A RESEARCH AGENDA FOR MODELS

TO PLAN AND SCHEDULE MANUFACTURING SYSTEMS

Chacko Abraham*
Brenda Dietrich*
Stephen Graves†
William Maxwell**
Candace Yano††

WP 1689-85

May 1985
Revised, July 1985

* Thomas J. Watson Research Center, I.B.M.

† Sloan School, Massachusetts Institute of Technology

** School of OR & IE, Cornell University

†† IOE Department, The University of Michigan

I. Introduction

In this report we propose a research agenda for the development of models to plan and schedule discrete-parts manufacturing systems. The report was "commissioned" by a set of industrial researchers, academic researchers, and Dr. William Spurgeon of the National Science Foundation (see Acknowledgments for details). Its purposes are to articulate research areas of national importance in the planning and scheduling of discrete-parts manufacturing systems, and to provide funding agencies and researchers with a related agenda.

By "planning and scheduling" we mean the acquisition and coordination of the means of production:

- capital resources

- human resources

- parts and materials

- tools and accessories

- information

The research agenda would provide models that give fundamental principles, insights and/or computer-based algorithms that can be used interactively and iteratively by a decision-maker.

The design of a manufacturing system has typically entailed four phases:

- Product

- Process

- Facilities

- Operational Controls

We address the development and use of models for the design of Operational Controls, namely planning and scheduling. These models address economic and marketplace objectives of efficient operation and customer delivery satisfaction rather than product and process performance objectives. As such, they are often called "macro" engineering models, in contrast to the "micro" engineering models used for design of product, process, and local workplace (which are the major focus of CAD/CAM/CAE models).

This report presents a framework of decisions (Section II), an assessment of past research and current practice (Section III), and a research agenda (Section IV) based upon this foundation.

The framework has four hierarchical levels:

- manufacturing system planning

- production planning

- flow planning

- scheduling

The first three levels use models to make decisions about the acquisition (manufacturing system planning - capital resources and human resources; production planning - "raw" parts, materials, and human resources; flow planning - parts and materials by stage of completion) of some of the means of production. Due to lead times involved in acquisition, plans at the first three levels are made before production actually takes place. The fourth level is characterized by the need to coordinate or merge all five means of production in order to produce and distribute a part, subassembly, or finished product.

In discussing past research and current practice and in setting a research agenda, we limit ourselves to the last three levels of the decision framework (i.e. exclude manufacturing systems planning). We do this to limit the scope to tactical and operational decisions. But by this we do not intend to imply that the tactical/operational issues of planning and scheduling are of more importance or relevance than the strategic issues that arise in manufacturing system planning. Indeed, the research opportunities for model development may be the greatest for manufacturing system planning.

Existing research results are classified into broad categories, and the nature of past research on models for planning and scheduling is assessed in Section III. Our intent here is not to be comprehensive, but rather to give a brief background for the research agenda.

The need for new models, in terms of the framework of Section II, is identified in Section IV, by a comprehensive set of research topics that address the full gamut of the planning and scheduling activity.

II. Framework for Decision Making

The four-level framework is an extension of the framework proposed by Maxwell, Muckstadt, Thomas, and van der Eecken (1983) and closely resembles that of Hax and Meal (1975) and of Morton, Fox, and Sathi (1984). The framework is hierarchical in nature and presumes a corresponding hierarchical decision-making organization. We intend to identify or develop decision-support models for each level in the framework.

Manufacturing_system_planning specifies and organizes the manufacturing resources necessary to meet long-term production goals, often expressed in terms of production volumes and mixes. This includes the determination of equipment, labor and information requirements, and results in the facilities design.

Production_planning takes the facilities design as given, and sets the aggregate production rates or volumes, consistent with the capability of the manufacturing system, in order to satisfy aggregate demand in an economic manner. Production planning also includes the gross determination of reorder intervals for products and parts.

