

TECHNICAL NOTE

Optimal Control of an Assembly System with Multiple Stages and Multiple Demand Classes

Saif Benjaafar

Industrial and Systems Engineering, University of Minnesota, Minneapolis, Minnesota 55455, saif@umn.edu

Mohsen ElHafsi

The A. Gary Anderson Graduate School of Management, University of California, Riverside, California 92521, mohsen.elhafsi@ucr.edu

Chung-Yee Lee

Department of Industrial Engineering and Logistics, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, cylee@ust.hk

Weihua Zhou

Department of Management Science and Engineering, Zhejiang University, Hangzhou, China, larryzhou@zju.edu.cn

We consider an assembly system with multiple stages, multiple items, and multiple customer classes. The system consists of m production facilities, each producing a different item. Items are produced in variable batch sizes, one batch at a time, with exponentially distributed batch production times. Demand from each class takes place continuously over time according to a compound Poisson process. At each decision epoch, we must determine whether or not to produce an item and, should demand from a particular class arise, whether or not to satisfy it from existing inventory, if any is available. We formulate the problem as a Markov decision process and use it to characterize the structure of the optimal policy. In contrast to systems with exogenous and deterministic production lead times, we show that the optimal production policy for each item is a state-dependent base-stock policy with the base-stock level nonincreasing in the inventory level of items that are downstream and nondecreasing in the inventory level of all other items. For inventory allocation, we show that the optimal policy is a multilevel state-dependent rationing policy with the rationing level for each demand class nonincreasing in the inventory level of all nonend items. We also show how the optimal control problem can be reformulated in terms of echelon inventory and how the essential features of the optimal policy can be reinterpreted in terms of echelon inventory.

Subject classifications: production and inventory control; systems with multiple echelons; inventory rationing; Markov decision processes; make-to-stock queues.

Area of review: Manufacturing, Service, and Supply Chain Operations.

History: Received October 2006; revisions received April 2008, December 2008, December 2009, April 2010; accepted June 2010.

1. Introduction

Multistage assembly, one might argue, is the most common process with which physical goods are produced nowadays. Despite their pervasiveness, assembly systems are notoriously difficult to analyze and manage. Few results exist on how to optimally control production and inventory in assembly systems, especially when there is variability in either demand or production times. The difficulty of identifying optimal policies is in part due to the multidimensionality of the problem (multiple items, multiple production facilities and, in some cases, multiple demand classes). Moreover, there are interdependencies between these various dimensions: (1) demands for different items are correlated and (2) the production of one item depends on the availability of other items. These interdependencies require that production and inventory decisions be

coordinated among the different items. Coordination is, however, challenging when items have different production times that may be stochastic, different inventory holding costs, or reside in different locations within the assembly process. This can be further compounded when there are demand classes with different priorities. In that case, deciding whether or not to satisfy a demand from a particular class must take into account the ability to fulfill demand from other classes in the future.

In this paper, we consider an assembly system with multiple production and inventory stages and multiple demand classes. The system consists of multiple facilities, each producing a single item in batches of variable size, one batch at a time. Batch production times at each facility are stochastic, exponentially distributed, with production rates that can vary from item to item. Items can be produced in a make-

to-stock fashion in anticipation of future consumption. Each item, other than the end product, has one successor item but possibly multiple predecessors. One unit of each of the predecessor items is needed for the production of a successor item. Demand for the end product arises from different demand classes. Orders from each demand class take place continuously over time according to a compound Poisson process. Demand classes differ in the corresponding shortage penalties incurred if orders are not immediately satisfied from inventory. At any point in time, the system manager must decide (1) how many units, if any, of each item to produce (production decision) and (2) should a customer order be placed, how many units of that order, if any, to satisfy from available inventory (inventory allocation decision).

We formulate the problem as a Markov decision process (MDP) and characterize the structure of the optimal policy. We show that the optimal production policy for each item is a state-dependent base-stock policy, where the state of the system is specified by the vector of inventory levels for items. We show that the base-stock level for each item is nonincreasing in the inventory level of downstream items and nondecreasing in the inventory level of all other items. Moreover, the base-stock level of an item becomes less sensitive to the inventory levels of items the further downstream these items are. For inventory allocation, we show that the optimal policy is a multilevel *rationing* policy with state-dependent rationing levels. These rationing levels determine whether an order from a class is fully or partially fulfilled or not fulfilled at all. The rationing level for each class is nonincreasing in the inventory levels of all items (other than the end product).

