

HIERARCHICAL INTEGRATION OF PRODUCTION PLANNING
AND SCHEDULING

by

Arnoldo C. Hax and Harlan C. Meal

May 1973

656-73

**HIERARCHICAL INTEGRATION OF PRODUCTION PLANNING
AND SCHEDULING**

Arnoldo C. Hax and Harlan C. Meal

ABSTRACT

This paper describes the development of a hierarchical planning and scheduling system for a multiple plant, multiple product, seasonal demand situation. In this hierarchical structure, optimal decisions at an aggregate level (planning) provide constraints for detailed decision making (scheduling).

Four levels of decision making are used: *first*, products are assigned to plants (using mixed-integer programming), making long-term capacity provision and utilization decisions; *second*, a seasonal stock accumulation plan is prepared (using linear programming), making allocations of capacity in each plant among product Types within which the products have similar inventory costs; *third*, detailed schedules are prepared for each product Family (using standard inventory control methods for items grouped for production since the Items in a Family share a major setup), allocating the type capacity among the product Families in the Type; *fourth*, individual run quantities are calculated for each Item in each Family, again using standard inventory methods.

The approximations used and the procedures developed are described in sufficient detail to guide a similar application. We also discuss the problems encountered in implementation and the approach used to resolve these problems. Finally, we estimate the costs and benefits of this system application.

Hierarchical Integration of Production Planning and Scheduling

Arnoldo C. Hax* and Harlan C. Meal**

I. INTRODUCTION

The objective of the present paper is to provide a framework for a hierarchical decision making approach in which the aggregate results of capacity planning provide constraints for detailed scheduling decisions. In order to illustrate, with specific examples, the design and implementation issues in such a hierarchical system, we describe the actual application of this approach to the development of a production planning and scheduling system for a firm producing many products in several plants.

We will first describe the characteristics of the production and distribution problems presented in the firm where the system design was carried out. Then, we will comment on the general nature of hierarchical planning and scheduling systems, and justify the specific approach followed. Subsequently, we will discuss the details of each of the components of the overall system. Finally, we will describe the implementation efforts, the difficulties encountered, how these were overcome and what costs and benefits can be derived from the system operation.

II. SITUATION

In order to protect the anonymity of this manufacturing firm we will describe it as a process manufacturing situation analogous in some ways to a chemical plant or a steel mill. There are some minor assembly operations; but for our purposes, these can be treated as though they were mixing operations in a batch process chemical plant.

1. Multiple Plants. There are four plants, geographically separated so as to service reasonably separated market areas in the north, south, east and west. The combination of manufacturing and transportation costs indicates that some items should be made in only one plant, some in two, and so on. The assignment of product to plants is an important problem faced by the firm.

2. Seasonal Demand. The customer demand for this set of products is seasonal, with three distinct seasonal demand patterns. Some products can be characterized as winter season and some as summer season, with significant differences in the size of these two markets. A third seasonal pattern is neither winter nor summer, but shows distinct variations through the year.

* Massachusetts Institute of Technology, Sloan School of Management

** Arthur D. Little, Inc.

3. Items Grouped for Production. There are natural groupings of products in manufacturing. A major setup cost is incurred whenever a particular set of products is run, while minor setup costs are required in switching production from one of these products to another. Moreover, production of any of the group in a plant requires a fixed capital investment (tooling cost) which then permits production of the remainder of the group without additional tooling cost.

4. Level Production. To a greater extent than is common in most manufacturing operations, there is a strong incentive to maintain a nearly level manufacturing rate. This arises from two primary considerations: a) The capital cost of equipment is very high compared with the cost of shift premiums for labor, and the system normally operates three shifts on a five-day week. Occasionally a sixth and a seventh day are worked. b) The labor union is very strong and demands that production levels be kept nearly constant throughout the year. Thus, even though there may be a tradeoff between investment in capital equipment and investment in seasonal inventories, the question is largely moot, since the labor union is strongly opposed to substantial swings in work force level.

The planning of inventories, particularly seasonal stocks, is difficult in this situation. There is a need to build seasonal stocks, but very often the time at which one should build them does not seem to be reasonable, given the immediate demand problems. Also, as in most seasonal stock situations, it is hard to decide in which items to build seasonal stocks.

The list of symptoms exhibited by the production and distribution activities in this company was almost a classic set:

1. Poor customer service
2. High inventory
3. High production cost.

The poor customer service was exhibited in two ways, which are related. Some promised shipping dates were missed on make-to-order items for which a specific promise was given to the customer. In addition to these, shortages of stock items had risen to a level which caused the Sales Department a good deal of concern. These two were related because attempts to correct the former led to expediting and short runs, which resulted in reduced production output and stock shortages.

While high inventory in the face of poor customer service seems to be an inconsistency, it is a common occurrence. Poor control leads to excessive seasonal stock accumulation in some items and shortages in others. The excess inventories cannot be used to improve the service in the short items.

The high production costs are primarily a consequence of runs which are uneconomically short with consequent high setup cost and low productivity associated with the high frequency of changeover. These, in turn, resulted from the high runout rate and the consequent need to produce a small amount of each of many items in order to clear up some of the backorders created. This situation often arises when a simple order point, order quantity system is used to control inventories in an environment characterized by limited capacity and strong seasonalities. At the beginning of the peak season, when the demand starts to increase, many items trigger the order point simultaneously, creating a surge in manufacturing load and thus forcing the normal order quantities to be reduced because of the limited production capacity available. At the end of the season, if no seasonal stock limits are built into the order quantity procedures, the last production run will exceed the demand requirement to the end of the season; this leads to an inventory which is inactive until the beginning of the new season.

The system to be described here was developed to help management solve these problems. The center of the difficulties seemed to lie in the planning of aggregate production levels, particularly in allocating available production capacity among several product types with differing seasonal demand patterns, and in the subsequent detailed scheduling of each item belonging to a product type. The result of these efforts is an integrated production planning and scheduling system.

III. STRUCTURE OF THE PLANNING AND SCHEDULING SYSTEM

We decided upon a hierarchical system, one which makes decisions in sequence, with each set of decisions at an aggregate level providing constraints within which more detailed decisions must be made. We did this because we found that we could not, with available analytic methods and data processing capability, develop an optimization of the entire system. However, even if the current state of the art allowed solution of a detailed integrated model of the production process we would have rejected that approach because it would have prevented management involvement at the various stages of the decision-making process. A model that facilitates overall planning can only become effective if it helps in establishing at the various organizational levels subgoals which are consistent with the management responsibilities at each level. The model should allow for corrections to be made to these subgoals by the managers at each level, and for coordination among the decisions made at each level. This is the essential characteristic of hierarchical planning.

