Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler eds.

A SIMULATION-BASED APPROACH FOR INVENTORY MODELING OF PERISHABLE PHARMACEUTICALS

Ana R. Vila-Parrish
Julie Simmons Ivy
Russell E. King

Edward P. Fitts Department of Industrial & Systems Engineering
North Carolina State University
Raleigh, NC 27695-7906, U.S.A.

ABSTRACT

Pharmaceutical expenditures are increasing for hospital systems nationwide. We model the inventory and ordering policies for perishable drugs in the setting of an inpatient hospital pharmacy. We consider two stages of inventory: raw material and finished good (e.g. intravenous). We use a two-phased approach to explore policy structures that could be implemented in the hospital pharmacy. We develop a policy which is based on the idea that hospitals can improve both costs and patient demand fulfillment by using knowledge of patient mix to guide their drug inventory and preparation decisions. We compare this policy to a simpler stationary base stock policy. The policies are evaluated on the basis of (1) shortage cost, (2) outdating cost (expirations), and (3) holding cost through a range of cost scenarios.

1 INTRODUCTION

Hospital pharmacies throughout the United States are experiencing inventory problems that result in waste and shortages that affect patient outcomes due to delayed procedures and drug substitutions (Fox and Tyler 2003; Baumer et al. 2004; Gillerman and Browning 2000; Weigner 2001). While some drug shortages are uncontrollable (e.g. due to a natural disaster), improper inventory management can result from the clinical versus procurement expertise of those managing the inventory (Alverson 2003). While the hospital's pharmacists have expertise in the efficacy and treatment protocols of drugs that are administered to patients, they are also tasked with managing, ordering, and producing the drugs that are ultimately dispensed to patients. The pharmacists serve as the gate keepers of drug distribution by ensuring the accuracy and appropriateness of prescribed medications.

Serious consequences of traditional pharmacy purchasing include missed contract compliance, excess inventory levels, frequent stock-outs and costly deliveries, workflow interruptions and expensive rework, and increased health system labor requirements (Alverson 2003). These problems are exacerbated by the large number (typically more than 2,000) of perishable drugs in the hospital, the manual inventory stock keeping and daily order entry.

Pharmacists must make decisions regarding their inventory levels and how and when to produce drugs in response to, or in anticipation of, patient demand. Frequently, these decisions are made by measuring drug utilization from historical data and devising a common inventory level (measured in days of inventory) which makes the process easy to manage but not efficient. However, there is additional information available that pharmacists are currently not using to make inventory decisions. Drug demand is a function of the patient's condition and the prescribing protocols of the physician. The patient's condition provides a forecast of their drug need during their length of stay. The consumption of a particular drug on a given day is conditional on the mix of patients in the hospital that require this drug and may be for more than one unit of drug. Since drug usage changes over time and is not known with certainty, the daily demand is dynamic and stochastic.

2 RESEARCH OBJECTIVES

We consider an inpatient hospital pharmacy which stocks pharmaceutical drugs in two stages, namely raw material and the finished good (i.e. a dispensed form such as intravenous, IV). Our research objective is to use patient condition information to determine the appropriate inventory level in each stage per period that minimizes wastage and holding cost while maximizing timely access to required (preferred) drugs.

Policy development consists of two phases. First, we model the problem as a Markov decision process (MDP).

978-1-4244-2708-6/08/$25.00 ©2008 IEEE

We use the MDP to represent drug demand as a function of patient condition and accordingly determine the appropriate drug inventory levels. From the results of the MDP, we gain some structural insights on the form of the optimal policy for a simple problem. Since the state space increases rapidly as a function of the number of patient types and the number of periods until drug expiration, the MDP becomes computationally intractible for large-scale problems. Therefore, in the second phase we use simulation to evaluate inventory policies characterized from the first phase for a variety of realistic problem instances.

3 LITERATURE REVIEW

Problems related to the management of perishable inventories arise in different sectors including blood banks, foodstuffs, chemicals, and drugs where demand is stochastic (see Nahmias 1982 for a complete review). Analysis of these problems introduces complexity due to the dimensionality of the inventory age vector. We focus our discussion on perishable inventory literature for products with a fixed lifetime and periodic review. Although we acknowledge there is substantial literature for deteriorating products with continuous review (refer to Weiss 1980; Graves 1982; Schmidt and Nahmias 1985; Liu and Lian 1999) we focus on the literature which assumes a periodic review since it is most appropriate for our problem setting.