Flow_planning is the disaggregation of the production plan. The flow plan determines the actual production batches based on the reorder intervals from the production plan. Furthermore, flow planning specifies the time flow of these batches through their process steps in a way that is consistent both with the production plan (i.e. the aggregate production rate) and with the resource constraints and demand requirements at each process step. Flow

planning also includes the determination and location of protection stocks (e.g. WIP) necessary to buffer the flow plan from disruptions.

Scheduling is the implementation of the flow plan and results in the real-time sequencing and coordination of the production activities. At this level, there is a very short horizon over which there are explicit constraints on all resources; typically, the primary objective is to try to execute the flow plan as closely as possible.

As an illustration of the decision framework, consider an MRP system. The manufacturing_system_plan has resulted in a specific production facility. Rough-cut capacity planning and the generation of the master schedule corresponds to the production plan. The explosion of the master schedule into time-phased requirements is the flow_plan. And the shop floor control system accounts for the scheduling activity.

Each level has a set of closely related decisions and input data requirements that are different only with respect to time horizons and level of aggregation. For instance, a capacity plan uses monthly or yearly data over a horizon of several years or more, for highly aggregated groups of products. At the other extreme, a schedule deals with individual orders or batches and in time units of seconds, minutes, or perhaps hours.

Table 1 indicates, for each of the four levels

(1) possible objectives or criteria,

(2) control variables,

(3) relevant costs,

(4) constraints,

(5) model robustness considerations,

(6) future start of planning horizon,

(7) length of planning horizon, and

(8) planning time period.

Table 1 should serve as an indicator of the typical factors and issues which should be incorporated into models at each level.

We have listed more than one criterion or objective for each of the four levels. Although most models in the literature have a single objective, "real" problems have layers of objectives. Generally, the first and foremost objective is economic feasibility with respect to a set of salient constraints. The traditional objective of cost minimization (or in the case of a manufacturing system plan, profit maximization) may be secondary or tertiary.

Three items of note in this table are as follows: (1) We include under scheduling the costs of information acquisition (both for initialization and for on-line monitoring), and implementation (e.g., software costs). While these costs may be difficult to estimate, they can be significant, and therefore should be considered (at least implicitly) in models. (2) The availability of tools and accessories, limitations of the material handling equipment, and maintenance time, are constraints that are not typically included in production planning and scheduling models. (3) The hierarchical nature of the constraints is emphasized. That is, the flow plan must be consistent over the short term with the production plan, and the schedule need be

consistent with the flow plan over the scheduling horizon.

The robustness considerations in Table 1 are perhaps the most important new dimension that must be considered in models for any of the four levels of plans. Quantifying this "robustness", and developing models to measure it are one of the major intellectual challenges for the development of meaningful models.

For example, traditional manufacturing system planning assumes that both equipment capacities and the nature of the products to be produced are known, even if product volumes and costs are uncertain. More often than not, however, equipment capacities are not known, because the equipment is new and untested and/or because the actual output of a piece of equipment is sensitive to the scheduling and inventory policies, product mix, and interactions with other equipment in the plant. In addition, equipment purchased to produce a known initial set of products is often used to produce different products within a time frame shorter than the economic life of the equipment. Thus, it is desirable to incorporate both known changes and uncertain elements in the planning to ensure that both the average rate of return and its variance are acceptable.

TABLE 1

LEVEL	OBJECTIVES OR CRITERIA	CONTROL VARIABLES	RELEVANT COSTS	CONSTRAINTS
Manufacturing System Plan	Versatility Strategic Market Positioning Return on Investment Return on Equity	Equipment choice & layout Assignment of production families to production facilities	Equipment purchase costs Operating costs	Product mix Product volumes
Production Plan	Cost minimization Feasibility	Aggregate production rate Workforce level Reorder intervals	Workforce level & change Setups Inventory (cycle and seasonal) Subcontracting	Workforce level change Process capabilities Bottleneck Resource
Flow Plan	Cost minimization Service level	Actual batch sizes Time-phased equipment loading Protection stock	Inventory Expediting Setups Reworks	Consistency with Production Plan Maintenance Plan Tools and Accessories Due dates
Schedule	Cost minimization Meet due dates	Sequencing of activities Rescheduling to disruptions	Information acquisition Implementation Work-in-process inventory Setups	Consistency with Flow Plan Material Handling Capability Space Due Dates Data Handling

