ANALYTICAL EVALUATION OF HIERARCHICAL PLANNING SYSTEMS

M.A.H. Dempster
Balliol College, Oxford, England

M.L. Fisher
University of Pennsylvania, Philadelphia, PA, USA

L. Jansen, B.J. Lageweg, J.K. Lenstra
Mathematisch Centrum, Amsterdam, The Netherlands

A.H.G. Rinnooy Kan
Erasmus University, Rotterdam, The Netherlands

PREFACE

Previous research in the System and Decision Sciences Area at IIASA has involved study of both hierarchical systems and optimization techniques. More recently, SDS research in optimization has concentrated on stochastic systems. This paper — the first in a series by an international group of researchers — fuses these strands of research.

The central observation of the authors is that practical hierarchical planning has a natural temporal structure which typically involves a sequence of decisions at an increasing level of detail and with increasingly accurate information. Such problems may be modeled by multistage stochastic programs which, due to their complexity, and in parallel with real-life procedures, require the use of approximations and heuristics for near-optimal solution.

In this paper the general modeling approach and a simple example are discussed. The hope is held out that these models will eventually serve as an analytical yardstick against which the increasingly popular computer-based hierarchical planning systems may be measured.

All of the authors are active in the development of computer software for planning in various environments, so that in a very real sense this is theoretical research stemming from practice.

M.A.H. DEMPSTER

Analytical Evaluation of Hierarchical Planning Systems

M. A. H. DEMPSTER

Balliol College, Oxford, England

M. L. FISHER

University of Pennsylvania, Philadelphia, Pennsylvania

L. JANSEN, B. J. LAGEWEG, J. K. LENSTRA

Mathematisch Centrum, Amsterdam, The Netherlands

A. H. G. RINNOOY KAN

Erasmus University, Rotterdam, The Netherlands

Hierarchical planning systems have become popular for multilevel decision problems. After reviewing the concept of hierarchical planning and citing some examples, we describe a method for analytic evaluation of a hierarchical planning system. We show that multilevel decision problems can be nicely modeled as multistage stochastic programs. Then any hierarchical planning system can be measured against the yardstick of optimality in this stochastic program. We demonstrate this approach on a hierarchical system that can be shown to be asymptotically optimal for a job shop design/scheduling problem.

MANY OPERATIONS management planning and control problems require a series of decisions over time at an increasing level of detail. For example, there are at least two distinct decision making levels in most production operations. At the lowest level, detailed production scheduling decisions determine who will do a particular job on what machine and when. Considerations at this level include minimizing setups and meeting due dates. At a higher level, aggregate planning decisions are made concerning hiring and layoffs, overtime, production levels for product groups, ordering of raw materials, and setting due dates. The time horizon for aggregate decisions ranges from several months to one year. At the time aggregate decisions are made, much detailed information is not known with certainty. This may include future product demand, job processing times, machine breakdowns, worker availability,

and raw material availability. In addition, other details are deliberately ignored at the aggregate level. For example, the sequence dependent nature of setups is usually ignored, and product groups are used rather than individual stock-keeping units.

Hierarchical planning systems are becoming increasingly popular for multilevel decision problems. General discussions of hierarchical planning can be found in Bitran and Hax (1977), and Bradley et al. ([1977] pp. 212–213). A hierarchical system uses separate mathematical program-ming models to make the decisions at each level. The solution of a higher level model creates some of the constraints for the model below it. Of course, there is nothing new about using interacting optimization models. What seems to be new in the hierarchical approach is the explicit emphasis on the linkages between the models and on designing all models in the system simultaneously so that they fit well together.

There are two fundamental reasons for using a hierarchical approach.

1. *Reducing complexity.* Breaking a problem into subproblems is a standard method for simplifying the solution process. A tenet of hierar-chical planning is that this partitioning can be done so that the interaction effects between subproblems are acceptably weak.

