

A Decomposition Approach for a Capacitated, Single Stage, Production-Inventory System

Ganesh Janakiraman¹
IOMS-OM Group
Stern School of Business
New York University
44 W. 4th Street, Room 8-160
New York, NY 10012-1126.

John A. Muckstadt
286 Rhodes Hall
School of Operations Research and Industrial Engineering
Cornell University
Ithaca, NY 14853.

July 2003

Subject Classification: Inventory/Production: Policy - optimality of state dependent modified base-stock policies, Markov modulated demands and lead times. Inventory/Production: Planning Horizon - Finite and Infinite, discounted and average cost.

¹This author was a Visiting Assistant Professor in the School of ORIE at Cornell University when this work was conducted.

Abstract

We provide a new proof of the optimality of modified base-stock policies for capacitated single item, single location systems. A key idea in the proof is that the system can be decomposed into C subsystems having unit capacity, where C is the capacity per period and that an optimal policy for the system can be found by managing each subsystem independently and optimally.

1 Introduction

Consider a production and inventory system, in which we manage a single item, that has the following characteristics. Time is divided into periods of equal length. At the beginning of each period, an order is placed on an outside supplier. The processing capacity of this supplier is limited to C units per period, which limits the amount ordered. (Alternately we can also interpret this capacity as a contracted upper bound on order sizes.) Once an order is placed, it is completed and available in inventory in no more than $m - 1$ periods. The exact length of this lead time is a random variable. Furthermore, the sequence in which orders are received from the supplier corresponds to the sequence in which they are placed on the supplier. Subsequent to the time the order is placed in a period, orders are received from the supplier; customer demands are then observed. Demand and lead time processes are governed by an exogenous Markov Chain. All excess demand is backordered. At the end of each period, both holding and backorder costs are incurred.

The purpose of this paper is to present a new decomposition technique to prove that a modified base-stock policy is optimal for managing this system. That is, in each period there exists a desired order-up-to level, and an order is placed to bring the stock level up to this value if there is sufficient capacity available to do so. Otherwise, C units are ordered. This result extends earlier work on capacitated systems with deterministic lead times. The significance of the decomposition technique is not in this simple extension because we also demonstrate how traditional dynamic programming methods can be used to achieve the same result. The significance is the potential applicability of the proof technique employed here in more complicated inventory models.

The remainder of the paper is organized as follows. We provide a summary of related papers in section 2. Some notation and preliminaries are developed in section 3. The main optimality result is proved in section 4 for both finite (sections 4.1 and 4.2) and infinite planning horizons (section 4.3). More specifically, in section 4.1, we show that the system can be decomposed into C independent subsystems, and in section 4.2 we characterize the optimal policy for a subsystem as one which leads to a modified base-stock policy in the original system. In section 4.3, the optimality result for the infinite horizon case is discussed and the well known newsvendor formula is extended to these systems when the performance measure is the average cost per period over an infinite horizon. In the same section, we also demonstrate how existing results from the literature can be extended using the traditional dynamic programming method for these systems.

2 Literature Survey

The optimality of base-stock policies for uncapacitated systems has been derived under various assumptions by Karlin and Scarf (1958), Clark and Scarf (1960) and Chen and Song (2001), among several others. The optimality of modified base-stock policies for capacitated

systems was first established by Federgruen and Zipkin (1986a) and Federgruen and Zipkin (1986b) with the assumption of stationary demands. This result was extended to the case of non-stationary demand processes in Aviv and Federgruen (1997), Kapuscinski and Tayur (1998) and Scheller-wolf and Tayur (1997). The proofs of these results for capacitated systems hold when lead times are deterministic. Parker and Kapuscinski (2001) extend these results to a special class of two echelon, capacitated serial systems by introducing and proving the optimality of a new class of policies called “modified, echelon base-stock policies”.

All the papers mentioned above, with the exception of Chen and Song (2001) who use a lower-bounding argument, use dynamic programming to prove the optimality of base-stock policies. However, in Muharremoglu and Tsitsiklis (2001), the paper that is most closely related to our work, a novel alternate proof technique has been developed. They refer to this new approach as the “single-unit decomposition” approach. They consider uncapacitated serial systems with Markov modulated demand and lead times under the assumptions that lead times do not cross and excess demand is backordered. They establish the optimality of state dependent echelon base-stock policies for these systems. The crux of their work is the decomposition of the original inventory control problem into a countable collection of subproblems, each of which consists of one unit of inventory and demand.

In this paper, we use a similar decomposition approach to prove the optimality of state dependent modified base-stock policies for single location capacitated systems with Markov modulated demands and lead times. Our proof is based on decomposing the system into C subsystems, where C is the capacity, measured in units, in each period. Each of the C subsystems consists of a countable number of units and customers and has unit capacity.