LEVEL	START OF FUTURE PLANNING HORIZON	LENGTH OF PLANNING HORIZON	TIME PERIOD	MODEL ROBUSTNESS CONSIDERATIONS
Manufacturing System Plan	6 months to 2 years	1 to 10 years	years	Changing product mix and volumes over time Uncertainty in estimates of costs and capacities
Production Plan	1 to 3 months	6 to 18 months	months/quarters	End-of-horizon effects Rolling horizon implementation
Flow Plan	1 to 4 weeks	1 to 6 months	days/weeks	Information deficiency Uncertainty of parameter estimates and distributions
Schedule	present to 1 week	hours to 4 weeks	continuous	Quality of information about state of system Responsiveness to change Resource Unavailability

III An Assessment of Past Research and Current Practice

A detailed classification and review of the literature on the planning and scheduling of manufacturing systems would be a tremendous undertaking of far greater scope than is necessary here. However, some brief comments and observations on past research and on current practice can guide our research agenda.

There are a few broad categories into which one can classify the past research. As one might expect, there has been a tendency to associate problem types with certain methods of analysis, e.g. combinatorial, probabilistic, or experimental. In particular, there has been significant research in the areas of static scheduling (combinatorial), dynamic scheduling (experimental), production lot-sizing (optimization), production planning (optimization), and protection stock positioning (probabilistic).

Progress in static scheduling is disappointing when one considers the insights gained relative to the number of published papers. Up to the mid 1970's the holy grail of solving an easily posed problem kept researchers busy until the establishment of the essential computational intractibility of these combinatorial problems. The voluminous literature has limited actual application to scheduling of manufacturing, due to the fatal flaw of addressing a static problem with error-free information. Furthermore, we have often defined the scheduling problem exclusive of any production plan, when in fact, scheduling should be the implementation of the production plan over the scheduling horizon (e.g. 1 day to 4 weeks). One important exception is the

work on the Travelling Salesman problem, which is often embedded within the details of the process planning of a series of operations on the same item on a machine or is a consequence of the changeover time from one item to another on a machine.

Comparing dispatching rules for the dynamic scheduling of manufacturing operations was fashionable in 1958-1965. This research was experimental, always involving simulation and usually involving hypothetical manufacturing situations. The resulting ratio-type rules based upon due-dates and processing times, are now used routinely in the design and operation of many manufacturing facilities and have had an impact upon how jobs are dispatched in a job shop. There has been very little focus, however, on what many practitioners view as the dynamic scheduling problem: the quotation of a reasonably realizable due-date for a new potential shop order. Past analyses have assumed that a job due-date is not subject to negotiation by manufacturing management, yet this negotiation is the crux of the marketing-manufacturing interface.

Models for production planning and lot-sizing have met with some success in terms of their applicability, especially those that have been tailored to specific classes of manufacturing problems. Notable in this category are the works of Geoffrion and Graves (1976), Caie and Maxwell (1981), Jaikumar (1974) and Hax and Meal (1975). However, there has been less success with more generic production planning problems. In particular, we cite the finite-horizon, periodic, known-demand problem as examined by Wagner and Whitin (1958, single product without resource

constraints) and by Manne (1958, multiple products, resource constrained, single production stage). When this problem structure is extended to include not only resource constraints and many products, but also many stages of production and a dynamic rolling horizon with possible WIP inventories, there are modelling demons that many have faced, few have recognized, and none have conquered. The first modelling demon is that WIP inventories, destined for assembly, cannot be in backlog status before the assembly. The second modelling demon is the artificial division of the time availability of a resource into time brackets of fixed capacity.