Moreover, each hierarchical level has its own characteristics including the type of manager in charge of controlling the execution of the plan, the scope of the planning activity, the level of aggregation of the required information (and the form in which the information should be disaggregated when transferred to lower levels) and the time horizon of the decision. The lower one gets in the hierarchy, the narrower is the scope of the plan, the lower is the management level involved, the more detailed is the information needed, and the shorter is

the planning time horizon. Each level of planning has its own objectives and constraints in which decisions have to be made. It is only natural, therefore, that a system designed to support the overall planning process should correspond to the hierarchical structure of the organization.

Finally, as Emery [8] points out, when a high-level plan significantly restricts the options available at lower levels the plan becomes centralized. This can only be justified if centralization improves the overall performance of the organization, by recognizing broader objectives which cannot be perceived at lower organizational levels. This is particularly true when the degree of interaction existing among subunits is critical, as is usually the case when dealing with production and transportation decisions in a multiplant-multiproduct corporation.

In spite of the importance that this hierarchical approach has in production planning, very few integrated solutions to this problem have been reported. Most of the published efforts have concentrated on analysis of individual components of the overall problem. Although it is theoretically possible to develop iterative procedures which converge to an optimum final plan, by sequential adjustment of lower level actions [1], this approach is not computationally feasible.

Prior to designing the hierarchical system we are about to describe, we considered the attractiveness of using other approaches, primarily those of Landon and Terjung [13], Connors [5], Zangwill [17], and Zoller [18]. We found, however, that those approaches were either difficult to implement or they were based on a given level of aggregation, ignoring the problems associated with detailed scheduling. We decided, therefore, to construct our own hierarchical system, making full use of the idiosyncracies of our particular problem. In taking this pragmatic approach we recognized there was a risk of arriving at planning and scheduling decisions that were not optimal. We are confident, however, that the reasons for isolating and linking manageable portions of the overall decision were sound enough to prevent major deviations from optimality and that we reached our goal of maintaining a simple design, relatively easy to implement. We will now describe the modeling effort to show why this is the case.

To decompose the overall problem we examined the extent to which the various planning and scheduling are coupled. If two sets of decisions were independent we could totally separate them in structuring the hierarchy of decisions.

Starting at the most detailed level, we found that Items sharing a major setup cost had to be scheduled jointly. If these Items are grouped into a Family, the production cost is substantially reduced relative to independent scheduling. Thus, the scheduling decisions for the Items in a Family are strongly coupled. On the other hand, the coupling between the schedules of Items in *different* Families is very weak. Items produced in one plant are produced with the same equipment and therefore compete for capacity, but otherwise are completely decoupled.

We also found that the scheduling decisions for a Family in one time period are strongly coupled to the scheduling decisions for the same Family in other time periods. This is a consequence of the need to accumulate seasonal stock in most Families of Items, due to the competition among Families for scarce capacity. If this were not so the Families could be scheduled independently and no seasonal stock would have to be accumulated. Furthermore, since we have batch production there is a potential coupling between the Family run length and the seasonal stock accumulation.

The latter coupling was removed by finding the optimum seasonal stock accumulation pattern for all Families simultaneously, under the assumption that the unit cost of production in accumulating seasonal stocks was independent of run length. In effect, we ignored run length economics in developing the Family seasonal stock accumulation pattern. However, we treated explicitly the coupling among Families in competing for capacity and the coupling among Family schedules in different time periods.

Ignoring run length economics in accumulating seasonal stocks is valid if the seasonal stock accumulation quantities do not influence the unit cost of production. This can be accomplished by neglecting the seasonal peak in calculating the economic run length for each Family and then extending the run length each time a batch is run to accumulate the seasonal stock needed without incurring an extra setup charge.

We found that the approaches which integrate production planning and economic lot sizes (see, for example, [15] and [7]) are impractical because they lead to evaluation of a large number of production sequencing alternatives, present difficulties in aggregating and disaggregating information, and are computationally expensive for a problem of this size.

Families were aggregated into Types in performing the seasonal planning since this made the seasonal planning computations much simpler and also facilitated the development of the Family and Item schedules. Families which have the same seasonal pattern and the same production rate (measured by inventory investment produced per unit time) are indistinguishable from a seasonal planning point of view. If the cost to carry the inventory is the same for two Families in a Type, only the total accumulation for the two Families (or all Families in the Type) need be considered in developing the plan. The schedules of the Families in the Type are coupled only in that the total of the Family schedules must add up to the Type total as given by the seasonal plan.

Having several Families in which a given seasonal stock can be accumulated simplifies short term scheduling. It reduces the danger of too many Items running out at the same time and allows a general lengthening of runs relative to the constant demand economic run length, without inventory penalty.

Thus, the coupling between Family run lengths and the Family seasonal plans is removed by setting up a seasonal plan which assumes no incremental setups are incurred to accumulate seasonal stock and then accumulating the seasonal stock by extending runs which were scheduled to meet the current demand.

Finally, the schedule for a Family in one plant may be coupled to the schedules for the same Family in the other plants where the same Families can be produced. These schedules will be coupled only if it is desirable to produce in one plant to satisfy demand in another plant's territory and if trans-shipment is cheaper than seasonal stock accumulation as a means of dealing with the shortage. Since, in this case, it is always cheaper to carry the inventory than it is to ship the product, the plants are decoupled and may be scheduled independently. If there were a net *annual* shortage in one plant and an excess in another, the territories should be redefined or capacity adjusted by transfer of equipment.

The assignment of Families to plants for production is independent of the plant production schedules for the same reason. We found the most economic set of production locations for each Family and then used those locations, adjusting plant capacity as needed in the annual planning process. In this way we accumulate the seasonal stock without interplant shipments.

This decoupling process leads to the following hierarchy of decisions, as shown in Figure 1.

1. Assignment of Items and Families to Plants. The process starts with the forecasts of the annual demand in each plant territory. This information, as well as the additional capital investment to be incurred in each plant for the production of a given Family and the interplant transportation costs, is input to the Plant/Product Assignment Subsystem. This determines the plant locations at which each Family should be manufactured. This assignment is done annually.

2. Seasonal Planning. Once each Family has been assigned to a plant or plants, a monthly demand forecast for each product Type is computed for each plant. These forecasts, together with the available production capacity at each plant, are input to the Seasonal Planning Subsystem which defines a monthly production plan and a seasonal inventory accumulation strategy for each product Type at each plant location. The Seasonal Plan is updated monthly to take account of changes in actual inventory position and changes in forecast.

