firms have a policy of servicing each customer from a single distribution center for as many products as is economically feasible. This is done as a convenience to the customer, to simplify accounting, communication, and marketing functions, to obtain maximum economies of scale for outbound freight, and to achieve greater consistency in order cycle times. The main exceptions to this policy occur when plant direct shipments are more economical for some product-customer combinations, and when it is more economical to concentrate in relatively few locations the inventories of slower-moving products or products requiring expensive specialized handling or storage facilities. To reflect this policy properly a model must explicitly include the multiple choice aspect of assigning each customer to a single distribution center (subject, of course, to the possible exceptions just mentioned).

6. *Shipments to Customers: Preserve the Identity of Originating Plants*

In some industries much freight moves from plant to customer under a "storage-in-transit" privilege which leads to significantly lower freight rates than the simple sum of plant-to-warehouse and warehouse-to-customer rates. This requires keeping track of the plant origin of transit commodities all the way to the customer. This capability is also useful for dealing with other difficult problem features such as deteriorating commodities for which total transit time must be controlled, and products for which the price or demand structure is dependent upon the identity of the originating plant.

7. *Various Desired Constraints on Distribution System Configuration*

Often there are additional configuration constraints to be taken into account, as mentioned under restriction R6 in the introduction, in order for a distribution system design to be valid and acceptable to management. A deliberate effort should be made to discover and incorporate all such constraints at the outset of a study, lest they come to light later as an obstacle to managerial acceptance. Configuration constraints are also necessary as a technical device to properly select among alternative major projects for expanding or contracting existing warehouses.

Some comments may be in order concerning some of the possible problem features *not* selected for inclusion on this list. One example is multiple time periods, which permit the long-term dynamic evolution of a distribution system to be studied. There are two reasons why multiperiod models have been of limited usefulness in the past. One is the greatly increased difficulty of obtaining the necessary data; a five period model requires five times as much data as a one period model (which is already a major undertaking). The result is that multiperiod models tend to use grossly scaled first period data for all subsequent periods in order to reduce the data gathering cost, with the consequence that model realism beyond the first period is in serious question. Also, there are inherent difficulties of reliably forecasting costs and demands farther into the future even on the most generous of project budgets. The second inhibiting factor is the sheer computational difficulty of solving multiperiod models reliably and inexpensively enough to make them a truly useful tool. The usual alternative to an explicit multiperiod model is to run a single period model with one or more cost-demand scenarios reflecting years beyond the target year of the study (which is usually two to three years in the future). These "snapshots" are then taken into consideration when interpreting the more detailed analysis of the target year. This snapshot approach may be inadequate, however, if the demand or cost structure is changing rapidly in time. It then may be necessary to formulate a model with several time periods. In this case it is useful to know that there is a trick by which any computational method designed for a single period model with features 1, 3, 4 and 7 can be made to solve its multiperiod counterpart. This can be done by a mathematical trick based on the idea of replicating all product groups and warehouse locations as many times as there are time periods, and adding side constraints to make everything behave properly. Obviously this can be practical only if the number of periods is fairly small (e.g., two to four).

Another problem feature omitted from the list is that of explicitly distinguishing between alternative modes of transportation and the different unit rates that apply to shipments of different weight and composition. The reason is that, for planning purposes, these aspects all can be taken into account in the data gathering phase of the study and distilled into appropriately weighted unit transportation costs for each product group over each allowable link (plant-to-warehouse, plant-to-customer, warehouse-to-customer). The weighting procedure takes into account the ordering habits and delivery requirements of the individual customers and the firm's manner of shipping to its warehouses. The best mode or mix of modes would be selected for each link. The unit transportation rates in and out of a warehouse can be made dependent on its annual throughput, if desired, by providing several alternative size ranges for a given facility (this can be done if problem features 3, 4 and 7 are provided for).

Finally, plant location and plant expansion planning are not included in the list of key problem features. This is because most distribution system planning studies either take production (or other supply) facilities and their capacities as given, or there are few enough alternatives of this type that they can be studied as individual cases. (In this event management may well *insist* on studying the alternatives on an individual basis.) Virtually any model and computational method can be used to examine cases individually (provided, however, that the results achieved for each case are sufficiently accurate that comparisons among them is meaningful). When the number of cases is too large to study individually, several of the available methods to be discussed below are extensible to resolve plant location and expansion questions simultaneously with distribution center location questions.

Essential Criteria for a Useful Computational Method

Even the most realistic model is useless unless a computational method is available that can "solve" it with reasonable efficiency for the desired insights and conclusions. When a model represents a complex distribution system with numerous configuration and shipping alternatives, this can be a monumentally difficult task.[2] Yet it must be accomplished economically and in a manner that will win the confidence of management. A computational method should meet the following three criteria in order to be fully satisfactory.

1. *It Should Truly Optimize*

It is important to draw a distinction between a computational method that is truly optimizing and one that is not. For a truly optimizing method one can specify in advance of a computer run an arbitrarily tight tolerance on the error within which the solution must approximate the true global mathematical optimum. This tolerance can be expressed in terms of a number of dollars or a certain percentage over the true minimal cost. The ability to specify extremely tight tolerances does not mean that this will necessarily be done in every run, but there are many occasions where this *must* be done in order to reach meaningful conclusions. If such a tolerance cannot be specified in advance, with guaranteed results, the method is not a truly optimizing one.

