

OPTIMAL POLICIES FOR MULTI-ECHELON INVENTORY PROBLEMS WITH BATCH ORDERING

FANGRUO CHEN

Graduate School of Business, Columbia University, New York, New York 10027, fc26@columbia.edu

(Received August 1996; revisions received June 1997, January 1998, May 1998; accepted May 1998)

In many production/distribution systems, materials flow from one stage to another in fixed lot sizes. For example, a retailer orders a full truckload from a manufacturer to qualify for a quantity discount; a factory has a material handling system that moves full containers of parts from one production stage to the next. In this paper, we derive optimal policies for multi-stage serial and assembly systems where materials flow in fixed batches. The optimal policies have a simple structure, and their parameters can be easily determined. This research extends the multi-echelon inventory theory in several ways. It generalizes the Clark-Scarf model by allowing batch transfers of inventories. Rosling (1989) shows that assembly systems can be interpreted as serial systems under the assumption that there are no setup costs. We show that the series interpretation still holds when materials flow in fixed batches which satisfy a certain regularity condition. Finally, Veinott (1965) identifies an optimal policy for a single-location inventory system with batch ordering. This paper generalizes his result to multi-echelon settings.

1. INTRODUCTION

In many production/distribution systems, materials flow from one stage to another in fixed lot sizes. For example, a retailer orders a full truckload from a manufacturer to qualify for a quantity discount; a factory has a material handling system that moves full containers of parts from one production stage to the next. In this paper, we derive optimal policies for two multi-echelon problems where materials flow in fixed batches.

The first problem is a serial system with N stages. Customer demand arises periodically at stage 1, stage 1 orders from stage 2, 2 from 3, etc., and stage N orders from an outside supplier with unlimited stock. The demands in different periods are independent and identically distributed. When demand exceeds the on-hand inventory at stage 1, the excess is backlogged. Each stage can only order a fixed, stage-specific quantity or multiples thereof. The base quantities at different stages are coordinated in the sense that they satisfy an integer-ratio constraint, i.e., the base quantity at stage i is an integer multiple of the base quantity at stage $i - 1$. The production-transportation leadtimes from one stage to the next are constant. The system incurs holding and backorder costs. The objective is to minimize the long-run average total cost in the system.

For the above serial problem, the optimal policy for each stage is a reorder-point policy; i.e., whenever its echelon stock (the inventory position of the subsystem consisting of the stage and all its successor stages) falls to or below a reorder point, order a minimum integer multiple of the base quantity to increase the echelon stock to above the reorder point. Note that this is precisely the echelon-stock (R, nQ) policy where R is the reorder point and Q the base quantity. It can also be seen as a generalized base-stock policy; i.e.,

each period an order is placed to keep the echelon stock within an interval of length Q . (Thus, a *generalized* base-stock policy orders up to an interval, while an *ordinary* base-stock policy orders up to a point.)

The second problem considered in this paper is an assembly system that produces a single end item from several components. The assembly operation consists of multiple steps involving a total of N distinct items (components, subassemblies, and the end item). The assembly structure is a tree; its root represents the end item and its leaves are the components. Customer demand is for the end item only, with complete backlogging. The demands in different periods are also independent and identically distributed. The production batch for each item is restricted to be an item-specific base quantity or multiples thereof. The procurement leadtimes for the components as well as the production leadtimes for the subassemblies and the end item are all constant. The objective is to minimize the long-run average holding and backorder costs in the system.

Under one condition, the above assembly system is equivalent to a serial system with N stages. This requires rank ordering the N items of the assembly system according to their total leadtimes. The total leadtime for item i is the sum of the leadtimes of the item and all its successor items. Thus the end item is item 1, and the item with the longest total leadtime is item N . Now consider an N -stage serial system where stage i is associated with item i for all i . If the base quantities in this serial system satisfy the integer-ratio constraint; i.e., the base quantity of item i is an integer multiple of the base quantity of item $i - 1$ for all i , then the assembly system is equivalent to the serial system. Under this condition, the echelon-stock (R, nQ) policy, when modified slightly, remains optimal for the assembly system.

Subject classifications: Inventory/production: multi-stage/item, stochastic, fixed lot-sizes, optimal policies. Manufacturing: assembly systems.
Area of review: LOGISTICS AND SUPPLY CHAIN OPERATIONS.

The above serial model is a generalization of the Clark-Scarf model that effectively assumes a base quantity of one for every stage. It is perhaps fair to say that the Clark-Scarf model is now well understood: Clark and Scarf (1960) show that the base-stock policy is optimal in a finite-horizon setting, Federgruen and Zipkin (1984) extend this result to the infinite-horizon case with both discounted and average costs, and Chen and Zheng (1994b) provide a simple way to show the optimality of the base-stock policy. This paper suggests that the base-stock policy, when modified to accommodate the base order quantities, is still optimal when every stage orders in batches.

Compared with serial systems, assembly systems with stochastic demand have attracted relatively less attention in the literature. Schmidt and Nahmias (1985) characterize an optimal policy for a system where two components are assembled into one end item. Rosling (1989) shows that a general assembly system (with any number of items) is equivalent to a serial system, and thus the Clark-Scarf result prevails. Both papers assume zero setup costs. This paper extends this equivalence between assembly and serial systems to the batch-ordering case.

Many have studied the (R, nQ) policy. For single-location inventory systems, Hadley and Whitin (1961) show that the inventory position is uniformly distributed; Veinott (1965) shows that the policy is optimal when the order quantity is restricted to be integer multiples of Q ; and Zheng and Chen (1992) provide an optimization algorithm and a sensitivity analysis. On the other hand, the policy has been proposed as a reasonable heuristic for serial systems (e.g., De Bodt and Graves 1985, Axsater and Rosling 1993, and Chen and Zheng 1994a) as well as distribution systems (e.g., Axsater 1993a, Cachon 1995, and Chen and Zheng 1997). However, the (R, nQ) policy has received much less attention in the context of assembly systems, except a brief mention in Axsater and Rosling (1993). To this literature we make two contributions: a generalization of Veinott's optimality result to multi-echelon systems and a demonstration that a heuristic policy widely studied in the literature is actually optimal under some plausible scenarios. (If demand is for a single unit at a time, then the (R, nQ) policy degenerates into an (R, Q) policy whereby each order is exactly Q units. We refer the reader to Axsater (1993b) for a review of studies on the (R, Q) policy in various multi-echelon systems.)

Finally, we all know that it is extremely difficult to characterize an optimal policy for a multi-echelon, stochastic inventory system with setup costs at all echelons (see Clark and Scarf 1962 for an example). The existence of setup costs implies that replenishment must be carried out in batches. If we ignore the setup costs but insist that every stage order in fixed quantities, then we have a new formulation of the problem. This paper demonstrates that under this new formulation, it is sometimes tractable to identify an optimal policy. (Although fixed order quantities may not be a perfect substitute for the setup costs, they can accommodate aspects that are not captured in the setup costs, e.g., the convenience of standardized shipments.)

The rest of the paper has four sections. Section 2 states a key observation and uses a simple example to illustrate the methodology to be used throughout the paper. Section 3 characterizes an optimal policy for the serial system. Parallel to §3, §4 deals with assembly systems. Section 4 extends the results to continuous-time models, discusses a key assumption made for assembly systems, and poses a few open questions.

2. PRELIMINARIES

This section sets the stage for the later ones. We first present an observation that is useful in characterizing optimal policies for inventory problems with batch ordering. This observation is then used to re-derive an existing optimal policy for a single-location problem with batch ordering. Our purpose is to familiarize the reader with the methodology used in this paper.

Let $\mathfrak{I}$ be the set of integers and $\mathfrak{R}$ the set of real numbers.

LEMMA 1. *Let $G(\cdot): \mathfrak{I} \rightarrow \mathfrak{R}$ be a function and Q a positive integer. Define $\bar{G}(y) = \sum_{x=1}^Q G(y+x)$, $y \in \mathfrak{I}$. Suppose that $\bar{G}(y)$ is quasiconvex and minimized at $y=R$ which is a finite integer.*

(i) *For any fixed $z \in \mathfrak{I}$, $G(z+xQ)$ is quasiconvex in $x \in \mathfrak{I}$. Let x_z be the unique integer so that $R+1 \leq z+x_zQ \leq R+Q$. Then $G(z+xQ)$, as a function of $x \in \mathfrak{I}$, is minimized at $x=x_z$.*

(ii) *For any $x \in \mathfrak{I}$, define*

$$O[x] = \begin{cases} x, & x \leq R+Q, \\ x-nQ, & x > R+Q, \end{cases}$$

where n is the largest integer so that $x-nQ > R$. Then $\sum_{x=1}^Q G(O[y+x]) = \bar{G}(\min\{R, y\})$ which is quasiconvex and nonincreasing in $y \in \mathfrak{I}$.