The infinite-horizon, continuous-time, constant-demand lot-sizing problem (EOQ for a single product) has seen a change in focus that may result in insights for process design (Maxwell and Muckstadt (1983)), and production planning (Roundy (1984)). As such, for the former, one can use these models to examine where process improvements (e.g. setup time reduction) will have the biggest impact on lot-sizing in a complex product structure. For the latter, the thrust of the models has been on determining economic reorder intervals as opposed to reorder quantities.

The positioning of stock to protect against demand uncertainty in a distribution system has a long history of models which give both insights and workable policies. Recent work has started to develop some insight on the positioning of stock in manufacturing to protect against demand uncertainty (e.g. Yano and Carlson, 1984). Yet, there has been essentially no work on the problem of positioning protection WIP, in space or time, to cope

with lead time uncertainty, or supply uncertainty in quantity or timing.

Whereas the research literature can be categorized by problem type and/or analysis methodology, current practice is often characterized in terms of 'philosophies' or software products. Manufacturing Resource Planning (MRP), Just in Time production (JIT), and Optimized Production Timetables (OPT) are representative of the philosophies and associated computer software products that have captured the attention of corporate executives in their quest to maintain competitive advantage in domestic manufacturing. Competitiveness hinges on broader issues of tax and trade policies, labor skill classifications, manufacturing management, and technology advantage, in addition to the tactical issues of manufacturing control. Given, however, that manufacturing managers are being instructed to investigate such things as JIT and OPT, we find it compelling and necessary to comment on the scientific basis for these techniques.

Just in Time (JIT) production is a _philosophy_ that calls for reducing work-in-process inventory (WIP) to aid process improvement and reduce process variability. Unfortunately, it has been misinterpreted by some as a _method_ that can achieve zero or minimal WIP with a lot size of one. There are no models or theory to help one to achieve the JIT goals, and in particular, to help to determine when and where to maintain this minimal inventory. For instance, in the literature on the Toyota production system one can find only two operational equations and these are Little's Law and Maxwell's Saturation Law (1985).

The Japanese have recognized, from the point of view of manufacturing control, that engineering of the setup often has higher payoff than engineering of the per unit run time (faster tool cutting times, tools with longer life, etc). They have also developed a method of focusing engineering effort on bottlenecks, rather than making local improvements in productivity that may not make the system as a whole more efficient or effective.

Commercial software packages for MRP, OPT, RESULT, SPEED, EXJIT, (and a host of others) simultaneously address production planning, flow planning, and scheduling. Yet developing a schedule for future operations given conditions on inputs, requirements for outputs, and resource constraints over time is a gory instance of the so-called two-point boundary value problem of control theory. The initial point is a time vector of the inputs -- initial inventory levels and schedules of incoming receipts of parts and materials. The final point is a time vector of the outputs -- due-dates of completed products to be shipped. The boundaries are the time vector of resource availability.

This instance of the two-point boundary value problem is computationally intractable. Often a "solution" is possible if one relaxes one of the three constraints: the initial point, the final point or the boundary value. MRP systems choose to ignore the boundary value, OPT chooses to ignore the final point, and EXJIT chooses to ignore the initial point. As such, this software reflects our lack of success at solving a monolithic model for production planning and scheduling.

In summary, there is a disparity between the insights,

orientation and methodologies of past research and the possible scientific bases of current practice. There is a substantial need to focus research on the proper issues so that the next generation of practice has a scientific basis for each of the levels of planning.

V   Research Agenda

In the following we indicate and discuss the research needs in operational control, according to the categories used in the previous sections. First, though, we comment upon some requirements that must go hand-in-hand with the specification of this research agenda, as well as mention some generic issues that do not fit into the decision categories.

While there has been extensive research in production planning and scheduling, we lack a unifying framework within which to position and evaluate this research. In particular, there is a need for a taxonomy of manufacturing environments that can be used to define and describe production planning/scheduling problems. Abraham and Dietrich (1985) have proposed an extensive taxonomy as part of the effort in developing this reserach agenda; in Table 2 we give an abridged version of this taxonomy. Such a taxonomy is essential to move our field forward in a managed and systematic fashion. Furthermore, this taxonomy could be the first step in developing a generic data-base structure that would capture the information requirements for the range of problems and manufacturing environments. This data-base structure would be the basis for communicating and transferring research results into practice.