The results presented in this paper can be extended in a variety of ways. In the full-length version of this paper (Benjaafar et al. 2008), we show that our analysis can be extended to (1) systems where the decision criterion is the average cost instead of the expected discounted cost, (2) systems with backorders instead of lost sales (in the case of a single class), and (3) systems where production times have a generalized Erlang distribution, so that the production time of an item consists of a series of stages, with each stage being exponentially distributed. In each case, we show that the optimal policy retains the same structure and the same properties described in Theorem 1. In the case of systems with Erlang production times, we describe additional properties regarding the effect of the production stage for each item on the optimal production and order fulfillment policy.

Literature dealing with optimal control policies for assembly systems is relatively limited.¹ Rosling (1989) provides a characterization of the optimal policy for a multistage assembly system with fixed assembly times. He shows, under some conditions on initial inventory, that the assembly system reduces to an equivalent serial system. The optimal policy for this serial system is known from Clark and Scarf (1960) to be an *echelon* base-stock policy

(the echelon inventory of an item consists of its own local inventory and the inventory of all downstream items). Chen and Zheng (1994) extend the result of Rosling to systems with continuous review and compound Poisson demand, and Chen (2000) extends it to systems with batch ordering. To the best of our knowledge, there are no known results for the structure of the optimal policy for systems with exogenous and *stochastic* assembly lead times.

The paper that is most closely related to ours is Benjaafar and ElHafsi (2006). They study a single-stage assemble-to-order (ATO) system with exponentially distributed component production times, Poisson demand, and multiple demand classes. They assume that assembly of the end product is instantaneous and, therefore, only components are held in stock. They show that the optimal component production policy is a state-dependent base-stock policy and the optimal inventory allocation is a multilevel rationing policy with state-dependent rationing levels. The problem we treat in this paper is different in that we consider a multistage system, allow for positive assembly time for the end product, and for the end product to be held in inventory ahead of demand. Consequently, the optimal policy has features that are absent in the single-stage ATO problem, including the difference in the effect of downstream and upstream inventories on production decisions and the sensitivity to the specific location of these inventories in the assembly process. The approach used to characterize the structure of the optimal policy relies on induction, as it does in the case of ATO, to prove that the optimal cost function satisfies certain properties. However, the number of properties in our case is significantly larger and involves notions of *modularity* more general than simple convexity or submodularity. Moreover, in contrast to the ATO case, proving the validity of these properties is less direct than those used for the ATO case and requires the introduction of additional conditions whose validity then implies the desired properties (see the proof of Lemma 1 in the electronic companion that is part of the online version at <http://or.journal.informs.org/>).

2. Problem Formulation

We consider an assembly system with multiple stages, multiple items, and one end product. The system consists of m production facilities, each producing a different item. Items may correspond to *starting components*, *intermediates*, including subassemblies, or the *end product*. Starting components are produced from material supplied by an external source, whereas intermediates are produced from other items that are themselves produced internally. We refer to the set of items needed to produce item k , $k = 1, \dots, m$, as $P(k)$, the set of *predecessors* of item k , where $P(k) = \emptyset$ if k is a starting component. We consider *pure* assembly systems where each item is needed for the production of exactly one other item. We refer to this other item as the *successor* item. The exception is the

end item, item 1, which does not have a successor. We use the notation $SS(k)$ to refer to the successor of item k for $k = 2, \dots, m$. Let $SS^r(k)$ denote the item obtained by r successive applications of the successor operator SS to item k . Hence, for every item k , $k = 2, \dots, m$, there exists an $r(k)$ such that $SS^{r(k)}(k) = 1$. In other words, successive applications of the operator SS eventually lead to the end item. We refer to any item that can be obtained via r applications of the successor operator SS , $r = 0, \dots, r(k)$, as being on the successor *path* from item k to item 1, and denote the set of such items as $S(k)$, with $SS^0(k) = k$. Special cases of the types of systems we consider include single-stage systems, serial systems, and two-stage assembly systems where components are produced in the first stage and assembled into the end product in the second stage.

Items are produced in batches of variable size one batch at a time, where batch size is a decision variable. We do not place an upper bound on the size of batches. However, the analysis can be easily extended to the case where such a bound is in effect. To produce one unit of an item, one unit from each of its predecessor items is needed. Upon production completion, items are placed in inventory. Items in inventory incur a holding cost per-unit per unit time (more about this later). Batch production times, regardless of batch size, are independent and exponentially distributed with mean μ_k^{-1} for item k . (As described in the full-length version of the paper, the analysis can be extended to more general production time distributions such as the Erlang distribution.) Each facility can hence be viewed as a single server with batch service and finite service rate μ_k . Batch processing of this kind is common in practice and includes processes such as heating, chemical treatment, certain semiconductor manufacturing processes, and robotic assembly, among many others.