PROOF. (i) Take any $z \in \mathfrak{I}$. Consider $G(z+xQ)$ as a function of $x \in \mathfrak{I}$. Note that

$$\begin{aligned} G(z+(x+1)Q) - G(z+xQ) \\ = \bar{G}(z+xQ) - \bar{G}(z+xQ-1). \end{aligned} \quad (1)$$

Take any $x < x_z$. Thus $z+xQ \leq R$. From (1) and the quasiconvexity of $\bar{G}$, $\bar{G}(z+xQ) - \bar{G}(z+xQ-1) \leq 0$ or $G(z+(x+1)Q) \leq G(z+xQ)$ for all $x < x_z$. Similarly, $G(z+(x+1)Q) \geq G(z+xQ)$ for all $x \geq x_z$. Therefore, $G(z+xQ)$ is quasiconvex in x and is minimized at $x=x_z$.

(ii) If $y \leq R$ then $O[y+x] = y+x$ for $x=1, \dots, Q$, which implies $\sum_{x=1}^Q G(O[y+x]) = \bar{G}(y)$. Now take any $y > R$. Note that for each $x=1, \dots, Q$, $O[y+x]$ is a unique and different point in $\{R+1, \dots, R+Q\}$. Thus $\{O[y+x], x=1, \dots, Q\} = \{R+1, \dots, R+Q\}$ or $\sum_{x=1}^Q G(O[y+x]) = \bar{G}(R)$. Combining these two cases, $\sum_{x=1}^Q G(O[y+x]) = \bar{G}(\min\{R, y\})$. Because $\bar{G}(\cdot)$ is quasiconvex and R is its minimum point, $\bar{G}(\min\{R, y\})$ is quasiconvex and nonincreasing in y . $\square$

Next, we use the above result to re-derive an optimal policy characterized by Veinott (1965). Consider a single-location, periodic-review inventory problem with independent and identically distributed demands. Assume that the demands only take integer values. When demand exceeds the on-hand inventory, the excess is backlogged. Each order placed must be a positive integer multiple of Q , which is itself a fixed positive integer. Let $G(y)$ be the (conditional) expected cost in a period given that the inventory position at the beginning of the period after ordering is y . (The inventory position is the on-hand inventory plus outstanding orders minus backorders.) The planning horizon is infinite, and the objective is to minimize the long-run average total cost. (The same approach can be used for the infinite-horizon, discounted case or the finite-horizon case with or without discounts.)

Assume that $\bar{G}(y)$, $y \in \mathfrak{I}$, is quasiconvex and it is minimized at $y = R$, a finite integer. Thus Lemma 1 applies. To determine an optimal policy for the above inventory problem, we first show that the long-run average cost of any feasible policy is bounded below by a constant. We then construct a feasible policy that actually achieves the lower bound. Thus the constructed feasible policy is optimal.

Consider any feasible policy. Take any period t , and let y_t be the inventory position at the beginning of period t after ordering. This period's expected cost is $G(y_t)$. From Lemma 1, $G(y_t) \geq G(y'_t)$ where $y'_t \stackrel{\text{def}}{=} y_t + nQ$ where n is the unique integer (positive or otherwise) so that $y'_t \in \{R+1, \dots, R+Q\}$. The long-run average cost of the feasible policy must be greater than or equal to the long-run average value of $G(y'_t)$. To determine the latter, consider the stochastic process $\{y'_t\}$. Let D_t be the demand in period t . Note that under any feasible policy,

$$y_{t+1} = y_t - D_t + mQ$$

for some nonnegative integer m . Because $y_t - y'_t$ and $y_{t+1} - y'_{t+1}$ are both integer multiples of Q ,

$$y'_{t+1} = y'_t - D_t + m'Q,$$

where m' is an integer. Moreover, given y'_t and D_t , the value of m' is unique because $R+1 \leq y'_{t+1} \leq R+Q$. Because the demands in different periods are independent, $\{y'_t\}$ is a Markov chain. And it has a finite state space, i.e., $\{R+1, \dots, R+Q\}$. When the demand distribution satisfies some mild conditions, the steady-state distribution of the Markov chain is uniform. In this case, the long-run average value of $G(y'_t)$ is

$$\frac{1}{Q} \sum_{y=R+1}^{R+Q} G(y),$$

which is a lower bound on the long-run average cost of any feasible policy. (One sufficient condition for the uniform distribution is that the demand in each period takes on the value 1 with positive probability. This is indeed a mild condition satisfied by most demand distributions. Let

us make this assumption for the sake of brevity. However, the uniform distribution—possibly over a subset of the state space—holds for general demand distributions. The results in this paper can be easily extended to general demand distributions; see the appendix.)

The above lower bound can be achieved if the inventory manager follows the (R, nQ) policy. It operates as follows. At the beginning of each period, if the inventory position is at or below R , the manager orders a minimum integer multiple of Q to raise the inventory position to above R ; otherwise, no order is placed. It is easy to verify that under this policy, $y'_t = y_t$ for all t and thus the long-run average cost is equal to the lower bound. As a result, the (R, nQ) policy is optimal. This is precisely the policy characterized by Veinott (1965) for the above single-location inventory problem.

Note that our assumption—i.e., $\bar{G}(\cdot)$ is quasiconvex—is slightly more general than Veinott's original assumption that $G(\cdot)$ is quasiconvex. The above analysis suggests that the optimality of (R, nQ) policies is a direct consequence of the shape of $\bar{G}(\cdot)$, not $G(\cdot)$. When linear holding and backorder costs are incurred in each period, it is well known that $G(\cdot)$, thus $\bar{G}(\cdot)$, is quasiconvex. The same still holds for certain types of nonlinear backorder costs, see Chen and Zheng (1993).

3. SERIAL SYSTEMS

In this section, we characterize an optimal policy for serial inventory systems with batch ordering. The methodology is essentially the same as the one used in the previous section for single-location systems.

Consider a serial inventory system with N stages where customer demand arises at stage 1 only, stage 1 replenishes its inventory from stage 2, 2 from 3, etc., and stage N from an outside supplier with unlimited stock. (The outside supplier is also called stage $N+1$.) Time is divided into equal intervals called periods. We assume that the demands in different periods are independent, identically distributed, discrete random variables. When the customer demand exceeds the on-hand inventory at stage 1, the excess is backlogged. Each order placed by stage i must be a positive integer multiple of a stage-specific base quantity, Q_i , $i = 1, \dots, N$. These base quantities represent the basic units of transportation (e.g., a full truckload) or production (e.g., a pallet). We assume that they satisfy the following integer-ratio constraint:

$$Q_{i+1} = n_i Q_i, \quad i = 1, \dots, N-1, \quad (2)$$

where n_i is a positive integer. The transportation/production leadtimes from one stage to the next are fixed. The system incurs holding and backorder costs. The objective is to minimize the long-run average total cost in the system.

Define the *echelon inventory level* at stage i to be the inventories on hand at stages $1, \dots, i$ plus inventories in transit to stages $1, \dots, i-1$ minus backorders at stage 1. In short, the echelon inventory level at stage i is the *net* inventory level in the subsystem consisting of stages $1, \dots, i$. Define

the *echelon inventory position* at stage i to be the echelon inventory level at stage i plus inventories in transit to stage i .

We assume that all the replenishment activities in a period occur at the beginning of the period. At stage $i > 1$, they occur in the following sequence: An order from stage $i - 1$ is received, an order is placed with stage $i + 1$, a shipment is received from stage $i + 1$, and a shipment is sent to stage $i - 1$. For stage 1, order placement occurs at the beginning of the period, while customer demand arrives during the period. Let t^- be the beginning of period t after all the replenishment activities in the period have taken place. Let t^+ be the end of period t (after demand occurrence). Define

$IP_i(t)$ = echelon inventory position at stage i at t^- ,

$IL_i^-(t)$ = echelon inventory level at stage i at t^- ,

$IL_i(t)$ = echelon inventory level at stage i at t^+ ,

$B(t)$ = backorder level at stage 1 at t^+ .

We assume that the system starts with a plausible initial state, i.e., the initial on-hand inventory at stage $i + 1$ is a nonnegative integer multiple of Q_i , $i = 1, \dots, N - 1$. This is reasonable because the order size by stage i is always an integer multiple of Q_i , and thus there is no incentive to keep a fraction of Q_i at stage $i + 1$. This initial condition, together with the integer-ratio condition (2) that implies that both the in-flow and the out-flow at stage $i + 1$ are integer multiples of Q_i , suggests that the on-hand inventory at stage $i + 1$ is always an integer multiple of Q_i . Because the on-hand inventory at stage $i + 1$ at t^- is $IL_{i+1}^-(t) - IP_i(t)$,

$$IL_{i+1}^-(t) - IP_i(t) = mQ_i, \quad i = 1, \dots, N - 1, \quad (3)$$

where m is a nonnegative integer.

Let $g_i(y)$ be the expected one-period cost incurred at stage i in period t , given $IP_i(t) = y$. The expected total cost in the system in period t , given the echelon inventory positions at all stages, is thus

$$\sum_{i=1}^N g_i(IP_i(t)). \quad (4)$$

Later, we will provide specific expressions for the $g_i(\cdot)$ under a plausible holding-backorder cost structure.

We proceed to establish a lower bound on the long-run average value of (4). It is convenient to express (4) in a time-shifted manner. Let L_i , a nonnegative integer, be the lead-time from stage $i + 1$ to i . Let M_i be the total leadtime at stage i , i.e., $M_i = \sum_{j=1}^i L_j$. Take any period t . Write IP_i for $IP_i(t - M_i)$, $i = 1, \dots, N$, and IL_i^- for $IL_i^-(t - M_{i-1})$, $i = 2, \dots, N$. The following is (4) shifted in time:

$$\sum_{i=1}^N g_i(IP_i). \quad (5)$$

Clearly, (4) and (5) have the same long-run average value.