Perishable inventory research to date has made many simplifying assumptions deemed necessary to make the problem tractable but limit its applicability to the pharmacy inventory problem. For example, Veinott (1960) and Dave (1979) consider perishable products with known demand. Those that model the dynamic nature of the demand function assume that it is essentially independent of the system (exogenously defined) and stationary (Nahmias 1975; Brodheim et al. 1976; Cohen 1976; Chazan and Gal 1977; Schmidt and Nahmias 1985; Lian and Liu 1999b; Katagiri and Ishii 2002). The use of non-stationary demand functions is rare but Haijema *et al*. (2007) considered a seven day discrete demand function. Our approach is most closely related to Haijema (2007) as they use a combined optimization by dynamic programming and simulation in the context of a blood bank inventory problem. Our research addresses a significant gap in the inventory literature by defining the stochastic demand function as a property of the system, instead of an exogenously defined random variable, and a significant gap in the pharmacy literature by relating patient condition to drug demand.

The hospital setting is unique because there is an interrelationship between the patient's condition, the patient's demographics, and drug utilization. Patient condition impacts drug demand, however patient response to drugs is uncertain and may also affect drug demand. Drug availability may impact patient condition, resulting in increased hospital stay, and thus impact demand. We incorporate and model this interrelationship between the demand function and the inventory management of perishable goods. The patient's duration of hospitalization stay is stochastic and therefore results in further uncertainty in expected demand.

4 MDP METHODOLOGY

Consider, for sake of discussion, a hospital unit with two patient types. In general we assume there are N patient types, but as an example we will define a system with two patient types. Each patient type has an associated demand profile for a single drug. Because there is uncertainty associated with the patient's condition and response to treatment, the patient's condition changes stochastically during their stay. If the patient's condition changes such that their demand for the drug changes we assume that their "type" also changes.

STATE DEFINITION: Our MDP state definition involves two components: the number of patients in the system of each patient type and the two-stage inventory position. First, we define the patient component of the state. Let type (1) represent a patient in the unit with highly variable demand with a large mean and type (2) represent a patient in the unit with demand that has lower variability and mean. A third patient state is defined to be absorbing and represents the patient's discharge from the hospital. A graphical representation of this Markov chain is shown in figure 1. The patient can arrive to this hospital unit and be characterized as either type 1 or type 2 with rates λ_1 and λ_2, respectively. During the patient's hospitalization they may stay in the same condition, change condition, or transition out of the system. A probability transition matrix, P, specifies the probability that a patient changes type or leaves the system given their current type classification. For example, a patient which is considered to be type 1 at time t changes to type 2 by time t+1with probability p_{12}. We define a q_{it} as the number of patients of type i in the system in period t and $Q_t = \{q_{it}\}$ is the patient vector in period t.

Patient Demand is modeled as a discrete non-negative random variable D. We define a stationary probability mass function, $f_j(d) = P(D = d \mid j)$ which is the probability that demand (in units) equals a discrete value d for patients in state j. We assume that maximum demand is a finite number M.

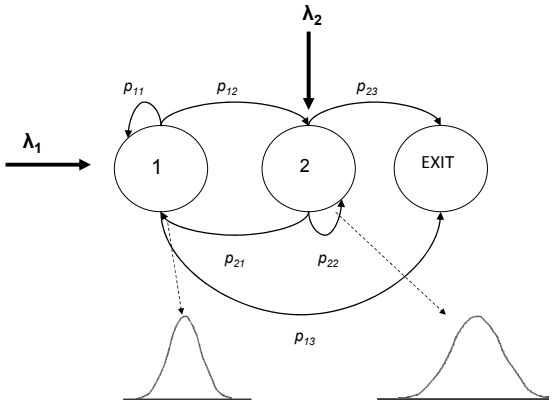

Figure 1: Markov chain for a system with two patient types (1) and (2) and third state which represents patient's discharge.

Second, we define the inventory component of the state space. As previously mentioned, we assume the drug takes on two forms: raw material and finished good or dispensed form (achieved by additional processing), the inventory vector must account for both forms. We define $\bar{I}_t = (i^r_{1t},...,i^r_{mt},i^f_{1t},...,i^f_{jt})$ as the inventory level of raw material, i^r_{kt} , in period t with remaining life of k periods (k = 1,.., m) and finished goods, i^f_{lt} , prepared in period t with remaining life of l units (l = 1,..j) where $j < m$. Order, inventory, and processing quantities are a function of both patients and current inventory levels therefore the state space of our MDP is multi-dimensional, $(\bar{I}_t, \bar{Q}_t)$.