Demand for the end product arises from n different demand classes. Demand from class l , $l = 1, \dots, n$, takes place continuously over time according to an independent compound Poisson process with rate λ_l . This means that interarrival times between consecutive orders from each class are exponentially distributed with mean $1/\lambda_l$. Orders may be for multiple units, with the number of units requested by each order being a random variable. More specifically, each arriving order from class l requires t units with $t = 1, \dots, q_l$, where q_l can be arbitrarily large. The probability that an order requests t units is p_l^t , independently of the size of previously placed orders from the same class or from other classes. Demand for the end product from any class can be satisfied only if there is positive inventory available for the end product. Units from any order that are not satisfied from on-hand inventory are considered lost (or must be expedited through other means such as overtime or outsourcing to a third party). A unit demand from class l that cannot be immediately fulfilled from stock incurs a *lost-sale* cost c_l per unit, which can

vary from class to class. Without loss of generality, we assume $c_1 > c_2 > \dots > c_n$.

Because the lost-sale costs can be different for different classes, it may not always be optimal to satisfy demand from a class even if there is on-hand inventory for the end product. In fact, it might be more desirable not to satisfy demand from a class, or to satisfy it only partially, in order to reserve the available inventory for future demand from a more important class (i.e., one with a higher lost-sale cost). Consequently, each time an order is placed, the system manager must decide on how many units of an order, if any, to satisfy from on-hand inventory.

In addition to deciding on how much of each order to fulfill, the system manager must decide at any point in time how many units of each item to produce. If an item is not currently being produced, this means deciding whether or not to initiate production. If the item is currently being produced, this means deciding whether or not to modify the amount being produced. Modifying the amount being produced involves either interrupting the production of some units that are currently being produced or initiating the production of additional ones. Note that an item can be produced only if there is at least one unit available in inventory for each of its predecessor items. If the production of a unit is interrupted, it can be resumed the next time the production of that unit is initiated (because of the memoryless property of the exponential distribution, resuming production from where it was interrupted is equivalent to initiating it from scratch). We assume that there are no costs associated with interrupting production. This assumption is not restrictive because, as we show in Theorem 1, it turns out that it is never optimal to interrupt production of a unit once it has been initiated.

In our model, we assume that order interarrival times and production times are exponentially distributed. This assumption is made in part for mathematical tractability because it allows us to formulate the control problem as an MDP enables us to describe the structure of an optimal policy. It is also useful in approximating the behavior of systems where variability is high. This assumption is consistent with previous treatments of production-inventory systems; see, for example, Benjaafar and ElHafsi (2006) and the references therein. Our treatment is more general than the one found in these papers because we allow for demand to be of variable order size and for production to be of variable batch size.

The state of the system at time t can be described by the vector $\mathbf{X}(t) = (X_1(t), \dots, X_m(t))$, where $X_k(t)$, $k = 1, \dots, m$, is a nonnegative integer denoting the on-hand inventory for item k at time t . We let $h(\mathbf{X}(t)) = \sum_{k=1}^m h_k(X_k(t))$, where h_k is an increasing convex function, denote the inventory holding cost rate when the state of the system is $\mathbf{X}(t)$. Note that because of the possibility of interrupting production, it is not necessary to include in the state description the number of units of each item currently being produced. Furthermore, because both order

interarrival times and production times are exponentially distributed, the system is memoryless and decision epochs can be restricted only to times when the state changes (i.e., the completion of an item or the fulfillment of an order). The memoryless property allows us to formulate the problem as an MDP and to restrict our attention to the class of Markovian policies for which actions taken at a particular decision epoch depend only on the current state of the system.

In each state, the system manager must decide on how many units of each item to produce and on how many units of an order from a particular customer class to satisfy should one arise. A control policy π specifies, for each state $\mathbf{x} = (x_1, \dots, x_m)$, an action that indicates how many units of each item to produce and how many units of each order from each class to satisfy from inventory should such orders arise (while the system is in state \mathbf{x}). Hence, for each state, a policy specifies two types of decisions: a production decision and an inventory allocation (or *rationing*) decision.

In evaluating a policy π , we consider the associated expected discounted cost (the sum of inventory holding and lost-sales costs) over an infinite planning horizon. Our analysis can be easily extended to the average cost case (also over an infinite planning horizon); see Benjaafar et al. (2008) for details. We choose to present the discounted cost case here because it forms the basis for proving results for the case of average cost. We denote by $v^\pi(\mathbf{x})$ the expected discounted cost obtained under a policy π and a starting state $\mathbf{x} = (x_1, \dots, x_m)$. Our objective is to choose a policy π^* that minimizes this cost.

