Operations Research

In This Issue

IN THIS ISSUE

Remarks from the Editor

I am pleased to present this Focus Issue on Manufacturing Logistics. Unlike a special issue, where new papers are solicited that deal with a theme, this issue was created by pulling together previously accepted papers from all topical areas of the journal. Thus, this focused issue achieves two goals. First, it creates a volume that presents state-of-the-art thinking on manufacturing logistics. Second, it reduces the backlog of accepted papers thereby reducing the publication delay—i.e., the time from acceptance to publication—which had been as long as 24 months. Our goal is, through the use of this supplement, to reduce the publication delay to 12 months over the next year. I hope that you enjoy this focused issue and I look forward to your comments and suggestions.

PATRICK T. HARKER, Editor
Philadelphia, PA

A Stock Control Problem

Balancing the load of a stock is one the fundamental problems in modern production planning. This problem can be demonstrated by trucks which are delivering and accepting various quantities of a good at a warehouse. A truck is only dispatched if sufficient quantities of the required good and sufficient stock size to store all of the delivered goods are available. The following natural problems arise in this context: a schedule of the trucks must be found where a sufficient amount of the required good is available in stock. The preferred schedule is the one which keeps the maximum stock capacity needed to store all the goods as small as possible. In "The Stock Size Problem," H. Kellerer, V. Kotov, F. Rendl, and G. J. Woeginger present several algorithms which deliver efficient approximate solutions for this problem.

Combining Make-to-Stock with Make-to-Order Production

Production systems are often categorized as either "make-to-order" or "make-to-stock," but the nature of product demands and product characteristics often necessitate a mixed mode of production, where some products are produced to inventory while others are produced only upon customer demand. In a typical production environment, 20% of the product categories account for 80% of total demand volume, and the remaining 20% of the volume is made up of 80% of the product varieties. While it may be beneficial to hold inventory to provide quick, reliable service for items with high demand, it can become prohibitively expensive to do so for all products in the mix. It may be economical to serve items with high, stable demand from inventory and produce items with low, highly variable demands in a make-to-order fashion. In "A Multiclass Hybrid Production Center in Heavy Traffic," V. Nguyen presents an analysis of a production system that makes multiple product types, some made-to-stock, some made-to-order. The analysis leads to an effective procedure for approximating performance measures such as fill rates and average inventory levels.

How Good are Stationary Policies?

Firms usually prefer to operate in a stationary manner in the sense that the same set of activities is repeated, for example, each month. Inventory managers are no exceptions. Therefore a replenishment policy is preferred if it leads to a stationary flow of material through the system, everything else being equal. However, everything else is often not equal; there may exist a more elaborate, replenishment policy that achieves lower cost. In "Stationary Policies in Multi-Echelon Inventory Systems with Deterministic Demand and Backlogging," F. Chen aims to assess the cost effectiveness of stationary policies relative to the minimum cost achievable among all feasible policies. It is found that stationary policies can be quite costly in the worst case. But a simulation study shows that such a worst case rarely occurs and that stationary policies perform very well on average.

Planning Requirements when Forecasts Change

Many discrete-parts manufacturing firms plan their multistage-manufacturing operations using the logic of materials requirements planning (MRP) systems. A key input to an MRP system is a forecast of demand. The planning system then converts the demand forecast into time-phased material and production requirements under the assumption that the forecast is accurate and will not change. Inevitably, though, the forecast does change, and the firm must revise the requirements plan. In the operations research literature, there are very few attempts to model how MRP-based planning systems function in realistic environments, namely a world with changing and inaccurate forecasts. In "A Dynamic Model for Requirements Planning with Application to Supply Chain Optimization," S. C. Graves, D. B. Kletter and W. B. Hetzel provide a tractable model that determines production and inventory requirements for a multistage operation operating under MRP control, subject to a dynamic and stochastic forecast process. The authors also report on the successful application of this model, a reduction of the value of the inventory by more than 20%, to a supply chain for film manufacturing at Eastman Kodak.

Controlling Yields in Fiber-Optic Cable Manufacturing

Many models of manufacturing with uncertain yields focus on establishing safety stocks, which are used to supplement production quantities when a failure occurs. In fiber-optic cable manufacturing, uncertainties in the process result in cables that are shorter than the length planned. When cables are being made to order, extra lengths of cable cannot be stored separately and added to a cable that is too short. If the planners add enough length to every cable, no cables will be rejected because of length, but this extra length is costly to manufacture. Given the chance to specify the effective yield of the production process, what amount of "safety length" should the planner add to an order? Should thicker cables (containing more individual fibers) have more "safety length?" Should longer cables have more "safety length?" In "A Subassembly Manufacturing Yield Problem with Multiple Production Runs," P. Nandakumar and J. Rummel develop a model for such a yield problem, characterize the behavior of the model, and provide efficient solution heuristics.

Limited Customer Willingness to Wait

Increasingly, management is becoming more concerned about providing an appropriate level of inventory for products carried by companies. Current thinking has it that whenever there is no stock on hand, a customer either will leave immediately or will wait indefinitely, each with negative consequences to the company. However, in reality, customers may show some propensity to wait for a period of time for satisfaction before deciding to leave. This willingness to wait for a while may have a considerable impact on management decisions regarding an inventory stocking policy. In "An $(S - 1, S)$ Inventory System with Fixed Shelf Life and Constant Lead Times," D. Perry and M.J.M. Posner consider a fairly general "willingness to wait" distribution. The analysis is carried out in the framework of certain inventory policies with constant replenishment times and for products having fixed shelf life or obsolescence time. The rather simple solution results that are generated should assist managers in determining better stocking decisions, leading to potentially large cost savings.