DECISIONS: Each period (e.g. each day) there are multiple decisions that must be evaluated. At the beginning of period t, we observe the inventory level $\bar{I}_t$ and the expected patient demand which is a function of the patient mix, $\bar{Q}_t$. Thus, we observe the state of the system $(\bar{I}_t, \bar{Q}_t)$. We determine the drug order quantity, x_t, and how many units of raw material, r_t, to convert to the dispensed form for use in period t, where $r_t \in R$, $x_t \in X$: where R,X is the set of feasible process/order quantities (alternatives) that minimizes the total expected cost. We assume that the finished goods are received immediately after an order has been placed while the raw material has a one period delay.

The sequence of each period's events proceed as follows: the state of the system $(\bar{I}_t, \bar{Q}_t)$ is observed, a finished goods order is placed (if needed) and arrives instantaneously, an order for raw material is placed (if needed), the day's actual demand is fulfilled, then at the end of the period, any unused drug which outdates in this period is depleted from inventory, any departing patients leave the system, the raw material ordered arrives, and the resulting cost (described below) is calculated. The timeline of events are shown in Figure 2 for a system with two patient types, raw material with infinite shelf life, and finished goods inventory with a two period shelf life (i.e., finished goods ordered in period t expires at the end of period t+ 1).

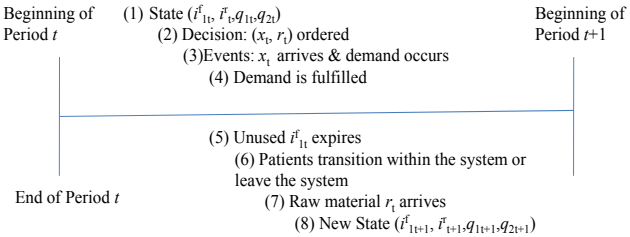

Figure 2: Timeline of decisions and events for the MDP model

OBJECTIVE FUNCTION: The objective is to minimize cost assuming no backlogging of demand. Instead, we assume that the patient demand will be fulfilled that day by procuring the drug from another channel (e.g. other hospital). Two shortage costs are incurred which are dependent on the stage of inventory which has the stockout. The stockout cost, per unit, for the finished good form of the drug is defined as b_2. In addition to b_2, the cost of stocking out of both the finished and raw good imposes an additional cost of b_1 . Lastly, there is a per unit cost, w, for inventory which expires. The single period cost minimization is defined by equation (1).

$$G_t(\bar{I}_t,\bar{Q}_t) = \min_{\substack{0\le x_t\le i^r_t \\ r_t\ge 0}} \begin{bmatrix} b_1 \max\{0, \sum_{i=1}^{N}\sum_{d=0}^{M}(q_{it}f_i(d) - x_t - \sum_{l=1}^{j} i^r_{t-l} - \sum_{k=1}^{m} i^f_{t-k})\} + \\ b_2 \max\{0, \sum_{i=1}^{N}\sum_{d=0}^{M}(q_{it}f_i(d) - x_t - \sum_{k=1}^{m} i^f_{t-k})\} + \\ w\max\{0, i^f_{t-m} - \sum_{i=1}^{N}\sum_{d=0}^{M} q_{it}f_i(d)\} \end{bmatrix} \quad (1)$$

The MDP is solved using backwards recursion and the recursive relationship is defined in equation (2).

$$C_t(\bar{I}_t,\bar{Q}_t) = \min_{r\ge 0, 0\le x\le i1} E_{d|q_i}\left[G_t(\bar{I}_t,\bar{Q}_t) + C_{t+1}(\bar{I}_{t+1},\bar{Q}_{t+1})\right] \quad (2)$$
$$t = 0,...,T$$

NUMERICAL SOLUTION OF MDP: MDP experimentation was conducted for a small problem instance with a single patient whom can change between two patient types, infinite raw material shelf life, and a two period finished goods shelf life. As reflected in equation

(1), the number of units of raw material in stock imposes an upper bound on the number of units of finished goods that can be processed. The optimal inventory policy for the finished good is dependent on the patient type and is constrained by the total raw material available. The optimal inventory policy at time t for the raw good is dependent on the expected patient mix at time t+1. From the results of the MDP we observe that the policies have a base stock structure where the base stock levels are dependent on the raw and finished goods inventories and the patient mix. The dependence on patient mix motivates the development of the base stock levels for the adaptive policy that we develop and evaluate via simulation in Section 5.