As we will see in the next section, our model permits linear ordering, holding and backorder costs whereas the earlier work on capacitated, single-stage systems permitted convex holding and backorder costs. Our approach is limited in that sense. For the sake of completeness, we also describe briefly how the traditional dynamic programming approach for inventory systems can be extended to prove the optimality of modified base-stock policies for a stochastic lead time model with convex holding and backorder costs.

The lead time model we employ is based on the model first introduced by Kaplan (1970) and subsequently redefined and simplified by Nahmias (1979). With the same model, Erhardt (1984) demonstrated the optimality of (s, S) policies for uncapacitated, single-stage systems with a fixed order cost.

3 Notation and Definitions

Let us start by presenting the notation and key definitions that we will use throughout the paper. Initially assume the planning horizon of the system consists of T periods, numbered $t = 1, 2, \dots, T$ in that order. In section 4.3, we examine the infinite horizon case. Assume

there is an exogenous finite-state, ergodic Markov Chain s_t that governs the lead time and demand processes. The state of s_t is observed at the beginning of each period t . The transition probabilities for the Markov Chain s_t are assumed to be known. Furthermore, given s_t , the probability distribution of d_t , the demand in period t , is known.

The lead time process evolves as follows. There is a random variable ρ_t , whose distribution is determined completely by s_t , that specifies the least “age” of orders that will be delivered in period t at location 1. This means all *orders that are outstanding* at the beginning of period t (after order placement in t) and were placed in period $t - \rho_t$ or earlier are delivered in period t . We assume that the sample space of the random variable ρ_t is $\{0, 1, 2, \dots, m - 1\}$, since the maximum lead time of an order is $m - 1$ periods. Zipkin (1986) noted that this lead time model does not permit arbitrary marginal distributions for the lead time even when $\{\rho_t\}$ is an i.i.d. sequence of random variables. He showed that the failure rate of this distribution has to be non-decreasing and has to approach 1 in the limit. From a practical point of view, this model is appropriate when the supplier sequences the orders on a first-come, first-serve basis and processes a random number (independent of the size of the waiting orders), possibly zero, of these orders every period.

We consider each unit of demand as an individual customer. Suppose at the beginning of period 1 there are v_0 customers waiting to have their demand satisfied. These customers are indexed $1, 2, \dots, v_0$ in any order. All subsequent customers are indexed $v_0 + 1, v_0 + 2, \dots$ in the order of the period of their arrivals, arbitrarily breaking ties among customers that arrive in the same period.

Next, we define the concept of the *distance of a customer* at the beginning of any period. (See figure 1.) Every customer who has been served is at distance 0; every customer who has arrived, placed an actual order, but who has not yet received inventory, is at distance 1; all customers arriving in subsequent periods are said to be at distances 2, 3, \dots corresponding to the sequence in which they will arrive. Distances are assigned to customers that arrive in the same period in the same order as their indices. This ensures that customers with higher indices are always at “higher” distances. Also, note that there is exactly one customer at each of the distances 2, 3, \dots .

Next, we define the concept of a *location of a unit* at the beginning of a period before order placement. (See figure 1.) There are $m + 2$ possible locations at which a unit can exist. If the unit has been used to satisfy a customer’s order, the location of this unit is 0. If it is part of the inventory on hand, it is in location 1. If it has been ordered from the supplier n periods ago ($1 \leq n < m$) but is still in the pipeline, due to its lead time, it is in location $m + 1 - n$. If the unit has not been ordered from the supplier, it is in location $m + 1$.

At the beginning of period 1, we assign an index to all units in a serial manner, starting with units at location 1, then location 2, \dots , location $m + 1$, and arbitrarily assign an order

to units present at the same location. We assume a countably infinite number of units is available at the supplier, that is, location $m + 1$, at the beginning of period 1. However, at most C of these units can be ordered in any period.

j and k are used to denote the indices of units as well as customers. Denote the distance of customer j at the beginning of period t by y_{jt} and the location of unit j at the beginning of period t by z_{jt} . Let \mathcal{S} refer to the entire system with all the units and all the customers and the capacity constraint of C units per period.

The state of the system at the beginning of period t is given by the vector $\mathbf{x}_t = (s_t, (z_{1t}, y_{1t}), (z_{2t}, y_{2t}), \dots)$.

Next, we explain the sequence of events in period t . (Though redundant at this point, we repeat the phrase “in \mathcal{S} ” for the sake of conciseness at a later stage in the paper.)