True optimization requires, of course, the ability to deal *simultaneously* with all decision variables identified by the model. None of the questions posed at the outset of this paper can be answered independently of the others. To attempt to take the problem apart and deal with the pieces separately, or to consider only a plausible subset of all possible decision alternatives, is to risk obtaining a suboptimal solution of unknown quality.

It should be stressed that our use of the term "optimizing" is meant to be *relative* to the degree of realism achieved by the model which is adopted. Every model involves approximation and simplification. It would be folly to suppose that the mathematical optimum of any model necessarily indicates the "best" executive decision. We do claim, however, that the ability to compute mathematically optimal solutions of even an inexact model is indeed necessary as a means toward the ultimate objective of finding and justifying the best executive decisions. This point is fairly subtle, but we hope that the arguments to follow will clarify it.

The most obvious reason why it is important to use a method that truly optimizes is that otherwise one may miss significant opportunities for savings. Such opportunities can easily amount to many times the total cost of the entire study when a firm's distribution costs amount to a few million dollars per year or more. It is often argued that simpler, nonoptimizing methods are likely to perform just as well as truly optimizing methods when the model data contain substantial errors. The fallacy of this is that unbiased random errors tend to cancel one another and preserve the *relative ranking* of alternative solutions. Computational experiments have confirmed this effect. It follows that if one solution is better than another under the estimated data,

---

[2] For instance, there are more than seven million ways of selecting between five and ten distribution center locations from a list of just twenty-five possibilities!

then it will probably still compare favorably under the (unknown) true data. Thus the use of truly optimizing methods is warranted in order to obtain the maximum savings even when the available data contain (unbiased) errors. Of course, no approach, computer-assisted or otherwise, can be immune from the possible deleterious effects of serious *biased* errors in the data.

The prudent staff analyst or consultant will keep in mind that to overlook an opportunity for significant savings is to risk loss of credibility to the point of professional suicide. Consider the consequences of using a suboptimal computational method that leads, by chance, to a recommendation that management can beat manually by a significant amount (perhaps by massaging the recommendation itself). This actually happened on one occasion to a well-known consulting firm doing a $75,000+ project for a large shipper, with predictable disastrous results for the client-consultant relationship.

Another persuasive reason for using an optimizing method, one that is often overlooked, is that otherwise it is meaningless to attempt comparative runs to study important "what if . . .?" questions and others of vital interest to management (see criterion 3 below). Consider what might happen if there is even 1 percent wobble in the quality of the computer "solution" and management asks what will happen if a certain existing warehouse is kept open instead of being closed as dictated by the solution. Keeping the warehouse open might actually leave the true optimal costs essentially unchanged, but there is a 50-50 chance that the comparison will report a penalty (as high as 1 percent) for keeping it open. Thus management could easily arrive at the false conclusion that it is necessary to close this warehouse.

The final argument we wish to cite in favor of an optimizing method is that such methods usually lend themselves more readily to tracking down the sources of the inevitable data input errors. The reason is that input errors are revealed and diagnosed via their influence on the model's output. This influence is much easier to trace back through the precision logic of an optimizing method than through the "mushier" innards of most nonoptimizing methods which attempt to do more than simply add up costs. Needless to say, thorough input validation is absolutely essential.

2. *Computer Costs Should Be Moderate*

In the past, computations for truly optimizing methods have acquired the reputation of being quite expensive. The result has been that these methods often have been used in a suboptimal fashion in order to conserve computer time, or have been passed over altogether in favor of less expensive approximate techniques. We saw in the previous section that this may well be a false economy. Nevertheless, management will probably still continue to economize on the computer budget as a way of keeping total planning study costs down. It is fortunate, as we shall see in the following section, that some of the most recently developed computational methods offer great improvements in efficiency and may actually be *less* expensive to use than available approximate methods (which usually rely on brute force rather than the sophistication accruing from technological progress).

In gauging the cost of an optimizing technique it is customary to estimate how much the full optimization run will cost, and then add some fraction or small multiple of this amount to account for the data validation phase. This is a mistake. To carry out a study properly and take full advantage of the laboriously built distribution model, a much larger multiple should be used (perhaps on the order of twenty-five). The reason is that nearly all the real benefit occurs in the "secondary run" phase of the study, rather than in the phase leading to the initial full optimization. This point is discussed at length under the next criterion.

3. *The Computational Method Should Facilitate Multiple Secondary Optimization Runs*

It is important to make a distinction between the *primary* or *base case* optimization run, in which the full distribution system design problem is solved for the first time using best estimates of all data, and the many subsequent *secondary* optimization runs done to address specific questions of concern to management. A common misconception is to consider the primary optimization run as the principal objective of the study. It is not. It merely serves to establish a leading contender for subsequent analysis and criticism from all relevant points of view. That is where the secondary runs come in; they provide the capability of answering many of the questions arising during this phase of critical analysis and review.

There are many types of useful secondary runs. The following general categories will serve for illustrative purposes, though each particular study will pose its own unique requirements.