3. Scheduling of Families. The Item and Family Run Length Subsystem develops nominal minimum run quantities for all the Items and Families, and the Overstock Limit Subsystem imposes upper limits on the accumulation of seasonal stock in any one Item, based upon the risks of understock and overstock, and the associated costs at the end of each season. The run quantities and overstock limits are evaluated using standard procedures which will not be discussed in the present paper (see, for example, references [9] and [14]).

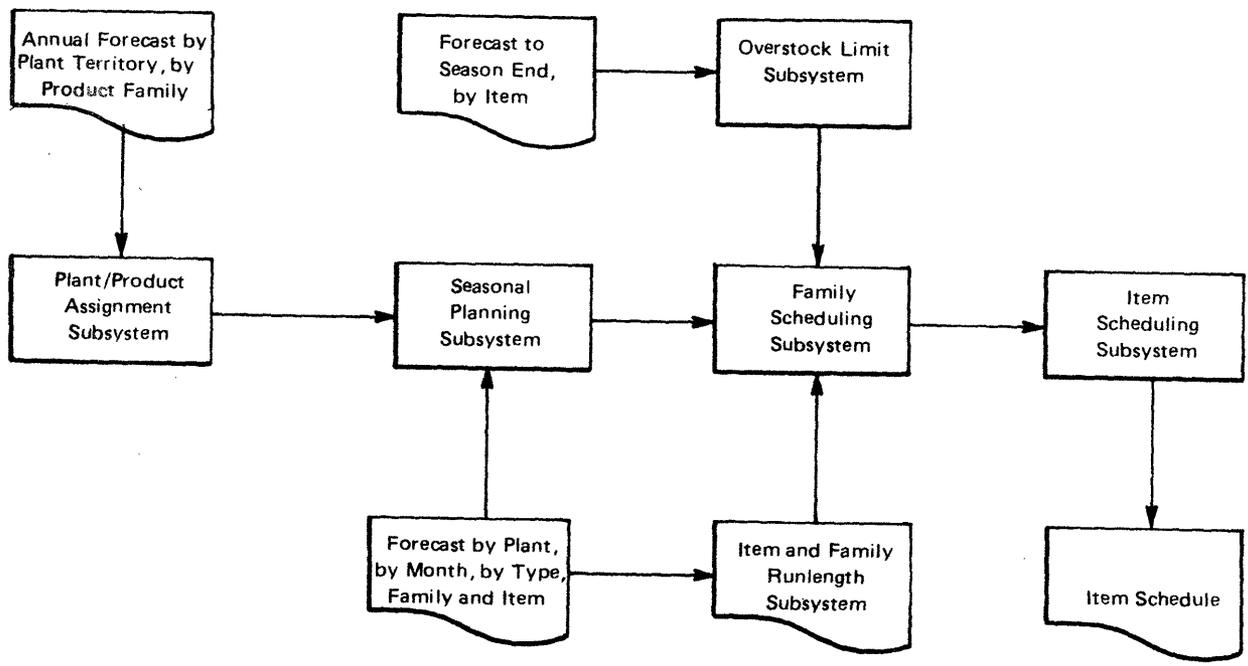


FIGURE 1 DECISION SEQUENCE IN PLANNING AND SCHEDULING

The heart of the production planning and control system is the Family Scheduling Subsystem. This Subsystem uses the results of the Seasonal Planning, Item and Family Run Length and Overstock Limit Subsystems to determine the run quantities to be produced for each Family in the coming month. This system is used monthly as soon as data are available from the preceding month and prepares a schedule for the following month.

4. Scheduling of Items. The Item Scheduling Subsystem allocates the Family production run quantity among the Items in the Family. Runout times are equalized, as this provides best service for the inventory investment, and Overstock Limits are observed. This system is used monthly, immediately following preparation of the Family Schedule.

It is clear that the proposed system requires several different kinds of forecasts. These were prepared using standard techniques discussed in [3], [4] and [6].

IV. PLANT/PRODUCT ASSIGNMENT SUBSYSTEM

When an Item Family can be manufactured at more than one plant, one must first decide where that product should be assigned for manufacture. The analysis must balance the cost of interplant freight and handling against the incremental capital investment cost required to set up the proper facilities, including tools, for the manufacture of the product in question. This decision is most critical for a new product, when it is necessary to determine the initial capital investments to be made, but it is also important for existing products because the variable manufacturing costs differ for each plant location and are subject to change. Furthermore, demand patterns may change, making it desirable to transfer production facilities from one plant to another.

It would have been desirable to develop a multiperiod, multiproduct model to take seasonality into account in the allocation of Item Families to the various plants and to impose capacity constraints on all the products. Such a model, which is very easy to formulate, gives rise to a very large mixed-integer program that cannot be solved within the existing state-of-the art. A capacity-free model offers several advantages. First, it provides the "ideal" distribution of plant capacity for the whole manufacturing process and, thus, it constitutes an effective tool for long-range capacity expansion policies. Second, it permits quantification of the desirability of transferring equipment from one plant, where it currently stands, to another plant, which represents the optimum location for processing the product. Finally, a capacity-free model allows each Item Family to be treated independently, since the only interaction among Items is their competition for scarce capacity.

The following notation will be used to describe the Plant/Product Assignment model in mathematical terms.

C_m = Incremental capital investment required to produce the Item Family at plant location m .

g_{mn} = cost of making one unit of the Item Family at plant m and transferring it to plant n for delivery to the final customer. If $m = n$, the transfer cost is zero.

F_m = forecast annual demand in plant territory m .

x_m = $\begin{cases} = 1, & \text{if the Family is assigned to plant } m; \\ = 0, & \text{otherwise} \end{cases}$

y_{mn} = production at plant m to fill demand in plant n territory

The model has as its objective the minimization of the total cost of capital, production and transportation.

$$\text{Cost} = \sum_m C_m x_m + \sum_m \sum_n g_{mn} x_m y_{mn}$$

subject to constraints:

$$\sum_m x_m y_{mn} = F_n \quad \text{for all } m, n$$

$$y_{mn} \geq 0$$

$$x_m = 0, 1$$

If the number of plants is small, it is possible to enumerate all the alternatives given by the possible values of the variables x_m and proceed to solve a sequence of trivial linear programming problems that can be computed manually.

If there is a large number of plants to consider, the computational methods suggested by Manne [16] have proven to be quite successful. After all Item Families have been processed, we can verify whether capacity constraints for each plant have been violated or not. If the aggregate solution is not feasible, reassignments should be made based on the reassignment cost penalties (i.e., increase in cost as a consequence of changing from the best solution to a second or third alternative) and the capacity constraints at each plant. A procedure for doing this has been suggested by Hax [7]. At this point, the model can be expanded into a multiperiod model to include seasonalities. In the application described here, this step was not required; the seasonalities were explicitly handled in the Seasonal Planning Subsystem.