(1) \mathbf{x}_t is observed. **(2)** Units in \mathcal{S} at location 2 become available in inventory, that is, location 1. Units in \mathcal{S} at locations 3, 4, \dots, m move to the next location. Ordering decisions are made for units in \mathcal{S} at location $m + 1$. Formally, $u_{jt} \in \{0, 1\}$ is decided for all $j \in \mathcal{S}$ such that $z_{jt} = m + 1$. Unit j is ordered (we will use “released” and “ordered” interchangeably) *if and only if* $u_{jt} = 1$. Every unit ordered in this period moves to location m . Number of units released from the supplier is $q_t = \sum_{j \in \mathcal{S}: z_{jt} = m+1} u_{jt}$. For *capacity feasibility*, q_t is C or less. **(3)** ρ_t is realized; if $\rho_t \leq m - 2$, all units in \mathcal{S} at locations 2 through $m - \rho_t$ arrive from the supplier, and are at location 1. **(4)** Demand d_t is realized. That is, customers in \mathcal{S} at distances 2, 3, $\dots, 2 + d_t - 1$ all arrive and are by definition at distance 1. Customers in \mathcal{S} currently at distances $2 + d_t, 3 + d_t, \dots$ move d_t steps towards distance 1. **(5)** Units on-hand in \mathcal{S} and waiting customers in \mathcal{S} are matched to the extent possible. That is, as many waiting customers are satisfied as possible and as many units on hand are consumed as possible. *We assume that units and customers in \mathcal{S} at location 1 and distance 1, respectively, are matched in a first-come, first-serve order based on the indices, starting from the lowest index.* **(6)** h dollars are charged per unit of inventory on hand (at location 1) in \mathcal{S} and b dollars are charged per waiting customer (at distance 1) in \mathcal{S} . Clearly, only one of these costs will be incurred in any period. We assume $b > 0$. This ensures that if the inventory position is negative, then the optimal policy will be to increase the inventory position to some non-negative level, if possible.

Note: Though we have not mentioned purchase costs in the model, linear purchase costs payable at the time of receipt of inventory can easily be accommodated.

The performance measure under consideration is the expected sum (discounted or undiscounted) of costs over the T period planning horizon. A set of mappings, one for every t , from \mathbf{x}_t to (u_{jt}) is called a *policy*. A *feasible policy*, is one that satisfies the capacity constraint. That is, $\sum_{j: z_{jt} = m+1} u_{jt} \leq C$ for all t and \mathbf{x}_t . A *monotone policy* is one that satisfies the constraint $u_{jt} \geq u_{(j+1)t}$ for all j, t and \mathbf{x}_t . In words, a monotone policy always releases

a lower indexed unit no later than a higher indexed unit. Similarly, we define a *monotone state* to be one where lower indexed units are in the same or lower indexed locations. That is, $z_{kt} \leq z_{jt}$ if $k \leq j$. Next we state a lemma with some facts about monotone policies. The proofs are trivial and hence omitted.

Lemma 1 (i) *For every feasible policy, we can construct a monotone, feasible policy which attains the same cost in every period along every sample path. Consequently, the class of monotone policies contains an optimal policy.* (ii) *When a monotone policy is used in every period, since customer demands are satisfied based on the indices and orders do not cross, no unit other than j can satisfy customer j 's demand. Thus, unit j and customer j are matched when monotone policies are used.* (iii) *When a monotone policy is used in every period, \mathbf{x}_t is a monotone state for all t .*

From now on, our attention is restricted to monotone states when analyzing the system \mathcal{S} without any loss of generality. In the next section, we develop a proof of the optimality of modified base-stock policies for periodic review, single location capacitated systems of the type we have described.

4 Optimality of Modified Base-stock Policies

In this section, we first show that the system can be decomposed into C subsystems of unit capacity. Subsequently, we prove that each subsystem can be managed optimally by using a “critical distance” policy. When the same “critical distance” policy is used to manage each subsystem, the system follows a modified base-stock policy. A discussion about infinite horizon results, extending the traditional dynamic programming approach to these systems and a computational approach are presented in the final subsection.

4.1 Decomposition of the System into C Subsystems

As mentioned earlier, Muharremoglu and Tsitsiklis (2001) provided the basis and the motivation for this work. Let us summarize the basic steps of their proof methodology for uncapacitated systems. First, they recognized that the cost of the system is the sum of the costs of each of the unit-customer pairs. Second, they observed that each of these pairs can be controlled independently and optimally and that the resulting policy is optimal for the entire system. Third, they examined an individual unit-customer problem and showed that the optimal policy is a “critical distance” policy (*this terminology was not used in their paper*): Release a unit if and only if the corresponding customer is closer than a critical distance. Last, they recognized that operating each unit-customer pair using a critical distance policy produces an echelon base-stock policy in the original system.

Unfortunately, in capacitated systems, we cannot operate each unit-customer pair independently. In the capacitated system, the capacity constraint links the “release” decisions of the units. Consequently, we propose an alternate decomposition which we describe below.

Definition 1 *Subsystem w , represented by \mathcal{S}_w , $1 \leq w \leq C$, refers to the subset of unit-customer pairs with indices $w, w + C, w + 2C, \dots$. Each subsystem has a unit capacity.*

The intuitive reason for defining a subsystem in this way is the fact that when a monotone policy is used in \mathcal{S} , unit j can be affected by the *capacity constraint* in any period only if unit $j - C$ has still not been released. This provides a natural connection between unit j and unit $j - C$ for any j .

