

# A MANAGEMENT SUPPORT SYSTEM FOR DISTRIBUTION PLANNING

ARTHUR M. GEOFFRION

*Graduate School of Management  
University of California, Los Angeles*

GLENN W. GRAVES

*Graduate School of Management  
University of California, Los Angeles*

SHAO-JU LEE

*Department of Management Science  
California State University  
Northridge, California*

## ABSTRACT

This paper provides a status report on the authors' long-term work in the area of distribution planning. Technical challenges relating to large-scale optimization continue to be stimulating, but an increasingly large share of our attention has been absorbed by the need to imbed optimization techniques within a complete system containing the many other capabilities needed to facilitate successful practical application. We call a complete system of this sort a "management support system." We describe such a system in some detail, with an emphasis that deliberately excludes the technical aspects of optimization. In this way we hope to accommodate the broadest possible audience.

No new research results are presented. The exposition is non-mathematical and should be accessible to anyone interested in modern approaches to distribution planning. A modest amount of supporting technical detail, however, is given in an appendix.

## RÉSUMÉ

Cet article est un rapport sur le statut du travail à long terme des auteurs dans le domaine de la planification de la distribution. Les défis techniques se rapportant à l'optimisation à grande échelle continuent d'être stimulants, mais une part grandissante de notre attention a été absorbée par le besoin d'introduire des techniques d'optimisation dans un système complet contenant les nombreuses autres capacités requises pour faciliter l'application pratique avec succès. Nous appelons un système complet de cette sorte un "système de support de gestion". Nous décrivons en détail un tel système, en insistant délibérément sur l'exclusion des aspects techniques de l'optimisation. De cette façon nous espérons accommoder la plus large audience possible.

Aucun résultat nouveau n'est présenté. L'exposé est non-mathématique et devrait être accessible à quiconque s'intéresse aux approches modernes de la planification de la distribution. Un modeste ensemble de détails techniques est tout de même donné en appendice.

*A management support system* is composed of a model, supporting data files, a solver to exercise the model, and interface facilities between these things and the manager's world, all integrated and organized so as to support

managerial decision-making within a prescribed domain on a continuing basis.<sup>(6)</sup> ODS is such a system for strategic and tactical distribution planning problems. Originally designed for a particular food industry application<sup>(2,4)</sup> it has evolved through many generations of changes in the course of numerous applications over the years.

The current version of this system is described in a non-technical fashion. The description is non-mathematical and should be accessible to anyone interested in modern approaches to distribution planning.

Section A describes the kinds of target issues with which ODS is designed to deal and reviews the management support system concepts used to organize much of the remaining exposition. Section B outlines the modelling options and associated data elements, with details deferred to appendix 1. Section C gives a simple intuitive interpretation of the decomposition technique by which optimization is accomplished. Section D describes the elaborate interface facilities needed for productive man-machine communication. Section E places ODS in the larger context of a typical planning project, and a concluding section comments briefly on practical applications to date and prospects for the future. A modest amount of technical detail is given in appendix 2.

#### A. SYSTEM OBJECTIVES

ODS aims to provide an efficient optimization-based management support system applicable to a wide variety of distribution planning issues at an annual level of detail.

Table 1 presents some typical target issues for which the system is intended. These are divided into three categories. The focus of *network rationalization issues* is the classical question of warehouse location and the many important corollary questions which this entails. All these questions are addressed simultaneously and the answers are obtained in each optimization run. The "*What if ...*" questions, by contrast, must be addressed individually on a case-by-case basis. For example, three alternative plant capacity expansion proposals necessitate the analysis of three separate cases. Finally, *sensitivity issues* require a series of optimization runs where a single factor (like the level of customer service) is varied systematically.

A more extensive discussion of table 1 and related matters is given elsewhere.<sup>(5)</sup>

It is important to recognize the difference between a fixed-installation system – one designed for a specific application in a single organization – and a system designed for many applications in a variety of organizations. The latter type of system, of which ODS is an example, must be designed to function in a much less stable and more demanding environment characterized by a constant stream of inexperienced users, frequent

TABLE 1  
Typical Target Issues

---



---

<b>I. NETWORK RATIONALIZATION ISSUES</b>
1. How many DCs (warehouses) should there be?
2. Where should the DCs be located?
3. What is the proper ownership and design option for each DC?
4. What is the proper size of each DC?
5. How should inventory be deployed?
6. How much of each product should come from each source?
7. What is the proper use of contract suppliers?
8. What is the proper allocation of supplies to DCs and customers?
9. Which customers should get which products from which DC?
10. How to resolve direct sourcing versus warehousing versus crossdocking choices for serving individual customers?
11. What is the relative profitability of customers and products?
<b>II. "WHAT IF" QUESTIONS</b>
A. Environmental issues
1. Impact of possible changes in demand structure
2. Impact if deregulation or higher fuel costs for transportation
3. Impact on a plant or DC of transportation strike, natural disaster, weather closure, energy shortage, etc.
B. Business decision and policy issues
1. Evaluation of plant capacity expansion proposals
2. Plant location studies
3. Impact of introducing a new product line or discontinuing one
4. Pricing policy analysis
5. Split delivery policy analysis
6. Expanding into new market regions
7. Transportation policy analysis
8. Outbound freight consolidation program evaluation
9. Evaluation of DC capacity expansion or mechanization proposals
10. Inventory policy comparisons
11. Impact of combining autonomous divisional distribution systems
12. Evaluation of alternative distribution echelon structures
13. Implementation priority analysis
<b>III. SENSITIVITY ISSUES</b>
1. Quantify the tradeoff between distribution cost and customer service
2. Quantify the variation of least distribution cost with the number of DCs
3. System sensitivity analysis with respect to any factor, e.g., inflation by cost category
4. Influence of demand change over time

---

requests for new system capabilities, and continual technical challenges. Survival in such an environment demands a fast pace of evolutionary improvement. ODS is quite different from what one would expect of a fixed-installation system, and it is safe to predict that the ODS of tomorrow will exhibit still other evolutionary adaptations.

As an optimization-based management support system ODS has four parts: a *modelling framework*, a *data file structure*, an *optimizer*, and *interface facilities* between these and system users.

A *modelling framework* is a conceptual context and a menu of options within which a variety of specific models can be tailored to represent a variety of