V. SEASONAL PLANNING SUBSYSTEM

The Seasonal Planning Subsystem determines the production and seasonal stock accumulation for the product Types to be manufactured at each plant so that production and inventory costs are minimized, within the fixed constraints of the existing production facilities and the forecast demand requirements. Inventories are the logical way to smooth the production requirements created by the seasonal and fluctuating demand of the various product types, since it is uneconomical, and at times unfeasible, under these conditions to match perfectly the demand and production patterns. We used a linear programming model to accumulate seasonal stock, minimizing the total cost of inventory and production. Similar models have been developed by Bowman [2], and Hanssmann and Hess [10].

Planning Horizon and Number of Time Periods.

The total time horizon that should be considered in planning seasonal stock accumulation must be long enough to cover at least one seasonal cycle for all the products involved. On the other hand, it should be short enough to make the model computationally feasible and to allow demand forecasts to be reasonably accurate. In the problem under consideration, a time horizon of fifteen months met these two criteria.

The fifteen-month time horizon was divided into seven time periods; the first three are one month each, and the remaining four are one quarter (i.e., three months) each. The model was updated monthly for operational purposes; therefore, the aggregation of time periods did not lose any significant information, but it reduced substantially the computational time required to solve the model.

Product Aggregation.

The Seasonal Planning Subsystem was concerned exclusively with planning production of product Types. In a given product Type, we included all Items having a similar seasonal demand pattern and a similar inventory holding cost per hour of production. Considerations affecting individual run lengths and overstock limits for each Item were taken into account by other subsystems, as explained previously. These subsystems are described in later sections.

Regular and Overtime Production Capacities.

It is important to make a distinction between regular and overtime production because of the increase in labor cost when working overtime. The model properly balanced the tradeoff existing between the overtime cost penalty and the inventory carrying cost to determine the optimum production plan.

Demand Forecast.

A basic input to the subsystem was the demand forecast for each product Type during each time period covered by the planning time horizon. First, we replanned production monthly, updating the model once a month, in order to find the revised optimal production patterns as newer and more accurate demand forecasts were obtained. Second, by discounting the cost elements considered in the model, we gave more weight to the data describing earlier time periods, presumably known more accurately, than to the data describing more distant time periods. Third, the inventory at the end of each time period was required to be at least equal to the safety stock of the product Type. This safety stock is the sum of safety stocks of the Items belonging to the product Type [9].

Cost Components.

The objective of the model is the minimization of the total cost involved in accumulating seasonal stock. We included the cost of overtime production and the cost of holding inventory for each product Type. Regular production labor costs had to be absorbed by the firm, so they were treated as fixed costs. Backorder costs and hiring and firing costs can be easily introduced in this formulation [10]. The backorder problem was handled by a service policy statement and work force leveling was handled outside the model also, as described in the problem statement. Thus, these cost elements were not incorporated in the model.

Mathematical Formulation of the Seasonal Planning Model.

To describe the Seasonal Planning model (ignoring discounting and aggregated technological constraints) in precise mathematical terms, we will adopt the following notation:

R_{it}	=	hours of regular production of product Type i to be scheduled during time period t
O_{it}	=	hours of overtime production of product Type i to be scheduled during time period t
$(R)_t$	=	total hours of regular production available during time period t
$(O)_t$	=	total hours of overtime production available during time period t
I_{it}	=	inventory of product Type i on hand at the end of time period t (units)
r_i	=	Production rate for product Type i (units/hr)
D_{it}	=	forecast demand of product Type i during time period t (units)

- $(CO)_{it}$ = cost of overtime production of product Type i during time period t (\$/hour)
 $(CI)_{it}$ = inventory holding cost for product Type i during time period t (\$/unit-period)
 SS_{it} = safety stock required for product Type i at end of time period t (units).

Using this notation, the Seasonal Planning model has as its objective the minimization of the total of regular and overtime production cost and inventory holding cost given by

$$\text{Total Cost} = \sum_i \sum_t (CO)_{it} O_{it} + \sum_i \sum_t (CI)_{it} I_{it}$$

We find the best seasonal stock accumulation plan by minimizing this cost expression, subject to the following constraints:

$$\begin{aligned} \sum_i R_{it} &\leq (R)_t && \text{for all } t \\ \sum_i O_{it} &\leq (O)_t && \text{for all } t \\ r_i(R_{it} + O_{it}) - I_{it} + I_{i,t-1} &= D_{it} && \text{for all } i, t \\ R_{it} &\geq 0 && \text{for all } i, t \\ O_{it} &\geq 0 && \text{for all } i, t \\ I_{it} &\geq SS_{it} && \text{for all } i, t. \end{aligned}$$

The last constraint introduces lower bounds to the ending inventory at each time period to incorporate safety stock requirements.

VI. ITEM FAMILY SCHEDULING SUBSYSTEM

The Seasonal Planning Subsystem determines the production hours to be used for each product Type. To set the Family run quantities correctly and preserve the optimality of the seasonal stock accumulation plan, the Family Scheduling Subsystem must schedule just enough production in the Families in the product Type to use all the time available to the Type. It must also accumulate the required seasonal stock without incurring setup costs just for that purpose, and insure that customer service standards are met. Finally, the Family Scheduling Subsystem must avoid stock accumulations which exceed individual Item over-stock limits. Within these constraints and because of the way in which product Types are defined, it is immaterial which Families belonging to the product Type we choose to run.

Accordingly, we use a system which, *first*, determines which Families in the Type must be run in the scheduling interval in order to meet Item service requirements; *second*, sets initial Family run quantities so as to minimize cycle inventory and changeover costs; *third*, adjusts the run quantities of the Families scheduled to use all the production time available to the Type, while, *fourth*, observing Item overstock limits.

Two general methods are available for accumulating the seasonal stock needed in each product Type. We can increase the run length of each Family which must be run in order to meet service requirements, or we can include more Families, going “deeper into the run-out list,” making minimum or nominal runs of each Family. We chose the first method since this allows us to save on setup costs, gaining longer runs with no inventory penalty. Further, we selected an allocation logic which balances inventories relative to overstock limits, to minimize the problems of end-of-season carryover stocks.

A Family must be produced in the schedule interval if the available stock A_{kj} of any Item k in Family j , is less than forecasted demand of the Item during the schedule interval plus the desired amount of safety stock S_{kj} . S_{kj} is established using standard safety stock methods. The minimum or base run length Q_j is set for each Family using standard methods for items grouped for production. The overstock limits, L_{kj} , are set using standard newsboy methods.* The cost of understock is equal to the changeover cost associated with having to make an extra run between now and the end of the season; the cost of overstock is equal to the cost of carrying stock over to the next season. Using these together with the distribution of uncertainty in Item forecast to the end of the season yields the desired overstock limit L_{kj} .