2. *Coping with uncertainty.* It is important to realize that the decisions at the various levels in the planning process need to be made at different points in time. For example, aggregate planning decisions are made early enough to implement plans for hiring/layoff, raw materials acquisition, etc. On the other hand, a decision to assign a particular job to a specific machine can be postponed until the instant before the job begins proc-essing. This is important in light of the fact that much data at the detailed level is uncertain at the time aggregate decisions are made. If detailed and aggregate decisions were combined in a single giant optimization model, as is sometimes proposed, the detailed decisions would be made earlier than necessary and hence would be based on less reliable forecasts of the uncertain data. The hierarchical approach postpones the detailed decisions as long as possible so that they can be based on more timely and hence more accurate data.

A third advantage often attributed to hierarchical systems is that they parallel the hierarchical organization of most firms. While this is certainly an important consideration, we believe that hierarchical planning orga-nizations, as well as hierarchical planning systems, are a response to the nature of the problems being solved, and to the need to reduce complexity and respond to uncertainty cited above.

Past work in hierarchical planning has consisted mainly of building clever systems. This paper is an expository work that reviews many of the hierarchical systems that have been developed and then considers the question of how one decides whether one system is better than

another. To obtain an objective standard of performance, we introduce a stochastic programming model that encompasses the entire multilevel decision process. At each level we model lower levels accurately, but with stochastic parameters. The objective at each level is to minimize current costs plus the expected objective value of an optimal lower level solution. The performance of a hierarchical planning system can be studied analytically by determining how close it comes to optimality in this stochastic programming model.

In Section 1 we provide a rich set of examples of hierarchical planning systems. Section 2 describes the analytic evaluation approach in detail. In Section 3 we illustrate this approach by analyzing a hierarchical system for a simplified prototype of one of the examples given in Section 1. This hierarchical system is shown to have the desirable property of being asymptotically optimal as the number of jobs in the problem approaches infinity. This result provides insight into why and when the hierarchical approach works well.

1. EXAMPLES OF HIERARCHICAL PLANNING SYSTEMS

Examples are provided of hierarchical systems for four types of problems. Specific implementations of each type are also referenced.

1.1. Aggregate/Detailed Production Scheduling

This example has already been discussed at some length in the introduction. The aggregate planning problem is usually modeled as a multiperiod linear program with a planning horizon of about one year. The basic decision variables are target production and inventory levels by period for aggregated groups of products. These target levels are fed to the detailed production scheduling model which is concerned with scheduling the actual production of each stock-keeping unit over a relatively short horizon. This detailed model is usually an integer program solved by a heuristic. It may be run much more often than the aggregate model.

Hax and Meal (1975), and Hax and Golovin (1978a, b) have designed and successfully implemented a system of this type. They distinguish three increasingly aggregated product units: items, families and types. At the type level, they use a linear program to set long range target production levels that minimize production, overtime, and inventory carrying costs. At the family level, the production target for a type is allocated to the families within that type considering setup costs. At the item level, item production is planned over a short planning horizon using recent demand data.

Jaikumar (1974) has developed a system for the Booth Fisheries Division of Consolidated Foods that has many novel features and has

achieved impressive economic results. At the aggregate level, a linear program is used to determine levels of production, marketing promotions, and raw material acquisitions that maximize sales revenue minus the cost of inventory holding, overtime, hiring/layoff, production and raw materials. This model has a 26-week planning horizon with time periods varying from one to six weeks. The detailed model is an integer program with a planning horizon of one week and time periods of two hours. This model allocates the production of particular products to individual production lines.

A novel feature of Jaikumar's system is the use of a dual method to link the aggregate and detailed models. Dual variables on the production capacity, raw material and manpower constraints in the aggregate model are fed to the detailed model and used in the objective function to cost out usage of those resources.

1.2. Job Shop Design/Scheduling

This problem is concerned with specifying the number of machines of various types to have in a job shop and the scheduling of work on those machines to minimize the total of machine costs and some job based measure of performance such as average job tardiness. The first algorithmic approach to this problem was given by Fisher (1969).