I. Production Description

    A. Process Complexity
       1. Stage Description
       2. Machine Redundancy
       3. Machine Flexibility
       4. Machine Utilization
       5. Process Sequences
       6. Time Constraints Among Machines
       7. Set-up Requirements/Cost vs. Quantity Curve
       8. Operator Requirements

    B. Product Complexity
       1. Volume and Number of Part Types
       2. Product Mix
       3. Assembly/Subassembly Production
       4. Batch Regrouping
       5. Batch Contents
       6. Batch Size

    C. Materials Handling Complexity
       1. Nature of Materials Handling System
       2. Flexibility
       3. Storage Capacity at Float Buffers (WIP)
       4. Storage Capacity at Bank Buffers (parts and raw materials)
       5. Transportation and Storage of Bulk and Replacement Parts

    D. Scheduling Criteria
       1. Scheduling Cost Constraints
       2. Scheduling Objectives/Requirements

II. Problem Specification

    A. Requirements Generation
       1. Type of Shop
       2. Use of Vendors

    B. Data
       1. Product Demand
       2. Process Data: process rates for each machine
       3. Machine Failure Rates at Each Stage
       4. Machine Repair Times
       5. Process Yields
       6. Materials Handling Service Rates
       7. Materials Handling Failure Rates
       8. Materials Handling Repair Times

TABLE 2: ABRIDGED TAXONOMY

The research agenda is predicated on the hierarchical
decision framework discussed earlier. However, this framework
is clearly not the only way nor necessarily the right way to
structure decision making. Furthermore, with this framework we
suggest that a hierarchical system of models is needed to support
decision making, i.e. a model(s) for production planning, linked
to a model(s) for flow planning, linked to a model(s) for
scheduling. Thus, we rule out the consideration of a monolithic
model that simultaneously addresses the entire range of planning
and scheduling decisions (e.g. solves the two-point boundary value
problem). The basis for this is our belief that not only is it
not possible to solve such monolithic models, but also that it may
not be necessary, since the proper hierarchical system can give
decisions of comparable quality. But this is primarily
conjecture. It would be valuable to find more definitive evidence
for, or support of, the quality of decision-making from a
hierarchical system.

The construction of a hierarchical system of models results
in decomposing the set of decisions to different levels in the
hierarchy. Inherent in this decomposition is the need to link the
models at the different levels to produce good overall decisions.
Hax and Meal introduced the idea of having a model from one level
impose a constraint on the model at the next lower level (e.g. the
flow plan imposes a constraint on the schedule), but there has
been little consideration of how decisions at one level affect the
decisions at the next higher level. In particular, what type of
information, perhaps "dual" values or shadow prices, should be fed

back to the model at the next higher level? This is an important research topic upon which we comment further in the discussion of research needs for each decision level.

PRODUCTION PLANNING: AGGREGATE PRODUCTION SMOOTHING

Aggregate production smoothing is necessary when it is either not possible or not economic to match production to demand over all points of a planning horizon (e.g. 3 to 6 month). In such cases, the aggregate production rate is set to a smoothed average of the demand rate. As a result, aggregate inventory is built during periods of low demand in anticipation of subsequent periods of high demand. In this way, the production facility adapts to variability in the demand rate with a mixture of anticipatory inventory and smoothed adjustments to the production rate.

We discuss below four issues that need to be addressed by research on production smoothing:

(i)   How do we smooth production in a dynamic environment in which there may be significant uncertainties in both the demand forecasts and the production process? We need models that incorporate these uncertainties and recognize that the production plan is never frozen, but should be updated, for example by using a rolling horizon.

(ii)   How do we smooth production over multiple stages of
        production, e.g. parts fabrication and assembly?
        Models are needed that can coordinate production rates
        across multiple stages, and that can determine how much
        decoupling inventory is needed between stages and
        whether some stages are more suitable for carrying
        anticipatory inventory.