To establish the Family schedule we proceed as follows:

Define the runout time RT_j for each Family as the time at which the available inventory is expected to reach the safety stock level:

$$RT_j = \min_k [(A_{kj} - S_{kj})/F_{kj}],$$

where

A_{kj} = available stock of the k -th Item in Family j

S_{kj} = Item Safety Stock level required at the end of the schedule period to meet the service requirement

F_{kj} = forecast usage rate of the k -th Item.

* For these standard methods consult any good inventory control text such as [9] or [14].

In this case the schedule period is one month and the available stock is measured at the beginning of the prior month. Thus, the safety stock must cover the forecast errors during a two month interval. All those Families with RT_j less than the time until the end of the period being scheduled must be put into the schedule for that period. In the remainder of this subsystem description the index j will be restricted to the set of Families in the current schedule.

Define a trial Family run quantity

$$RQ_j = \sum_k RQ_{kj} = \sum_k \min [Q_{kj}, (L_{kj} - A_{kj})]$$

as the sum of the minimum runs for each Item in the Family.

We now sum the trial run quantities for all the Families in the Type and compare the sum with the total production available to the Families in the Type, $r_i(R_i + O_i)$ where R_i and O_i are the regular and overtime hours allocated to Type i and r_i is the production rate in units per hour.

If $\sum_j RQ_j < r_i(R_i + O_i)$, schedule RQ_j^* for each Family where

$$(1) RQ_j^* = \max \left\{ \sum_k (L_{kj} - A_{kj}), RQ_j + [r_i(R_i + O_i) - \sum_j RQ_j] \frac{\sum_k (L_{kj} - A_{kj})}{\sum_j \sum_k (L_{kj} - A_{kj})} \right\}$$

The seasonal stock accumulation amount for the j -th Family is the difference between RQ_j and RQ_j^* . That cannot exceed the amount required to reach the overstock limit for each Item in the Family, $\sum_k (L_{kj} - A_{kj})$. If overstock limits do not interfere (and they usually do not), the amount of production required in the Type, $r_i(R_i + O_i)$ in excess of the base run quantities $\sum_j RQ_j$, is allocated among the Families so that after allocation all Families will have the same fractions of their overstock limits, keeping inventories in good balance.

Notice that when the overstock limit does not interfere, RQ_j^* will be determined by the second term inside the bracket in equation (1), then

$$\begin{aligned} \sum_j RQ_j^* &= \sum_j RQ_j + [r_i(R_i + O_i) - \sum_j RQ_j] \frac{\sum_j \sum_k (L_{kj} - A_{kj})}{\sum_j \sum_k (L_{kj} - A_{kj})} \\ &= r_i(R_i + O_i) \end{aligned}$$

as required.

If $\sum_j RQ_j \geq r_i(R_i + O_i)$, schedule RQ^* for each Family where

$$RQ_j^* = r_i(R_i + O_i) RQ_j / \sum_j RQ_j$$

This expression adjusts run quantities downward if the total production time allocated to the Type is insufficient to produce base run quantities of all the Families which need to be run in the schedule interval. Here we also have

$$\Sigma_j RQ_j^* = r_i(R_i + O_i) \Sigma_j RQ_j / \Sigma_j RQ_j = r_i(R_i + O_i)$$

When the overstock limits do interfere, it may be impossible to produce all the assigned quantity $r_i(R_i + O_i)$ with those Families belonging to Type i which have been put into the current schedule (because their runout times were less than the time to the end of the schedule period). In this case, one should go deeper into the runout list and add additional Families belonging to the product Type. The new Families should be run up to their maximum allowable quantities, given by $\Sigma_k (L_{kj} - A_{kj})$, until the product Type production requirements are met.

No seasonal stock is being accumulated. This simply reflects the realities of the end game as the peak season draws to a close. The Family scheduling logic is shown in flow chart form in Appendix A.

VII. ITEM SCHEDULING SUBSYSTEM

The Item Scheduling Subsystem determines the production quantities for each Item so that the total of the quantities equals the Family schedules. In doing this we must continue to observe the overstock limits and attempt to maximize the customer service which can be obtained with the stock which will be produced. We do this by equalizing the expected runout times for the Items in the Family. This maximizes the expected time until the first Item in the Family runs out. Such a runout requires scheduling the Family again and therefore should be deferred as long as possible, within the constraints of the overstock limits and the total Family production quantity available.

The Item run quantity which accomplishes this is

$$(2) \quad RQ_{kj}^* = S_{kj} - A_{kj} + F_{kj} [RQ_j^* + \Sigma_k (A_{kj} - S_{kj})] / \Sigma_k F_{kj}$$

where

RQ_{kj}^* = desired run quantity for k -th Item belonging to Family j

F_{kj} = forecast demand rate for k -th Item

RQ_j^* = $\Sigma_k RQ_{kj}^*$ = Family run quantity, from the Family Scheduling System

A_{kj} = available inventory of the k -th Item

S_{kj} = desired safety stock of the k -th Item at the end of the schedule interval.

Notice that the new runout time for Item k will be

$$RT_k = (RQ_j^* + A_{kj} - S_{kj})/F_{kj}$$

and, by expression (18), this is equal to

$$RT_k = [RQ_j^* + \sum_k (A_{kj} - S_{kj})] / \sum_k F_{kj}$$

and this is constant for every Item k . This equalizes the expected runout time for all the Items in the Family. Moreover, adding each side of equation (2) with respect to k gives us:

$$\sum_k RQ_{kj}^* = RQ_j^*$$

and, therefore, guarantees that the total amount scheduled for the Family, RQ_j^* , has been allocated among the Items belonging to that Family.

The resulting run quantities must be tested for negativity and against the overstock limits for each Item. If the Item run quantity does not lie between these limits it is set to zero or to the overstock limit, as appropriate. The normalizing constant, $[RQ_j^* + \sum_k (A_{kj} - S_{kj})] / \sum_k F_{kj}$, is then increased (by eliminating the Item from the summation) for the remaining Items until the total Family run quantity has been scheduled. A logical flow chart of the Item Scheduling System is shown in Appendix B.

We did not allocate the Family production run among Items in the same way we allocated the Type production time among Families. The Type allocation should build a balanced seasonal stock and avoid carryover stocks. The Family allocation, within the constraint of a specified seasonal stock accumulation, should defer the next requirement to run the Family as long as possible. For this reason we allocate the Family production run among Items so as to equalize expected Item runout times.