The sequence of events in \mathcal{S}_w are the same as steps (1)-(6) in section 3 with the following obvious modifications: \mathcal{S} is replaced by \mathcal{S}_w and C is replaced by 1. Note that we still assume that \mathbf{x}_t , the information about the entire system \mathcal{S} , is available when managing \mathcal{S}_w .

For subsystem \mathcal{S}_w , a policy is *monotone* if unit j is released no later than unit $j + C$ for any unit j in \mathcal{S}_w . A policy for \mathcal{S}_w is *feasible* if it never releases more than one unit in a period. Note that the class of monotone policies is optimal to each subsystem \mathcal{S}_w and these policies ensure that unit j is matched with customer j .

We now show that the subsystems can be optimally managed independent of each other even though the lead time and demand processes of different subsystems are *not stochastically independent* and that these policies, when combined, form an optimal policy for \mathcal{S} . Let us first define $\mathbf{x}_t^w =_{def} (s_t, (z_{wt}, y_{wt}), (z_{(w+C)t}, y_{(w+C)t}), \dots)$, that is, the information in \mathbf{x}_t that pertains to \mathcal{S}_w .

Theorem 2 *For any monotone state \mathbf{x}_t , the optimal expected discounted (undiscounted) cost in periods $t, t + 1, \dots T$ for system \mathcal{S} equals the optimal expected discounted (undiscounted) cost in periods $t, t + 1, \dots T$ for the group of subsystems $\{\mathcal{S}_w\}$. \mathcal{S}_w can be optimally managed using \mathbf{x}_t^w instead of \mathbf{x}_t . Furthermore, when each \mathcal{S}_w is managed optimally using \mathbf{x}_t^w in periods $t, t + 1, \dots T$, the resulting policy is optimal for the entire system, \mathcal{S} .*

Proof : A *feasible* policy for subsystem \mathcal{S}_w can be constructed from any *feasible, monotone* policy in \mathcal{S} by implementing the (u_{jt}) actions suggested by the latter policy on the elements of \mathcal{S}_w . Similarly, a *feasible* policy for \mathcal{S} can be constructed from any set of *feasible* policies for $\{\mathcal{S}_w\}$ by combining these policies as follows: for every unit $j \in \mathcal{S}$ implement the u_{jt} action suggested by the policy for the subsystem to which j belongs. Furthermore, note that the cost incurred by \mathcal{S} in any period is the sum of the costs incurred by the units and customers belonging to the C subsystems. Combining these three observations with the optimality of the class of *monotone* policies in \mathcal{S} proves the first statement.

Next, notice that the cost incurred in \mathcal{S}_w in period t depends only on \mathbf{x}_t^w , and the probabilities necessary to describe the transition from a state \mathbf{x}_t^w to \mathbf{x}_{t+1}^w depend only on the actions in \mathcal{S}_w and the information in \mathbf{x}_t^w . This proves the second statement.

The last statement in the theorem is a direct consequence of the first two statements. \square

It is now clear that examining a subsystem's optimal policy is useful in determining the optimal policy for the entire system. In the next section, we derive the structure of the optimal policy for a subsystem.

4.2 Optimal Policy Structure for a Subsystem

Before examining an individual subsystem, we first observe that all subsystems are identical in the sense that (i) they have identical cost structures and (ii) given a state (\mathbf{x}_t^w) and a fixed operating policy for a subsystem, the stochastic evolution of the subsystem is independent of the index w . Consequently, the optimal policy(ies) is(are) identical across all subsystems.

Next, we develop some necessary preliminaries about optimal policies for the subsystems. Consider only the class of monotone policies for the subsystems, which contains at least one optimal policy. Therefore, in any period t and any \mathcal{S}_w , the state \mathbf{x}_t^w is monotone, that is, $z_{wt} \leq z_{(w+C)t} \leq \dots$. Consecutive units in location $m+1$ belonging to \mathcal{S}_w have indices that differ by C . Let n_{wt} be the lowest index, that is, unit n_{wt} is the only candidate for being released in period t in \mathcal{S}_w . Let us now examine the information that is actually required to manage \mathcal{S}_w optimally. Consider a given t , s_t and n_{wt} . Monotonicity implies that all units indexed below n_{wt} in \mathcal{S}_w have already been released from location $m+1$. Consequently, the expected costs associated with all these units and the corresponding customers are sunk; that is, these costs are the same for all policies from period t onward. Therefore, having information about the locations (distances) of units (customers) in \mathcal{S}_w with indices below n_{wt} is unnecessary. The location of all units with indices higher than n_{wt} is $m+1$. It is now clear that (s_t, n_{wt}) and the distances of all customers in w with indices n_{wt} and higher is sufficient for this subsystem. Even this information turns out to be more than needed, as we will see next.