5 SIMULATION EXPERIMENTS

In order to explore more realistic problems we turn to simulation and compare the cost performance measures under various parameters. We develop two inventory policies one which is based on the MDP solution and another which uses a fixed base stock level that is independent of patient mix. We describe these policies in more detail in subsequent sections.

5.1 Simulation Model Parameters

We consider a system with three patient types that have unique demand characteristics (i.e. mean, variance). Figure 3 shows the Poisson demand distributions for each patient type used for this experiment and Table 1 summarizes the demand distribution associated with each patient type. We chose a mean demand which decreases from patient type 1 to 3. The patient types are ordered such that type one represents a patient whose condition is more severe while type three's condition is least severe.

Table 1: Patient type and demand parameters

Patient Type	Mean (λ)
1	10
2	4
3	1

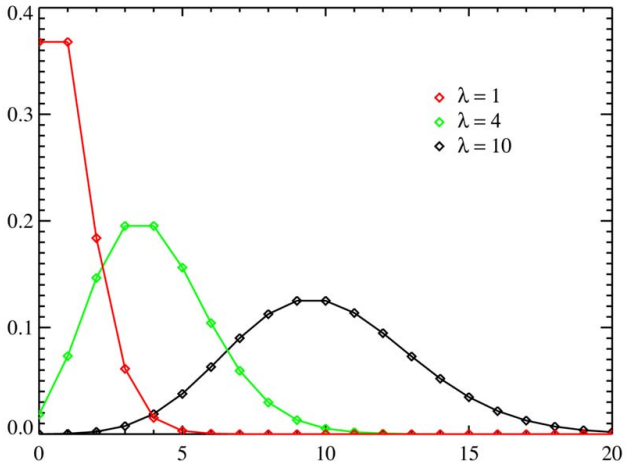

Figure 3: Demand distributions associated with each patient type

In addition to demand characteristics for each patient type, we define a set of transition probabilities. These probabilities are the chance that a patient's classification changes from one period to the next period. This is important in the model, as in the MDP, since a change in patient type corresponds to a change in patient demand. In the simulation these probabilities are used to formulate a policy, called the Adaptive Policy (AP), which uses the transition probabilities as a forecast of demand for the single drug in the next period. The transition probabilities, given that the patient does not depart from the system, are outlined in Table 2. In addition to the transition probabilities, we define arrival rates by patient type and the probability that the patient exits the system. Arrival and departures parameters are defined in Table 3. Patients are assumed to arrive at the same rate regardless of patient type however the probability of departure increases from patient type 1 to patient type 3. As before, we chose these probabilities of departure to correspond with the patient condition. A patient whose condition is severe (type 1), at period t, has a probability of zero of being discharged in period t+1.

Table 2: Patient transition probabilities given that the patient does not depart from the system

Patient Type	1	2	3
1	0.5	0.3	0.2
2	0.2	0.6	0.2
3	0.05	0.3	0.65

Table 3: Arrival rates and probability that a patient departs by patient type

Patient Type	Arrival Rate	Prob(Exit)
1	4/day	0
2	4/day	0.2
3	4/day	0.8

We define three cases for our simulation and policy experimentation according to the parameters summarized in Table 4. The Base Case represents a drug cost structure with a larger finished good expiration cost than finished good stockout cost. This scenario represents a drug that is expensive and thus expiration is costly. The stockout cost represents the per unit cost of processing additional units of finished goods outside of the initial daily batch. The stockout cost of the raw good is the highest cost because if this shortage occurs the hospital must procure the drug from the vendor (or another hospital) and pay a premium. Case 2 assumes that we incur a holding cost penalty on the raw material while Case 3 represents a less expensive drug than the Base Case. The changes from the base case are highlighted in Cases 2 and 3. The parameters in Tables 1-3 were held constant throughout the experiments.