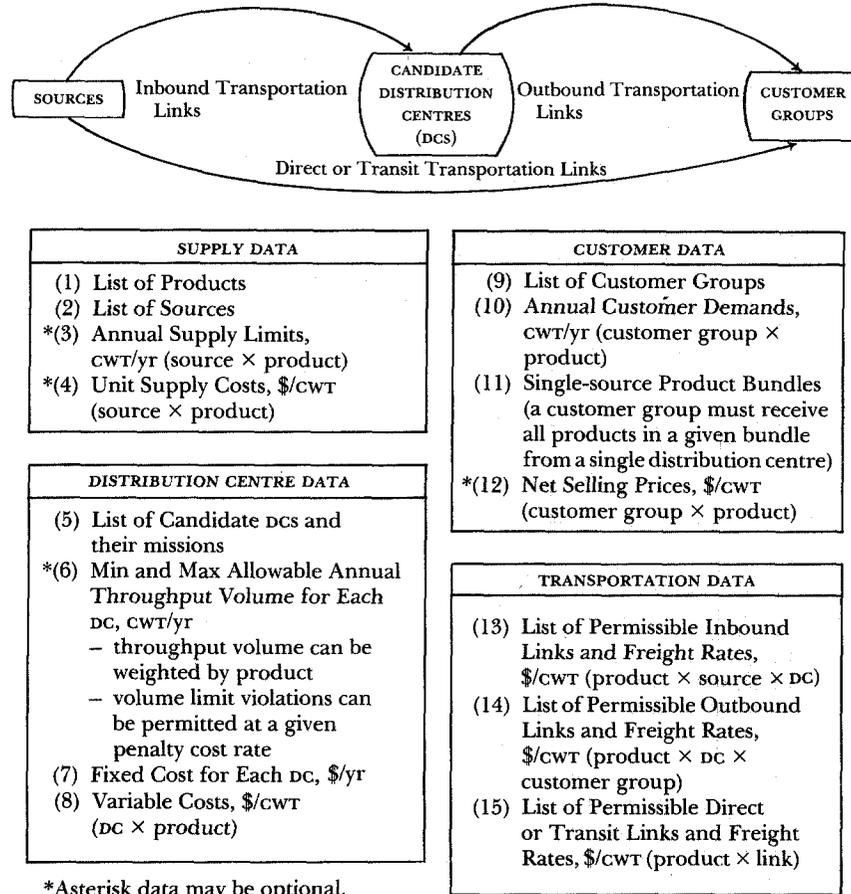


FIG. 1. Modelling framework and data element checklist (see appendix 1 for further details).

specific real distribution systems. A *model* is a careful description of a real system and of the relationships by which, under trial assumptions, alternative managerial choices are believed to influence pertinent system performance measures.

A *data file structure* is a set of guidelines specifying, for the data elements required by a model, how the data are to be organized and represented in computer-readable form.

An *optimizer* is a computer routine for exercising a model so as to find those particular managerial choices which, under trial assumptions, yield the best value for a designated performance measure.

*Interface facilities* serve two functions: to facilitate creation and maintenance of data files, and to facilitate communication between distribution managers (or supporting staff) and the management support system as a functioning whole.

To describe ODS it suffices to describe each of the four parts individually. This is done in the next three sections.

## B. MODELLING FRAMEWORK AND DATA ELEMENTS

Figure 1 depicts the overall modelling framework of ODS and indicates the associated data elements. Briefly, multiple products from (possibly) capacitated sources flow through (possibly) capacitated candidate distribution facilities to meet customer demands. Direct shipment may be allowed, and customers may stipulate single-sourcing requirements. The fundamental task of the optimizer is to find the distribution system design that achieves the desired level of customer service at minimum cost within all necessary constraints and desired management policies.

### *Determine*

- how many and which of the candidate distribution centres to open
  - what size they should be
  - which customers they should be assigned to serve for which products
  - what all the transportation flows should be throughout the system
  - how much of each product should be shipped from each source
- so as to minimize the sum of all costs associated with*
- supply
  - transportation
  - distribution centres
  - inventory
  - system reconfiguration
- subject to all appropriate constraints on*
- available supply
  - allowable throughput of the distribution centres
  - demands to be met
  - single-sourcing of customers
  - customer service.

Proper use of this optimization capability and of the modelling options implicit in figure 1 enables all the issues of table 1 to be addressed. Further details are given in appendix 1.

The kinds of specific models that can be built within this framework include the most widely studied “warehouse location” models appearing in the literature of management science / operations research.

Two comments concerning the generality of this modelling framework are in order. The first concerns the fact that it is stated at an *annual* level of detail, and hence does not explicitly address operating decisions that must be made in a short time frame (like vehicle scheduling). Nevertheless, an annual model can be designed to take detailed account of the cost / service implications of operating policies, can be used to evaluate alternative operating policies in depth, and is fully compatible with a number of analytical models of operating policies in areas such as inventory control, expediting, and mode selection.

The second comment is that the modelling framework actually embraces a still broader range of practical distribution systems than figure 1 appears to cover. This is due in part to the existence of numerous modelling devices for converting seemingly different problem features back to standard form and in part to special software features not described explicitly here. Such modelling devices and software features permit accommodation, for instance, of economies of scale in warehousing and transportation, centralized stock point selection, crossdocking choices, and a number of seemingly more general echelon structures (raw material flows, interplant transfer, master-satellite warehouse, etc.).

The data elements listed in figure 1 are assembled in computer-readable form as an Input Data File. Organization and formats have been designed to facilitate manual preparation and visual inspection.

### C. SOLUTION BY DECOMPOSITION

The ods optimizer uses a decomposition technique of advanced design,<sup>(4)</sup> but it is possible to give an intuitively meaningful interpretation of it as a *trial-and-error procedure with learning*. This should (1) dispel some of the needless mystery that often surrounds sophisticated optimization techniques, and (2) clarify some of the advantages inherent in this approach.

Distribution systems existed long before the advent of management science in the 1940s. How were they designed? By trial-and-error, of course, as pictured in figure 2: select a trial system configuration, evaluate its performance, select another trial configuration based on the results of previous trials, and so on, until the improvements expected from new trial configurations appear small relative to the trouble of further calculations. A "trial system configuration" was usually an assumed set of locations for distribution facilities. An "evaluation" of performance usually meant assigning customers to the closest facilities, assigning facilities to the cheapest plants, adjudicating any capacity violations or other constraints in an ad hoc manner, and then calculating various cost and customer service measures.

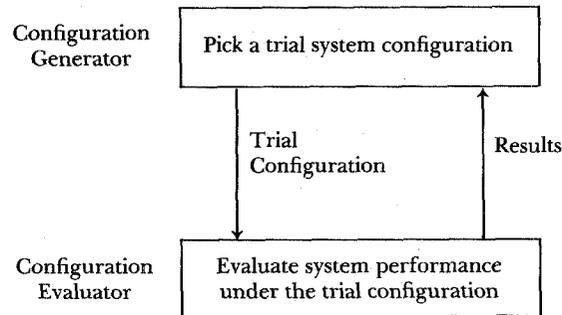


FIG. 2. A simple trial-and-error approach to distribution system configuration design.

The chief merit of this time-honoured trial-and-error approach is its understandability. Among its major shortcomings are the following.

1. It can consider only a small fraction of all possible alternatives.
2. It is very tedious, with many tempting but dangerous short-cut opportunities (e.g., regionalize, marginalize, linearize, omit).
3. Sensitivity analysis is impractical.

With the advent of linear programming about three decades ago, it became possible at least partially to computerize the evaluation step of the trial-and-error approach, since an important part of this step is to find least cost transportation flows meeting prescribed supply and demand constraints. This approach, still in use today, helped somewhat to mitigate the disadvantages of the completely manual approach, but all three of the shortcomings mentioned still remain.

The decomposition approach employed can be viewed as the ultimate refinement of the line of attack sketched above in that it automates the selection of trial configurations in a manner that guarantees true optimization. It eliminates the first and second shortcomings and makes significant progress on the third.

More precisely, the decomposition approach can be schematized as in figure 3. The main apparent differences between figures 2 and 3 are as follows.

- (A) A trial system configuration now includes customer group assignments as well as which facility locations are open.
- (B) Evaluation of a trial configuration now separates into a simple independent evaluation subproblem for each product (in fact, each subproblem is a network flow problem of classical transportation type).
- (c) The results returned from the evaluation step are of a particular type dubbed "Cost Estimating Relationship."