VIII. IMPLEMENTATION

We will divide the implementation activities into two parts – technical and tactical. The technical part refers to the transformation of the system logic into computer programs and input and output systems and procedures. The tactical part is concerned with the development of management familiarity with the system and an effective reduction of the system to practice.

Technical Parts

Surprisingly little difficulty was encountered in the technical parts of the system. We believe this is the result of what some would call a profligate use of computer test time and

programmer time. The major computer systems are the Seasonal Planning Subsystem, the Family and Item Run Length Subsystem, and the Monthly Family and Item Scheduling Subsystem. We first set up the LP for seasonal planning by hand and used the OMEGA code available from University Computing for use on their Univac 1108. In all but the earliest versions, we used a matrix generator code which we wrote for the IBM 360/40 at Arthur D. Little, Inc. We transmitted the matrix to UCC, where the LP was run as a remote entry batch job on the Univac 1108.

The client then programmed the matrix generator for his own IBM 1130. He also ran the LP on the 1130 using the IBM Mathematical Programming System (MPS). This was initially intended to be a pilot version. The plan was to convert it to MPS on the client's 360/40 or use a service bureau. At present, they plan to use the 1130 indefinitely for this job. It provides a great deal of flexibility in use, allows rapid turnaround time and is not expensive, if figured at marginal cost. Most importantly, the programmer analyst is very involved in the problem and is close enough to the computer to resolve his own difficulties.

The Family and Item Run Length Subsystem was programmed three times. It was first programmed on a time-shared console by us. This provided for "conversational" mode development of the logic. It was then programmed on the 1130 by the client. This provided pilot testing and some first-order changes. In fact, there was one major reprogramming of the system after some significant problems were uncovered. So, in a sense, there were two 1130 versions. Finally, it was programmed in a "final" version on the 360/40.

The scheduling system has been programmed only once, on the 1130.

This may sound like a lot of programming and reprogramming. It has not, however, been an expensive strategy. By knowing that initial versions would be modified and reprogrammed, we could then take a more casual view of the requirements and specifications and could get that job done quickly and cheaply. We could then see which aspects of the system needed most change and which changes were most important. This evaluation could be made with results in hand, so the discussions about program specification and design were concrete rather than abstract.

All in all, we believe that we saved programming time and cost by taking this iterative approach to the problem. We could, by running pilot versions of the logic, pave the way for the client and make his work not only easier but more certain to succeed. When he programmed the problem, he could concentrate on system problems, not logic or algorithms, since they had been worked out earlier.

The bulk of the implementation effort went into development of the data base. As is usually the case in firms with primarily manual information and control systems, the existing records contained many errors — obsolete part numbers, incorrect production rates, old costs, etc. Development of a system which would accurately maintain the data file and getting the file cleaned up initially was the major barrier to implementation.

Cleaning up this file had major side benefits. The cost files are now accurate and a better base for planning. Capacity estimation is also much more accurate. Perhaps most important is the development of a common language for production and distribution, essentially eliminating production and shipping errors as a consequence of incorrect part numbers.

Tactical Parts

The tactical implementation of a system of this type is primarily a question of developing management familiarity with the system – what it does and why, and how to make it achieve a desirable result.

A primary problem has been the rigidity of the system, or put differently, the discipline that such a system imposes. In particular, the seasonal planning system will not accept inconsistent sets of requirements and capabilities. If the time available to complete the required production is insufficient, the system will not produce a “best efforts” result; it just produces a printout which says there is no feasible solution.

This kind of system characteristic can be a problem to a manager who is accustomed to “dealing with those problems when we get to them.” Such a manager wants to avoid deciding which of two (or twenty) competing items will be produced with scarce production capacity, until faced with the fact of having both items run out. The manager hopes this will not happen and that he will have made an economic choice by having less than the expected required capacity available.

With the new system, the manager must choose a total production quantity he is prepared to accept; the system will then develop schedules such that the risk of simultaneous runout of two or more items is kept manageably low. This does not require the kind of long-term forecast accuracy or willingness to make blind commitments as may appear at first. The planning model is updated monthly in a routine way and can be updated any time a change large enough to justify such updating has occurred. Each such updating is capable of accepting new production requirements statements and new capacities, as well as reflecting what has happened thus far, via beginning inventories.

A second problem has been to develop an understanding of the difference between a forecast and a plan. For most manufacturers, a forecast of demand is what they plan to make, even though they recognize the forecast is uncertain. Often a seasonal stock accumulation plan is just a level production plan, by item, to reach the annual forecast.

In an explicit planning system of the sort described here, the plan at the beginning of the season may differ from the forecast because of capacity limitations. It certainly does not lead to a level plan by item. Thus, a manager must develop a new working style in using this planning system.

A third problem arises because of the need to decide to accumulate seasonal stocks before forecast uncertainties are resolved. Once the amount to be produced has been established, the seasonal planning system will find the most economic production pattern. It is not easy, however, to decide how much to produce, given forecast uncertainties. Although it is possible to incorporate costs of lost sales in such a model in order to treat this problem explicitly, this has not been done – primarily to avoid the additional complexity, but also because our understanding of these costs does not justify such a quantitative treatment.

The seasonal planning approach used here deals with these uncertainties in a heuristic way, by replanning each month with an updated forecast. While this produces a well-behaved production plan, it is unsettling to a manager who sees his plan greater than the forecast one month and less the following month. As was pointed out above many managers prefer to adjust the plan to match the forecast and make no attempt to deal with uncertainty until faced with the actual requirement. When requirements are a long way off in time, it seems as though an excess of such commitments could be accepted since some of them will certainly not materialize. It has been difficult for management to accept the fact that production resources must be committed to these items long in advance of need, because of seasonal planning; they have, therefore, been slow in providing the basic commitment decisions.

Costs and Benefits

This system required the efforts of a project director, about half time, for a period of about two years. He had a full-time system analyst and about half-time help from a production manager assigned to the project. The total cost of the development, including consultants' fees was \$150-200,000.

We cannot say yet how large the benefits will be. The primary benefit will be smoother production with fewer emergency interrupts. This will save money both in setup costs and in improved service. A secondary saving will be reduction in inventory carrying costs. These together are expected to amount to more than \$200,000 each year in each of the plants.

Acknowledgement

The project director for the manufacturing firm and his colleagues provided invaluable assistance in describing the nature of the operation to us and in reviewing the system as it developed. Unfortunately, they must remain anonymous. D. J. Smith and A. S. Freedman of Arthur D. Little, Inc., were responsible for the development of the Plant/Product Assignment model and the Family and Item Run Length model, respectively.