Let $D(t)$ be the customer demand in period t , a discrete random variable. Let $D[t_1, t_2]$ be the total demand in periods

$t_1, \dots, t_2 - 1$. For $i = 2, \dots, N$, write D_i for $D[t - M_i, t - M_{i-1}]$. Because the demands in different periods are independent, $D_2, \dots, D_N$ are independent. The following are well-known inventory balance equations:

$$IL_i^- = IP_i - D_i, \quad i = 2, \dots, N. \quad (6)$$

Moreover, from (3),

$$IL_{i+1}^- - IP_i = mQ_i, \quad i = 1, \dots, N - 1, \quad (7)$$

where m is a nonnegative integer.

Define a sequence of functions recursively as follows. Let $G_1(y) = g_1(y)$, $y \in \mathfrak{I}$. For $i = 1, \dots, N$, assume that $G_i(\cdot)$ is defined and that $\bar{G}_i(y) \stackrel{\text{def}}{=} \sum_{x=1}^{Q_i} G_i(y + x)$ is quasiconvex and minimized at $y = R_i$, a finite integer. Thus $G_i(\cdot)$ satisfies the conditions of Lemma 1. Define $O_i[\cdot]$ as in Lemma 1 by replacing R and Q with R_i and Q_i respectively. Define

$$G_{i+1}(y) = g_{i+1}(y) + EG_i(O_i[y - D_{i+1}]),$$

$$y \in \mathfrak{I}, \quad i = 1, \dots, N - 1.$$

The above assumptions will be verified later under a specific, plausible cost structure.

LEMMA 2. $G_i(IP_i) \geq G_i(O_i[IL_{i+1}^-])$, $i = 1, \dots, N - 1$.

PROOF. Let $O_i[IL_{i+1}^-] = y$. Thus $y \leq R_i + Q_i$. From the definition of $O_i[\cdot]$ and (7), $y - IP_i$ is an integer multiple of Q_i . From Lemma 1(i), $G_i(IP_i + xQ_i)$, as a function of $x \in \mathfrak{I}$, is quasiconvex and minimized at $x = x^*$, where x^* is the unique integer with $R_i + 1 \leq IP_i + x^*Q_i \leq R_i + Q_i$. Therefore, the lemma follows if $y = IP_i + x^*Q_i$. Otherwise, if $y < IP_i + x^*Q_i$ thus $y \leq R_i$, then $y = IL_{i+1}^-$ and (7) implies that $IP_i \leq y$. In this case, the lemma follows because $G_i(IP_i + xQ_i)$ is nonincreasing in x for $x < x^*$. $\square$

We next derive a lower bound on the long-run average value of (5). Note that

$$\begin{aligned} C &\stackrel{\text{def}}{=} E \left[\sum_{i=1}^N g_i(IP_i) | IP_N \right] \\ &= E \left[\sum_{i=2}^N g_i(IP_i) + G_1(IP_1) | IP_N \right] \\ &\geq E \left[\sum_{i=2}^N g_i(IP_i) + G_1(O_1[IL_2^-]) | IP_N \right] \\ &= E \left[\sum_{i=3}^N g_i(IP_i) + g_2(IP_2) + G_1(O_1[IP_2 - D_2]) | IP_N \right] \\ &= E \left[\sum_{i=3}^N g_i(IP_i) + G_2(IP_2) | IP_N \right], \end{aligned} \quad (8)$$

where the inequality follows from Lemma 2, the first equality from the definition of $G_1(\cdot)$, the second equality from (6), and the third equality from the definition of $G_2(\cdot)$ and

the fact that D_2 is independent of IP_2 and IP_N . Repeating the above procedure, we have

$$C \geq G_N(IP_N). \quad (9)$$

We can think of IP_N as the inventory position of a single-location inventory system. As we have shown in the previous section, the long-run average value of $G_N(IP_N)$ is bounded below by

$$C^* \stackrel{\text{def}}{=} \frac{1}{Q_N} \sum_{y=R_N+1}^{R_N+Q_N} G_N(y).$$

THEOREM 1. C^* is a lower bound on the long-run average costs of all feasible policies in the serial system.

The lower bound in Theorem 1 can be achieved. Consider the following feasible policy. Whenever the echelon inventory position of stage i , $i = 1, \dots, N$, is at or below R_i , stage $i + 1$ ships an integer multiple of Q_i to stage i to bring the echelon inventory position of stage i into the interval $\{R_i + 1, \dots, R_i + Q_i\}$; and if the on-hand inventory at stage $i + 1$ is insufficient, then ship as much as possible. This policy is called an echelon-stock (R, nQ) policy. (An alternative way to describe the policy is to allow each upstream stage to backlog orders from the downstream stage in case of a stockout. Under the echelon-stock (R, nQ) policy, every stage orders to bring its echelon stock, which is its echelon inventory position plus its outstanding orders that are backlogged at the upstream stage, into the critical interval.)

THEOREM 2. For the serial system, it is optimal to use the echelon-stock (R_i, nQ_i) policy at stage i , $i = 1, \dots, N$. The minimum long-run average cost is C^* .

PROOF. It suffices to show that the echelon-stock (R, nQ) policy achieves a long-run average cost equal to C^* . Because stage i follows an (R_i, nQ_i) policy based on its echelon inventory position, $IP_i = O_i[IL_{i+1}^-]$ in the long run. Therefore, Lemma 2 becomes an equality. Consequently, both (8) and (9) are equalities. From the single-location model considered in the previous section, the (R_N, nQ_N) policy at stage N achieves the lower bound C^* . $\square$

The optimal (R, nQ) policy is easy to compute. The reorder points R_i are the minimum points of N quasiconvex functions $\tilde{G}_i(\cdot)$. The computational procedure is bottom-up: First compute R_1 , which is used to determine $G_2(\cdot)$ thus $\tilde{G}_2(\cdot)$; then compute R_2 , so on and so forth. This is reminiscent of the Clark-Scarf (1960) model where the optimal base-stock levels are also determined sequentially. A key observation from the Clark-Scarf model is that an N -stage serial system can be decomposed into N newsboy problems if a proper *induced-penalty cost* is charged to every upstream stage for not meeting an order from the downstream stage. The same observation holds here. Suppose we charge

the following induced-penalty cost to stage $i + 1$:

$$G_i(O_i[IL_{i+1}^-]) - \min_{m \in \mathfrak{Z}} G_i(IL_{i+1}^- + mQ_i), \quad (10)$$

which, under the (R_i, nQ_i) policy at stage i , is equal to

$$G_i(IP_i) - \min_{m' \in \mathfrak{Z}} G_i(IP_i + m'Q_i).$$

The above expression has an intuitive interpretation: The first term is the *actual* cost at stage i , while the second term is the *minimum* cost at stage i if stage $i + 1$ has ample stock. The difference can thus be attributed to the lack of inventory at stage $i + 1$. Under this induced-penalty cost, stage $i + 1$'s expected cost, given its echelon inventory position $IP_{i+1} = y$, is

$$\tilde{G}_{i+1}(y) \stackrel{\text{def}}{=} g_{i+1}(y) + E \left\{ G_i(O_i[y - D_{i+1}]) - \min_{m \in \mathfrak{Z}} G_i(y - D_{i+1} + mQ_i) \right\}.$$

Because $\sum_{x=z+1}^{z+Q_i} \min_{m \in \mathfrak{Z}} G_i(y - D_{i+1} + x + mQ_i) = \tilde{G}_i(R_i)$ for all z, y and D_{i+1} (Lemma 1),

$$\sum_{x=1}^{Q_{i+1}} \tilde{G}_{i+1}(y + x) = \tilde{G}_{i+1}(y) - n_i \tilde{G}_i(R_i).$$

Because $\tilde{G}_{i+1}(y)$ is quasiconvex and minimized at $y = R_{i+1}$, $\sum_{x=1}^{Q_{i+1}} \tilde{G}_{i+1}(y + x)$ is also quasiconvex and minimized at the same point. Now assume that stage $i + 2$ has ample stock. The problem facing stage $i + 1$ is exactly the same as the single-location problem analyzed in §2 with a one-period cost function $\tilde{G}_{i+1}(\cdot)$. Consequently, the optimal policy at stage $i + 1$ is the (R_{i+1}, nQ_{i+1}) policy with a long-run average cost

$$\frac{\tilde{G}_{i+1}(R_{i+1}) - n_i \tilde{G}_i(R_i)}{Q_{i+1}} = \frac{\tilde{G}_{i+1}(R_{i+1})}{Q_{i+1}} - \frac{\tilde{G}_i(R_i)}{Q_i}.$$


The sum of the above cost across all stages is C^* . Therefore, the induced-penalty costs (10) decompose the N -stage problem into N single-stage problems.