Difference (A) is occasioned by the desire to confine all of the combinator-

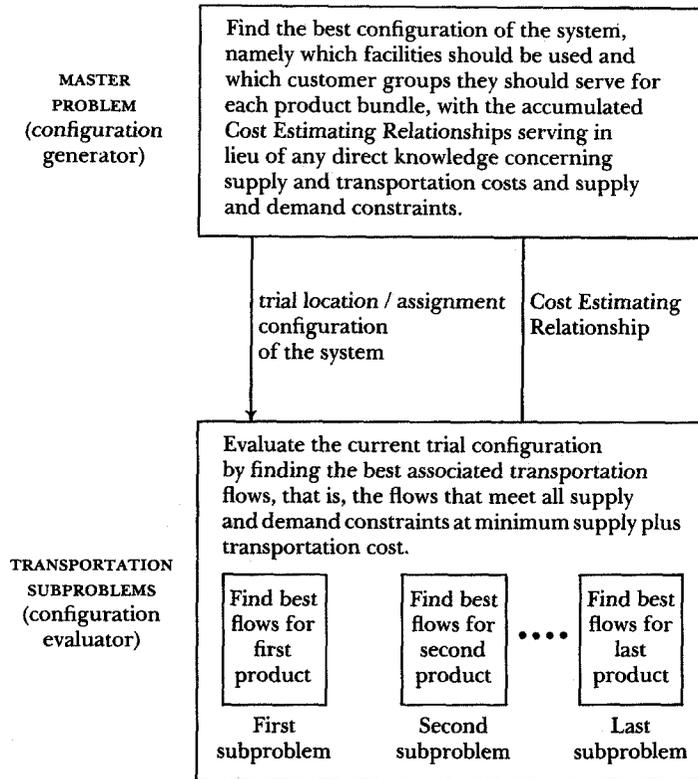


FIG. 3. The decomposition approach.

ial aspects of system design to the so-called Master Problem. This in turn makes possible the simplification described in (B), because, when locations and assignments are held fixed temporarily, there is nothing to make the best choice of one product's transportation flows depend on the best choice of another product's transportation flows. Prior to fixing the locations and assignments, there are the following interdependencies between the transportation flows of the different products:

- They compete for limited distribution centre throughput capacity.
- A distribution centre fixed cost is levied only once when at least one product flows through.
- The single-source product bundling rules must be obeyed.

These interdependencies are absent at the evaluation step. As for (C), each Cost Estimating Relationship is a readily available byproduct from the Transportation Subproblems and furnishes a lower estimate of the supply and transportation cost consequences of departing from the trial

system configuration for which its parent Transportation Subproblems were solved. The cumulative effect of these Cost Estimating Relationships in the Master Problem is to help it "learn" enough so that eventually it can produce a truly optimal system configuration and demonstrate that no better configuration exists.

The ping-ponging back and forth between Master Problem and Subproblems stops when a Termination Test indicates that further iterations are not worthwhile. The Termination Test is based on the fact that a lower bound on optimal total system cost is produced each time a Master Problem is solved (it is a "relaxation" of the full problem), while an upper bound on optimal total system cost is produced each time the Transportation Subproblems are solved. When the lower and upper bounds are "close enough" to each other, say within 1/4 of 1%, the best total system design found thus far must be within the same percentage of the true mathematical optimum. Some further technical details are given in appendix 2.

In summary, the decomposition scheme is a mathematically sound procedure for solving the problem as a whole with the help of any available mechanisms for the Master Problem and the Transportation Subproblems.

Several inherent advantages of the decomposition approach are evident from the preceding discussion. Perhaps the most important is that the approach is quite insensitive to the number of products (provided this number does not significantly influence the number of major iterations, which has in fact been the empirical case). The size of the Master Problem is completely independent of the number of products, and the Transportation Subproblems are solved *sequentially* by product. Thus primary computer storage requirements are essentially independent of the number of products and computing time grows only linearly with this number.

A second advantage is that the evaluation of manually prepared system configurations is quite convenient. Such configurations usually arise from interesting "What if ...?" modifications of results obtained from previous runs. It is simply necessary to feed such a configuration to the Transportation Subproblems.

A third advantage is that the Cost Estimating Relationship corresponding to the optimal configuration yields, with negligible extra work, a table of Customer Reassignment Penalties that is quite useful for interpreting just how sensitive the optimal configuration is to specific alterations of the service areas. For instance, the penalty table might show that the annual cost penalty is at least \$1,360 for switching the Orlando customer group from the Miami distribution centre to the Savannah distribution centre. It would require extensive additional calculations to

obtain such penalties with any conventional (non-decomposition) optimization approach. See appendix 2 for technical details.

#### D. INTERFACE FACILITIES

The usefulness of a management support system depends heavily on the quality and completeness of its interface facilities, particularly those supporting the creation, maintenance, and revision of the model and its data. This fact led to the development of a model / data management module called the Preprocessor, which runs independently of the Optimizer module. Figure 4 shows the relation between the two modules and the division of labour between them. Note that the Preprocessor accepts either an Input Data File (in creation mode) or an old Problem File (in edit mode), depending on whether the problem being prepared for optimization is a "new" one or a revised version of an "old" one.

Three design goals have been paramount: (1) correctness of the problem being solved, (2) proper documentation and control of the problem being solved, and (3) user convenience and productivity. These goals are reflected in the details of Preprocessor and Optimizer functions summarized below.

The main *functions of the Preprocessor* are:

1. read the Input Data File and insert missing values using prescribed data extension rules (extensive default, replication and generation capabilities substantially reduce the physical size of the Input Data File);
2. perform various error checks, internal consistency tests, and diagnostic routines, with explicit messages printed whenever difficulties are suspected (about fifty types of messages are possible); automatic recovery procedures are necessary to avoid premature termination of a Preprocessor run when errors are encountered;
3. implement certain modelling options (e.g., criteria for outbound link inclusion);
4. print an extensive *Problem Summary*, with various options, for reference purposes;
5. generate a *Problem File* suitable for direct access by the Optimizer;
6. perform various editing functions on any existing Problem File so that users can correct input errors or modify data in preparation for desired Optimizer runs; editing functions include deletion, insertion, or replacement of any data entry and scaling of groups of data in various ways.

The main *functions of the Optimizer* are:

1. read the Problem File and accept certain additional user commands such as locking open or closed selected distribution centres;