Let us define y_{wt} to be $y_{n_{wt}}$, the distance of customer n_{wt} corresponding to the next unit waiting to be released in \mathcal{S}_w . Since unit n_{wt} is still at location $m+1$, customer n_{wt} cannot be at distance 0. Assume $y_{wt} > 1$, that is, customer n_{wt} has not yet arrived. This implies that all customers with higher indices have also not arrived and that the subsequent customer in w is at distance $y_{wt} + C$, the next one at $y_{wt} + 2C$ and so on. That is, if $y_{wt} > 1$, all other required information about customer distances is already known. Let us now assume that $y_{wt} = 1$. In this case, customer n_{wt} has arrived and it is possible that some subsequent customers have also arrived. However, since customer n_{wt} has arrived, it is optimal to *release* unit n_{wt} since $b > 0$. Consequently, any information about other customer distances is unnecessary in determining the optimal decision for \mathcal{S}_w in period t . It is now clear that (s_t, n_{wt}, y_{wt}) is a minimally sufficient information vector to manage \mathcal{S}_w optimally in period t using a monotone policy. Furthermore, since all units are identical and all subsystems are identical, we can use a more compact information vector (s_t, y) where $y = y_{wt}$. We define $R_t^*(s_t, y) \subseteq \{0, 1\}$ as the set of optimal decisions for unit n_{wt} at time t if the state of the exogenous Markov Chain is s_t and if y_{wt} is y , where 0 refers to holding the unit back whereas

1 refers to ordering or releasing the unit in period t .

Next, we show that there is a “critical distance” policy that is optimal for a subsystem. We need the following Lemma to prove this fact. The lemma states that if it is (uniquely) optimal for \mathcal{S}_w to release unit n_{wt} in period t when the system is in the Markovian-state s_t and customer n_{wt} is at a distance $y + 1$, then it would be optimal to release it if the customer were any closer.

Lemma 3 $R_t^*(s_t, y + 1) = \{1\}$ implies that $R_t^*(s_t, y) \supseteq \{1\}$.

Proof: The proof is by contradiction. Assume the statement is not true. That is, there exists t , s_t and y such that $R_t^*(s_t, y + 1) = \{1\}$ and $R_t^*(s_t, y) = \{0\}$. Consider a monotone state \mathbf{x}_t in \mathcal{S} and the corresponding states \mathbf{x}_t^w for the subsystems such that (a) the exogenous Markov Chain is at state s_t , (b) there is a w in $\{1, 2, \dots, C - 1\}$ such that y_{wt} is y and $y_{(w+1)t}$ is $y + 1$ and (c) $n_{(w+1)t}$ is $n_{wt} + 1$. Consider some monotone policy for \mathcal{S} . Monotonicity implies that this policy would choose one of the following three pairs of actions for units n_{wt} and $n_{wt} + 1$: (a) release both n_{wt} and $n_{wt} + 1$, (b) hold both n_{wt} and $n_{wt} + 1$ and (c) release n_{wt} and hold $n_{wt} + 1$.

By assumption, cases (a) and (c) are suboptimal for \mathcal{S}_w while cases (b) and (c) are suboptimal for \mathcal{S}_{w+1} . This implies that any monotone policy for \mathcal{S} is suboptimal for at least one of subsystems w and $w + 1$. So, any monotone policy for \mathcal{S} has a higher expected cost than the sum of the optimal costs for the C subsystems from period t onwards, which is the same as the optimal cost for \mathcal{S} by Theorem 2. This implies that no monotone policy can be optimal for \mathcal{S} , which contradicts our earlier assertion about the optimality of some monotone policy. Therefore, our assumption about $R_t^*(s_t, y)$ and $R_t^*(s_t, y + 1)$ is invalid. \square

We use this Lemma to develop the notion of a “critical distance” policy. Let us define

$$y^*(t, s_t) \stackrel{\text{def}}{=} \max\{y : R_t^*(s_t, y) \supseteq \{1\}\}.$$

$y^*(t, s_t)$ is defined in such a way that it is optimal to release unit n_{wt} if and only if customer n_{wt} is at a distance of $y^*(t, s_t)$ or closer. This distance $y^*(t, s_t)$ is the “critical distance” at time t and Markovian state s_t for every subsystem. Consider the policy

$$R_t(s_t, y) = \{1\} \text{ if and only if } y \leq y^*(t, s_t).$$

Policy R_t is an optimal policy for every subsystem. The next observation we make is that when policy R_t is used in period t for every subsystem, the resulting policy for the original system \mathcal{S} is a modified order-up-to policy with an order-up-to level of $y^*(t, s_t) - 1$. This can be shown either using an algebraic proof or a more intuitive argument which we provide as our proof.