Inventory as a Cure for the Supply Chain Bullwhip Effect

Fluctuation in consumer demand is a fact of life in most industries. What is not as well understood is how stock-ordering policies within the supply chain can act to amplify these fluctuations, although it is generally assumed that inappropriate inventory management can contribute to this phenomenon. This so-called "Bullwhip Effect" has been observed widely and can give rise to significant costs due to swings in both inventory and backorder levels. In "The Stabilizing Effect of Inventory in Supply Chains," M. P. Baganha and M. A. Cohen introduce a model that examines the role of supply chain inventory, at intermediate echelons, as a stabilizer of demand variability. Our main results include characterization of conditions under which inventories act as a cure for the bullwhip effect. We also discuss mechanisms to promote demand smoothing through improved inventory management policies.

Limited Buffers Impact Scheduling Rules

Almost all of the previous literature on deterministic scheduling models assumes that buffers of infinite capacity for holding raw materials and finished goods are available at the start and end of the production process. A variety of efficiency considerations, however, suggest the need to limit buffer size. The implications for finding good scheduling rules are considerable. In "Scheduling with Finite Capacity Output Buffers," N. G. Hall, M. E. Posner, and C. N. Potts consider a variety of classical scheduling problems in the context of limited output buffer capacity. In several cases, they provide fast, optimal solution procedures. In other situations, they provide results which show that such procedures are unlikely to exist. In "Scheduling with Finite Capacity Input Buffers," the authors consider a variety of classical scheduling problems in the context of limited input buffer capacity and provide similar results to the output buffers case. As a consequence of these works, the relationship between buffer design and scheduling rules in manufacturing systems is clarified.

Cyclic Schedules Help in Plant Management

Quoting and meeting consistent and short lead times for customer orders, consisting of a set of products, is an important aspect of industrial plant management. This challenge is heightened because of randomness in demand timing, quantity and mix, as well as capacity constraints. An essential part in achieving this goal, therefore, is to be able to schedule bottleneck resources that produce a set of products in a predictable manner. Inventory levels must also be consistent with the capacity utilization. Motivated by problems from a laminate plant, in "Managing Stochastic Multi-Product Systems: Models, Measures, and Analysis," R. Anupindi and S. Tayur study the use of cyclic schedules by analyzing them using a simulation based optimization method, and report an actual implementation.

Insight into an Inventory Puzzle: Order Crossover

Traditional one-cycle inventory models of stochastic lead times often adopted the contradictory assumption that lead times were independently and identically distributed, but that there would be, at most, one order outstanding. That assumption, which led to an upper bound on the true cost, was necessary to avoid the complications that arise when orders may arrive out of sequence. In "An Inventory Model with Order Crossover," X. X. He, S. H. Xu, J. K. Ord, and J. C. Hayya present a multicycle, constant unit demand, exponentially distributed leadtime model to derive a tight upper bound on the total cost and to gain insight into the effects of order crossover. They find that the crossover model may yield a significantly lower total cost figure than the one-cycle model. The results provide some insight into the effects of order crossover on the performance of inventory systems, a problem that has puzzled analysts in the past.

Meaningless Solutions to Optimization Problems

Many practical problems can be modeled as optimization problems. The study of such problems is a focus of a major part of operations research. Optimization problems often involve objective functions that are based on scales of measurement. While a major focus of OR analysis involves sensitivity of conclusions to perturbations in data, it is rare for an investigator to ask whether an admissible change of scale of measurement (such as a change of unit or zero point) might affect the conclusion that a given solution is optimal. In "On the Meaningfulness of Optimal Solutions to Scheduling Problems: Can an Optimal Solution be Non-optimal," N.V.R. Mahadev, A. Pekeč, and F. S. Roberts point out that one cannot conclude that a given solution to a scheduling problem is optimal as that solution might change if an admissible change of scale takes place. If the conclusion can depend upon the accident of a particular choice of unit or zero point, they question whether it is meaningful to ask for an optimal solution. These points are made by considering the problem of finding an optimal schedule for jobs on a single machine. This schedule has penalties for both tardy and early arrivals and penalties that are measured using scales that may be changed by admissible transformations. The authors formulate and analyze this particular scheduling problem for four specific penalty functions under a variety of different assumptions. However, the main point of the article is a message about operations research theory and practice in general, a message that extends well beyond the particular application to scheduling.

Two-Echelon Inventory Systems with Batch-Ordering

When distributing products over a large geographical region, it is common to use a two-echelon inventory system with a central warehouse and a number of retailers close to the customers. The retailers replenish from the central warehouse. In the same way, the warehouse replenishes its stock from an outside supplier. In "Evaluation of Installation Stock Based (R, Q)-Policies for Two-level Inventory Systems with Poisson Demand," S. Axsäter considers such a system where all facilities apply general installation stock batch-ordering policies. The retailers are different and face Poisson demand. He presents a method for evaluation of holding and shortage costs that is exact for two retailers. The technique is based on the solution of a much simpler problem with one-for-one ordering policies. The solution for two retailers is used as an approximation in cases of more than two retailers.

Focusing on Setup Times

In many production or communication systems, a single resource, such as a machine or transmission line, is shared by many types of customers. Different customer types access the resource in a cyclic order. An important factor in the performance of such shared resource systems is the time spent on setups. This is often the focus of the KAIZEN, or continuous improvement, approach to improving system performance. However, considering mean setup times alone is not sufficient to improve performance. As shown in an earlier paper by Sarkar and Zangwill, reducing mean setup times may in fact lead to larger mean work-in-progress and waiting times. Hence, entire distributions must be considered. In "Multiclass Production Systems with Setup Times," R. Righter and J. G. Shanthikumar show that reducing setup times in the convex sense, stochastically reduces work-in-progress and waiting times. They also show that the joint distributions of various performance measures across time are invariant under changes in mean setup times, as long as the total mean setup time remains constant.