(iii)   Inherent in production smoothing is the notion of
        aggregation: of products, of time and of resources. We
        need to determine how to aggregate products for
        production planning, and how to disaggregate an
        aggregate production plan. We need to understand the
        proper choice of time period for planning and how that
        impacts the disaggregation of the production plan.
        Finally, we need to determine how to aggregate
        resources for production planning, and especially how
        to model and measure capacity.

(iv)   Aggregate production smoothing models usually do not
        reflect lot-sizing considerations, since they deal in
        terms of aggregate products for which the definition of
        a "lot" is ambiguous. Yet any disaggregation of an
        aggregate plan implies a specification of production
        lots. We need to characterize the impact of lot-sizing
        on production smoothing, and to refine production
        smoothing models accordingly.

PRODUCTION PLANNING: LOT SIZING AND REORDER INTERVALS

There is an enormous and varied literature on lot sizing that
has focused primarily on the tradeoff between the setup cost for a
production lot and its inventory holding cost. Various models
examine this tradeoff for both constant and time-varying demand,
for multiple production stages or complex product structures (e.g.
an assembly product), and when there are constraints on production
capacity. Further research should address the following issues:

(i) How should we schedule production lots (reorder
intervals) within a production plan? While advances
continue to be made on multi-item lot-sizing problems
with capacity constraints, we need to understand how to
link lot-sizing with the setting of aggregate
production rates. Whereas an aggregate plan seems to
impose a constraint on lot-sizing, it is not clear how
to feed back the results from a lot-sizing model to the
aggregate planning problem. Furthermore, production
planning is a dynamic activity; this suggests that we
should focus more on the _timing_ of production setups,
rather than on the _sizing_ of production lots.

(ii) There is increasing evidence that various uncapacitated
lot-sizing and reorder interval models are exceedingly
robust; that is, simple heuristics can provide very
good solutions. The next step is to understand how

this robustness extends to more complex problems where demand is time varying and/or where constraints exist.

(iii) New manufacturing technologies have created new classes of lot-sizing problems. One example is a production facility where the setup configures the facility with either a set of components or a set of tools, which in turn defines the family of products that can be processed on the production facility for this setup (e.g. automated assembly equipment for electronic modules and CNC machine tools). Lot-sizing now entails choosing the families, scheduling their setups, and then determining how individual product requirements are to be met within the schedule of family setups.

(iv) The choice of lot size typically depends on the setup time. For production equipment with sequence-dependent setup times, one cannot separate the lot-sizing problem from the sequencing problem, which in its most general terms, has the structure of a Travelling Salesman problem. In practice, one often simplifies the sequencing problem (and, thus, the lot-sizing problem) by identifying and exploiting a hierarchical structure that determines the setup times. This suggests research into the categorizing of manufacturing setup-time structures encountered in practice, and the developing of efficient algorithms for various cases.

(v)   The choice of lot sizes affects the flow of production
        in a complex operation such as a job shop. How do WIP
        requirements depend upon the choice of lot sizes, and
        how can this be incorporated into the lot-sizing
        decision? Also, what is the impact on the production
        flows from either process changes or product changes
        that permit different lot-sizing strategies (e.g.
        smaller lots from quicker setups)?

FLOW PLANNING: PLANNED LEAD TIMES

A planned lead time is the time that one allows for a
production "step": the smallest unit of production activity used
in planning and scheduling. For instance, in a job shop the
production step might be an operation on a job at a work station,
and the planned lead time is the expected time, both waiting and
in process, that a job will spend at the work station. In this
context a job requires, for completion, a series of steps at a
series of work stations.

In many production environments, planned lead times play a
very important role: they permit one to decompose a production
process into a series of steps or stages, and thus can be the
basis for a flow plan. Thus, planned lead times are used for flow
planning and material management in MRP systems, for sequencing
and control in shop floor control systems, and often for quoting
"promise dates" for customer orders. Furthermore, in MRP systems,

the establishment of planned lead times effectively determines the WIP level. Yet, there is virtually no literature that attempts to prescribe how to set and how to manage these critical control parameters.