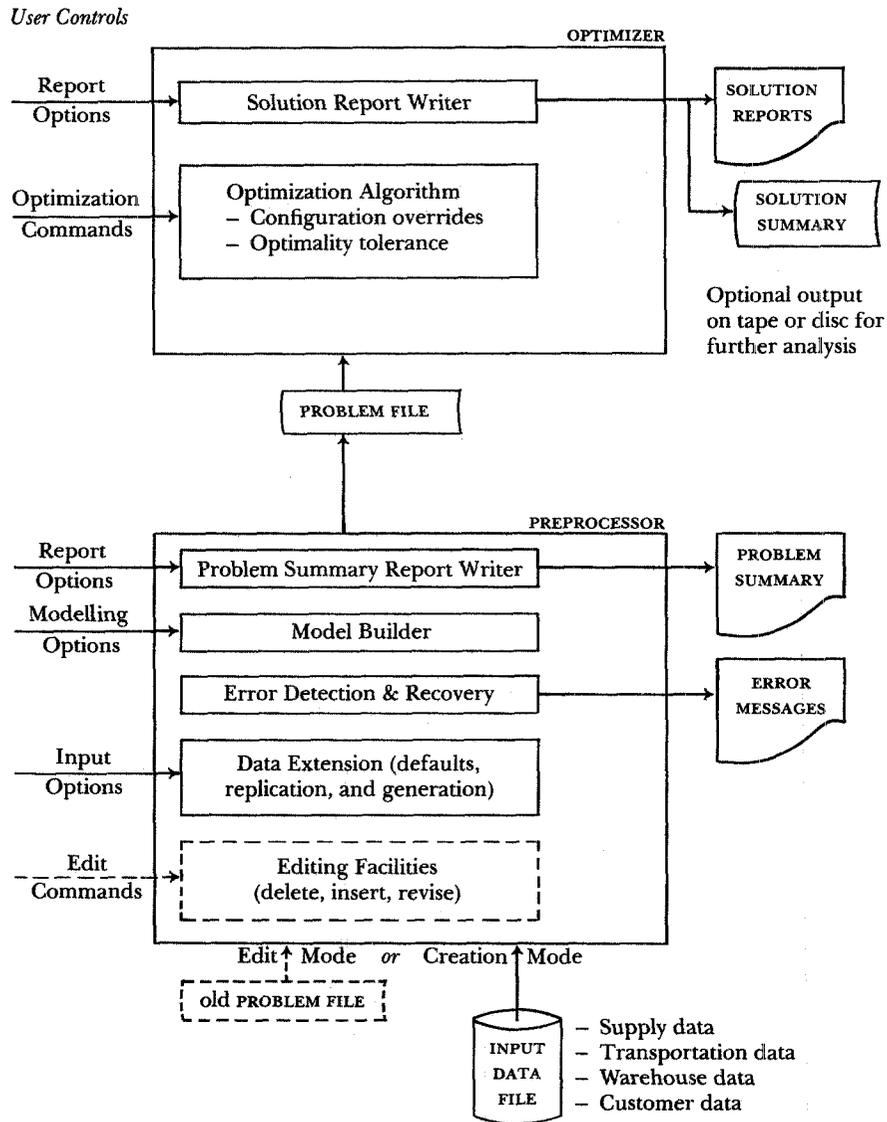


FIG. 4. System overview.

2. determine an optimal distribution system design via the decomposition technique;
  3. print comprehensive reports, with various options, describing the optimal distribution system design (see below).
- The reports produced by the Optimizer are designed to be read

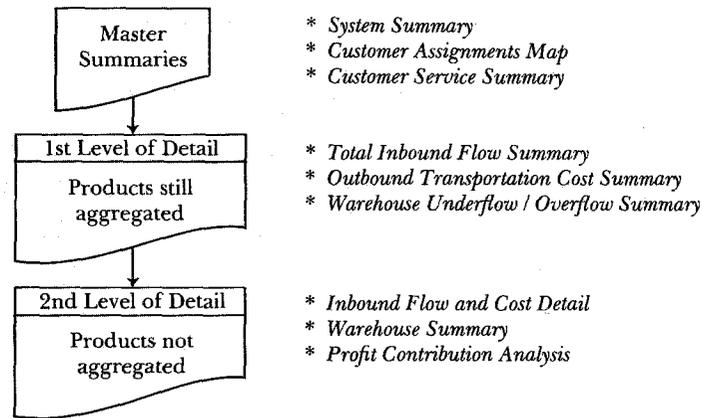


FIG. 5. Hierarchy of solution reports.

directly by distribution managers. The principal reports fall into a three-level hierarchy as shown in figure 5.

It suffices for illustrative purposes to give details on just one of these reports. The *System Summary* gives annual totals over all products as shown in figure 6.

Not all reports are of the traditional tabular or columnar type. The *Customer Assignments Map* shows the locations of open distribution centres and all customer groups superimposed on a computer-generated map of the United States in such a way that the optimal service area of each DC can be seen at a glance. And the *Customer Service Summary* is presented graphically in the form of a histogram.

The *Customer Reassignment Report*, mentioned at the conclusion of section c, falls between the first and second levels of detail (products are aggregated by bundles as explained in data element 11). It gives estimates of the cost increases that would be incurred by reassigning customers to distribution centres other than the one specified in the optimal system design. This report makes it possible to examine quickly, and without the need for a computer, any moderate variant of the "optimal" system configuration produced by any given Optimizer run.

#### E. STAGES OF A TYPICAL PLANNING PROJECT

The stages of a typical ODS-based distribution planning project have much in common with those of any planning project that makes use of a management support system. They are summarized in figure 7.

Time moves from left to right in this diagram, although the horizontal widths of the boxes representing the various stages are not proportional

Warehouse	Total Throughput	Fixed Cost	Variable Cost	Production Cost	Inbound Cost	Outbound Cost	Total Transportation Cost
name	quantity	xxx	xxx	xxx	xxx	xxx	xxx
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
name	quantity	xxx	xxx	xxx	xxx	xxx	xxx
	total	total	total	total	total	total	total
		Total Warehouse Costs		Total Production plus Transportation Costs			
		Grand Total System Costs					

FIG. 6. System summary.

to the time duration of each. Notice also that two activity streams are identified, corresponding to: (1) the "lead" members of the project team who assume responsibility for the success of the endeavour, and (2) the "supporting" members who carry out various specific tasks under the general direction of the lead persons. Although the six stages tend to occur in the sequence indicated, there is in fact a good deal of overlap (indicated by the slanted boundaries), anticipation, and even cycling (as between the third and fourth stages).

#### Stage 1: Project mandate

The objectives of a project and the guidelines within which they are to be pursued are set forth in an executive-level mandate. Toward this end it is appropriate to:

- A) make a realistic *situation analysis* of the distribution function – its role and mission in the company, synopsis of present system and operating policies, environmental trends and influences, problems and opportunities;
- B) decide on the specific *target issues* of interest (cf. table 1) and define the allowable *scope* of possible system changes;
- C) select the quantitative *performance measures* by which the desirability of any proposed changes will be judged (typically relating to cost, investment, customer service, and system flexibility and vulnerability to disruption);
- D) elicit any *managerial edicts* (assumptions, constraints, and policies) to be honoured (e.g., only public warehousing will be permitted);

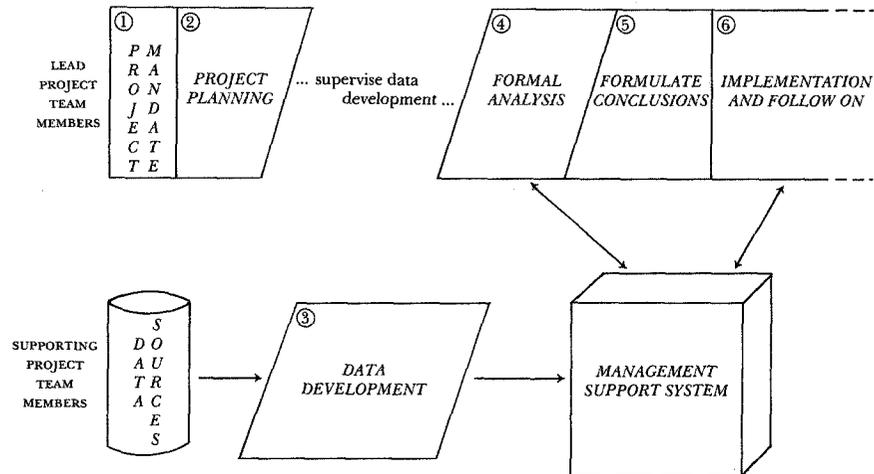


FIG. 7. Typical stages of a planning project.