The continuous-time control problem described above can be transformed into an equivalent discrete-time control problem by recognizing that control decisions can be restricted to event epochs coinciding with either production completions or order arrivals. The problem can be further simplified using the standard procedure of introducing a uniform transition rate between decision epochs given by $\beta = \sum_{l=1}^n \lambda_l + \sum_{k=1}^m \mu_k$ and by rescaling time so that $\alpha + \beta = 1$, where $0 < \alpha < 1$ is the discount rate. The detailed steps involved in this uniformization procedure are described in the online appendix; see also Bertsekas (1995, Chapter 7) and Porteus (2001, Chapter 8) for further discussion. It can then be shown that the optimal control problem can be restated as a discrete-time stochastic dynamic program whose cost function $v^* \equiv v^{\pi^*}$ satisfies the following optimality equation:

$$v^*(\mathbf{x}) = h(\mathbf{x}) + \sum_{k=1}^m \mu_k T_k v^*(\mathbf{x}) + \sum_{l=1}^n \sum_{t=1}^{q_l} p_l^l \lambda_l T_l^t v^*(\mathbf{x}), \quad (1)$$

where the operators T_k and T_l^t are defined as

$$T_k v(\mathbf{x}) = \min_{0 \leq q \leq \min_{i \in P(k)} x_i} v(\mathbf{x} + q(\mathbf{e}_k - \mathbf{e}_{P(k)})), \quad (2)$$

and

$$T_l^t v(\mathbf{x}) = \min_{0 \leq u \leq \min(t, x_l)} \{v(\mathbf{x} - u\mathbf{e}_l) + (t - u)c_l\}, \quad (3)$$

where \mathbf{e}_k is the k th unit vector of dimension m and $\mathbf{e}_{P(k)} = \sum_{j \in P(k)} \mathbf{e}_j$. In the optimality Equation (1), the first term refers to the cost incurred while in state \mathbf{x} whereas the remaining terms refer to the expected cost to go. This expected cost to go reflects all the possible states to which the system could transition, and the associated costs. The transition to a particular state depends on the decisions made while the system is in state \mathbf{x} . The operator T_k is associated with decisions about how many units of item k to produce, and the operator T_l^t is associated with decisions about how many units of an order of size t from customer class l to satisfy. Note that it is optimal to produce q units of item k if $v^*(\mathbf{x} + q(\mathbf{e}_k - \mathbf{e}_{P(k)})) \leq v^*(\mathbf{x} + q'(\mathbf{e}_k - \mathbf{e}_{P(k)}))$ for all $q' \neq q$ provided there are at least q units of on-hand inventory for each of item k 's predecessors. Similarly, it is optimal to satisfy u units from an order of size t from class l when the system is in state \mathbf{x} if $v(\mathbf{x} - u\mathbf{e}_l) + (t - u)c_l \leq v(\mathbf{x} - u'\mathbf{e}_l) + (t - u')c_l$ for all $u' \neq u$ provided there are at least u units of on-hand inventory for the end product.

3. The Structure of the Optimal Policy

In this section, we characterize the structure of an optimal policy. In order to do so, we will show that the optimal value function $v^*(\mathbf{x})$ for all states \mathbf{x} satisfies certain properties. We then show that these properties imply a specific rule for the optimal action in each state. First, let us introduce the following difference operators for functions v defined on \mathbf{Z}_+^m following difference operators for functions v defined on \mathbf{Z}_+^m (where \mathbf{Z}_+ is the set of nonnegative integers and \mathbf{Z}_+^m is its m -dimensional cross product):

$$\Delta_j v(\mathbf{x}) = v(\mathbf{x} + \mathbf{e}_j) - v(\mathbf{x}), \quad \text{for } j = 1, \dots, m,$$

and combinations of such operators, including

$$\Delta_{i,j} v(\mathbf{x}) = \Delta_i \Delta_j v(\mathbf{x}) = \Delta_j v(\mathbf{x} + \mathbf{e}_i) - \Delta_j v(\mathbf{x}),$$

where $\mathbf{x} = (x_1, \dots, x_m)$ and the variables x_j are state variables associated with an assembly system per the description in the previous section. Note that the order in which the differences are taken does not matter, i.e., $\Delta_{i,j} v(\mathbf{x}) = \Delta_{j,i} v(\mathbf{x})$. For notational convenience, we also define

$$\Delta_{j-l} v(\mathbf{x}) = v(\mathbf{x} + \mathbf{e}_j - \mathbf{e}_l) - v(\mathbf{x}),$$

and

$$\Delta_{k-P(k)} v(\mathbf{x}) = v(\mathbf{x} + \mathbf{e}_k - \mathbf{e}_{P(k)}) - v(\mathbf{x}).$$

Let \mathcal{V} denote the set of real-valued functions v defined on \mathbf{Z}_+^m that satisfy the following conditions for $i = 1, \dots, m$:

- A1:** $\Delta_{i-P(i),j}v(\mathbf{x}) \geq 0, j \in S(i);$
- A2:** $\Delta_{i-P(i),j}v(\mathbf{x}) \leq 0, j \notin S(i);$
- A3:** $\Delta_{i-P(i),j-l}v(\mathbf{x}) \geq 0, j \in S(i), j \neq l, l \in S(j), l \neq j;$
- A4:** $\Delta_{i,j}v(\mathbf{x}) \geq 0, j \in S(i);$
- A5:** $\Delta_{i-P(i),j-P(j)}v(\mathbf{x}) \leq 0, j \neq i; \text{ and}$
- A6:** $\Delta_1v(\mathbf{x}) \geq -c_1.$

These conditions, which describe the monotonicity (nondecreaseness or nonincreaseness) of the difference operators along several dimensions, are related to the notions of convexity, concavity, supermodularity, and submodularity; see, for example, Topkis (1998). In Lemma 1 below, we show that the optimal value function v^* satisfies these conditions. When applied to the optimal cost function, these conditions have simple intuitive interpretations and imply a specific structure for the optimal policy:

- Condition A1 implies that producing item i is economically less desirable when the inventory level of item j is higher and item j is on the successor path of item i .
- Condition A2, in contrast, implies that producing item i is economically more desirable when the inventory level of item j is higher and item j is not on the successor path of item i .
- Condition A3 implies that the economic value of producing an additional unit of item i does not increase with a simultaneous increase in the inventory level of item j and a decrease in the inventory level of item l , for $j \in S(i)$, and $l \in S(j)$.
- Condition A4 specifies that the economic value of having one more unit of item i does not increase as the inventory level of item j increases for $j \in S(i)$. In other words, the optimal cost function is supermodular in x_i and x_j . The special case of $i = j$ corresponds to convexity in x_i .
- Condition A5 specifies that the economic value of producing any item is nondecreasing with the production completion of any other item.
- Condition A6 implies that it is always preferable to reduce inventory by one unit rather than reject a unit demand from class 1 and incur the corresponding lost-sales cost c_1 .

LEMMA 1. *If $v \in \mathcal{V}$, then $Tv \in \mathcal{V}$, where $Tv = h(\mathbf{x}) + \sum_{k=1}^m \mu_k T_k v(\mathbf{x}) + \sum_{l=1}^n p_l^q \sum_{i=1}^q p_i^l \lambda_i T_i^l v(\mathbf{x})$. Furthermore, the optimal cost function v^* is an element of \mathcal{V} . That is, v^* satisfies conditions A1–A6.*

The proof of Lemma 1 and all other subsequent results in this paper can be found in the online appendix. We are now ready to state the main result of this paper. First, let us introduce the notation $\mathbf{x}_{-k} = (x_1, \dots, x_{k-1}, x_{k+1}, \dots, x_m)$ to denote an $(m - 1)$ dimensional vector with elements corresponding to the on-hand inventory levels of items $j \neq k$. Also, we define $s_k^*(\mathbf{x}_{-k}) = \min\{x_k \geq 0 \mid \Delta_{k-P(k)}v^*(\mathbf{x}) \geq 0\}$, and $r_l^*(\mathbf{x}_{-1}) = \min\{x_1 \geq 0 \mid \Delta_1v^*(\mathbf{x}) \geq -c_1\}$.

THEOREM 1. *There exists an optimal stationary policy that can be specified as follows.*

- *The optimal production policy for item $k, k = 1, \dots, m$, is a base-stock policy with a state-dependent base-stock level $s_k^*(\mathbf{x}_{-k})$, such that it is optimal to produce a quantity that brings the inventory level of item k as close as possible to $s_k^*(\mathbf{x}_{-k})$, but without exceeding it, if $x_k < s_k^*(\mathbf{x}_{-k})$, and not to produce at all otherwise. That is, the optimal production quantity is given by*

$$\min\left\{(s_k^*(\mathbf{x}_{-k}) - x_k)^+, \min_{t \in P(k)}(x_t)\right\}.$$

- *The optimal inventory allocation policy is a multilevel rationing policy with a vector of state-dependent rationing levels $\mathbf{r}^*(\mathbf{x}_{-1}) = (r_1^*(\mathbf{x}_{-1}), \dots, r_n^*(\mathbf{x}_{-1}))$ such that it is optimal to satisfy as many units as possible of an order from class $l, l = 1, \dots, n$, as long as the inventory level does not drop below $r_l^*(\mathbf{x}_{-1})$. That is, the optimal number of units to satisfy from an order of size t from class l is given by $\min\{x_l - r_l^*(\mathbf{x}_{-1}), t\}$ if $x_l \geq r_l^*(\mathbf{x}_{-1})$; the entire order is rejected otherwise.*

The optimal policy has the following additional properties:

PROPERTY 1 (P1). $s_k^*(\mathbf{x}_{-k})$ is nonincreasing in x_j if $j \in S(k)$ and $j \neq k$ and is nondecreasing in x_j if $j \notin S(k)$.