Theorem 4 *The optimal policy for \mathcal{S} is to release as many units as possible to raise the inventory position to $y^*(t, s_t) - 1$ at time t when in Markovian state s_t and the planning horizon consists of T periods. That is, a state dependent modified order-up-to or base-stock policy is optimal for the entire system when the planning horizon is finite.*

Proof: Let us use the policy R_t on each subsystem since we know this is an optimal policy for the entire system. In any monotone state \mathbf{x}_t and the corresponding states \mathbf{x}_t^w for the subsystems, the units that are candidates for release in each subsystem bear consecutive labels. (More formally, a simple inductive argument can be used to show that there is a cyclic sequence starting from some v such that $n_{vt} = n_{(v-1)t} + 1 = \dots = n_{1t} + (v - 1) = n_{Ct} + v = n_{(C-1)t} + (v + 1) = \dots = n_{(v+1)t} + (C - 1)$.) By construction, the corresponding customers who have not arrived are in consecutive distances. Unit n_{wt} is released if customer n_{wt} is at distance 1, in other words, backordered, or in any of the subsequent $y^*(t, s_t) - 1$ distances. This is the same as ordering up to $y^*(t, s_t) - 1$ if the inventory position is below $y^*(t, s_t) - 1$ and the inventory position plus C exceeds or equals $y^*(t, s_t) - 1$. If the former condition fails, all C subsystems release a unit and if the latter condition fails, none does. This is the same as a modified-base-stock policy with target $y^*(t, s_t) - 1$. \square

4.3 Infinite Horizon

When the discount factor is strictly smaller than 1, all the analysis presented in the paper thus far holds for an infinite horizon. In other words, the optimality of modified base-stock policies for the infinite horizon, discounted cost model can be proved using our decomposition approach. Two questions remain: Does the optimality result hold when the holding and backorder costs are convex but non-linear? Does the optimality result hold for the infinite horizon, average cost model? We discuss these issues briefly using the traditional dynamic programming approach and results from Erhardt (1984) for uncapacitated, single-stage systems with stochastic lead times.

Erhardt (1984) considers an uncapacitated, single-stage inventory system with the same stochastic lead time model as ours. He allows a fixed ordering cost in addition to a linear ordering cost, both payable at the time of the receipt of inventory. Convex holding and shortage costs are charged. Expression (3) in his paper is a dynamic programming formulation for the finite horizon model, whose recursion is identical to that of a problem with zero lead times. (Note that his “ $\sum_{i=0}^m \alpha^i l_i g_n(i, y)$ ” is a convex function of y since the single period holding and backorder cost function is convex.) This dynamic programming recursion thus established the optimality of (s, S) policies, hitherto known for deterministic lead time models, to these stochastic lead time models. Observe that Erhardt’s expression (3) with this recursion remains valid for the capacitated system we are interested in with the following additions: (i) the capacity constraint, (ii) an additional state variable for the exogenous Markov Chain and (iii) the dependence of the probabilities l_i on the state of the Markov Chain. (Erhardt did not use a Markov modulated model; in his model, demands

are independently distributed and the random variables corresponding to our (ρ_t) 's are independent and identically distributed.) In other words, the dynamic program for the Markov modulated, capacitated inventory model with stochastic lead times can be cleverly rewritten as a dynamic program for a similar Markov modulated system with zero lead times. It is known that the optimal policy for this system is a state-dependent, modified base-stock policy for the finite horizon, infinite horizon discounted and infinite horizon average cost models. For example, Scheller-wolf and Tayur (1997) discuss Markov modulated models and Aviv and Federgruen (1997) study a periodic model with convex holding and backorder costs.

A key formula for computing the optimal base-stock level, developed by Tayur (1992) for stationary, capacitated systems with deterministic lead times, can also be extended to a stationary version of our model with stochastic lead times. We will assume the following: (a) $\{d_t\}$ is a sequence of i.i.d. random variables, (b) $\{\rho_t\}$ is a sequence of i.i.d. random variables, (c) $\{d_t\}$ and $\{\rho_t\}$ are independent, (d) the planning horizon is infinite and the performance measure is the long run expected cost per period, (e) a stationary modified base-stock policy is used and (f) $E(d_t) < C$.

The shortfall in a period is the amount by which the inventory position (after ordering) is below the base-stock level. The key idea in Tayur (1992) is that the shortfall process in an inventory system is related to the “water level” process in dams. These processes, which are subsets of a wider class of processes called stochastic storage processes, have been extensively studied by Prabhu (1998).