- E) consider the *capabilities and limitations of ODS* and any other likely analytical tools; consider the current implications of maintaining a distribution planning system beyond the completion of the seminal project;
- F) finalize a written *project mandate*, in the light of all the above, that is specific and prioritized as to expectations and project scope and realistic as to attainability within the allocated resources.

#### Stage 2: Project planning

This is a master planning activity to work out how the project will fulfil its mandate. Activities include:

- A) design the basic model within the available modelling framework;
- B) identify data requirements and available data sources;
- C) develop a strategy for data development and aggregation (definitions, working assumptions, criteria, estimation and sampling procedures);
- D) specify requirements for any supplementary computer programs needed to facilitate data development, and for any necessary modifications to ODS;
- E) decide on an approach to model validation;
- F) develop a preliminary schedule of stage 4 computer runs and what these require in terms of Input Data File generation and editing;
- G) create an organization, assign responsibilities and due dates for the major tasks, and establish suitable project control procedures;
- H) prepare a written comprehensive Project Plan and secure consensual approval.

Project planning is a high-level activity requiring a wide range of technical and managerial skills.

*Stage 3: Data development*

The data development process enters its detail phase when the Project Plan has been completed. The "raw materials" are the many possible sources of data both within the firm and without. The "end product" is an Input Data File acceptable to the Preprocessor (usually more than one, and often with variants in a form suitable for digestion via editing). The principal activities are:

- A) prepare any computer programs needed for data development;
- B) data acquisition (manual acquisition, computerized extraction from internal sources, and procurement from external sources);
- C) aggregation, summary, analysis, and synthesis of the necessary data elements (see fig. 1);
- D) verification of data elements;
- E) reformat data elements and assemble as an Input Data File.

These activities are guided by the Project Plan, but a considerable amount of tactical maneuvering usually occurs, owing to unforeseen circumstances.

*Stage 4: Formal analysis*

Verification of individual data elements will for the most part have been done during the data development process. Now it is necessary to gain assurance that the various elements fit properly when joined together. One or more ODS validation runs are required.

Analysis of the project's target issues begins in earnest once the validation exercises have been concluded successfully. ODS is used as an analytical problem-solving and insight-generating tool. Some runs will have been planned in advance, but the need for others will become evident only as the analysis proceeds.

There should be a relentless attempt to *explain* the findings that emerge at this stage. It is not enough just to know *what*; *why* is equally important, because it is the key to genuine understanding and insights into the firm's distribution system.

*Stage 5: Conclusions*

The findings developed at the previous stage are now interpreted in light of considerations outside the scope of the formal analysis – current organizational priorities and politics, environmental trends, modelling uncertainties, and the like. Specific prioritized recommendations and an implementation plan are worked out, documented, and approved by the necessary authorities.

Stages 4 and 5 are closely intertwined, the distinction being mainly one of problem-solving posture: objective analysis and research, on the one hand, versus judgmental synthesis and pragmatism on the other.

Project documentation, which began with the written Project Mandate at stage 1, now reaches its culmination. The Project Plan will have been revised and expanded continually throughout the third and fourth stages. Now conclusions and recommendations are set forth, along with their associated supporting arguments. Care should be taken to make these arguments truly substantive rather than just defences of methodology. Care should also be taken to capture the newly emerged system understanding and insights, which, whether or not they impact the final recommendations, are a valuable project legacy.

*Stage 6: Implementation and follow-on*

The process of turning the project's recommendations into reality depends largely on traditional managerial skills. In most cases it will be found that ODS has a useful role to play as a tool for monitoring the progress of the implementation plan and adapting it to changing circumstances. This requires, of course, that ODS be maintained during the implementation period.

A management support system based on ODS has many of the attributes of a capital asset. It is capable of evolving over time and finding uses far beyond the original project that spawned it (see table 1 and<sup>(5)</sup>).

## F. RESULTS

The first application in the early 1970s of the original version of ODS was reported in a Harvard Business Review article.<sup>(2)</sup> The system has evolved continually since then and has been applied successfully to such industries as automotive parts, chemicals, consumer products, food products, health care products, household appliances, industrial supplies, mining, paper products, and wholesale distribution.<sup>(6,7)</sup> Applications to several other industries are currently in progress. Virtually all the sample issues listed in table 1 have been addressed in various applications.

Such applications involve up to approximately 100 products, 100 sources, 100 distribution centres, 400 customer groups, and 30,000 customer group / bundle → distribution centre assignment choices. The typical industrial application has 2–50 products, 5–20 plants, 30–60 candidate distribution centres, 125–300 customer groups, and 3000–15,000 binary choices. Solutions to within a few tenths of 1% of global optimality usually require 2–5 major iterations, 300–900 kilobytes of main storage, and a few minutes of CPU time on an IBM 3033.

The system also served as the primary analytic tool in a three-year

high-level Department of Defense study aimed at reorganizing the entire defence logistics system for the continental United States (all four armed services and the Defense Logistics Agency).<sup>(7)</sup> This application gave rise to considerably larger problems, in the neighbourhood of 15,000–35,000 binary choices. Many such problems have been solved in about half an hour and the equivalent of several megabytes of main storage on an IBM 370/168.

Some concluding observations are now offered on the future of management support systems for distribution planning. First, the continuing importance of distribution planning is guaranteed by the sheer dollar magnitudes involved – some 20% of the GNP. Second, distribution is capturing an increasing share of upper management attention because of rapid wage and freight rate inflation, the high capital cost of carrying inventory, oil market uncertainties, and the momentous swing toward transportation deregulation. Management support systems are undeniably valuable in dealing with such concerns. Third, computer-based data sources are gaining ground in most companies, thereby facilitating the development of new management support systems. Fourth, recent technical advances promise continued evolution of systems like the one we have described. The successor to ODS, already largely in place, will have a modelling framework of substantially enlarged scope and extensive portable data preparation facilities to facilitate the always-tedious data development chores. Finally, the professional arts of building and using such systems are gradually being systematized (e.g.,<sup>(8)</sup>). Greater attention to this long-neglected province can be expected to improve the cost-effectiveness of tomorrow's management support systems for distribution planning.

#### APPENDIX I: MODELLING AND DATA ELEMENT OPTIONS

There is considerable discretion concerning how the various data elements and modelling options are used in particular applications. For instance, the application described in<sup>(2)</sup> omitted or did not use the full flexibility allowed by elements 3, 6, 8, 11, and 12.

It is assumed for illustrative purposes that all physical quantities are measured in hundredweight (CWT).

##### *Supply data*

1. A list of *products* (syn. product groups)

This is the final level of aggregation of a much larger number of individual stock-keeping units.

2. A list of *sources* (syn. plants, contract manufacturers, co-packers, etc.)

It is usually both permissible and desirable to combine sources of supply

that happen to be geographically proximate to one another into a single "source."

3. A table of *lower and upper limits (CWT/yr) on annual source shipments by product*

Upper limits (which can be infinity) permit the realities of currently installed plant capacity to be taken into account. Lower limits (which can be zero) permit at least a break-even level of operations to be maintained when the economics of distribution otherwise might dictate unacceptable shifts in the sources of supply. For validation purposes it may be useful temporarily to set both lower and upper limits equal to historical supply levels.

4. A table of *unit supply costs (\$/CWT) for each product at each source*  
Landed raw material costs should not be overlooked as a component of supply cost. For products procured from outside sources supply costs should be interpreted as procurement costs F.O.B. supplier.