It is interesting to note that different induced-penalty costs have been used in the literature for multi-echelon inventory problems with fixed ordering costs (thus batch transfers). For example, one can charge the following induced-penalty cost to stage $i + 1$:

$$G_i(\min\{s_i, IL_{i+1}^-\}) - G_i(s_i),$$

where s_i is some constant. This type of induced-penalty cost has been used (a) by Clark and Scarf (1962) to decompose the dynamic program for a two-stage serial system, (b) by Chen and Zheng (1994b) to establish lower bounds on minimum costs in various multi-echelon systems, and (c) by Chen and Zheng (1998) to derive approximate cost functions that lead to near-optimal control parameters in serial systems. Note that in (a) and (b), s_i is obtained by solving a single-location (s, S) model, while in (c), $s_i = R_i$. The

Figure 1. The G-function and two types of induced-penalty functions.



induced-penalty costs in (10), however, lead to the optimal policy. Figure 1 depicts these two types of induced-penalty costs in a serial system with $R_1 = 8$ and $Q_1 = 5$.

We pause here to note that it is not a novel idea to use echelon-stock (R, nQ) policies in serial systems; see, e.g., De Bodt and Graves (1985), Axsater and Rosling (1993), Chen and Zheng (1994a, 1998). What we have shown here is that this class of policies is actually *optimal* when the material flow is regulated by fixed lot sizes.

For the remainder of this section, we show that the assumptions made earlier about $G_i(\cdot)$, $i = 1, \dots, N$, are satisfied under a plausible holding-backorder cost structure. Suppose the system incurs linear holding and backorder costs. Let

H_i = installation holding cost at stage i per unit per period,
 h_i = echelon holding cost at stage i per unit per period,
 $= H_i - H_{i+1} > 0$, with $H_{N+1} = 0$,
 p = backorder penalty cost (at stage 1) per unit per period,
 $p > 0$.

(At first glance, it seems restrictive to assume $h_i > 0$ for $i = 1, \dots, N$. It is not. If $h_N = H_N = 0$, then we can keep infinite inventory at stage N and treat the stage as the new outside supplier of the system. If $h_i \leq 0$ for some $i < N$, then no inventories should be kept at stage $i + 1$. In other words, any shipment to stage $i + 1$ should be sent to stage i directly. This effectively eliminates stage $i + 1$. Continue in this fashion as long as there is a stage with a nonpositive echelon holding cost rate. This will lead to a serial model with positive echelon holding cost rates at all stages.)

We follow the convention of assessing holding and back-order costs based on the period-ending inventory levels. Note that the installation (on-hand) inventory at stage i at t^+ , $i \geq 2$, can be written as $IL_i(t) - IL_{i-1}(t)$, and the installation (on-hand) inventory at stage 1 at t^+ is $IL_1(t) + B(t)$. (Thus the inventories in transit to stage $i - 1$ are part of the installation inventory of stage i . This is standard.) Charging H_i for each unit of installation inventory at stage i , $i = 1, \dots, N$, and p for each unit of customer backorders, we have the following holding and backorder costs in period t :

$$\begin{aligned} & \sum_{i=2}^N H_i [IL_i(t) - IL_{i-1}(t)] + H_1 [IL_1(t) + B(t)] + pB(t) \\ &= \sum_{i=1}^N h_i IL_i(t) + (p + H_1)B(t). \end{aligned}$$

Call $h_1 IL_1(t) + (p + H_1)B(t)$ the cost at stage 1 and $h_i IL_i(t)$ the cost at stage i , $i = 2, \dots, N$. Because $IL_1(t + L_1) = IP_1(t) - D[t, t + L_1]$, where $D[t_1, t_2]$ denotes the total demand in periods $t_1, \dots, t_2$, the expected one-period cost at stage 1, given $IP_1(t) = y$, is

$$\begin{aligned} g_1(y) &= E[h_1(y - D[t, t + L_1]) \\ &\quad + (p + H_1)(y - D[t, t + L_1])^-], \end{aligned}$$

where $(x)^- = \max\{-x, 0\}$. Similarly, because $IL_i(t + L_i) = IP_i(t) - D[t, t + L_i]$ for $i = 2, \dots, N$, we have

$$g_i(y) = E[h_i(y - D[t, t + L_i])], \quad i = 2, \dots, N.$$

We next verify the assumptions made earlier about $G_i(\cdot)$; i.e., $\bar{G}_i(\cdot)$ is quasiconvex and is minimized at a finite point. Let μ_1 be the expected value of $D[t, t + L_1]$ and define

$$G_1^d(y) = h_1(y - \mu_1)^+ + (p + H_2)(y - \mu_1)^-, \quad y \in \mathfrak{I}.$$

For $i = 2, \dots, N$, let μ_i be the expected value of $D[t, t + L_i]$ and define

$$\begin{aligned} G_i^d(y) &= h_i(y - \mu_i)^+ + (p + H_{i+1})(y - \mu_i)^- \\ &\quad - (h_2 + \dots + h_i)\mu_i, \quad y \in \mathfrak{I}, \end{aligned}$$

where μ is the expected demand in one period. Note that $G_1^d(y)$ can be interpreted as the holding and backorder costs in an EOQ model (with backorders) with holding cost rate h_1 , backorder cost rate $p + H_2$, and leadtime demand μ_1 , given that the inventory position is y . Similar interpretations can be made for $G_i^d(\cdot)$, $i = 2, \dots, N$.

LEMMA 3. $\bar{G}_i(\cdot)$ is convex and $G_i(y) \geq G_i^d(y)$, $\forall y \in \mathfrak{I}$, $i = 1, \dots, N$.

PROOF. It is clear that $G_1(\cdot)$ is convex and $G_1(y) \geq G_1^d(y)$, $\forall y \in \mathfrak{I}$ (Jensen's inequality). Thus the lemma holds for $i = 1$. Now take any $i \geq 1$ and suppose the lemma holds for i . We next show that it also holds for $i + 1$.

By definition,

$$\begin{aligned}\bar{G}_{i+1}(y) &= h_{i+1} \sum_{x=1}^{Q_{i+1}} (y + x - \mu_{i+1} - \mu) \\ &\quad + \sum_{x=1}^{Q_{i+1}} EG_i(O_i[y + x - D_{i+1}]).\end{aligned}$$

(Recall that D_{i+1} is the total demand in periods $t, \dots, t + L_{i+1} - 1$ and its mean is μ_{i+1} . So the mean of $D[t, t + L_{i+1}]$ is $\mu_{i+1} + \mu$.) The first term on the right side is clearly convex in y . From Lemma 1 (ii), we can write the second term as

$$\begin{aligned}&\sum_{z=0}^{n_i-1} \sum_{x=1}^{Q_i} EG_i(O_i[y + x + zQ_i - D_{i+1}]) \\ &= \sum_{z=0}^{n_i-1} E\bar{G}_i(\min\{R_i, y + zQ_i - D_{i+1}\}),\end{aligned}$$

which is convex in y because R_i is a minimum point of $\bar{G}_i(\cdot)$. In sum, $\bar{G}_{i+1}(\cdot)$ is convex.

It remains to show that $G_{i+1}(y) \geq G_{i+1}^d(y)$ for all $y \in \mathfrak{I}$. First, note that if $y \leq R_i + Q_i$ then $O_i[y] = y$; otherwise, if $y > R_i + Q_i$ then $O_i[y] = y - nQ_i$ for some positive integer n and thus $O_i[y] < y$. Combining these two cases and noting that $(x)^-$ is a nonincreasing function, we have

$$(O_i[y])^- \geq (y)^-. \quad (11)$$

On the other hand, from the definition of $G_i^d(y)$, we have

$$\begin{aligned}G_i^d(y) &\geq (p + H_{i+1})(y - \mu_i)^- - (h_2 + \dots + h_i)\mu \\ &\geq (p + H_{i+1})(y)^- - (h_2 + \dots + h_i)\mu,\end{aligned}$$

where the second inequality follows because $\mu_i \geq 0$ and $(x)^-$ is a nonincreasing function. The inductive assumption, the above inequality, and (11) lead to

$$\begin{aligned}G_i(O_i[y]) &\geq G_i^d(O_i[y]) \geq (p + H_{i+1})(y)^- \\ &\quad - (h_2 + \dots + h_i)\mu.\end{aligned}$$

Using the above inequality in the definition of $G_{i+1}(y)$, we have

$$\begin{aligned}G_{i+1}(y) &\geq h_{i+1}(y - \mu_{i+1} - \mu) + (p + H_{i+1})E(y - D_{i+1})^- \\ &\quad - (h_2 + \dots + h_i)\mu \\ &\geq h_{i+1}(y - \mu_{i+1}) + (p + H_{i+1})(y - \mu_{i+1})^- \\ &\quad - (h_2 + \dots + h_{i+1})\mu \\ &= h_{i+1}(y - \mu_{i+1})^+ + (p + H_{i+2})(y - \mu_{i+1})^- \\ &\quad - (h_2 + \dots + h_{i+1})\mu \\ &= G_{i+1}^d(y),\end{aligned}$$

where the second inequality is Jensen's. $\square$

Lemma 3 implies that $\bar{G}_i(\cdot)$ is quasiconvex and minimized at a finite point because $\lim_{|y| \rightarrow +\infty} G_i^d(y) = +\infty$, $i = 1, \dots, N$. This completes the section.


4. ASSEMBLY SYSTEMS

This section is parallel to §3. We first describe an assembly system where the materials flow in fixed batches. We then introduce several key assumptions, among which are several properties of the one-period cost function. Based on these assumptions, we establish a lower bound on the long-run average costs of all feasible replenishment policies. We then present a feasible policy that achieves the lower bound. Finally, we verify the assumed properties of the one-period cost function under a specific holding-backorder cost structure.