Table 4: Simulation experiment parameters

Parameters	Base Case	Case 2	Case 3
Holding Raw	0	1	0
Expired FG	10	10	1
Stockout FG	5	5	5
Stockout Raw	20	20	20
Raw Shelf Life	infinite	infinite	infinite
FG Shelf Life	2 days	2 days	2 days

The simulation replication length was 1000 days and executed for 40 replications in order to collect and analyze the cost performance measures. The discrete event simulation utilized some of the simulation libraries developed by Telford (2008) using C#.

5.2 Inventory Policies and Results

For the cases outlined in Table 4, two types of policies are evaluated. The first is a fixed base-stock policy with parameters (r, x), called the "Fixed Policy" (FP). The FP is independent of the actual patients admitted in the hospital and is based on historical daily demand. That is we find a fixed base stock level, for each inventory stage that is implemented regardless of actual patient mix. FP is easy to implement (can be based on historical averages) and therefore attractive to pharmacy inventory managers. We found the "best" FP base stock level by running our simulation over a range of candidate levels and selecting the best performing. Second, we evaluate an inventory policy with base stock levels based on the number and type of patients admitted which we call the Adaptive Policy (AP). As found in the structure of the optimal solutions using the MDP, the Adaptive Policy also uses knowledge of patient mix to process finished goods in advance of demand and to order raw material by forecasting next period's demand. Specifically, the order up to level for the finished good in period t takes into account the patient mix in period t. Different base stock levels for the finished good were evaluated based on the expected demand plus an additional quantity which was proportional to the standard deviation. Similarly, the range of candidate stock levels for the raw material was evaluated based on the expected demand in period t+1 (which is based on the patient mix in period t).

The base stock levels for the AP may be more complicated to derive than the FP as it requires more knowledge of patient demand characteristics to determine the nonstationary base stock levels. The FP could be determined based on historical drug demand data collected by the pharmacy. We are interested in comparing the performance of these two policy types. The mean total cost (holding cost + stockout cost + expiration cost) and 95% confidence interval for both ordering policies, using the best parameters, for each case are shown in Table 5.

Table 5: Results for the base case for each policy type: Adaptive Policy and Fixed Policy

	Base Case		Case 2		Case 3	
	AP	**FP**	**AP**	**FP**	**AP**	**FP**
Mean	1524	2657	45427	115996	1524	1689
Std. Dev	182	508	1854	19252	182	188
Upper 95% CI	1582	2819	46020	122154	1582	1749
Lower 95% CI	1465	2494	44834	109838	1465	1629

As we can see from the results in Table 5, the Adaptive Policy has a statistically lower total cost than the Fixed Policy for all three cases. The policies behaved as expected as the problem's cost structure changed. For example, comparing the policy parameters of the Base Case to Case 2 for the Fixed Policy, it is more cost effective to decrease the amount of raw material held in inventory. With the change in expiration cost in Case 3, the total cost is lowered by raising the Fixed Policy's finished goods base stock level.

6 CONCLUSIONS AND FUTURE WORK

We develop a MDP-based model for determining optimal inventory policies for an in-patient hospital pharmacy. We develop a simulation model for comparing the results of a fixed inventory policy to our adaptive policy. The simulation experiments conducted in this research allowed us to

explore larger problem instances and quickly evaluate different policy structures. First, we can observe that the difference in cost performance of the two policies vary depending on the scenario's cost structure. Second, the Adaptive Policy outperformed the Fixed Policy with the defined experiment set.

For future work we want to explore the impact of setup or fixed production costs on our policy selection. These costs may be appropriate in some pharmacy environments such as those with compounding facilities. As noted in Section 5, we held the patient demand, arrival, and transition paremeters constant while varying the costs. If the patient demand is stationary throughout their stay then perhaps the Fixed Policy would be equivalent to the Adaptive Policy. Therefore a mix of policies may be appropriate, when considering the entire portfolio of pharmaceuticals, which depend on drug cost, demand variability, and demand uncertainty. Varying the parameters which represent the patient characterics are of further interest to us and is the subject of future work.

REFERENCES

Alverson C. (2003-11-01). Beyond purchasing--Managing hospital inventory. Managed Healthcare Executive. Retrieved on 2008-01-01. <http://www.managedhealthcareexecutive.com/mhe/article/articleDetail.jsp?id=75802>

Brodheim, E., C. Derman and G.P. Prastacos. 1975. On the Evaluation of a Class of Inventory Policies for Perishable Products Such as Whole Blood. Mgmt. Sci. 21, 1320-1325.