"What If . . .?" Runs, spawned by management's need to explore a variety of issues such as the effect of forcing certain warehouses to be open or closed, altering warehouse service areas to conform more closely to current marketing zones, negotiating new rail rates, examining the impact of substantially higher truck rates due to increased fuel costs, closing down certain production lines at certain plants, and implementing alternative marketing strategies or channels of distribution. The ability to answer such questions quickly and confidently is not only essential to a successful distribution planning study, but may well justify maintaining the model in a "ready" status indefinitely as a quick response tool to cope with questions arising from continually changing conditions.

Sensitivity Analysis Runs, in which possible errors or changes in the data are explored in order to determine how "robust" the model's implications are. For instance, several sets of estimated demand data may be run corresponding to pessimistic, expected, and optimistic forecasts. Or one may examine the implication of changes in assumptions regarding the future cost of capital and inflation rate.

Trade-off Analysis Runs, which can develop the trade-offs between total cost and the quality of customer service, between capital expenditure and annual operating cost reductions, between energy consumption and total cost, etc. The cost/customer service trade-off curve is probably among the most effective tools for decision making in the elusive area of setting appropriate customer service levels.

Capacity Valuation Runs, which assess the marginal value of increased capacity of various types at plants and warehouses.

Implementation Priority Analysis Runs, which identify the specific changes in the distribution system (especially opening and closing warehouses) that account for the largest potential savings. This helps management to decide which changes to implement first, which to do later, and which are too marginal in value to be worth the organizational upset of doing at all.

These and other types of secondary runs[3] are necessary to answer properly the primary questions addressed by a distribution planning study and to capture the full benefits inherent in a comprehensive distribution model. They must earn sufficient credibility and enthusiasm so that management will be willing to implement the major findings of the study and maintain the model for future use.

The crucial role played by secondary runs has not always been fully appreciated. Consider the following recent statement by a respected management consultant of wide experience:

> . . . A highly complex computer program to solve the problem is fine. But I prefer something simpler. The reason is that the decision to open or close a distribution center will be made by the President of your company, or someone fairly high up in management. It's not enough for you to say to this decision-maker, "The computer says to do thus and so." As a corporate officer who takes his responsibility seriously, before he makes a decision he wants to understand the situation. He wants to know the *reasons why* he should open or close a distribution center. He wants to know how his decision will affect service and cost. He may want to ask some "What if . . . ?" questions.
>
> I prefer a simple method of analysis because it's easier then for the decision-maker to understand, to know the reasons why, to make his own evaluation of the results that will follow from his decision.[4]

The whole point of the five types of secondary runs described above is precisely to serve such executive responsibilities. The wise executive will insist above all that the case for each recommendation must be *persuasive;* prop-

---

[3] See, for instance, Geoffrion and Graves [16], Sec. 4.2.

[4] See Hupp [20], pp. 466-7.

erly specified secondary runs using a comprehensive computerized model now provide the capability to develop recommendations of unprecedented persuasiveness. Such a tool may actually be *simpler* for the executive to use than a more elementary approach, because it can quickly provide answers to more comprehensive questions than could otherwise be asked.

It is not feasible to carry out multiple secondary runs, however, unless the computational method meets certain requirements. An important one discussed previously is that there must be very tight control over the degree of optimality so that comparing different runs is meaningful. A second obvious requirement is that the computer cost of executing a secondary run must be very low, for it is typically desirable to do many of them, perhaps on the order of twenty to fifty.[5] This implies that the cost of a secondary run should be kept below $100 or so. Additionally, it follows from the large number of secondary runs that they should be easy to set up (have convenient data management facilities), execute reliably without the need for manual intervention during the solution process, and be provided with report generators that display the results in compact form clearly understandable by analysts and managers alike. This will conserve scarce technical manpower resources and encourage maximum interest and participation by management.

Evaluation of Alternative Available Computational Methods

Criteria have now been established by which to judge the realism of a model and the effectiveness of a computational method. It should come as no surprise that, generally speaking, these criteria tend to conflict with one another; the more realistic the model, the more difficult it is to devise an effective computational method to solve it. This is the main trade-off to be kept in mind when evaluating the merits of alternative models and methods for a given application.

In this section we attempt to describe and evaluate the capabilities of the major model/methods presently available from the published literature and from various consulting firms and software vendors. A brief tabular summary of our conclusions for the main optimizing methods is given at the very end of the paper.

1. *Decomposition Method*

The decomposition method of Geoffrion and Graves [16] is among the newest and so far the only method which accommodates all seven problem features and amply meets all three computational criteria. It is referred to as a "decomposition" method because one of the primary technical principles used

---

[5] The most comprehensive use of secondary runs known to the author was a study of the U.S. Air Force logistics system involving 215 separate runs (Figgens and Thompson [14])! These runs explored different scenarios for the future overseas distribution of military forces, different levels of future materiel demand, and different assumptions on transportation costs, delivery requirements, and distribution center costs.

in its design involves the decomposition of the full multicommodity problem into a series of simpler single commodity problems. This principle, quite unique among the available alternative methods, permits vastly improved computational performance by comparison with the previous state-of-the-art. Additional efficiency is gained by other innovations of interest only to the technical specialist.