Consider the following assembly system. Several components are purchased from outside suppliers. These components are assembled into intermediate products, or subassemblies, which are further assembled into other subassemblies, so on and so forth until a finished product (or end item) is produced. We assume that the system produces only one end item. The system has N distinct items: the components, the subassemblies, and the end item. The assembly structure is a tree where the root represents the end item, the leaves are the components, and each item (except the end item) has exactly one successor. We assume that the component suppliers have unlimited stock and each order for a component arrives after a fixed leadtime. The production leadtimes for the other items (subassemblies and the end item) are also fixed. The customer demand is for the end item only, and the demands in different periods are independent, identically distributed, discrete random variables. When the demand exceeds the on-hand inventory of the end item, the excess demand is backlogged and filled as soon as the finished-good inventory becomes available. The system incurs holding and backorder costs. The objective is to minimize the long-run average total cost in the system.

The items are numbered $1, \dots, N$ with item 1 being the end item. Let $s(i)$ = the immediate successor of item i , with $s(1) = 0$, $S(i)$ = the set consisting of item i and all its successors, $P(i)$ = the set of the immediate predecessors of item i . Note that $S(1) = \{1\}$ and $P(i) = \emptyset$ if i is a component. We define the units of the items so that one unit of item i requires exactly one unit of item j for all (i, j) with $j \in P(i)$.

Let L_i , a fixed nonnegative integer, be the leadtime for item i , $i = 1, \dots, N$. If item i is a component, then L_i is the procurement leadtime; otherwise, if item i is a subassembly or the end item, then L_i is the number of periods required to produce the item. Following Rosling (1989), we define M_i to be the total leadtime associated with item i , i.e., $M_i = \sum_{j \in S(i)} L_j$. Thus M_i is the total assembly time required to produce item i and all its successors. For convenience, we number the items so that M_i is nondecreasing in i and if item i is a successor of item j then $i < j$. (Thus the end item is item 1.) Define $l_i = M_i - M_{i-1}$, $i = 1, \dots, N$, with $M_0 = 0$. Clearly, $l_i \geq 0$ for all i . Figure 2 depicts an assembly system with five items. Items 3, 4, and 5 are components, and item 2 is a subassembly. Associated with each item are a circle and a triangle: The circle marks the source of the item, the triangle represents the stockpile of the item, and the length of the

Figure 2. An assembly system.

arrow linking the two represents the assembly/procurement leadtime. The circle associated with item i is connected to the triangle associated with item 1 through a chain of arrows. The total length of these arrows is the total leadtime associated with item i . Note that the items are indexed according to their total leadtimes.

A unique feature of our assembly system is that each order for item i must be a positive integer multiple of a fixed base quantity, Q_i , which is a positive, item-specific integer, $i = 1, \dots, N$. These base quantities represent the basic units of transportation or production. The following assumption is critical:

$$Q_{i+1} = n_i Q_i, \quad i = 1, \dots, N - 1, \quad (12)$$

where n_i is a positive integer. Therefore, we require that the base quantities for the items do not decrease with their total leadtimes. This is a strong assumption. Note that to some extent the base quantities are a consequence of the setup costs, which are not modeled explicitly here. A larger setup cost leads to a larger base quantity. Therefore, the assumption is plausible when a longer total leadtime is associated with a larger setup cost. A special case of the above assumption is $n_i = 1$ for all i , i.e., all the items share a common base quantity. This is true in some production environments. (In §5, we suggest a solution for assembly systems that do not satisfy (12).)

We next define key variables to describe the inventory state of the system. The following terminology is useful. When a unit of item i is being assembled to produce a unit of item j , we say that one unit of item i is *leaving* and one unit of item j is *arriving*. The units of a component that are in transit from an outside supplier are also called *arriving*. The units of an item that are neither leaving nor arriving are the *on-hand* units of the item. Therefore, the components and subassemblies can have arriving, on-hand, and leaving units, while the end item can only have arriving and on-hand units. For example, in Figure 2, suppose we take one unit of item 4 and one unit of item 5 and start assembling them into one unit of item 2. Then we say the following: One unit of item 4 is leaving, one unit of item 5 is leaving, and one unit of item 2 is arriving. To determine the arriving, on-hand, and leaving units associated with an item, one can visually stand at the triangle associated with the item and determine the inbound, on-hand, and outbound units, respectively. The *installation inventory* of an item is its on-hand units plus leaving units. The *echelon inventory level* of item i is the sum of the installation inventories of the items in $S(i)$ minus the customer backorders. The *echelon inventory position* of item i is the echelon inventory level of the item plus the arriving units of the item. As for the serial model considered in the previous section, we assume that all the replenishment decisions in a period, i.e., how many units of each item to produce/acquire, are made at the beginning of the period, and that customer demand arrives during the period. Similarly, write t^- for the beginning of period t after all the replenishment decisions in the period have been made, and t^+ for the end of the period after customer demand. We use $IL_i(t)$ to denote the echelon inventory level of item i at t^+ , $IL_i^-(t)$ to denote the echelon inventory level of item i at t^- , and $IP_i(t)$ to denote the echelon inventory position of item i at t^- . One can use Figure 2 to visualize these inventory variables. Consider item 3. As mentioned earlier, the circle associated with item 3 is connected to the triangle associated with the end item through a chain of arrows. Think of this chain as a pipeline filled with inventories. Then $IP_3(t)$ is the total net inventory in the pipeline. The echelon inventory level of item 3 is, however, associated with a *shorter* pipeline, one that starts from the triangle associated with item 3.

Note that the echelon inventory level of an item does not include any units of the item that are arriving. Following Chen and Zheng (1994b), we define the *extended echelon inventory level* of item i to be the echelon inventory level of the item plus those arriving units of the item that have completed the first l_i periods of the assembly/procurement operation. Note that the leadtime to produce/acquire item i is L_i periods and $L_i = M_i - M_{s(i)} \geq M_i - M_{i-1} = l_i$. Therefore, if we start producing/acquiring one unit of item i in period t , then this unit becomes part of the echelon inventory level of item i in period $t + L_i$, but it becomes part of the *extended* echelon inventory level of item i earlier, in period $t + l_i$. Clearly, if $s(i) = i - 1$, then $L_i = l_i$ and thus item i 's extended echelon inventory level is the same as its echelon inventory

level. This is always true for the end item. For example, in Figure 2, we have $s(i) = i - 1$ for $i = 1, 2$ and $s(i) < i - 1$ for $i = 3, 4, 5$. Therefore, the extended echelon inventory levels of items 3, 4, and 5 are different from their echelon inventory levels. As mentioned above, there is a pipeline connecting the circle associated with item 3 with the triangle associated with item 1. The extended echelon inventory level of item 3 is the total net inventory in a *truncated* pipeline that begins at the dotted line. Let $EL_i(t)$ be the extended echelon inventory level of item i at t^- . The following inventory balance equations are useful:

$$EL_i(t + l_i) = IP_i(t) - D[t, t + l_i], \quad i = 1, \dots, N, \quad (13)$$

and

$$IL_i(t + L_i) = EL_i(t + l_i) - D[t + l_i, t + L_i], \quad i = 1, \dots, N. \quad (14)$$

We assume that the system begins with a plausible initial state. Let the first period be period 0. Note that the initial on-hand inventory of item i (> 1) is $IL_i^-(0) - IP_{s(i)}(0)$, which satisfies

$$IL_i^-(0) - IP_{s(i)}(0) = m_i Q_{s(i)}, \quad \forall i \neq 1, \quad (15)$$

where m_i is a nonnegative integer. This assumption is reasonable because each order for item $s(i)$ is an integer multiple of $Q_{s(i)}$, and thus there is no incentive to keep an on-hand inventory of item i that is a fraction of $Q_{s(i)}$. We also assume that the system is *balanced* initially. For all $i < N$, if $i + 1 \notin P(i)$ then

$$EL_{i+1}(0) - IP_i(0) = m'_i Q_i, \quad (16)$$

where m'_i is an integer (positive or otherwise). This assumption has the following intuitive interpretation. Consider the assembly system depicted in Figure 2. Take $i = 2$. Clearly, $i + 1 = 3 \notin P(2) = \{4, 5\}$. As noted earlier, the extended echelon inventory level of item 3 is the total net inventory in a truncated pipeline from the dotted line to the triangle associated with item 1. On the other hand, the echelon inventory position of item 2 is the total net inventory in the pipeline from the circle associated with item 2 to the triangle associated with item 1. The left side of (16) is the difference between the total inventories in these two pipelines. (Note that the two pipelines have the same length.) When the difference is an integer multiple of Q_2 , it is possible for the system to initiate a production batch of item 2 so as to balance the inventories in the two pipelines. It is meaningful to try to balance the inventories in the different pipelines, because these inventories will eventually “merge,” on a one-to-one basis, into a single stream of finished goods. Also note that when all the items share a common base quantity, (15) implies (16). The above initial conditions lead to the following two lemmas.