If unit costs are not constant with respect to volume, marginal costs should be used that are valid for the plausible range of volumes to be examined by the model (cf. element 3). Actually, it is not the absolute magnitudes of these costs that matter, but rather their relative magnitudes; for each product, it is sufficient to estimate the *difference* between unit costs at the different plants or suppliers. It follows that supply costs can be ignored entirely for any product supplied from just one source.

In some applications marginal costs increase over the volume range of interest, owing to diseconomies of scale in manufacturing, premium overtime costs, buy-out rather than make, etc. A good broken-line approximation to these cost nonlinearities can usually be achieved by the simple expedient of replicating the source several times and using different unit costs and capacities for each. For instance, the first replicate could have a unit cost and capacity corresponding to regular time production, while the second replicate would have a higher unit cost and a capacity equal to what could be produced in overtime.

*Distribution centre data*

5. A list of *candidate distribution centres* (syn. warehouse, plant DC, depot, cross-dock) and their missions

Candidate distribution centre locations are specified as to city and mission (which products it can carry, whether public or private, degree of mechanization, etc.). Include promising new locations as well as current ones. More than one candidate DC can be given for a city if mission options are to be selected among. Particular candidate DCs can be "locked open" or "locked closed" if desired.

Some distribution centres may be full-line and others only partial-line. Designating partial-line facilities permits the system designer to concentrate slow-moving products in fewer locations, or to reflect the regional

character of some products, or to take account of the partial commonality of two or more partially autonomous distribution systems. See element 11 for further discussion.

6. For each distribution centre a *lower and upper limit (CWT/yr) on total annual DC throughput*, summed over all products, to be satisfied if it is open.

At the user's discretion:

- a. throughput limits may be stated in terms of weighted volumes, with a different weighting factor for each product;
- b. violations of throughput limits may be permitted for a stipulated penalty cost (\$ per CWT or weighted throughput), with a separate lower and upper violation penalty for each facility.

The need for throughput limits arises from two different kinds of considerations: practical and economic. Nearly all facilities have a limited practical throughput range imposed by available physical structures and systems for storage and handling. The limits tend to be wider for non-existing facilities than for existing ones, of course, owing to the greater latitude of choice or design when a new facility must be acquired. Economic considerations, on the other hand, have to do with the need to restrict the size of a facility to the range over which the associated cost estimates are valid (viz., elements 7 and 8). For instance, a public warehouse may not be willing to contract for less than a certain minimum annual volume. If economies-of-scale or diseconomies-of-scale are substantial enough to require different fixed-plus-variable cost figures for different ranges within the throughput spectrum, the ability to impose throughput limits allows a standard trick for accomplishing this: simply introduce several possible versions of a given facility, each with its own throughput range and cost figures. The same trick also permits various volume-dependent aspects of inbound and outbound freight costs to be modelled more accurately.

Lower throughput limits can be zero and upper limits can be infinite. Since they apply only if a facility is open, a positive lower limit does not force the facility to be open; it just guarantees that if the facility is open at all, it will at least be large enough to be viable.

The following comments pertain to the two discretionary features, 6a and 6b.

- a. In some applications the total annual throughput cannot be related accurately enough to facility size because a hundredweight of one product burdens capacity differently than a hundredweight of another product. The answer is usually to state throughput limits in terms of *weighted* annual volume, with a suitable weighting factor used for each product. In this way one can convert throughput to cube, cases, line items issued, or any other convenient measure.
- b. Experience has shown that rigid throughput limits generally yield

less realistic models than when limits are made "elastic" by means of under / overflow cost penalties. These penalties may be managerially meaningful (e.g., the cost of overtime, extra shift, or temporary warehousing), or they may be introduced in order to explore the sensitivity of the model to the given throughput limits.

7. For each distribution centre a *fixed cost* (\$/yr) to be incurred if it is open

The fixed cost should include annual possession costs, if any, as well as fixed portions of operating, inventory, and administration costs. One-time reconfiguration costs and savings would be added in at an equivalent annualized figure corresponding to a suitable term of amortization.

Some components of this data element and the next may come from linear approximations to nonlinear cost curves, and thus will not correspond to any traditional accounting definitions.

8. For each distribution centre a *variable cost* (\$/cwt) for each product that it is permitted to carry

Variable costs by product can be different for different distribution centres, but this need not be the case. Storage, handling, and inventory carrying are the most common cost categories.

#### *Customer data*

9. A list of *customer groups*

This is the final aggregation of a much larger number of individual customers, typically done on the basis of geography and possibly mode preference (e.g., Atlanta area TL vs. Atlanta area LTL customers). Between 100 and 300 groups usually suffice to represent adequately the entire country, with each being small enough that there will be no incentive for different customers in the group to require different distribution centre assignments.

10. A table of *demands* (cwt/yr) giving the amount of each product demanded by each customer group

It is assumed that all forecast demands will be met.

11. A partition of products into one or more *bundles*, with the understanding that if a customer group is assigned to a particular distribution centre for a certain product in bundle  $x$ , that customer group must be assigned to that distribution centre for all products in bundle  $x$ ; the partition can be the same for every customer group or it can be different for different customer groups

The bundling concept is designed to avoid split-sourcing customers to the extent that this makes physical and economic sense. Single-sourcing is convenient administratively and tends to reduce small delivery shipments. The reasons that products may have to be partitioned into more than one bundle include:

- (a) some distribution centres may not be full-line;
- (b) some products may not lend themselves to physically sharing out-bound transportation.

For a given partition each bundle contains one or more products, and every product must be in exactly one bundle. It is permissible for there to be a single bundle comprised of all products.

As an illustration, a firm may decide to carry "A" and "B" products in all distribution centres but the slow moving "C" products in just three regional facilities. Then there would be two bundles, one composed of AS and BS and the other composed of CS. A customer would receive all AS and BS from the primary assigned DC, but all CS from one of the regional facilities.

12. A table of *net selling prices* (\$/CWT) for each product in each customer group

Net selling price is gross selling price less discounts, allowances, and commissions. This information cannot influence the best system design but does facilitate analysis of product profitability by customer group.

A comment similar to the last paragraph under element 14 also applies here.

#### *Transportation data*

13. For each product a list of *permissible inbound transportation links*; and a *unit inbound transportation rate* (\$/CWT) for each link

Links would be omitted when a source cannot supply the product at all or a distribution centre is not permitted to carry it. Additional omissions could be made on the basis of policy or obvious economics.

The rates normally are weighted averages reflecting the mix of modes and shipment sizes deemed likely to prevail. Pipeline inventory costs can be included if desired.

14. For each product a list of *permissible outbound transportation links* satisfying the desired level of customer service (optionally, the customer service level of each link can be designated); and a *unit outbound transportation rate* (\$/CWT) for each link

Links would be omitted when a distribution centre is not permitted to carry the product or when the origin and destination are too far apart to permit an acceptable degree of customer service.

Designating some convenient measure of customer service level for each included link (such as its length in multiples of fifty miles or its average transit time in days) facilitates study of the trade-off between distribution cost and customer service.

Because each customer group is to be single-sourced for each product, the annual transportation quantity can be presumed known on each outbound link. This knowledge of what the annual transportation

quantity would be over each link (if used) enables the annual transportation cost associated with each link (if used) to be estimated with great precision. Of course, it is advisable to take into account the mix of modes and shipment sizes deemed likely to prevail. Pipeline inventory costs can also be included if desired.