Geoffrion and Graves [16] include details of a successful application at a large food company with 14 plants, 17 product groups, 45 possible distribution center sites, and 121 customer zones. The discussion includes considerable material concerning the use of secondary runs. Typical computer runs consistently took on the order of one minute on an IBM 360/91, and are easily set up with the help of console interactive data management procedures.

The method is readily extensible to incorporate: (i) simultaneous plant location and capacity expansion decisions, (ii) a demand structure which depends on the proximity of customers to their assigned warehouses, and also (iii) arbitrary additional (linear) constraints on all flow variables (e.g., to achieve a more detailed representation of production or supply limitations).

2. *General Mixed Integer Linear Programming*

The only other way to accommodate all seven problem features is to use a general mixed integer linear programming system. The best available commercial systems are UMPIRE for the Univac 1108 [30], MPSX-MIP for the IBM 360 and 370 series [23], and OPHELIE MIXED for the CDC 6000 series [25]. They all use a branch-and-bound strategy.[6] These systems extend the already great power of linear programming, which is acknowledged by many to be the most versatile and useful management science tool yet devised, by allowing some decision variables to take on only integer values. This permits warehouse economies of scale and fixed charges (feature 4) to be incorporated, and also features 5 and 7 if necessary. These three features are beyond the scope of ordinary linear programming (for instance, handling warehouse fixed charges requires defining 0-1 "location" variables that are 1 when a warehouse is contemplated at a given site and 0 otherwise). The generality of these systems offers the ultimate in model flexibility even beyond the seven problem features we have singled out.

The very generality of these software systems is, however, also their greatest weakness when it comes to computational efficiency. They were not designed specifically for distribution planning models, and they fail to exploit the special structure of such models as a means of achieving the best possible computing times. Even when the so-called generalized upper bounds (GUB) option is used, and this option *should* be used when available to deal with customer demand constraints efficiently, computing times are likely

---

[6] See Geoffrion and Marsten [17].

to be about two orders of magnitude longer than for the best specialized methods. The existing systems can handle only a few dozen integer variables reliably,[7] which severely limits the number of warehouse locations and the number of economy-of-scale break points on the warehousing cost curve which can be handled in practice. Strict enforcement of the one-warehouse-source-per-customer requirement for problems of realistic size is not a practical possibility.

Thus true optimization will not always be practical. For many problems of realistic size it will be necessary to suboptimize (i.e., terminate the computations prematurely) in order to conserve computer time. The following two examples may be of interest. One firm reportedly ran a problem with about 10 plants, 30 commodities, 15 fixed distribution centers (location choices were not allowed), and 80 customer zones for one hour on a Univac 1108 under UMPIRE without even finding a feasible solution. The difficulty can be attributed to problem feature 5. Another firm[8] used the Bonner and Moore FMPS-MIP system to solve a relatively small problem with 4 plants, 4 products, 10 warehouse sites with 3 possible sizes for each, and 39 customers. The troublesome single warehouse per customer requirement was not imposed. The primary optimization run took forty-five minutes on a Univac 1108. Times like these severely limit the number of secondary runs that can be afforded (see the discussion of the third computational criterion).

In all fairness, it must be said that the Univac and other software systems of their class are reported to have successfully solved numerous practical distribution planning problems. However, there are virtually no available written reports to enable one to evaluate seriously what has been accomplished in terms of our criteria.

A second drawback of the general mixed integer linear programming systems is that they tend to be cumbersome to use. Specialized matrix generators and report writers must be acquired or developed in order to deal with the problem at hand. Moreover, it takes an experienced user to know the best settings of the many possible user controls for each run. These impediments run contrary to the third computational criterion.

There do exist at least two (proprietary) matrix generator/report writer packages specifically for distribution planning problems. One, dubbed MULTICOM, was built around OPHELIE MIXED for CDC computers by S.I.A. Limited.[9] It is not designed to accommodate problem features 5 or 6. The other, known as the DS/SD Generalized Distribution Model,[10] is primarily intended for use with IBM's MPSX system. It is not designed to accommodate problem feature 5, and resorts to a computational option known as "separable programming" when the number of integer variables due to prob-

---

[7] See Geoffrion and Marsten [17].

[8] See Bartakke *et al.* [3].

[9] See Elson [12] and SIA Depot Location [27].

[10] See Bender [4].

lem feature 4 exceeds the capabilities of MPSX (as is well known, the separable programming option leads to non-optimal solutions of uncertain quality). Convenience of use and flexibility of modeling are the main advantages of the MULTICOM and DS/SD packages. The main drawback is that their computational efficiency is limited by that of their host commercial software systems.

3. *Optimizing Methods for a Single Stage of Distribution and (Ordinarily) a Single Product*

A number of methods have been devised for relatively simple applications in which it suffices to focus on but a single product and a single stage of distribution. Generally speaking, these two simplifications permit improved computational performance by specially tailored methods. All of the methods discussed here treat warehouse fixed charges and economies of scale via a branch-and-bound approach.[11] As a matter of convenience we shall refer below to the single stage of distribution as though it consisted of warehouses and customers, though it could just as well consist of plants and warehouses or plants and customers (ignoring warehouses).