LEMMA 4. *The on-hand inventory of item i , which is $IL_i^-(t) - IP_{s(i)}(t)$ at t^- , is always a nonnegative integer multiple of $Q_{s(i)}$, $i = 2, \dots, N$.*

PROOF. Note that the on-hand inventory of item i is decreased by multiples of $Q_{s(i)}$ and increased by multiples of Q_i , which is itself an integer multiple of $Q_{s(i)}$ (see (12)). The lemma thus follows from (15). $\square$

LEMMA 5. *For all $i < N$ with $i + 1 \notin P(i)$, $EL_{i+1}(t) - IP_i(t)$ is always an integer multiple of Q_i .*

PROOF. Note that both $EL_{i+1}(t)$ and $IP_i(t)$ are simultaneously decreased by the customer demand. The former is increased by multiples of Q_{i+1} , and the latter is increased by multiples of Q_i . The lemma thus follows from the initial condition in (16) and the integer-ratio condition in (12). $\square$

As mentioned earlier, the system incurs holding and back-order costs. Suppose that these costs can be assessed in the following manner. Let $g_i(y)$ be the expected cost associated with item i in period t given $IP_i(t) = y$, $i = 1, \dots, N$. We assume that

$$g_i(y) \text{ is nondecreasing in } y, \quad i = 2, \dots, N. \quad (17)$$

The expected total cost in the system in period t , given the echelon inventory positions of all items, is thus

$$\sum_{i=1}^N g_i(IP_i(t)).$$

As in §3, it is convenient to introduce a time shift in the above expression. Take any period t . Write IP_i for $IP_i(t - M_i)$, $i = 1, \dots, N$. Write IL_i^- for $IL_i^-(t - M_{i-1})$ and EL_i for $EL_i(t - M_{i-1})$, $i = 2, \dots, N$. We proceed to establish a lower bound on the long-run average value of

$$\sum_{i=1}^N g_i(IP_i). \quad (18)$$

Define a sequence of functions recursively as follows. (Here further assumptions are made about the one-period cost functions.) Let $G_1(y) = g_1(y)$, $\forall y \in \mathfrak{I}$. For $i = 1, \dots, N$, assume that $G_i(\cdot)$ is defined and that $\bar{G}_i(y) \stackrel{\text{def}}{=} \sum_{x=1}^{Q_i} G_i(y+x)$ is quasiconvex and minimized at $y = R_i$, a finite integer. Thus $G_i(\cdot)$ satisfies the conditions of Lemma 1. Define $O_i[\cdot]$ as in Lemma 1 by replacing R and Q with R_i and Q_i , respectively. Define

$$G_{i+1}(y) = g_{i+1}(y) + EG_i(O_i[y - D_{i+1}]),$$

$$y \in \mathfrak{I}, \quad i = 1, \dots, N - 1,$$

where D_{i+1} denotes $D[t - M_{i+1}, t - M_i]$, the total customer demand in periods $t - M_{i+1}, \dots, t - M_i - 1$, for $i = 1, \dots, N - 1$.

We next introduce a sequence of variables that help establish a linkage between the assembly system and the serial system considered in the previous section. Let $IP^N = IP_N$. Suppose IP^i is defined, $i = 2, \dots, N$. Let

$$EL^i = IP^i - D_i. \quad (19)$$

Define $IP^{i-1} = \min\{IP_{i-1}, EL^i\}$ for $i = 2, \dots, N$. Now we have the $IP^1, \dots, IP^N$ and $EL^2, \dots, EL^N$. Since $IP_i \geq IP^i$ for all i by definition, we have from (13) and (19) that

$$EL_i \geq EL^i, \quad i = 2, \dots, N. \quad (20)$$

LEMMA 6. $IP^1 = IP_1$.

PROOF. Based on the above recursive definition of IP^i ,

$$\begin{aligned} IP^1 &= \min\{IP_1, EL^2\} \\ &= \min\{IP_1, IP^2 - D_2\} \\ &= \min\{IP_1, IP_2 - D_2, EL^3 - D_2\} \\ &= \dots \\ &= \min\{IP_1, IP_i - D_i - \dots - D_2, i = 2, \dots, N\}. \end{aligned}$$

It suffices to show that for any $i = 2, \dots, N$, $IP_i - D_i - \dots - D_2 \geq IP_1$. Take any $i = 2, \dots, N$. Let $S(i) = \{i_1, i_2, \dots, i_n\}$ for some positive integer n with $i = i_1 > i_2 > \dots > i_n = 1$. (The end item, or item 1, is always a member of $S(i)$.) Note that i_{m+1} is the immediate successor of i_m for $m = 1, \dots, n - 1$. Because

$$\begin{aligned} D_i + \dots + D_2 &= D[t - M_i, t - M_{i-1}) \\ &\quad + \dots + D[t - M_2, t - M_1) \\ &= D[t - M_i, t - M_1) \\ &= D[t - M_{i_1}, t - M_{i_2}) + D[t - M_{i_2}, t - M_{i_3}) \\ &\quad + \dots + D[t - M_{i_{n-1}}, t - M_{i_n}), \end{aligned}$$

we have

$$\begin{aligned} IP_i - D_i - \dots - D_2 &= IP_{i_1}(t - M_{i_1}) - D[t - M_{i_1}, t - M_{i_2}) \\ &\quad - D[t - M_{i_2}, t - M_{i_3}) \\ &\quad - \dots - D[t - M_{i_{n-1}}, t - M_{i_n}) \\ &= IL_{i_1}^-(t - M_{i_2}) - D[t - M_{i_2}, t - M_{i_3}) \\ &\quad - \dots - D[t - M_{i_{n-1}}, t - M_{i_n}) \\ &\geq IP_{i_2}(t - M_{i_2}) - D[t - M_{i_2}, t - M_{i_3}) \\ &\quad - \dots - D[t - M_{i_{n-1}}, t - M_{i_n}) \\ &\geq \dots \\ &\geq IP_{i_{n-1}}(t - M_{i_{n-1}}) \\ &\quad - D[t - M_{i_{n-1}}, t - M_{i_n}) \\ &= IL_{i_{n-1}}^-(t - M_{i_n}) \\ &\geq IP_1, \end{aligned}$$

where the inequalities follow because the echelon inventory level (at the beginning of a period) of an item is at least as large as the echelon inventory position of its immediate successor. $\square$

LEMMA 7. $EL^{i+1} - IP_i = m_i Q_i$, where m_i is an integer, $i = 1, \dots, N - 1$.

PROOF. We first prove the lemma for $i = N - 1$. By definition, $EL^N = IP^N - D_N = IP_N - D_N = EL_N$, where the last equality is from (13). If $N - 1 = s(N)$, then $EL_N = IL_N^-$, and thus $EL_N - IP_{N-1}$ is the on-hand inventory of item N , which is an integer multiple of Q_{N-1} (Lemma 4). Otherwise, if $N - 1 \neq s(N)$, then the lemma follows from Lemma 5.

Now suppose the lemma holds for $i + 1$. We next show that it also holds for i . Because $EL^{i+1} = IP^{i+1} - D_{i+1}$ and $IP^{i+1} = \min\{IP_{i+1}, EL^{i+2}\}$ by definition, it suffices to show that both $IP_{i+1} - D_{i+1} - IP_i$ and $EL^{i+2} - D_{i+1} - IP_i$ are integer multiples of Q_i . Because the lemma holds for $i + 1$, i.e., $EL^{i+2} - IP_{i+1}$ is an integer multiple of Q_{i+1} , which is itself an integer multiple of Q_i , it suffices to show that $IP_{i+1} - D_{i+1} - IP_i$ is an integer multiple of Q_i . Note that $IP_{i+1} - D_{i+1} = EL_{i+1}$ (see (13)). Now if $i + 1 \in P(i)$, then $EL_{i+1} - IP_i = IL_{i+1}^- - IP_i$ represents the on-hand inventory of item $i + 1$, which is an integer multiple of Q_i (Lemma 4). Otherwise, if $i + 1 \notin P(i)$, then $EL_{i+1} - IP_i$ is still an integer multiple of Q_i (Lemma 5). This completes the proof. $\square$

COROLLARY 1. $EL^{i+1} - IP^i$ is a nonnegative integer multiple of Q_i , $i = 1, \dots, N - 1$.

PROOF. Follows from Lemma 7 and the fact that $IP^i = \min\{IP_i, EL^{i+1}\}$. $\square$

Now we are ready to present the series analogy. Consider a serial system with N stages. The leadtime at stage i is l_i and the base quantity at stage i is Q_i , $i = 1, \dots, N$. The demand process is the same as in the assembly system. Equation (19) and Corollary 1 establish that the inventory variables IP^i and EL^i can be replicated in the serial system. That is, the echelon inventory position at stage i is IP^i , the echelon inventory level (at the beginning of a period) at stage i is EL^i , and $EL^{i+1} - IP^i$ is the on-hand inventory at stage $i + 1$. As a result, several results obtained in the previous section can be carried over here. In particular:

LEMMA 8. $G_i(IP^i) \geq G_i(O_i[EL^{i+1}])$, $i = 1, \dots, N - 1$.