PROPERTY 2 (P2). $s_k^*(\mathbf{x}_{-k} + \mathbf{e}_j - \mathbf{e}_{P(j)}) \geq s_k^*(\mathbf{x}_{-k})$ for $j \neq k$.

PROPERTY 3 (P3). $s_k^*(\mathbf{x}_{-k} + \mathbf{e}_j) \leq s_k^*(\mathbf{x}_{-k} + \mathbf{e}_l)$ for $j \in S(k)$ and $l \in S(j), l \neq j$.


PROPERTY 4 (P4). *It is never optimal to interrupt the production of an item once it has been initiated.*

PROPERTY 5 (P5). $r_l^*(\mathbf{x}_{-1})$ is nonincreasing in each x_j for $j \neq l$.

PROPERTY 6 (P6). $r_n^*(\mathbf{x}_{-1}) \geq \dots \geq r_1^*(\mathbf{x}_{-1}) = 0$.

For production, the results of the theorem show that it is determined for each item by a base-stock level that depends on the inventory level of all other items. Property P1 shows that the base-stock level for an item does not increase with an increase in the inventory level of items that are on its successor path; that is, inventory in *downstream* stages can be viewed as a substitute for inventory in *upstream* stages. On the other hand, the base-stock level for an item does not decrease with an increase in the inventory level of items that are not on its successor path; in this case, the inventories can be viewed as complements. For example, for items that are assembled together at a particular stage having less of one item reduces the need for producing the other items, because eventually one unit from each item will be needed. Property P2 shows that the base-stock level of an item does not decrease with the production completion of any other item. Property P3 implies that the base-stock level of any item is more likely to decrease with an increase in the inventory level of items on its successor path that are more closely located (i.e., items that are further upstream on the successor path). For inventory allocation, the results of the theorem show that different demand classes should

Figure 1. The impact of downstream inventory levels on the base-stock level.



Notes. $\mu_1 = \mu_2 = \mu_3 = 1$; $\lambda = 0.95$; $c = 200$, $h_1 = 2$, $h_2 = 1.5$, $h_3 = 1$, $q_l = b_k = 1$ for all k, l .

be treated differently, with each class assigned an inventory rationing level below which the demand from this class would be rejected in favor of reserving inventory for classes with higher lost-sales costs. Similar to the base-stock levels, the rationing levels depend on the inventory level of all items. Property P5 shows that the rationing level of each class is nonincreasing in the inventory level of each item. Property P6 shows that the rationing levels are ordered, with demand from class 1 always fulfilled as long as there is on-hand inventory.

In Figure 1, we illustrate the fact that the base-stock level of each item becomes less sensitive to the inventory level of other items the further downstream these items are (Property P3 in Theorem 1). In the example, we consider a serial system with three items and a single customer class, with item 3 being the starting item. As we can see, an increase in x_1 has significantly less of an impact on the base-stock level of item 3 than an increase in x_2 . This result shows that it is not sufficient to monitor only the sum of downstream inventory levels in making decisions about whether or not to produce. Instead, the specific contribution to total inventory from each item matters. This is different from results for other types of inventory systems, such as systems with fixed supply lead times, where an echelon base-stock policy with fixed echelon base-stock levels has been shown to be indeed optimal; see, for example, Rosling (1989).

4. An Echelon Inventory Reformulation

Although the optimal production policy for items is not an echelon base-stock policy with fixed echelon base-stock levels, the problem can still be reformulated in terms of

echelon inventory, and the optimal policy can be reinterpreted in terms of echelon inventory. In this section, we show how such a reformulation can be carried out and describe an alternative approach, using echelon inventory, to characterize the structure of the optimal policy. We do so for two reasons: (1) to gain additional insight into the structure of the optimal policy and (2) to show that the results obtained under the local inventory formulation can be more informative about the structure of the optimal policy than those obtained under the echelon inventory formulation.