The most general method in this category is FLAC (Facility Location with Additional Constraints), developed by Geoffrion and McBride [18]. It accommodates upper and lower limits on warehouse throughput and arbitrary additional (linear) constraints on transportation as well as location variables (this is more general than the seventh problem feature requires). In addition, the model can be optimized with or without the requirement that each customer be serviced by a single warehouse. All three computational criteria are met for problems of reasonable size. Problems of practical size (about 25 warehouses and 100 customers) typically are solved in less than one minute on an IBM 360/91.

Next in order of decreasing generality comes problems without lower limits on warehouse throughput, without the requirement that each customer be serviced by a single warehouse, and without configuration constraints. Numerous methods have been proposed for this class of problems, among which the most successful is currently that of Akinc and Khumawala [1]. It appears to meet all three computational criteria. The methods of Davis and Ray [8] and of Soland [28] are also noteworthy.

The final and simplest class of problems drops the warehouse throughput limits, leaving our fourth problem feature as the only one treated exactly. A great many methods have been proposed, among which the most frequently used are Efroymson and Ray's [10], Atkins and Shriver's FLP [2], and Spielberg's SPLT 1 [29]. Khumawala's method [21] seems to be the best, easily satisfying all three computational criteria.

It is worth noting that all of the methods above, although designed for

---

[11] See Geoffrion and Marsten [17].

problems with but a single product, can also handle more than one product if a simple trick is used: namely, replicate each customer as many times as there are products to be distinguished and assign an appropriate product demand to each. This trick essentially converts a multiple product problem to an enlarged but equivalent single product problem.

4. *POLIGAMI*

POLIGAMI is a software package for the CDC 6600 marketed by S.I.A. Ltd. for optimization of networks with fixed charges.[12] Facility location problems are perhaps the primary intended area of application.

POLIGAMI can accommodate problems with a single product, two stages of distribution, plant capacities and upper limits on warehouse size, and fixed charges and simple economies of scale for warehouses. It can handle multiple products (by the trick mentioned at the end of the previous subsection) if just one stage of distribution is distinguished. It does not handle lower limits on warehouse size, the supply of each customer from a single warehouse, the preservation of the identity of the plants whose production supplies each customer, or additional constraints on system configuration. On the other hand, it can explicitly handle some features beyond the seven we have identified, such as interwarehouse transfer.

Thus the first computational criterion (true optimization) holds for the class of problems indicated.

The second criterion of moderate computer cost appears to hold for moderate-sized problems of the class indicated. An analysis of the detailed nature of the computational scheme used by POLIGAMI reveals that the technology employed reflects the state-of-the-art during the mid-1960s and thus may be adequate for problems with as many as thirty to forty potential sites.[13]

The third computational criterion is probably met for problems of moderate size.

5. *CAPFLO*

This is a revised version of a software package developed earlier in England known as DEPLOC,[14] and is available through Haverly Systems Inc.[15] CAPFLO is basically designed to accommodate problems with a single product, two stages of distribution, plant capacities and upper limits on warehouse size, and fixed charges and economies of scale for warehouses. It can also

---

[12] See SIA Depot Location [27] and Poligami User's Manual [26].

[13] Basically it is a conventional branch-and-bound approach with pure transportation network flow bounding subproblems solved by the "out-of-kilter" method. Today's state-of-the-art would permit improvements via more sophisticated branching criteria, stronger bounding problem design, and use of the revolutionary new methods for solving the underlying network flow problems (see Geoffrion [15] for specifics).

[14] See Fieldhouse [13].

[15] See Haverly and Fieldhouse [19].

handle multiple products if there is just a single stage of distribution, and several types of system configuration constraints. It does not handle in general and in an exact fashion lower limits on warehouse size, the requirement that each customer must be serviced by a single warehouse, or preservation of the identity of the plants supplying each customer.

The first computational criterion of true optimization holds for applications having the basic features indicated.

The second criterion of moderate computer cost appears to hold for small-to-medium sized problems. Suboptimal termination is sometimes recommended for "practical problems"[16] presumably due to possibly excessive running times for the larger problems. The difficulty has to do with the technical design of the computational method.[17]

The degree of satisfaction of the third computational criterion (convenience for multiple secondary runs) depends on the particular problem features being used and the size of the problem being solved. The prescribed practice of manual intervention and broken runs[18] does not enhance user convenience.

6. *Casewise Linear Programming*

By the casewise application of linear programming we mean the use of ordinary linear programming to study a modest number of manually selected design alternatives. A "design alternative" must specify predetermined definite locations for all facilities. For each alternative the optimal transportation flows throughout the distribution system can be found provided all economies of scale are ignored. Thus this approach cannot deal with problem features 4, 5 or 7.

This approach has been widely used, for lack of anything better, since linear programming became a standard management planning tool in the 1950s. Obviously it cannot hope to arrive at an optimal solution except in those rare cases where the total conceivable number of plausible design alternatives is manageably small; yet in artful hands it may yield quite satisfactory solutions. Awkwardness of use for secondary optimizations is another serious drawback (for instance, multiple runs and manual intervention are required to evaluate a single revised demand scenario).

This approach must be regarded as obsolete except for the very simplest problems, or for environments where there is no convenient access to more appropriate software.

---

[16] See Haverly and Fieldhouse [19], p. 1.