In a hierarchical approach, the higher level decision is obviously how many machines to have, and the lower level decision is how to sequence jobs in the shop. Armstrong and Hax, and Shwimer (1972) have described hierarchical systems that use an integer programming model and a simulation with an embedded heuristic sequencing rule to make the higher level decision. The lower level decision is made with the heuristic sequencing rule.

1.3. Distribution System Design/Control

The higher level problem includes the distribution design questions of where to locate plants and warehouses, whether to expand capacity at existing plants and warehouses, whether to install automatic materials handling equipment, etc. The lower level problem concerns questions such as the allocation of customers to warehouses and the determination of commodity flows through the system.

Hax (1977) describes a system that he developed in the aluminum industry. The higher level problem is a linear program that is run a number of times for different cases. In addition to the standard distribution system design questions, this model is concerned with whether or not the firm should enter into long term product swapping contracts with

competitors. The lower level problem is concerned with which sources should supply a set of orders in hand. It is modeled as a linear program with a heuristic post-adjustment of the solution to eliminate order splitting.

Glover et al. (1979) describe a comprehensive production and distribution planning model developed for the Agrico Chemical Company. This is a network flow model explicitly concerned with the lower level decision of scheduling the flow of shipments from plants through the distribution system to warehouses. The model has also been successfully used in a case-study mode to analyze higher level decisions on capacity expansion and equipment modernization at the plants and warehouses.

Federgruen and Lageweg (1980) describe a distribution system developed for a producer of industrial gases in the Netherlands. At the higher level, the number and locations of national distribution centers and the associated territories are determined. At the lower level, the commodity flows are routed from distribution centers through regional depots to final customers. The system is modeled as a network flow problem and uses a vehicle routing algorithm to calculate routing costs.

1.4. Vehicle Routing/Scheduling

Many organizations operate vehicle fleets to deliver their products to customers. Frequently, for administrative convenience and other reasons, the fleet is scheduled using a fixed route system. In a fixed route system, particular customers are assigned to a vehicle and this assignment is revised infrequently (e.g., every 6–12 months). Daily customer orders are random. A customer orders on a given day with some probability, and the amount of any order is random. On any given day, a delivery schedule must be developed for the customers that have ordered for each vehicle. In this daily problem one can consider the possibility of delaying delivery of a customer order, or of using an alternative mode of delivery such as airfreight.

This problem suggests an obvious hierarchical system in which one model is used to assign customers to vehicles to form the fixed routes, and another model is used to solve the daily scheduling problem ("cluster first, route second"). Fisher and Jaikumar (1981) have developed a partitioning algorithm for vehicle routing which follows the hierarchical structure just outlined. Customers are assigned to vehicles using a generalized assignment model. Each vehicle is scheduled using a traveling salesman model with side constraints. This algorithm has been implemented at a large chemical company and produced significant economic benefits. Golden and Yee (1979) have also suggested a general framework for this problem.

2. A STOCHASTIC PROGRAMMING FRAMEWORK FOR EVALUATION OF HIERARCHICAL PLANNING SYSTEMS

Past work on hierarchical systems has concentrated on building clever systems, many of which were chronicled in the previous section. This paper is concerned with the question of how one evaluates the quality of the decisions produced by a hierarchical system. One is interested both in comparing different systems and in direct evaluation of a single system.

Some methods already exist for partially answering these questions. Different systems can be compared empirically by Monte Carlo simulation. In this approach, higher level models are run with forecasts of the uncertain lower level data. Lower level models are run with actual data values generated randomly by the Monte Carlo method. Hax and Golovin (1978b) have used this approach in evaluating their system for different settings of various parameters. One can also evaluate either by analytic or empirical methods the degree of optimality in the solutions to the submodel at each level.

All of these evaluation methods fail to answer the nagging question of how good a particular hierarchical system is when compared with an optimal system. To answer this we need a measure of optimality for the overall system, not just for each subproblem. Put differently, hierarchical systems are often called suboptimizing systems. If so, what is the optimization problem being suboptimized?