Let \hat{x}_k be the echelon inventory of item k , so that $\hat{x}_k = \sum_{t \in S(k)} x_t$. Then, because there is a one-to-one correspondence between each vector of echelon inventory $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_m)$ and each vector of local inventory $\mathbf{x} = (x_1, \dots, x_m)$, we can fully represent the state of the system in terms of echelon inventory. Using echelon inventory, the optimality equation in (1) can be rewritten as

$$v^*(\hat{\mathbf{x}}) = \hat{h}(\hat{\mathbf{x}}) + \sum_{k=1}^m \mu_k T_k v^*(\hat{\mathbf{x}}) + \sum_{l=1}^n \sum_{t=1}^{q_l} p_t^l \lambda_t T_t^l v^*(\hat{\mathbf{x}}), \quad (4)$$

where the operators \hat{T}_k and \hat{T}_t^l are defined as

$$\hat{T}_k v(\hat{\mathbf{x}}) = \min_{0 \leq q \leq \min_{i \in P(k)} (\hat{x}_i - \hat{x}_k)} v(\hat{\mathbf{x}} + q\mathbf{e}_k), \quad (5)$$

$$T_t^l v(\hat{\mathbf{x}}) = \min_{0 \leq u \leq \min(t, \hat{x}_1)} \{v(\hat{\mathbf{x}} - u\mathbf{e}) + (t - u)c_l\}, \quad (6)$$

with $\mathbf{e} = (1, \dots, 1)$ being an m -dimensional vector with ones in all positions, and

$$\hat{h}(\hat{\mathbf{x}}) = \sum_{i=1}^n h(\hat{x}_i - \hat{x}_{SS(i)}). \quad (7)$$

Equation (7) follows from the fact that $x_k = \hat{x}_k - \hat{x}_{SS(k)}$, which leads to

$$\hat{h}(\hat{\mathbf{x}}) = \sum_{i=1}^n h(\hat{x}_i - \hat{x}_{SS(i)}) = \sum_{i=1}^n h(x_i) = h(\mathbf{x}).$$

Based on the above reformulation, we can make the following important observations: (1) the decision to produce item k does not change the echelon inventory level of any other item and affects only the echelon inventory level of the item itself, and (2) the decision to fulfill demand from on-hand inventory reduces the echelon inventory level of all items by an equal amount. This formulation is strikingly similar to the formulation of a single-stage ATO system with m items, where each unit of demand is fulfilled by one unit of inventory from each item, and the production of an item affects only its own inventory level (see, for example, Benjaafar and ElHafsi 2006). Here, echelon inventory plays the same role as local inventory in the ATO case. In fact, we can show that features of the optimal policy can be shown using conditions similar to those used to show the structure of the optimal policy in the case of ATO.

In particular, we can show that (1) the optimal production policy is a state-dependent echelon base-stock policy by showing that the optimal cost function is convex in the echelon inventory level of each item (see the online appendix for a proof for this and all other results in this section). That is, $\Delta_{i,i}v(\hat{\mathbf{x}}) \geq 0$, which means that the economic value of increasing the echelon inventory of any item i , while keeping the echelon inventory of all other items fixed, is nondecreasing in the echelon inventory level of that item. If we define $\hat{s}_k^*(\hat{\mathbf{x}}_{-k}) = \min\{\hat{x}_k \geq 0 \mid \Delta_k v^*(\hat{\mathbf{x}}) \geq 0\}$, then, by virtue of the convexity of v^* in $\hat{\mathbf{x}}_k$, it is optimal to produce units of item k as long as the resulting echelon inventory level of item k does not exceed $\hat{s}_k^*(\hat{\mathbf{x}}_{-k})$. Furthermore, we can show that (2) the optimal echelon base-stock level of any item i is nondecreasing in the echelon inventory level of any other item by showing the optimal cost function is submodular in \hat{x}_i and \hat{x}_j (i.e., $\Delta_{i,j}v(\hat{\mathbf{x}}) \leq 0$ for all i and all $j \neq i$). In other words, echelon inventories are all complements of one another. Note that this is different from the results obtained with the local inventory formulation where the optimal base-stock level for an item can be either nondecreasing or nonincreasing in the inventory level of other items depending on whether or not these items are on the successor path.

Similarly, we can show that (3) the optimal order fulfillment policy is determined by state-dependent rationing levels by showing that $\Delta_{e,1}v(\hat{\mathbf{x}}) \geq 0$. If we define $\hat{r}_l^*(\hat{\mathbf{x}}_{-1}) = \min\{\hat{x}_1 \geq 0 \mid \Delta_e v^*(\hat{\mathbf{x}}) \geq -c_l\}$, then it is optimal to fulfill demand from class l as long as the echelon inventory level of the end item (which also corresponds to the local inventory of the end item) does not drop below $\hat{r}_l^*(\hat{\mathbf{x}}_{-1})$. More generally, we can show that (4) the optimal rationing level for class l is nonincreasing in the echelon inventory of any item by showing $\Delta_{e,j}v(\hat{\mathbf{x}}) \geq 0$ for all j . Finally, by showing that $\Delta_e v(\hat{\mathbf{x}}) \geq -c_1$, we can show that (5) it is always optimal to fulfill demand from class 1 as long as there is positive echelon inventory for all items, or equivalently there is positive on-hand inventory for the end item. Note that given that there is a one-to-one correspondence between each vector of echelon inventory $\hat{\mathbf{x}} = (\hat{x}_1, \dots, \hat{x}_m)$ and each vector of local inventory $\mathbf{x} = (x_1, \dots, x_m)$, the optimal policy specified in terms of echelon inventory, if interpreted in terms of local inventory, is the same as the one specified originally in terms of local inventory.