Baumer AM, Clark AM, Witmer DR, Geize SB, Vermeulen LC, Deffenbaugh JH. 2004. National survey of the impact of drug shortages in acute care hospitals. Am J Health Syst Pharm 61(19):2015-22.

Chazan D and Gal S. 1977. A markovian model for a perishable product inventory. Management Science 23(5):512-21.

Cohen MA. Analysis of single critical number ordering policies for perishable inventories. Oper Res 24(4; 4):726.

Dave U. 1979. On a discrete-in-time order-level inventory model for deteriorating items. J Oper Res Soc 30(4):349-54.

Fox ER and Tyler LS. 2003. Managing drug shortages: Seven years' experience at one health system. Am J Health Syst Pharm 60(3):245-53.

Gillerman RG and Browning RA. 2000. Drug use inefficiency: A hidden source of wasted health care dollars. Anesth Analg 91(4):921-4.

Graves SC. 1982. The application of queueing theory to continuous perishable inventory systems. Management Science 28(4):400-6.

Haijema R, vanderWal J, vanDijk N. 2007. Blood platelet production: Optimization by dynamic programming and simulation. Computers & Operations Research 34(3):760-79.

Katagiri,Hideki and Ishii,Hiroaki. 2002. Fuzzy inventory problems for perishable commodities. European Journal of Operational Research

Lian Z and Liu L. 1999b. A discrete-time model for perishable inventory systems. Annals of Operations Research 87:103-16.

Liu L and Lian Z. 1999a. (s, S) continuous review models for products with fixed lifetimes. Oper Res 47(1):150-8.

Nahmias S. 1975. Optimal ordering policies for perishable inventory--II. Oper Res 23(4; 4):735.

Nahmias S and Pierskalla WP. 1976. A two-product Perishable/Nonperishable inventory problem. SIAM J Appl Math 30(3):483-500.

Nahmias S. 1982. Perishable inventory theory: A review. Oper Res 30(4; 4):680-708.

Schmidt CP and Nahmias S. 1985. (S - 1,s) policies for perishable inventory. Management Science 31(6):719-28.

Telford, S.J. 2008. Discrete Event Simulation in C#. Master's Project Report, Integrated Manufacturing Systems, N.C. State University.

Veinott, A. F., Jr. 1960. Optimal Ordering , Issuing, and Disposal of Inventory with Known Demand, unpublished Ph.D. dissertation, Columbia University.

Weinger MB. 2001. Drug wastage contributes significantly to the cost of routine anesthesia care. J Clin Anesth 13(7):491-7.

Weiss HJ. 1980. Optimal ordering policies for continuous review perishable inventory models. Oper Res 28(2):365-74.

AUTHOR BIOGRAPHIES

ANA R. VILA-PARRISH is a PhD student in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. She previously spent four years at Dell Computers as a new product operations engineer before returning to graduate school. Her research considers the impact of the incorporation of uncertainty and variability throughout product lifecycles on inventory and logistics strategies. More recently, this framework has been applied to inventory and replenishment decisions in healthcare. Her email address is <arvila@ncsu.edu>.

JULIE SIMMONS IVY is an Assistant Professor in the Edward P. Fitts Department of Industrial and Systems Engineering and Fitts Faculty Fellow in Health Systems Engineering at North Carolina State University. She previously spent several years on the faculty of the Stephen

M. Ross School of Business at the University of Michigan. She received her B.S. and Ph.D. in Industrial and Operations Engineering at the University of Michigan. She also received her M.S. in Operations Research at Georgia Tech. Her research interests are mathematical modeling of stochastic dynamic systems as applied to health care, manufacturing, and service environments. The focus of her research is decision making under conditions of uncertainty with the objective of improving the decision quality. Her email address is <jsivy@ncsu.edu>.

RUSSELL E. KING is Professor and Director of Graduate Programs in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University (NCSU). At NCSU he also serves as the Director of the Furniture Manufacturing and Management Center and is associate faculty of the Operations Research Program and the Integrated Manufacturing and Systems Engineering Institute. He is a Fellow of the Institute of Industrial Engineers, a member of the NCSU Academy of Outstanding Teachers, and an Alumni Distinguished Graduate Professor. His research interests can be generally classified as including control and scheduling of production and logistics systems and demand activated manufacturing/sourcing systems. He received a Ph.D. in Industrial Engineering from the University of Florida. His email address is <king@ncsu.edu>.