[17] It is more of an enumerative search procedure than a true branch-and-bound procedure. The only bounding subproblems permitted are pure transportation flow problems in which *all* facilities are locked open or closed. This results in much weaker bounding capability than is usually available from an up-to-date branch-and-bound method, with consequent overreliance on sheer enumeration through branching.

[18] See Haverly and Fieldhouse [19], p. 1.

7. *Computer Simulation*

Elaborate computer simulation packages are available for studying the design and operation of physical distribution systems. There is DSS (Distribution System Simulator)[19] available from IBM, LREPS (Long-Range Environmental Planning Simulator)[20] available through Systems Research Inc., and several other simulation packages specifically for distribution. In addition, many firms have programmed their own simulation studies using specialized simulation languages like SIMSCRIPT and GPSS, or a more general computer language.

The main point to be made about computer simulation is that this approach is not intended to optimize in the true sense of the term, but rather to evaluate a fully specified hypothetical system in great detail. It can take detailed account of policies and activities relating to inventory replenishment, individual buying patterns of customers, order filling, redistribution, transportation, and so on, and produce a simulated daily history of such activities for a period of a year or more. At this level of detail true optimization is usually unthinkable, so great is the problem complexity and the number of decisions and policies that could be selected. The best that can be done is to evaluate a modest number of manually selected detailed design alternatives, perhaps on the order of a dozen or so.

Does this mean that computer simulation is not useful for addressing the questions posed at the outset of this article? Not at all. It can address these questions in a narrow but deep way if used alone, and in a more powerful way if applied in tandem with (after) one of the better optimizing approaches described earlier. This could be viewed as optimization in the service of simulation (use the optimizer to replace the manual guessing of good alternatives to be simulated microscopically); or it could be viewed as simulation in the service of optimization, according to whether one is an advocate of one approach or the other (the two camps tend to be strongly polarized). A still more ambitious undertaking would be to alternate between optimization and simulation, feeding the results back and forth: optimization would provide improved alternatives to be simulated, and the simulation results would point the way to improved accuracy of the optimization model.[21]

A fact of some practical importance is that simulation tends to be considerably more expensive than optimization.[22] This is so because of the greater level of detail involved in a simulation model. The much larger appetite for numerical data leads to greater data gathering costs, and computer

---

[19] See Connors *et al.* [7].

[20] See Bowersox *et al.* [6].

[21] See Nolan and Sovereign [24].

[22] Atkins and Shriver [2] assert that typical simulation project costs "may range from $50,000 to $100,000 or more." There is a minimum contractual obligation of $32,520 for the use of IBM's Distribution Systems Simulator package.

running times are usually quite long even for a single design alternative. Given that a firm is willing to accept the greater expense and effort associated with simulation, a little appreciated fact is that for a relatively small additional cost the project can be expanded to include optimization too. The reason is that the data collection effort required for simulation usually includes all of the data needed for optimization. The added advantages of optimization to help select the design alternatives to be simulated would seem to be well worth the modest extra cost in most applications.

8. *"Continuous" Location Methods*

All of the models considered elsewhere in this article have assumed that management can specify a list of candidate cities where it may be plausible to acquire a warehouse if one is not already there. This list is usually composed of cities where the firm now has facilities, competitors have facilities, and selected gateway cities plausibly positioned with respect to the major demand locations.

"Continuous" location models, on the other hand, do not utilize a preselected list of candidate cities. Rather, each warehouse location is permitted to vary continuously in a geographical sense (its longitude and latitude are variable). One of the most comprehensive accounts of such models and methods for solving them is given in Eilon, Watson-Gandy, and Christofides [11].

One serious limitation of continuous location models is the impossibility of achieving a realistic cost structure for transportation flows. Consider, for instance, plant-to-warehouse flows. For any particular warehouse location one can look up the unit transportation rates by mode, weight and product and combine these appropriately to achieve a suitable composite unit transportation cost. It is well known that such costs in many cases cannot be reliably predicted based on a knowledge of distance alone. Yet this is exactly what must be attempted in a continuous location model: all transportation costs must be expressed as explicit well-behaved functions of distance (no look-ups allowed).

A second serious limitation is the terrible mathematical difficulty of solving continuous location models optimally. When more than one warehouse or one product is involved it is usually a practical impossibility to find and verify a true solution.

There are other major limitations, but the ones already mentioned suffice to show that this approach is not a strong contender for use in distribution systems planning. However, the simplest versions may be useful as a means of generating plausible candidate locations for subsequent use by a more realistic discrete location model/method.

9. *Heuristic Methods*

The persistent difficulty of optimally solving realistic distribution planning models has led to the proposal of various so-called "heuristic" methods for attempting to find good solutions. Such methods attempt to emulate on a computer various elementary manual techniques for sorting through design alternatives. An exposition of the philosophy of the heuristic approach can be found in the well-known article by Kuehn and Hamburger [22]. One of the more recent applications of the heuristic approach is that by Drysdale and Sandiford [9].

The inability of heuristic methods to satisfy the first and third computational criteria seriously limits their practical usefulness. Fortunately, the greatly increased power of recent truly optimizing methods to solve large realistic models usually makes it unnecessary to fall back on heuristic methods. The best of today's optimizing methods are actually *faster* than many available heuristic methods for the same model, by virtue of their far greater sophistication. In such cases there can be no justification for using heuristics.