As we can see, using the echelon reformulation, we can retrieve the essential features of the optimal policy. Namely, there are state-dependent thresholds that determine whether or not production takes place and whether or not demand is fulfilled and that these thresholds are monotonic in the echelon inventory levels of individual items. However, the structural properties are less informative than those obtained with the local inventory formulation. In particular, the monotonicity results do not differentiate between the effect of upstream and downstream inventory (as described in Property P1 of Theorem 1). They also do not distinguish

between the effects of different inventory levels depending on how far downstream they are from a particular item (as described in Property P3 of Theorem 1). This is because under the echelon formulation, an increase in the echelon inventory level of an item, given that the echelon inventory levels of all other items remain the same, corresponds to the production completion of that item (or equivalently to an increase in the local inventory of the item that is simultaneous with a decrease in the inventory level of its predecessor items). The submodularity property $\Delta_{i,j}v(\hat{\mathbf{x}}) \leq 0$ is in fact equivalent to property A5 in the local inventory formulation. Thus, it is not surprising that the echelon base-stock level of an item is nondecreasing in the echelon inventory level of any other item (this corresponds to property P2 in Theorem 1). To construct echelon properties equivalent to A1, A2, and A3, which would then imply properties similar to those of P1 and P3, is more difficult because increasing the local inventory level of an individual item, while keeping the local inventory of all other items constant, corresponds (under the echelon formulation) to increasing the echelon inventory of all upstream items. This means that the conditions on the optimal cost function that would have to be verified are more complicated and require evaluating cost differences involving simultaneous increases in the inventory levels of multiple items. In short, a local inventory formulation appears to be more effective in settings such as ours where distinguishing the effect of individual local inventory levels is important. The needed conditions to show structure under an echelon formulation appear more difficult to construct and perhaps less intuitive to identify initially.

5. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Endnote

1. We refer the reader to the excellent review of literature on assembly systems in Song and Zipkin (2003) and also to Zipkin (2000, Chapter 8), Axsater (2006, Chapter 5), and Chaouiya et al. (2000) for discussion of important results.

Acknowledgments

The authors are grateful to Area Editor Jeannette Song, the associate editor, and two anonymous referees for many helpful comments and suggestions. The corresponding author is Weihua Zhou. Research of the authors has been supported by grants from the U.S. National Science Foundation, the Hong Kong Research Grants Council, and the Natural Sciences Foundation of China.

References

- Axsater, S. 2006. *Inventory Control*, 2nd ed. Springer, New York.
- Benjaafar, S., M. ElHafsi. 2006. Production and inventory control of a single product assemble-to-order system with multiple customer classes. *Management Sci.* **52**(12) 1896–1912.
- Benjaafar, S., M. ElHafsi, C. Y. Lee, W. Zhou. 2008. Optimal control of an assembly system with multiple stages and multiple demand classes—The unabridged version. Working paper, University of Minnesota, Minneapolis.
- Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control*, Vol. 2, 2nd ed. Athena Scientific, Nashua, NH.
- Chaouiya, C., G. Liberopoulos, Y. Dallery. 2000. The extended kanban control system for production coordination of assembly manufacturing systems. *IIE Trans.* **32**(10) 999–1012.
- Chen, F. 2000. Optimal policies for multi-echelon inventory problems with batch ordering. *Oper. Res.* **48**(3) 376–389.
- Chen, F., Y.-S. Zheng. 1994. Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* **40**(11) 1426–1443.
- Clark, A. J., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6**(4) 475–490.
- Porteus, E. L. 2001. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Palo Alto, CA.
- Rosling, K. 1989. Optimal inventory policies for assembly systems under random demands. *Oper. Res.* **37**(4) 565–579.
- Song, J. S., P. Zipkin. 2003. Assemble-to-order systems. T. de Kok, S. Graves, eds. *Supply Chain Management*. Handbooks in Operations Research and Management Science, Vol. 30. North-Holland, Amsterdam.
- Topkis, D. 1998. *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.
- Zipkin, P. H. 2000. *Foundations of Inventory Management*. McGraw-Hill, New York.