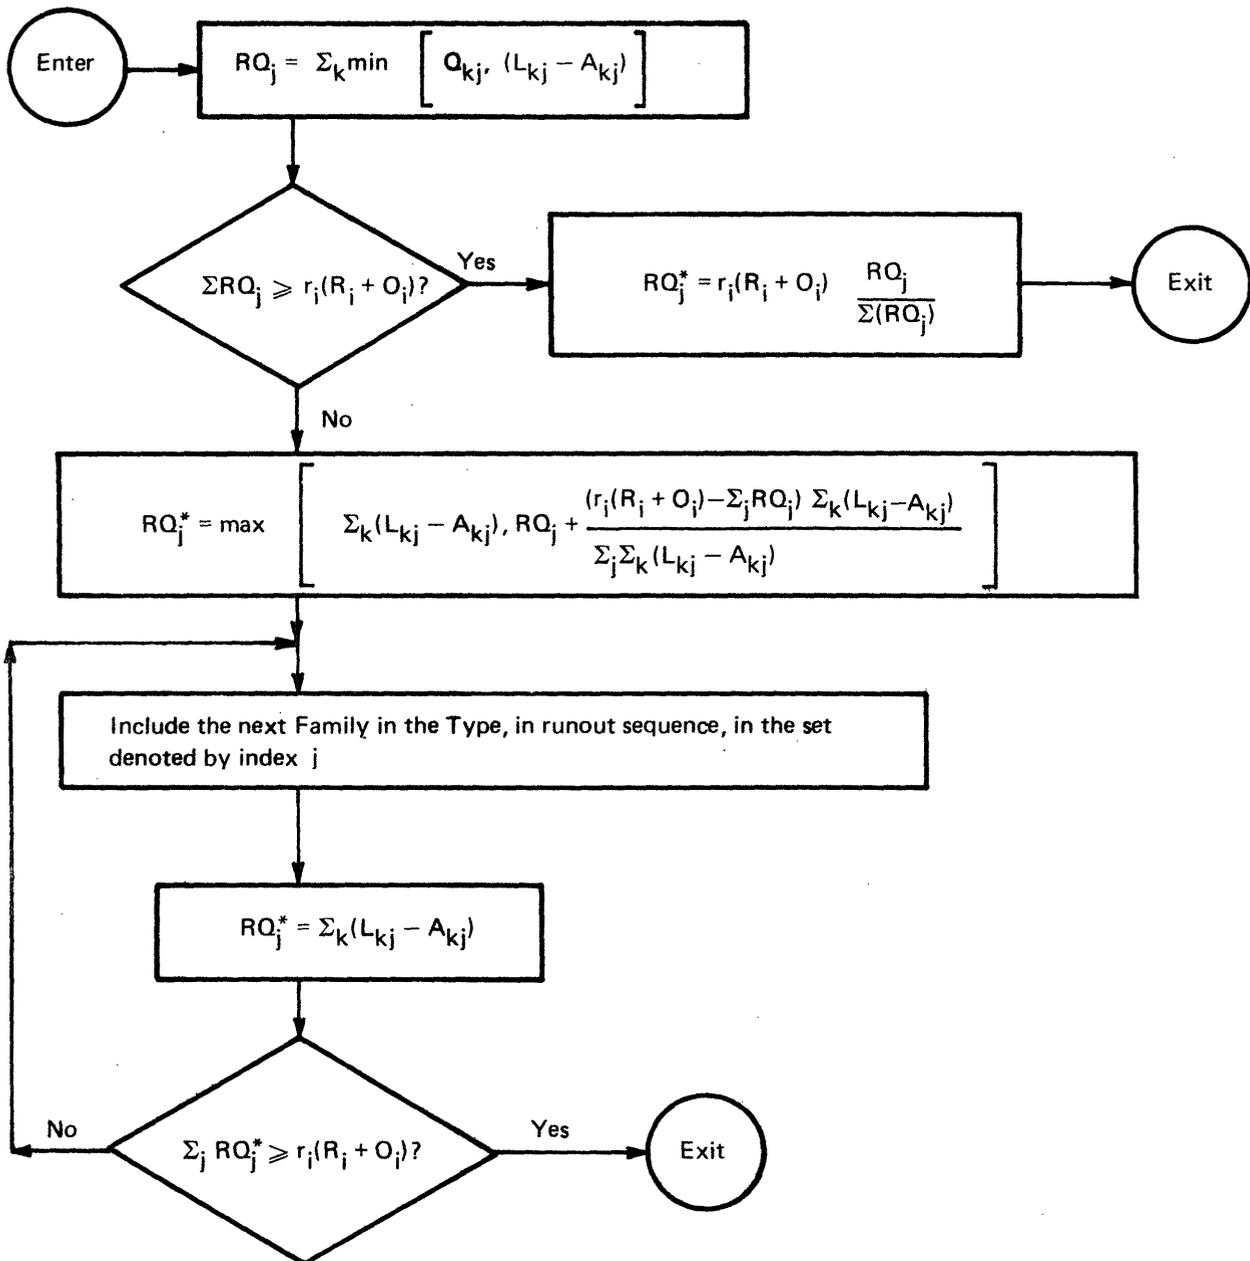
REFERENCES

1. Baumol, W. J., and T. Fabian, "Decomposition, Pricing for Decentralization and External Economies," *Management Science*, Vol. 11, No. 1, September 1964.
2. Bowman, W. H., "Production Scheduling by the Transportation Method of Linear Programming," *Operations Research*, Vol. 4, No. 1, February 1956.
3. Brown, R. G., "Smoothing, Forecasting and Prediction of Discrete Time Series," Prentice-Hall, 1962.
4. Chambers, J. C., S. K. Mullick, and D. D. Smith, "How to Choose the Right Forecasting Technique," *Harvard Business Review*, July-August 1971.
5. Connors, M. M.; C. Coray; C. J. Cuccaro, W. K. Green, D. W. Low, and H. M. Markowitz, "The Distribution System Simulator," *Management Science*, Vol. 18, No. 8, April 1972.
6. Draper, N. R., and H. Smith, "Applied Regression Analysis," John Wiley and Sons, 1966.
7. Dzielinski, B. P., C. T. Baker, and A. S. Manne, "Simulation Tests of Lot Size Programming," *Management Science*, Vol. 9, No. 2, January 1963.
8. Emery, J. C., "Organizational Planning and Control Systems," Macmillan, 1969.
9. Groff, G. K., and J. F. Muth, "Operations Management-Analysis for Decisions," R. D. Irwin, 1972.
10. Hanssmann, F., and S. W. Hess, "A Linear Programming Approach to Production and Employment Scheduling," *Management Technology*, Vol. 1, January 1960.
11. Hax, A. C., "Integration of Aggregate and Detailed Scheduling via Linear Programming," Forthcoming. Also, "Aluminum Products International (A) and (B)," Harvard Business School, 1972.
12. Iglehart, D. L., "Recent Results in Inventory Theory," *Journal of Industrial Engineering*, Vol. 17, January 1967.
13. Landon, L. S., and R. C. Terjung, "An Efficient Algorithm for Multi-Item Scheduling," *Operations Research*, Vol. 19, No. 4, July-August 1971.