In this writer's opinion the proper place of heuristic methods is in dealing with complex repetitive tasks of a tactical operational nature (e.g., production scheduling or vehicle routing) rather than of a strategic planning nature. Then "good solutions on the average" may suffice.

Conclusion

The introduction defined the general class of distribution planning problems to which this article is addressed. It also pointed out some of the shortcomings of conventional manual analysis by comparison with a computer-assisted study based on a comprehensive distribution system model. It follows that the latter approach should be employed whenever practicable.

Practicality criteria were developed in the next two sections. The first section identified and discussed seven advanced problem features that often need to be incorporated into a distribution planning model if it is to yield valid results and earn credibility in the eyes of management. The next section discussed the three criteria which determine the practical usefulness of any computational method for manipulating and solving such models. A general conclusion is that very stringent requirements are imposed on both models and computational methods if they are to achieve full success.

These requirements were applied in the third section to critically evaluate the leading available model/methods for distribution systems planning. A brief summary of how these alternatives stack up against the requirements is presented in Table 1 for those alternatives that are properly viewed as optimizing rather than suboptimizing. A number of suboptimizing alternatives were also discussed, including:

1. Casewise linear programming, which may be useful when there is but a small number of alternative configurations to be examined and little or no economies of scale;
2. Computer simulation, which is best for evaluating a few fully specified distribution design alternatives in great detail as regards their dynamic operating characteristics;

CRITERIA		MODEL/METHODS	Decomposition Method (Ref. 16)	General Mixed Integer Linear Programming (Refs. 23, 25, 30)
IMPORTANT PROBLEM FEATURES TREATED PROPERLY	1.	Multiple Products	YES	YES
	2.	Two Stages of Distribution: Plants, Warehouses and Customers	YES	YES
	3.	Capacities for Plants and Size Limits for Warehouses	YES	YES
	4.	Warehouse Economies of Scale and Fixed Charges	YES	YES
	5.	Each Customer Serviced by a Single Warehouse	YES	YES (but usually impractical)
	6.	Shipments to Customers: Preserve Identity of Originating Plant	YES	YES
	7.	Optional Restrictions on System Configuration	YES	YES
COMPUTATIONAL CRITERIA	1.	Capable of True Optimization	YES	YES
	2.	Probable Computational Efficiency for Problems of *Moderate* Size	GOOD	POOR TO ADEQUATE
	3.	Suitability for Multiple Secondary Optimization Runs: Fast, Reliable, Easy to Use	GOOD	POOR

Table 1 Summary Evaluation of Available Optimizing Model/Methods

<table>
<tr><th colspan="3">The Best Methods for a Single Stage of Distribution</th><th rowspan="2">POLIGAMI
(Refs. 26, 27)</th><th rowspan="2">CAPFLO
(Refs. 13, 19)</th></tr>
<tr><th>FLAC
(Ref. 18)</th><th>Akinc and Khumawala
(Ref. 1)</th><th>Khumawala
(Ref. 21)</th></tr>
<tr><td colspan="3">YES
(with the help of a simple trick)</td><td>YES
(In the absence of problem feature 2)</td><td>YES
(In the absence of problem feature 2)</td></tr>
<tr><td>NO</td><td>NO</td><td>NO</td><td>YES
(In the absence of problem feature 1)</td><td>YES
(In the absence of problem feature 1)</td></tr>
<tr><td>YES</td><td>YES
(upper limits only)</td><td>NO</td><td>YES
(upper limits only)</td><td>YES
(upper limits only)</td></tr>
<tr><td>YES</td><td>YES</td><td>YES</td><td>YES</td><td>YES</td></tr>
<tr><td>YES
(optional)</td><td>NO</td><td>NO</td><td>NO</td><td>DOUBTFUL</td></tr>
<tr><td colspan="3">← Not Applicable →</td><td>NO</td><td>NO</td></tr>
<tr><td>YES</td><td>NO</td><td>NO</td><td>NO</td><td>YES</td></tr>
<tr><td>YES</td><td>YES</td><td>YES</td><td>YES</td><td>YES</td></tr>
<tr><td colspan="3">← GOOD TO ADEQUATE →</td><td>PROBABLY ADEQUATE</td><td>PROBABLY ADEQUATE</td></tr>
<tr><td colspan="3">← GOOD TO ADEQUATE →</td><td>PROBABLY ADEQUATE</td><td>UNCERTAIN</td></tr>
</table>

3. Continuous location methods, which might be useful in the preliminary stages of a study to generate plausible candidate locations for new facilities; and

4. Heuristic methods, which have diminished appeal and justification now that powerful optimizing methods are available.

It is hoped that decision makers will find this evaluation framework useful for selecting the most suitable approach for a given application, and that researchers and software vendors will be stimulated to invest their future efforts in the most worthwhile directions. In this way, hopefully, computer-assisted distribution systems planning will more rapidly and widely achieve its full potential.

References

[1] Akinc, U., and Khumawala, B. M. *An Efficient Branch and Bound Algorithm for the Capacitated Warehouse Location Problem.* Working Paper, Graduate School of Business Administration, University of North Carolina, July 1974.