A little thought should make it clear that the answer to this question cannot be a deterministic mathematical programming model if we wish to accurately capture the uncertainty that exists at lower levels of the global decision problem. However, this phenomenon can be nicely modeled by a multistage stochastic program. The stochastic program would model lower levels accurately, but with stochastic parameters. The objective at each level is to minimize known costs at that level plus the expected objective value of an optimal lower level solution. For example, consider the form such a model would take for the example given in Section 1.2. At the time the job shop is designed, only probabilistic information is available on the jobs to be processed by the shop. A two-stage stochastic programming model of this problem would select the machine configuration of the job shop so as to minimize the cost of machines plus the expected cost (e.g., job tardiness or flow time) of operating the shop optimally with the given machine configuration. Similar stochastic programs can be created for all of the other examples given in Section 1. We note that all of these models would have integer variables at one or more of the levels.

We are not suggesting that any of these models should be solved optimally. Such an endeavor would be foolhardy given the reputation for

intractability enjoyed by both integer programming and stochastic programming. Rather, we suggest that a hierarchical system for a multilevel decision problem can be usefully viewed as a heuristic applied to a multistage stochastic programming model of the problem. Worst-case and probabilistic analysis techniques are now widely used in the study of integer programming heuristics (e.g., see Fisher [1980]). There is no reason why these same techniques could not be applied to stochastic programming heuristics. This would allow objective statements about how closely a particular hierarchical system approaches the ideal of optimality in the appropriate stochastic programming model.

Although the observations of this section are apparently straightforward, they have been the well-spring of some extensive analytic work on our part that we believe provides some interesting insights on the performance of hierarchical systems. In the next section we prove a result for a simplified prototype of the job shop design/scheduling example given in Section 1.2 and survey other results from a companion paper (Dempster et al. [1981]).

3. ANALYSIS OF A HIERARCHICAL SYSTEM FOR A JOB SHOP DESIGN/SCHEDULING PROBLEM

We are given n jobs to be processed on identical parallel machines. The problem is to decide how many machines to buy and how to sequence the jobs on the machines to minimize machine cost plus the maximum job completion time. Job processing times are independently distributed random variables whose values all become known with certainty at time $t = 0$ *after* a decision has been made on the number of machines to buy.

This problem is a simplified prototype of the job-shop design problem described in Section 1.2. In this section we will propose and analyze a hierarchical heuristic for this problem.

Define

$$c = \text{cost of a single machine,}$$
$$m = \text{number of machines to be bought,}$$
$$p_j = \text{processing time of job } j,$$
$$p = (p_1, \cdots, p_n),$$
$$C^*(m, p) = \text{earliest time at which all jobs are completed when scheduled optimally on } m \text{ machines, beginning at } t = 0, \text{ and with known processing times } p = (p_1, \cdots, p_n).$$

A tilde ($\sim$) under a variable will indicate that it is a random variable, and E will denote expected value.

The problem is to choose m prior to $t = 0$ to solve

$$Z^* = \min_m \{cm + EC^*(m, \underset{\sim}{p})\}. \tag{1}$$

Let m^* denote the optimal solution to (1). Then at time $t = 0$, sequence n jobs with known processing times on m^* machines to achieve the optimal completion time $C^*(m^*, p)$.

There is a natural two-stage hierarchical system for this problem in which the higher level problem is problem (1) and the lower level program is sequencing the jobs once processing times are known. It is not obvious how to solve (1) optimally. Determining $C^*(m, p)$ for fixed m and p is itself an NP-hard problem so finding $EC^*(m, p)$ as a function of m seems virtually impossible. We circumvent this difficulty with an idea that is fundamental to all of the hierarchical systems described in Section 1. In solving the higher level problem we suppress the combinatorial fine structure of the lower level problem by replacing $C^*(m, p)$ with P/m, where $P = \sum_{j=1}^{n} p_j$. The quantity P/m is a lower bound on $C^*(m, p)$ for any m and p. This bound is quite good if n is large and $p_{max} = \max_j\{p_j\}$ is sufficiently small. This replacement leads to the higher level problem.

$$\min_m\{cm + E\underset{\sim}{P}/m\}.$$

The derivative of the objective function is zeroed by $m = \sqrt{E\underset{\sim}{P}/c}$. Since m must be a positive integer, we choose $m^H \in \{\lceil\sqrt{(E\underset{\sim}{P}/c)}\rceil, \lfloor\sqrt{(E\underset{\sim}{P}/c)}\rfloor\}$ to minimize $cm + E\underset{\sim}{P}/m$ subject to $m^H \geq 1$. Here $\lceil a\rceil$ denotes the smallest integer not less than a and $\lfloor a\rfloor$ denotes the largest integer not greater than a.