14. Magee, J. H., and D. M. Boodman, "Production Planning and Inventory Control," McGraw-Hill, 1967.
15. Manne, A. S., "Programming of Economic Lot Sizes," Management Science, Vol. 4, No. 2, January 1958.
16. Manne, A. S., "Plant Location Under Economies-of-Scale – Decentralization and Computation," Management Science, Vol. II, No. 2, November 1964.
17. Zangwill, W. I., "A Deterministic Multiproduct, Multifacility Production and Inventory Model," Operations Research, Vol. 14, No. 3, May-June 1966.
18. Zoller, K., "Optimal Disaggregation of Aggregate Production Plans," Management Science, Vol. 17, No. 8, April 1971.

APPENDIX A

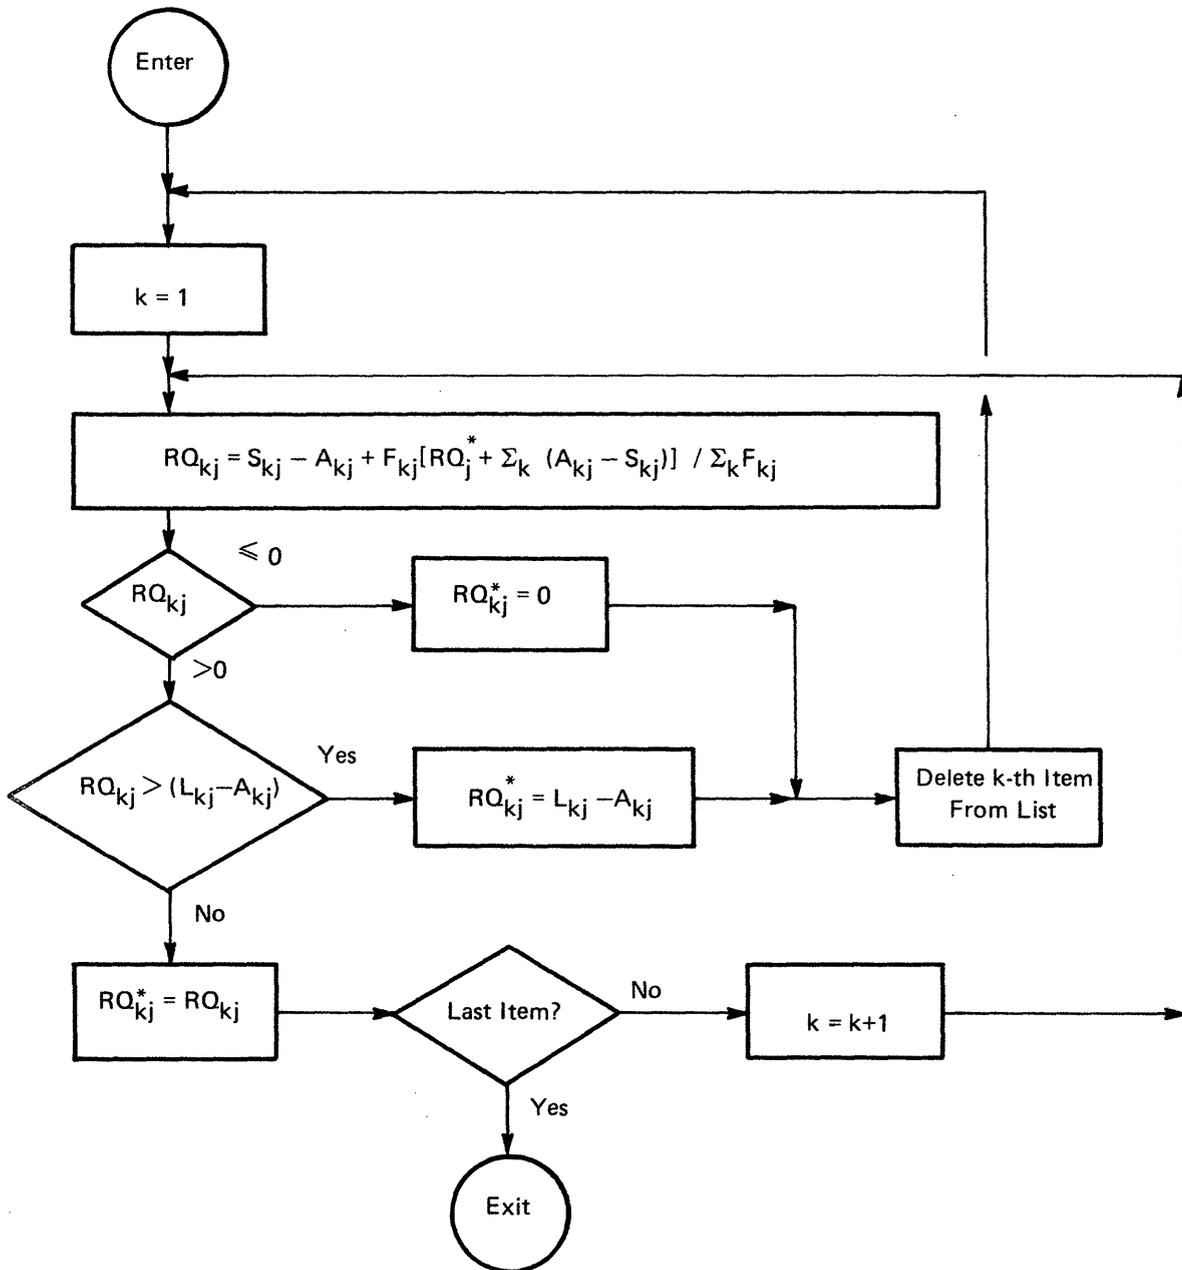
Logical Flow Chart for Family Scheduling Subsystem



Symbols are as defined in Sections V and VI. Initially the index j is limited to those Families in the Type which must be run in the period to meet service requirements.

APPENDIX B

Logical Flow Chart for Item Scheduling Subsystem



Initially the index k runs over all the Items in the Family. As Items are deleted as a consequence of having their run quantities set to zero or to the overstock limit, the list of Items identified by k is modified.