[2] Atkins, R. J., and Shriver, R. H. "New Approach to Facilities Location." *Harvard Business Review,* May-June, 1968, pp. 70-79.

[3] Bartakke, M. N.; Bloomquist, J. V.; Korah, J. K.; and Popino, J. P. "Optimization of a Multi-National Physical Distribution System." Sperry Rand Corporation, Blue Bell, Pa., presented at the 40th National ORSA Meeting, Anaheim, California, October 1971.

[4] Bender, P. S. "Principles of Physical Distribution Systems Configuration Design." Marketing Publications Incorporated, Washington, D.C., 1971.

[5] Bowersox, D. J. *Logistical Management.* New York: Macmillan, 1974.

[6] Bowersox, D. J.; Helferich, O. K.; Marien, E. J.; Gilmour, P.; Lawrence, M. L.; Morgan, F. W.; and Rodgers, R. T. *Dynamic Simulation of Physical Distribution Systems.* Michigan State University Business Studies, 1972.

[7] Connors, M. M.; Coray, C.; Cuccaro, C. J.; Green, W. K.; Low, D. W.; and Markowitz, H. M. "The Distribution System Simulator." *Management Science 18* 8 (April 1972).

[8] Davis, P. S., and Ray, T. L. "A Branch-Bound Algorithm for the Capacitated Facilities Location Problem." *Naval Research Logistics Quarterly 16* 3 (September 1969), pp. 331-344.

[9] Drysdale, J. K., and Sandiford, P. J. "Heuristic Warehouse Location— A Case History Using a New Method." *CORS Journal 7* 1 (March 1969), pp. 45-61.

[10] Efroymson, M. A., and Ray, T. L. "A Branch-Bound Algorithm for Plant Location." *Operations Research 14* 3 (May-June 1966), pp. 361-368.

[11] Eilon, S.; Watson-Gandy, C. D. T.; and Christofides, N. *Distribution Management.* New York: Hafner Publishing Company, 1971.

[12] Elson, D. G. "Site Location via Mixed-Integer Programming." *Operational Research Quarterly 23* 1 (March 1972), pp. 31-43.

[13] Fieldhouse, M. "The Depot Location Problem." University Computing Company, Ltd., London, presented at the 17th International Conference of TIMS, London, July 1970.

[14] Figgens, J. M., and Thompson, S. A. "A Methodology for Developing Sensitivity Analysis on Specified DOD Logistics Systems." Report GSA/SM/73-7, U.S.A.F. Air University, Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio, June 1973.

[15] Geoffrion, A. M. "A Guided Tour of

Recent Practical Advances in Integer Linear Programming." Working Paper No. 220, Western Management Science Institute, UCLA, September 1974.

[16] Geoffrion, A. M., and Graves, G. W. "Multicommodity Distribution System Design by Benders Decomposition." *Management Science 20* 5 (January 1974), pp. 822-844.

[17] Geoffrion, A. M., and Marsten, R. E. "Integer Programming Algorithms: A Framework and State-of-the-Art Survey." *Management Science 18* 9 (May 1972), pp. 465-491.

[18] Geoffrion, A. M., and McBride, R. D. "The Capacitated Plant Location Problem with Additional Constraints." Paper presented at the Joint National Meeting of AIIE, ORSA, and TIMS, Atlantic City, N.J., November 1972.

[19] Haverly, C. A., and Fieldhouse, M. "Warehouse Location Calculations." Notes distributed at a Mathematical Programming Project session at the SHARE XLII Meeting in Houston, March 1974.

[20] Hupp, B. W. "How Many Distribution Centers . . .?" In *Proceedings of the Eleventh Annual Conference* National Council of Physical Distribution Management, San Francisco, September 1973.

[21] Khumawala, B. M. "An Efficient Branch and Bound Algorithm for the Warehouse Location Problem." *Management Science 18* 12 (August 1972).

[22] Kuehn, A. A., and Hamburger, J. J. "A Heuristic Program for Locating Warehouses." *Management Science 9* 4 (July 1963), pp. 643-667.

[23] *Mathematical Programming System Extended, Mixed Integer Programming.* Manual SH20-0908, IBM Corporation, White Plains, New York.

[24] Nolan, R. L., and Sovereign, M. G. "A Recursive Optimization and Simulation Approach to Analysis with an Application to Transportation Systems." *Management Science 18* 12 (August 1972).

[25] *OPHELIE/LP Addendum: Mixed Integer Capability of OPHÈLIE/LP System.* Publication No. D0001507032, Control Data Corporation, Minneapolis, 1970.

[26] *Poligami User's Manual.* London: S.I.A. Ltd. (April 1970).

[27] *SIA Depot Location: POLIGAMI and MULTICOM.* Brochure circa 1971.

[28] Soland, R. M. "Optimal Facility Location with Concave Costs." *Operations Research 22* 2 (March-April 1974), pp. 373-382.

[29] Spielberg, K. "A Simple Plant Location Program—SPLT 1." IBM Contributed Program Library Writeup 360D-15.6.001, April 1968.

[30] *UMPIRE: Unified Mathematical Programming System Incorporating Refinements and Extensions: User's Guide.* Publication S00037-00-00, Computer Sciences Corporation, Los Angeles, 1970.