The lower level problem is the problem of sequencing n jobs with known processing times on m^H machines. In our hierarchical system, we simply solve this problem by list scheduling (*LS*). This heuristic assigns jobs in arbitrary order, placing each job on the machine that has the least processing already assigned. Let

$C^{LS}(m, p) =$ earliest time at which all jobs are completed when scheduled by list scheduling on m machines, beginning at $t = 0$, and with known processing times $p = (p_1, \cdots, p_n)$.

The overall value achieved by our hierarchical planning system is

$$Z^H = cm^H + EC^{LS}(m^H, \underset{\sim}{p}).$$

This value is compared with the optimal value in the following theorem.

THEOREM 1. $Z^H/Z^* \leq 1 + E\underset{\sim}{p}_{max}/2\sqrt{(cE\underset{\sim}{P})}$.

Proof. Consider an arbitrary schedule produced by the list scheduling rule, and let $\underset{\sim}{T}$ denote the latest time that all machines are occupied. Let k denote the index of a job completed last. For any instance of $\underset{\sim}{p}$ the following inequalities hold by the nature of list scheduling.

$$C^{LS}(m^H, \underset{\sim}{p}) \leq \underset{\sim}{T} + \underset{\sim}{p}_k \leq \underset{\sim}{P}/m^H + \underset{\sim}{p}_{max}.$$

Hence, for any instance of $\underset{\sim}{p}$, we have

$$cm^H + C^{LS}(m^H, \underset{\sim}{p}) \leq cm^H + \underset{\sim}{P}/m^H + \underset{\sim}{p}_{\max}. \tag{2}$$

Taking expectations gives

$$Z^H \leq cm^H + E\underset{\sim}{P}/m^H + E\underset{\sim}{p}_{\max} \tag{3}$$

$$\leq cm^* + E\underset{\sim}{P}/m^* + E\underset{\sim}{p}_{\max} \tag{4}$$

$$\leq Z^* + E\underset{\sim}{p}_{\max} \tag{5}$$

where (3) follows from (2), (4) from (3) and the fact that m^H minimizes $cm + E\underset{\sim}{P}/m$, and (5) from (4) and the fact that $C^*(m^*, \underset{\sim}{p}) \geq \underset{\sim}{P}/m^*$ for any instance of $\underset{\sim}{p}$.

Finally, we observe that

$$Z^* \geq \min_m \{cm + E\underset{\sim}{P}/m\} \geq 2\sqrt{cE\underset{\sim}{P}}. \tag{6}$$

Combining (5) and (6) establishes the theorem.

In Dempster et al. we show that, if the p_j have independent identical distributions with finite second moment, then the hierarchical planning system is asymptotically optimal in the sense that $\lim_{n\to\infty}(E\underset{\sim}{p}_{\max}/\sqrt{E\underset{\sim}{P}}) = 0$, and hence $\lim_{n\to\infty}(Z^H/Z^*) = 1$.

Under the same assumption we have for any $\epsilon > 0$

$$\lim_{n\to\infty} \text{prob.}\{(cm^H + C^{LS}(m^H, \underset{\sim}{p}))/(cm^* + C^*(m^*, \underset{\sim}{p})) \leq 1 + \epsilon\} = 1.$$

This is a strong result that says that with high probability the hierarchical planning system finds a near optimal solution for *any* instance of $\underset{\sim}{p}$.