Still greater accuracy can be achieved if desired by exploiting knowledge of how the demand for a given product at a given customer group breaks down by the various subproducts comprising that product. This permits the total freight cost over each outbound link to be built up at a greater level of detail than at first would appear possible with a given product aggregation.

15. For each product a list of *permissible direct or transit routes* and the associated *unit transportation rates* (\$/CWT); optionally, the customer service level of each route can be designated

Shipments under the storage-in-transit privilege cannot be modelled accurately using separate inbound and outbound rate files because the rate depends on the entire source  $\rightarrow$  distribution center  $\rightarrow$  customer *route*.

Another use of the transit feature involves perishables; only those routes would be included whose total journey times are acceptably short.

The comments under element 14 are also applicable here.

## APPENDIX 2: TECHNICAL DETAILS

This appendix gives a precise mathematical formulation of the problem, a technical overview of the decomposition technique, a technical explanation of the Customer Reassignment Penalties mentioned in section C, and comments on computational implementation.

### *Mathematical formulation*

Subscripts, variables, coefficients, and bundling structure are defined before giving a formal statement of the optimization problem.

#### *Subscripts*

$b$  indexes product bundles for a given customer group  
 $i$  indexes products  
 $j$  indexes sources  
 $k$  indexes candidate distribution centres  
 $l$  indexes customer groups

#### *Variables*

$v_k$  the amount of underflow (overflow) of  $v_{\min k}$  [ $v_{\max k}$ ], at a penalty rate of  $PL_k$  [ $PU_k$ ], for distribution centre  $k$ .  
 $x_{ijkl}$  the amount (CWT/yr) of product  $i$  flowing from plant  $j$  through distribution centre  $k$  to customer group  $l$ .

- $y_{klb}$  a binary variable indicating whether or not customer group  $l$  is assigned to distribution centre  $k$  for bundle  $b$
- $z_k$  a binary variable indicating whether or not distribution center  $k$  is selected for use

*Coefficients*

- $c_{ijkl}$  unit supply plus transportation cost (\$/CWT) associated with the flow  $x_{ijkl}$ ; equal to

$$c_{ij}^p + \begin{cases} c_{ijk}^{\text{in}} + c_{ikl}^{\text{out}} & \text{if transit does not apply} \\ c_{ijkl}^t & \text{if transit applies,} \end{cases}$$

where

- $c_{ij}^p$  unit supply cost
- $c_{ijk}^{\text{in}}$  unit inbound transportation and associated pipeline inventory costs
- $c_{ikl}^{\text{out}}$  unit outbound transportation and associated pipeline inventory costs
- $c_{ijkl}^t$  unit transit transportation costs
- $F_k$  fixed cost (\$/yr) associated with distribution centre  $k$
- $V_{ik}$  variable throughput cost (\$/CWT) for product  $i$  at distribution centre  $k$
- $D_{il}$  demand (CWT/yr) for product  $i$  at customer  $l$ ;  $D_{il} \geq 0$
- $\beta_i$  burden factor for product  $i$  to be used in calculating the weighted annual throughput of the distribution centres (e.g., the units of  $\beta_i$  would be  $ft^3/\text{CWT}$  if weighted throughput in cube is desired)
- $PL_k$  penalty rate (\$ per unit of violation) for weighted throughput less than  $v_{\text{MIN}_k}$  at distribution centre  $k$ ;  $PL_k \geq 0$
- $PU_k$  penalty rate (\$ per unit of violation) for weighted throughput greater than  $v_{\text{MAX}_k}$  at distribution centre  $k$ ;  $PU_k \geq 0$
- $SL_{ij}$  lower limit on the total shipment of product  $i$  by plant  $j$  (CWT/yr)
- $SU_{ij}$  upper limit on the total shipments of product  $i$  by plant  $j$  (CWT/yr);  $SL_{ij} \leq SU_{ij}$
- $v_{\text{MIN}_k}$  lower limit on the total annual weighted throughput of distribution centre  $k$  if it is open
- $v_{\text{MAX}_k}$  upper limit on the total annual weighted throughput of distribution centre  $k$  if it is open;  $v_{\text{MIN}_k} \leq v_{\text{MAX}_k}$

*Bundling structure*

- $B_l$  a set composed of the bundle indices defined for customer group  $l$  (often a single product partition type is used for all customer groups, in which case  $B_l = B$  for all  $l$ )

$I_l(b)$  a set composed of the product indices corresponding to bundle  $b$  for customer group  $l$ ; thus  $i \in I_l(b)$  means that product  $i$  is in bundle  $b$  for customer group  $l$ ; for any given  $l$ , the set

$$\bigcup_{b \in B_l} I_l(b)$$

consists of all product indices.

### Optimization problem

$$\text{Minimize } \sum_{v,x,y,z} \sum_i \sum_j \sum_k \sum_l c_{ijkl} x_{ijkl} + \sum_k \left[ F_k z_k + \sum_l \sum_{b \in B_l} \sum_{i \in I_l(b)} V_{ik} D_{il} y_{klb} \right] + \sum_k [\text{PL}_k \underline{v}_k + \text{PU}_k \bar{v}_k] \quad (1)$$

subject to

$$s_{L_{ij}} \leq \sum_k \sum_l x_{ijkl} \leq s_{U_{ij}}, \text{ all } ij \quad (2)$$

$$\sum_j x_{ijkl} = D_{il} y_{klb}, \quad \text{all } klb \text{ such that } b \in B_l \text{ and } i \in I_l(b) \quad (3)$$

$$\sum_k y_{klb} = 1, \quad \text{all } lb \text{ such that } b \in B_l \quad (4)$$

$$y_{klb} \leq z_k, \quad \text{all } klb \text{ such that } b \in B_l \quad (5)$$

$$v_{\text{MIN}_k} z_k - \underline{v}_k \leq \sum_l \sum_{b \in B_l} \sum_{i \in I_l(b)} \beta_l D_{il} y_{klb} \leq v_{\text{MAX}_k} z_k + \bar{v}_k, \text{ all } k \quad (6)$$

$$\left. \begin{array}{ll} \underline{v}_k \text{ and } \bar{v}_k \geq 0, & \text{all } k \\ y_{klb} = 0 \text{ or } 1, & \text{all } klb \text{ such that } b \in B_l \\ z_k = 0 \text{ or } 1, & \text{all } k \end{array} \right\} \quad (7)$$

$$x_{ijkl} \geq 0, \quad \text{all } ijkl \quad (8)$$

It is understood that all indicial summations and constraint enumerations run only over legitimate combinations of subscripts.

The objective function (1) is made up of three parts: a supply-transportation-pipeline inventory cost part, a distribution centre cost part, and a throughput violation penalty part. A net sales revenue term is not included because it would just be a constant in view of the assumption that all demands must be met exactly.

Constraints (2) enforce the limits  $s_L$  and  $s_U$ .

Constraints (3) enforce the correct relationship between the  $x$ s and  $y$ s: when  $y_{klb} = 1$ , they ensure that the proper amount of each product in bundle  $b$  flows to customer group  $l$  from distribution centre  $k$ ; and when  $y_{klb} = 0$ , they ensure that the corresponding  $x_{ijkl}$  flows are zero for all  $i$  in bundle  $b$ .

Constraints (4) specify that each customer group must be assigned to a single distribution centre for each bundle of products.

Constraints (5) ensure that customer group assignments can be made only to open distribution centres.