We need research to determine (i) how to set planned lead times for flow planning in a variety of production environments; (ii) how to integrate the establishment and use of these lead times with other components of the control process, i.e. lot-sizing, scheduling and production planning; and (iii) how to use these lead times to set "promise dates" for customer orders, and to act as an interface between the manufacturing and marketing organizations.

FLOW PLANNING: PROTECTION STOCK

A production process is subject to a variety of uncertainties: in the availability of a production resource (e.g. labor, equipment, material); in the output quantity and timing of a production process; and in the quantity and timing of customer demand. While clearly we will always benefit from a reduction or removal of any of these uncertainties, in lieu of this we still need to create and manage protection stocks that buffer the flow plan from disruptions caused by these uncertainties. Indeed, having models for determining how much protection stock is needed and where, may be an essential step in assessing the value of reducing the uncertainties. Major research questions exist with

regard to the size and location of protection stocks for virtually every production scenario of interest. Thus, while we have a good understanding of how to use finished-goods inventory to buffer against uncertainty in demand quantity, we do not have a comparable understanding for other types of uncertainties. Furthermore, when options exist for the location of a protection stock (e.g. raw materials vs. WIP vs. finished goods), we have neither a theory nor set of tools to address this question. Recently, progress has been made on the location of protection stocks (or safety stocks) in distribution systems; some of this insight might be transferred to or extended for production systems. More specifically, we list below key issues for research.

(i)   We need fresh modeling approaches to deal with production disruptions and delays from resource unavailability, (e.g. equipment failures) and from uncertainty in the process output. While there is an extensive literature on the physical sizing of inventory buffers for transfer lines, there is very little on the more general problem of protecting a flow plan from these disruptions.

(ii)   We need fresh modeling approaches for protection stock positioning for complex product structures. As one example, consider an assembly product with a multi-level bill-of-material. An important question is how

to spread the protection stock across the various levels of the bill-of material. This is complicated by the fact that any component shortage necessarily delays the completion of an assembly or subassembly. A second complication comes from component or subassembly commonality; that is, a component (subassembly) may go into several distinct assembly products.

(iii) A second type of product complexity is when a family of distinct products all have the same antecedent; that is, the first n production steps, say, are identical for all members of the family, and only at the n+1 step do we make a process choice that determines the identity of the product. Again, the key question is how to spread the protection stock across the production steps for the family of products. In this instance, the problem of protection stock positioning is one of determining the level of commitment for the protection stock.

(iv) We need to understand how to integrate the setting of protection stocks with production planning and scheduling. For instance, one might think that larger lot sizes (less frequent production) might require less protection stock; also, when setting protection stocks we need understand how the production plan creates these stocks, and how the production schedule will

replenish them after they are used.

(v)   We need to understand the relationship between
        protection stocks and planned lead times. Planned lead
        times may serve several purposes, one of which may be
        to provide a safety time to protect against uncertainty
        in the timing of supply and/or demand. This leads to
        protection stock, but disguised as WIP. We need
        understand how to use protection stocks vis a vis
        planned lead times to protect against timing
        uncertainty.

SCHEDULING

We conjecture that scheduling in the factory of the future
will primarily entail the generation of computerized Gantt Charts.
At any point in time there will be a detailed specification of
what activities are to be done when, and by what and/or whom, over
the immediate scheduling horizon. Implicit in this capability is
the assumption that shop status information will be available in
(near) real time and that sufficient computer power will be
available to regenerate rapidly one or more schedules, as needed.
With this in mind, we indicate important directions for research.

(i)   How do we integrate the scheduling activity with the
        planning activity? Key issues are how to create a

Gantt-Chart schedule that is consistent with the flow/production plan, and how to define the flow/production plan so that this is possible. To the extent that it may not be possible to schedule the plan with available resources, what is the cost or consequences of violating the plan? How can new information from the scheduling activity be fed back to the planning activity, and in what form?