These results are extended in Dempster et al. to the case where the jobs are to be processed on uniform parallel machines (i.e., with different speeds) and the problem is to minimize machine cost plus the maximum job lateness with respect to a common constant due date. Under appropriate assumptions, it also appears possible to extend this approach to the case where n is random.

What are the implications of these results? Although the model we have considered here is a simplified prototype of a real problem, we think these results provide insight into why and when the hierarchical approach is reasonable for more complicated problems. The hierarchical system given here makes the same assumptions at the higher level that one sees in more complicated systems. Namely, all jobs are replaced by the *aggregate* processing requirement and complicating details are omitted. Moreover, this hierarchical system has the theoretical properties we have discussed. More complicated systems work well because the instances for which the higher level assumptions are severely violated occur with decreasingly small probability as the problem grows larger.

ACKNOWLEDGMENTS

REFERENCES

ARMSTRONG, R. J., AND A. C. HAX. 1977. A Hierarchical Approach for a Naval-Tender Job-Shop Design. In *Applied Mathematical Programming*, Chap. 10, S. P. Bradley, A. C. Hax and T. L. Magnanti. Addison-Wesley, Reading, Mass.

BITRAN, G. R., AND A. C. HAX. 1977. On the Design of Hierarchical Production Planning Systems. *Decision Sci.* **8**, 28–55.

BRADLEY, S. P., A. C. HAX, AND T. L. MAGNANTI. 1977. *Applied Mathematical Programming*. Addison-Wesley, Reading, Mass.

DEMPSTER, M. A. H., M. L. FISHER, L. JANSEN, B. J. LAGEWEG, J. K. LENSTRA, AND A. H. G. RINNOOY KAN. 1981. Analysis of Heuristics for Stochastic Programming: Results for Hierarchical Scheduling Problems, Report, Mathematisch Centrum, Amsterdam.

FEDERGRUEN, A., AND B. J. LAGEWEG. 1980. Hierarchical Distribution Modelling with Routing Costs, Report BW 117, Mathematisch Centrum, Amsterdam.

FISHER, M. L. 1969. A Branch and Bound Algorithm for Optimizing with a Simulation Model, M.S. thesis, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass.

FISHER, M. L. 1980. Worst-case Analysis of Heuristic Algorithms. *Mgmt. Sci.* **26**, 1–17.

FISHER, M. L., AND R. JAIKUMAR. 1981. A Generalized Assignment Heuristic for Vehicle Routing. *Networks* **11**, 109–124.

GLOVER, F., G. JONES, D. KARNEY, D. KLINGMAN, AND J. MOTE. 1979. An Integrated Production, Distribution and Inventory Planning System, Research Report CCS 338, University of Texas, Austin.

GOLDEN, B., AND J. YEE. 1979. A Framework for Probabilistic Vehicle Routing. *AIIE Trans.* **11**, 109–112.

HAX, A. C. 1977. Integration of Strategic and Tactical Planning in the Aluminum Industry. In *Applied Mathematical Programming*, Chap. 6, S. P. Bradley, A. C. Hax and T. L. Magnanti. Addison-Wesley, Reading, Mass.

HAX, A. C., AND J. J. GOLOVIN. 1978. Hierarchical Production Planning Systems. In *Studies in Operations Management*, A. C. Hax (ed.). North Holland-American Elsevier, New York.

HAX, A. C., AND J. J. GOLOVIN. 1978. Computer Based Operations Management System (COMS). In *Studies in Operations Management*, A. C. Hax (ed.). North Holland-American Elsevier, New York.

HAX, A. C., AND H. C. MEAL. 1975. Hierarchical Integration of Production Planning and Scheduling. In *Studies in Management Sciences; Vol. 1. Logistics*, M. A. Geisler (ed.). North Holland-American Elsevier, New York.

JAIKUMAR, R. 1974. An Operational Optimization Procedure for Production Scheduling. *Comput. Opns. Res.* **1**, 191–200.

SHWIMER, J. 1972. Interactions between Aggregate and Detailed Scheduling in a Job Shop, Massachusetts Institute of Technology, Sloan School of Management, unpublished Ph.D. thesis.