Constraints (6) ensure that the total weighted throughput of the  $k^{\text{th}}$  distribution centre is between the desired limits (except for penalized violations) when it is open. These constraints are redundant, in view of (5), when  $z_k = 0$ . By making PL and PU sufficiently high, of course, VMIN and VMAX can be enforced rigidly, if desired.

Optimization problem (1)–(8) is a very large mixed integer linear program.

*The decomposition technique*

If the  $v$ ,  $\bar{v}$ ,  $y$ , and  $z$ -variables are all held fixed at values satisfying (4), (5), (6) and (7), the problem of finding the best values for the  $x$ -variables separates into an independent transportation-type subproblem for each product. The  $i^{\text{th}}$  subproblem can be written as follows, if we introduce the notation  $\bar{k}(l)$  for the index of the distribution centre to which customer group  $l$  is assigned for product  $i$  by the temporarily fixed  $y$ -array (see <sup>(4)</sup>):

$$\text{Minimize } \sum_{jl} c_{ij\bar{k}(l)l} x_{ij\bar{k}(l)l} \quad (9)$$

subject to

$$s_{L-ij} \leq \sum_l x_{ij\bar{k}(l)l} \leq s_{U-ij}, \quad \text{all } j: ij \text{ exists} \quad (10)$$

$$\sum_j x_{ij\bar{k}(l)l} = D_{il}, \quad \text{all } l, \quad (11)$$

$$x_{ij\bar{k}(l)l} \geq 0, \quad \text{all } j, l: ij \text{ exists.} \quad (12)$$

Problem (9)–(12) is the generic Transportation Subproblem. It is classical in form except for the lower limits on supply – a feature that, with the help of a standard trick, can be converted back to the classical transportation problem structure.

A technical device is used to make sure that the subproblems are always feasible: in effect, a dummy plant is introduced with a very high cost rate link to every customer group.

Suppose that all the Transportation Subproblems are solved with a given trial configuration. The consequent Cost Estimating Relationship referred to earlier is a linear function of  $y$ ,

$$f(y) = \alpha + \sum_k \sum_l \sum_{b \in B_l} a_{klb} y_{klb}, \quad (13)$$

with this property:

$f(y)$  is less than or equal to the sum of the optimal values of the

Transportation Subproblems for all  $y$ , with equality holding for the  $y$  at which it was generated.

A method is given in <sup>(4)</sup> for deriving closed form expressions for the coefficients of (13) in terms of the  $c_{ijkl}$ s,  $D_{il}$ s,  $s_{Lij}$ s,  $s_{Uij}$ s, and the dual variables associated with (10).

The standard form of the Master Problem after the Transportation Subproblems have been solved  $H$  times is:

$$\text{Minimize (1) with } \sum_{v,y,z} \sum_i \sum_j \sum_k \sum_l c_{ijkl} x_{ijkl} \text{ replaced by Max } \{f^1(y), \dots, f^H(y)\} \\ \text{subject to (4), (5), (6), (7).}$$

The motivation for the replacement term in the objective function is that the pointwise maximum of all available Cost Estimating Relationships is the tightest known lower estimate of the optimal supply plus transportation costs as a function of  $y$ . It follows that the optimal value of the Master Problem is a valid lower bound on the optimal value of (1)–(8).

In practice we use a variant of the standard form of the Master Problem, as explained in Section 2.3 of <sup>(4)</sup>. That article also contains a complete discussion of the decomposition procedure for a special case of (1)–(8) from which the general case can be obtained in a straightforward manner.

#### *Customer reassignment penalties*

From the defining property of the Cost Estimating Relationship (13) corresponding to the least cost configuration  $y^*$ , we obtain that

$$f(y) + \text{variable cost term of (1) at } y \\ - (f(y^*) + \text{variable cost term of (1) at } y^*) \quad (14)$$

must be less than or equal to the total cost penalty incurred (exclusive of fixed and throughput penalty costs) by changing customer group assignments from  $y^*$  to  $y$  and reoptimizing flows. Expression (14) can be rewritten as

$$\sum_k \sum_l \sum_{b \in B_l} \left( a_{klb} + \sum_{i \in I_l(b)} V_{ik} D_{il} \right) (y_{klb} - y_{klb}^*). \quad (14a)$$

The special case of greatest interest, and one which we tabulate routinely, is the one where some customer group  $l$  is switched from its optimally assigned distribution centre  $k^*$  to some other distribution centre  $k$  for product bundle  $b$ . Then (14a) specializes to

$$\left( a_{klb} + \sum_{i \in I_l(b)} V_{ik} D_{il} \right) - \left( a_{k^*lb} + \sum_{i \in I_l(b)} V_{ik^*} D_{il} \right) \quad (15_{klb})$$

This quantity is only a lower bound on the cost change induced by the reassignment, but in practice is often a remarkably accurate one. The penalties (15) are additive in case several reassignments are to be made simultaneously. Moreover, it is a simple matter to account for any changes in facility fixed costs or throughput penalty costs. Experience indicates that it is better to do such adjustments manually than to extend (14) and (15) to include these two cost categories.

#### *Computational implementation*

The two most important performance measures are primary computer storage and CPU run time requirements. Primary storage requirements are so minimized by fully dynamic reallocation techniques and by exploiting the inherent advantage of the decomposition approach that it does not require the data for more than one product to be in main storage at once. This means that a problem with fifty products will run in essentially the same amount of main storage as a one-product problem.

Three key factors control CPU times:

- 1 the number of major iterations required, that is, how many times the Master Problem must be solved;
- 2 how efficiently the Transportation Subproblems can be solved;
- 3 how efficiently the Master Problems can be solved.

Our experience concerning (1) is that just a few major iterations suffice for most problems. In fact, it can be shown that just one major iteration is required for any problem where the supply constraints are not binding (more precisely, just one Cost Estimating Relationship – no matter what the parent system configuration – enables the Master Problem to produce an optimal system configuration). Factor (2) offers no difficulty if a modern network code is used. ODS uses a version of the GNET code described in <sup>(1)</sup>. The primary challenge with regard to factor (3) has been to achieve a large-scale linear programming technology able to solve efficiently the LP relaxations of the Master Problem. The number of type (5) constraints often runs well into the thousands, but dropping them leads to an inordinate weakening of the LP relaxation. This has necessitated innovations in data structure design, “elastic” programming, and degeneracy control.

#### ACKNOWLEDGMENT

The authors gratefully acknowledge the valuable suggestions of Dr Richard F. Powers and Dr Tony J. Van Roy.

## REFERENCES

- (1) G. Bradley, G. Brown and G. Graves, "Design and implementation of large scale primal transshipment algorithms," *Management Science*, September 1977.
- (2) A. Geoffrion, "Better distribution planning with computer models," *Harvard Business Review*, July-August 1976.
- (3) A. Geoffrion, "Making better use of optimization capability in distribution system planning," *AIIE Transactions*, June 1979.
- (4) A. Geoffrion and G. Graves, "Multicommodity distribution system design by Benders decomposition," *Management Science*, January 1974.
- (5) A. Geoffrion and R. Powers, "Facility location analysis is just the beginning," *Interfaces*, April 1980.
- (6) A. Geoffrion and R. Powers, "Management support systems," *Wharton Magazine*, Spring 1981.
- (7) J. Hall, ed., "Management science update column," *Management Science*, 1980.

Copyright of INFOR is the property of INFOR Journal: Information Systems & Operational Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.