PROOF. Similar to Lemma 2. $\square$

This lemma can then be used to establish a lower bound on the long-run average value of (18). Note that

$$\begin{aligned} C &\stackrel{\text{def}}{=} E \left[\sum_{i=1}^N g_i(IP_i) | IP^N \right] \\ &= E \left[\sum_{i=2}^N g_i(IP_i) + G_1(IP^1) | IP^N \right] \\ &\geq E \left[\sum_{i=3}^N g_i(IP_i) + g_2(IP_2) + G_1(O_1[EL^2]) | IP^N \right], \end{aligned} \quad (21)$$

$$\begin{aligned}
&\geq E \left[\sum_{i=3}^N g_i(IP_i) + g_2(IP^2) + G_1(O_1[IP^2 - D_2]) | IP^N \right] \\
&= E \left[\sum_{i=3}^N g_i(IP_i) + G_2(IP^2) | IP^N \right],
\end{aligned} \tag{22}$$

where the first equality follows because $G_1(\cdot) = g_1(\cdot)$, $IP^1 = IP_1$ and $IP^N = IP_N$; the first inequality follows from Lemma 8; and the second inequality follows because $g_2(\cdot)$ is non-decreasing and $IP^2 \leq IP_2$. Repeating the above procedure,

$$C \geq G_N(IP^N) = G_N(IP_N),$$

the long-run average value of which is greater than or equal to

$$C^* = \frac{1}{Q_N} \sum_{y=R_N+1}^{R_N+Q_N} G_N(y);$$

see §2. Therefore, we have the following theorem.

THEOREM 3. C^* is a lower bound on the long-run average costs of all feasible policies in the assembly system.

The lower bound in Theorem 3 can be achieved. Consider the following feasible policy. Whenever the echelon inventory position of item i , IP_i , falls to or below R_i , produce the *maximum* integer multiple of Q_i so that IP_i does not exceed $R_i + Q_i$ and EL_{i+1} , the extended echelon inventory level of item $i+1$ (with $EL_{N+1} = +\infty$). Of course, the production quantity is constrained by the on-hand inventories of the immediate predecessors of item i . If before the production decision IP_i is already greater than or equal to EL_{i+1} , then no production for item i is initiated. Note that if item $i+1$ is a predecessor of item i , then IP_i can never exceed EL_{i+1} , which now becomes the echelon inventory level of item $i+1$. In this case, we need only the control parameters R_i and Q_i to make the replenishment decisions for item i . Otherwise, if item $i+1$ is *not* a predecessor of item i , then EL_{i+1} serves as an additional control parameter. Call the above policy the *modified echelon-stock* (R_i, nQ_i) policy.

THEOREM 4. For the assembly system, it is optimal to use the modified echelon-stock (R_i, nQ_i) policy to replenish item i , $i = 1, \dots, N$. The minimum long-run average cost is C^* .

PROOF. Under the above policy, $IP^i = IP_i$ and $EL^i = EL_i$ for all i in the long run. This is true for $i = N$ by definition. Because the order-up-to level for item $N-1$ never exceeds EL_N , $IP_{N-1} \leq EL_N$. Thus, $IP^{N-1} = IP_{N-1}$ by definition, which implies $EL^{N-1} = EL_{N-1}$, so on and so forth. The policy also implies $IP_i = O_i[EL_{i+1}]$. This follows because $EL_{i+1} - IP_i$ is a nonnegative integer multiple of Q_i by Corollary 1 and the fact that $IP^i = IP_i$ and $EL^{i+1} = EL_{i+1}$. With these observations, the inequalities in (21) and (22) both

become equalities. Finally, as shown in §2, the (R_N, nQ_N) policy for item N achieves the minimum long-run average value of $G_N(IP_N)$, C^* . Therefore, the modified echelon-stock (R, nQ) policy is optimal. $\square$

For the remainder of this section, we show that the assumptions made earlier about the one-period cost functions, $g_i(\cdot)$ for $i = 1, \dots, N$, are true under a plausible holding-backorder cost structure. The assumptions to be verified are (17) and that $G_i(\cdot)$ satisfy the conditions of Lemma 1. We consider the standard cost structure with linear holding and backorder costs. Let

H_i = installation holding cost for item i per unit per period,
 h_i = echelon holding cost for item i per unit per period,

i.e., $h_i = H_i - \sum_{j \in P(i)} H_j > 0$,

p = backorder cost for item 1 (end item) per unit per period,
 $p > 0$.

When each assembly operation is value adding and the holding costs are primarily the costs of capital tied up in the products, it is plausible to assume that every item has a positive echelon holding cost. If this is untrue—i.e., some items have nonpositive echelon holding costs—then one should follow an approach developed by Rosling (1989) to redefine the system so that every item has a positive echelon holding cost.

We follow the convention of assessing holding and backorder costs at the end of each period. At the end of period t , the installation inventory of item $i (> 1)$ is $IL_i(t) - IL_{s(i)}(t)$, the installation inventory of item 1 is $[IL_1(t)]^+$, and the customer backorder level is $[IL_1(t)]^-$. Charging H_i for each unit of installation inventory of item i and p for each unit of customer backorders, we have the total cost in period t :

$$\begin{aligned}
&\sum_{i=2}^N H_i [IL_i(t) - IL_{s(i)}(t)] + H_1 [IL_1(t)]^+ + p [IL_1(t)]^- \\
&= \sum_{i=1}^N h_i IL_i(t) + (p + H_1) [IL_1(t)]^-.
\end{aligned}$$

Define

$$\begin{aligned}
g_1(y) &= E[h_1(y - D[t, t + L_1]) \\
&\quad + (p + H_1)(y - D[t, t + L_1])^-], \quad y \in \mathfrak{I},
\end{aligned}$$

and

$$g_i(y) = E[h_i(y - D[t, t + L_i])], \quad y \in \mathfrak{I}, \quad i = 2, \dots, N.$$

From (13) and (14),

$$IL_i(t + L_i) = IP_i(t) - D[t, t + L_i], \quad i = 1, \dots, N.$$

Thus, given the echelon inventory positions IP_i , the expected total cost in a period in the system is $\sum_{i=1}^N g_i(IP_i)$. Clearly, $g_i(y)$ is increasing in y for $i = 2, \dots, N$. Thus (17) is satisfied. Following §3, one can easily show that the $G_i(\cdot)$ also satisfy the conditions of Lemma 1.

5. CONCLUDING REMARKS

The paper has so far focused on discrete-time models, but the results, also hold for continuous-time models with compound Poisson demand. We need to make only a few changes. Below, we highlight the necessary changes with the serial model.

Consider an N -stage serial system. Suppose that customers arrive at stage 1 according to a Poisson process with mean arrival rate λ . Each customer demands a random quantity, and the demand sizes of different customers are independent and identically distributed. Moreover, the demand sizes are independent of the arrival process. In short, the demand process is compound Poisson. An optimal policy for this system can be obtained by properly reinterpreting the notation in §3. Here the leadtimes L_i are allowed to be any nonnegative *real* numbers. The inventory variables are defined for any time epoch, and it is no longer necessary to distinguish between $IL_i^-(t)$ and $IL_i(t)$. The initial condition (3) is still required. The cost function $g_i(y)$ should be interpreted as the expected *rate* at which the costs at stage i accrue given its echelon inventory position y . The rest is essentially identical to the discrete-time case.

The integer-ratio assumption made earlier for assembly systems clearly limits the applicability of the results. When the integer-ratio constraint is violated, it is likely that the series analogy collapses. Below, we suggest an intuitively appealing approach to deal with this situation. Consider, for example, the simplest assembly system with three items. Items 2 and 3 are assembled into a final product, item 1. As before, the items are indexed by their total leadtimes. Thus item 3 has the longest total leadtime. Now suppose the base quantities associated with these items do not satisfy the integer-ratio constraint. In particular, assume that $Q_1 = 2$, $Q_2 = 8$, and $Q_3 = 4$. (Notice that the base quantities are all integer powers of two. The deterministic multi-echelon inventory literature, together with some evidence reported in Chen and Zheng 1998, suggests that power-of-two order quantities are often close to being optimal.) To determine a *reasonable* replenishment policy for this system, one can increase Q_3 to 8 so that the new base quantities satisfy the integer-ratio condition. For this new system, determine the optimal (modified) echelon-stock (R, nQ) policy as in §4. This policy is clearly feasible for the original system. In particular, each order for item 3 is now a multiple of 8 and thus is a multiple of 4. Our conjecture is that this policy is close to optimality. The intuition is that when a downstream item is replenished in multiples of 8, there does not seem to be any compelling reason to replenish an upstream item in fractions of 8. The same idea can be applied to general assembly systems that violate the integer-ratio condition.

We conclude this paper with a few open questions. We have shown the optimality of (R, nQ) policies only in serial and assembly systems. These results are unlikely to hold for distribution systems. In fact, no optimal policies exist for such systems; see Clark and Scarf (1962) for a brief explanation of the difficulties. However, it is still interest-

ing to see if (R, nQ) policies are close to being optimal. (On this question, Chen and Zheng 1997 provide some numerical evidence by comparing the performance of (R, nQ) policies with a lower bound in one-warehouse multi-retailer systems.) Another possible extension is nonstationary demands in the spirit of Veinott (1965). Finally, installation stock (R, nQ) policies pose yet another interesting question: Are they optimal when every stage has access only to the local inventory information?