(ii) Scheduling is a dynamic activity, and the scheduling problem is really one of understanding how to reschedule. It is important to create schedules that are, in some sense, robust to disruptions, robust to the absence or inaccuracy of status information, and are flexible to change. We need a priori measures of schedule robustness and flexibility that would allow one to determine what is a good schedule. We also need to understand the development and use of contingency schedules, as well as rescheduling in general. Finally, to the extent that inventory protects the schedule from disruptions, we need to know how to replenish or replace the protection stock as it is used.

(iii) In some production environments we may be able to set a production plan that permits the development of a "stable" flow plan. This "stable" flow plan then serves as an ideal or target on which to base the actual schedules; as shop disruptions force the actual schedule away from the target, we reschedule to try to bring the shop back to the target. We need to determine how to create "stable" flow plans for a production plan, and then how to use these flow plans for rescheduling.

(iv) In certain production environments there is a significant element of uncertainty associated with rework. Items that fail a test or inspection step may require a diagnostic step followed by a repeat of several previous production steps. This rework is often performed on the same production equipment used for primary production, and hence needs to be scheduled accordingly. Furthermore, when only a portion of a batch needs to be reworked, questions arise over whether to hold the good part of the batch for the completion of the rework or split the batch. We need a scheduling approach that reflects these realities.

(v) To the extent that scheduling is not a totally automated activity, then we need understand how the scheduling system should interface with a human

scheduler. What is the proper role of a human vis a vis a model in the scheduling activity? What is the best information to present to the human, and how should it be presented?

References

Abraham, C. and B. Dietrich 1985, "Discrete Manufacturing Taxonomy," working document, IBM Research, Yorktown Heights, N.Y.

Caie, J. P. and W. L. Maxwell 1981, "Hierarchical Machine Load Planning," Chapter 5 in Schwarz, L. B. (ed.), Multi-Level Production/Inventory Control Systems: Theory and Practice, North Holland, New York, N.Y.

Geoffrion, A. M., and G. W. Graves 1976, "Scheduling Parallel Production Lines with Changeover Costs: Practical Application of a Quadratic Assignment/LP Approach," Operations Research 24, 595-610.

Hax, A. C., and H. C. Meal 1975, "Hierarchical Integration of Production Planning and Scheduling," in Studies in Management Sciences, Vol. I, 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.

Manne, A. S. 1958, "Programming of Economic Lot-Sizes, Management Science 4, 115-135.

Maxwell, W. L. and J. A. Muckstadt 1983, "Establishing Consistent and Realistic Reorder Intervals in Production-Distribution Systems," Tech. Report No. 561, School of Industrial Eng. and Oper. Research, Cornell University, January.

Maxwell, W., J. A. Muckstadt, L. J. Thomas and J. VanderEecken 1983, "A Modeling Framework for Planning and Control of Production in Discrete Parts Manufacturing and Assembly Systems," Interfaces 13, 92-104.

Maxwell, W. 1985, personal communication.

Morton, T. E., M. Fox and A. Sathi 1984, "PATRIARCH: A Multilevel System for Cost Accounting, Planning, Scheduling," partial working document, Graduate School of Industrial Administration, Carnegie-Mellon University, May.

Roundy, R. 1984, "A 98%-Effective Lot-Sizing Rule for a Multi-Product, Multi-Stage Production/Inventory System," Tech. Report No. 642, School of Oper. Research and Industrial Eng., Cornell University, December.

Yano, C. A. and R. C. Carlson 1984, "Buffering against Demand
    Uncertainty in Material Requirements Planning Systems -
    Systems with Emergency Setups," Technical Report 84-13, Dept.
    of Industrial and Operations Engineering, University of
    Michigan, Ann Arbor, Michigan.

Wagner, H. M., and T. Whitin 1958, "Dynamic Version of the
    Economic Lot Size Model," Management Science 5, 89-96.