Assume that we are following a modified base-stock policy with parameter Y . In other words, in every period, we order enough to raise the inventory position up to Y if possible, or C otherwise. Define $V(t)$ to be the “shortfall” after ordering decisions are made in period t . Observe that $V(t)$ evolves according to the following recursive equation.

$$V(t+1) = (V(t) + d_t - C)^+ .$$

Define $I(t)$ and $N(t)$ to be the inventory position (after ordering) in period t and the net inventory at the end of period t , respectively. Let $D[t_a, t_b]$ be the cumulative demand in periods t_a, t_{a+1}, \dots, t_b . Note that $I(t)$ equals $(Y - V(t))$. Next, observe that all orders placed in period $\lambda_t =_{def} \max\{s - \rho_s : s \leq t\}$ or earlier arrive by t and all subsequent orders are outstanding at t . Therefore, $N(t)$ equals $(I(t - \lambda_t) - D[t - \lambda_t, t])$. Consequently, the cost incurred in period t is

$$h \cdot (N(t))^+ + b \cdot (N(t))^- ,$$

which is nothing but

$$h \cdot (Y - V(t - \lambda_t) - D[t - \lambda_t, t])^+ + b \cdot (Y - V(t - \lambda_t) - D[t - \lambda_t, t])^- .$$

Let $V(\infty)$ be the stationary shortfall random variable. The existence of the stationary distribution for the shortfall is known from Theorem 1 of Glasserman and Tayur (1994).

Since the distributions of the random variables λ_t and $D[t - \lambda_t, t]$ do not depend on t , we can suppress the parameters and write D instead of $D[t - \lambda_t, t]$. Furthermore, it can be shown that the distribution of D is the same as the marginal distribution of demand over a leadtime. (See for example the discussion on page 771 of Zipkin (1986) for a similar continuous time model.) The long run expected cost per period is then

$$h \cdot E[(Y - V(\infty) - D)^+] + b \cdot E[(Y - V(\infty) - D)^-] .$$

Minimizing this function gives the standard newsboy formula. The optimal base-stock parameter Y^* is computed as follows:

$$Y^* = \min\{Y : P(Y \geq V(\infty) + D) \geq b/(b + h)\} .$$

Glasserman (1997) and Roundy and Muckstadt (2000) discuss computationally efficient approximations for the distribution of $V(\infty)$.

5 Conclusions

In this paper, we develop a decomposition technique to analyze single-stage, capacitated systems with stochastic non-crossing lead times. Our proof methodology, an extension of the “single-unit, single-customer” approach developed recently by Muharremoglu and Tsitsiklis (2001) for uncapacitated serial systems, is based on a decomposition of the original system into C subsystems, where C is the capacity per period. Each subsystem consists of a countable number of units and customers and has unit capacity in each period. We prove that the optimal policy for a subsystem is a “critical distance” policy and that, when each subsystem is managed independently and optimally, the resulting policy is also optimal for the original system. This policy is shown to be a modified base-stock policy. Our approach is applicable to both the finite-horizon and infinite-horizon, discounted cost models. It is worth reiterating that the contribution of the paper is the proof technique and not the optimality result. This fact is exemplified by our discussion on transforming stochastic lead time models into zero lead time models using the work of Erhardt (1984), which implies that the optimality results that exist in the literature for capacitated systems with deterministic lead times easily carry over to stochastic lead time models with no order-crossing. Furthermore, it appears from our ongoing research that the decomposition approach developed here for single stage capacitated system can also be successfully used for a class of serial inventory systems with capacity constraints. This is an indication of the significance of this decomposition approach.

6 Acknowledgement

We are grateful to Professor Robin Roundy for carefully reading the paper and providing several important suggestions. We sincerely thank Professor Peter Jackson and two anonymous referees for providing valuable comments that improved the exposition of the paper.

References

- Aviv, Y. and Federgruen, A. (1997), “Stochastic Inventory Models with Limited Production Capacity and Periodically Varying Parameters,” *Probability in the Engineering and Informational Sciences*, 11, 107–135.
- Chen, F. and Song, J. (2001), “Optimal Policies for Multi-echelon Inventory Problems with Markov-modulated demand,” *Operations Research*, 49(2), 226–234.
- Clark, A. and Scarf, H. (1960), “Optimal Policies for a Multiechelon Inventory Problem,” *Management Science*, 6, 475–490.
- Erhardt, R. (1984, Jan.-Feb), “(s,S) Policies for a Dynamic Inventory Model with Stochastic Lead Times,” *Operations Research*, 32(1), 121–132.
- Federgruen, A. and Zipkin, P. (1986a), “An Inventory Model with Limited Production Capacity and Uncertain Demands I. The Average-Cost Criterion.,” *Mathematics of Operations Research*, 11(2), 193–207.
- Federgruen, A. and Zipkin, P. (1986b), “An Inventory Model with Limited Production Capacity and Uncertain Demands II. The Discounted-Cost Criterion.,” *Mathematics of Operations Research*, 11(2), 208–215.
- Glasserman, P. (1997), “Bounds and Asymptotics for Planning Critical Safety Stocks,” *Operations Research*, 45(2), 244–257.
- Glasserman, P. and Tayur, S. (1994), “The Stability of a Capacitated, Multi-Echelon Production-Inventory System Under a Base-Stock Policy,” *Operations Research*, 42(5), 913–925.
- Kaplan, R. (1970, March), “A Dynamic Inventory Model With Stochastic Lead Times,” *Management Science*, 16(7), 491–507.
- Kapuscinski, R. and Tayur, S. (1998, Nov.-Dec.), “A Capacitated Production-Inventory Model with Periodic Demand,” *Operations Research*, 46(6), 899–911.
- Karlin, S. and Scarf, H. (1958), “Inventory Models Of The Arrow-Harris-Marschak Type With Time Lag,” *Studies in the Mathematical Theory of Inventory and Production*, 155–178.
- Muharremoglu, A. and Tsitsiklis, J. N. (2001), “A Single-Unit Decomposition Approach to Multi-Echelon Inventory Systems,” Working Paper, MIT.
- Nahmias, S. (1979, Sept.-Oct.), “Simple Approximations for a Variety of Dynamic Leadtime Lost-Sales Inventory Models,” *Operations Research*, 27(5), 904–924.