APPENDIX

Consider the following Markov chain $\{Y_t\}_{t=1}^{\infty}$ with state space $S = \{R + 1, \dots, R + Q\}$, where R is an integer and Q a positive integer. The transition from one state to the next is determined by

$$Y_{t+1} = Y_t - d_t + mQ,$$

where $d_1, d_2, \dots$ are i.i.d. random variables that take on non-negative integer values and m is the unique integer (non-negative) so that $Y_{t+1} \in S$. This appendix is concerned with the steady state distribution of this Markov chain without any restriction on the distribution of d_t .

Let $f(\cdot)$ be the probability mass function of d_t . The one-step transition probability from state i to j is thus

$$a_{ij} = \sum_{n=0}^{\infty} f(nQ + i - j), \quad i, j \in S.$$

Let $\mathbf{P} = (a_{ij})$ be the transition matrix. Let $\pi = (\pi_1, \pi_2, \dots, \pi_Q)$. It is well known that the uniform distribution $\pi = (1/Q, 1/Q, \dots, 1/Q)$ is a solution to the balance equations $\pi = \pi\mathbf{P}$. If the Markov chain is *irreducible*, then the uniform distribution is the unique steady state distribution. (Hadley and Whitin 1961 consider this case, which holds if $f(1) > 0$, a condition mentioned in §2.) Below, we suppose that the Markov chain is *reducible*. In this case, the state space is partitioned into transient subsets and recurrent subsets.

Take any $i, j \in S$. Suppose j can be reached from i . Then there exist values of $d_1, d_2, \dots, d_n$ for some positive integer n so that

$$j = i - (d_1 + d_2 + \dots + d_n) + mQ$$

for some integer $m \geq 0$. That is, j can be reached from i in n steps. Now take any integer δ so that $i + \delta, j + \delta \in S$. Note that

$$(j + \delta) = (i + \delta) - \hat{d} + mQ,$$

where $\hat{d} = d_1 + d_2 + \dots + d_n$, the total demand in the n periods. Therefore $j + \delta$ can also be reached from $i + \delta$ in the same n steps. This observation shows that if states i and j communicate then $i + \delta$ and $j + \delta$, given both are in S , also communicate. Consequently, the states in any recurrent subset must be equally spaced; i.e., each recurrent subset can be expressed as $\{r + \Delta, r + 2\Delta, \dots, r + q\Delta\}$ for some integers r, Δ and q with $\Delta > 1$, $R - \Delta < r \leq R$ and q being

the largest integer so that $r + q\Delta \leq R + Q$. Note that there may be many recurrent subsets. Let the initial state of the Markov chain be in a particular recurrent subset S_0 . Clearly, the Markov chain will remain in S_0 forever. Let

$$S_0 \stackrel{\text{def}}{=} \{r + \Delta, r + 2\Delta, \dots, r + q\Delta\}.$$

We consider two cases:

Case 1. $q = 1$. In this case, there is only one recurrent state, i.e., $S_0 = \{r + \Delta\}$. The Markov chain remains in this state forever. Of course, the steady state distribution is uniform (over S_0). Moreover, one can show that d_t must be in multiples of Q . To see this, simply note that in order to make the one-step transition from state $r + \Delta$ back to itself, we must have $r + \Delta - d_t + mQ = r + \Delta$, $m \geq 0$ integer.

Case 2. $q \geq 2$. In this case, Q must be an integer multiple of Δ . This is easy to show. First, note that $r + 2\Delta$ can be reached from $r + \Delta$ in n periods for some $n \geq 1$. Let $\hat{d} > 0$ be the total demand in these n periods. Thus $r + 2\Delta = r + \Delta - \hat{d} + mQ$ for some positive integer m , or

$$\hat{d} = mQ - \Delta, \quad m \geq 1. \quad (23)$$

Now suppose the current state of the Markov chain is $r + q\Delta$. Following the same sample path for those n periods, the Markov chain arrives at the state $r + q\Delta - \hat{d} + m'Q$ for some nonnegative integer m' . Of course, this new state must still be in S_0 . Let it be $r + k\Delta \in S_0$. Thus

$$R < r + q\Delta - \hat{d} + m'Q = r + k\Delta < R + Q, \quad (24)$$

where the second inequality is strict because of the following. Suppose, to the contrary, the new state is $R + Q$. Therefore, $r + q\Delta = R + Q$. From $r + q\Delta - \hat{d} + m'Q = R + Q = r + q\Delta$, we have $\hat{d} = m'Q$. Using this in (23) leads to $(m - m')Q = \Delta$, a contradiction because both Δ and Q are positive integers, with $\Delta < Q$ because of $q \geq 2$. Combining (23) and (24),

$$q\Delta + m'Q = mQ - \Delta + k\Delta. \quad (25)$$

On the other hand, from (23),

$$r + q\Delta - \hat{d} + mQ = r + q\Delta + \Delta.$$

Note that from the definition of q and the fact that $\Delta < Q$,

$$R + Q < r + q\Delta + \Delta < R + 2Q.$$

Therefore,

$$R + Q < r + q\Delta - \hat{d} + mQ < R + 2Q.$$

Comparing the above inequalities with the ones in (24), we have $m = m' + 1$. Using this in (25) reveals that Q is an integer multiple of Δ .

We can further show that d_t is in multiples of Δ . To see this, consider the one-step transition from state $r + \Delta$. The next state is $r + \Delta - d_t + mQ = r + k\Delta$ for some integers m and k with $m \geq 0$ and $1 \leq k \leq q$. Because Q is an integer multiple of Δ , so is d_t . It is then straightforward to show that the steady-state distribution of the Markov chain is uniform over the states in S_0 .

Optimization and Optimality Issues

Consider the single-location inventory problem in §2. Recall that $\bar{G}(y)$ is assumed to be quasiconvex and is minimized at $y = R$. Suppose the Markov chain $\{y_t'\}$ is reducible. Then, from the above analysis, the long-run average value of $G(y_t')$ is

$$\frac{\Delta}{Q} \sum_{k=1}^q G(r + k\Delta),$$

where Δ , q , and r are integers with $\Delta > 1$, $q \geq 1$, and r being determined by the initial inventory position (if $q = 1$ set $\Delta = Q$). As in §2, this long-run average value is lower bound on the long-run average cost of any feasible policy. Moreover, the lower bound is achieved by the (R, nQ) policy (with the R defined above). Notice that for the reducible case, there are multiple optimal reorder points (the one defined above is one of them). In sum, the single-location results apply to general demand distributions. The same is true for the multi-echelon results because a multi-echelon system is eventually reduced to a single-location problem, and the lower-bounding arguments do not depend on the demand distribution.

ACKNOWLEDGMENT

The author thanks the associate editor and the referees for their insightful comments and suggestions that have significantly improved the paper.

REFERENCES

- Axsater, S. 1993a. Simple evaluation of echelon stock (R, Q) policies for two-level inventory systems. Working Paper, Lulea University of Technology, Sweden.
- . 1993b. Continuous review policies for multi-level inventory systems with stochastic demand. In *Handbook in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory*. S. Graves, A. Rinnooy Kan, and P. Zipkin (eds.) North Holland.
- , K. Rosling. 1993. Installation vs echelon stock policies for multilevel inventory control. *Management Sci.* **39** 1274–1280.
- Cachon, G. 1995. Exact evaluation of batch-ordering inventory policies in two-echelon supply chains with periodic review. Working Paper, Duke University.
- Chen, F., Y.-S. Zheng. 1993. Inventory models with general backorder costs. *Euro. J. Oper. Res.* **65** 175–186.
- , —. 1994a. Evaluating echelon stock (R, nQ) policies in serial production/inventory systems with stochastic demand. *Management Sci.* **40** 1262–1275.
- , —. 1994b. Lower bounds for multi-echelon stochastic inventory systems. *Management Sci.* **40** 1426–1443.
- , —. 1997. One-warehouse multi-retailer systems with centralized stock information. *Oper. Res.* **45** 275–287.
- , —. 1998. Near-optimal echelon-stock (R, nQ) policies in multi-stage serial systems. *Oper. Res.* **46** 592–602.
- Clark, A., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* **6** 475–490.

- , ——, 1962. Approximate solutions to a simple multi-echelon inventory problem. In *Studies in Applied Probability and Management Science*. K. J. Arrow, S. Karlin, H. Scarf (eds.). Stanford University Press, Stanford, CA.
- De Bodt, M., S. Graves. 1985. Continuous review policies for a multi-echelon inventory problem with stochastic demand. *Management Sci.* **31** 1286–1295.
- Federgruen, A., P. Zipkin. 1984. Computational issues in an infinite horizon, multi-echelon inventory model. *Oper. Res.* **32** 818–836.
- Hadley, G., T. Whitin. 1961. A family of inventory models. *Management Sci.* **7** 351–371.
- Rosling, K. 1989. Optimal inventory policies for assembly systems under random demands. *Oper. Res.* **37** 565–579.
- Schmidt, C., S. Nahmias. 1985. Optimal policy for a two-stage assembly system under random demand. *Oper. Res.* **33** 1130–1145.
- Veinott, A. 1965. The optimal inventory policy for batch ordering. *Oper. Res.* **13** 424–432.
- Zheng, Y.-S., F. Chen. 1992. Inventory policies with quantized ordering. *Naval Res. Logist.* **39** 285–305.