- Parker, R. and Kapuscinski, R. (2001), “Optimal Inventory Policies for a Capacitated Two-Echelon System,” Working paper, Yale School of Management.
- Prabhu, U. (1998), *Stochastic Storage Systems: Queues, Insurance Risk, and Dams*, New York: Springer.
- Roundy, R. O. and Muckstadt, J. A. (2000, Jan.), “Heuristic Computation of Period-Review Base Stock Inventory Policies,” *Management Science*, 46(1), 104–109.
- Scheller-wolf, A. and Tayur, S. (1997), “Reducing International Risk Using Quantity Contracts,” GSIA Working Paper, CMU, Pittsburgh, PA.
- Tayur, S. (1992), “Computing the Optimal Policy for Capacitated Inventory Models,” *Communications in Statistics: Stochastic Models*, 9, 585–598.
- Zipkin, P. (1986), “Stochastic Leadtimes in Continuous-Time Inventory Models,” *Naval Research Logistics Quarterly*, 33, 763–774.

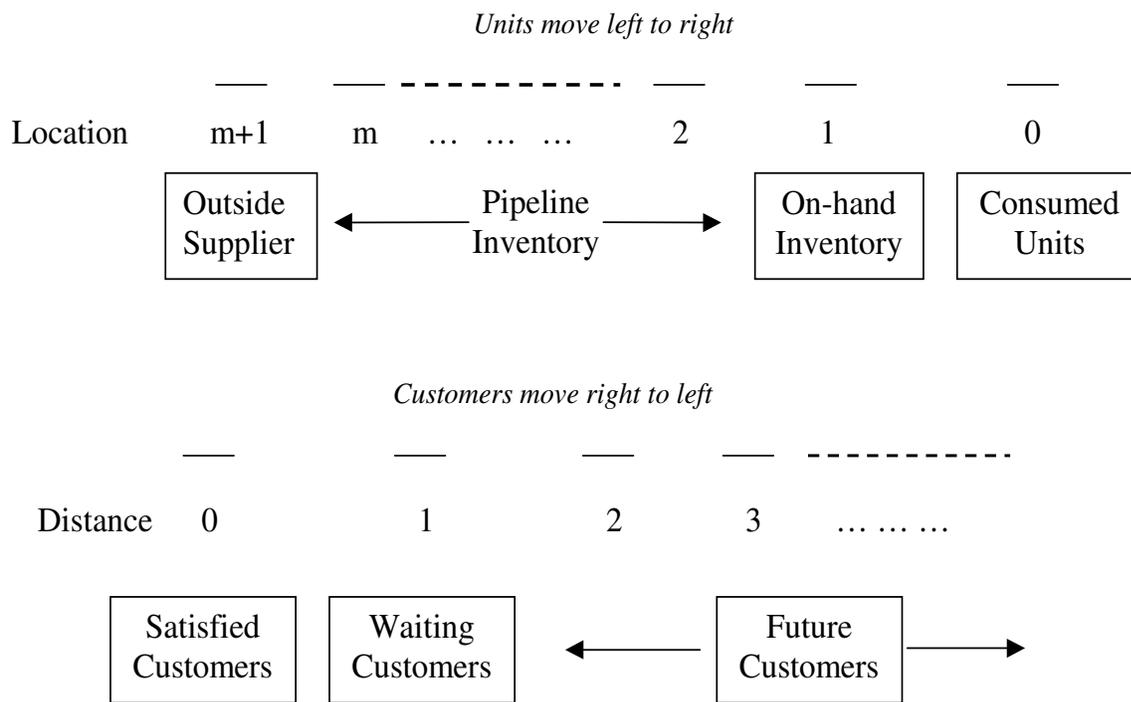


Figure 1: Locations of Units and Distances of Customers