A Decomposition Approach for A Class of Capacitated Serial Systems[1]

Ganesh Janakiraman [2]

IOMS-OM Group

Stern School of Business

New York University

44 W. 4th Street, Room 8-71

New York, NY 10012-1126.

John A. Muckstadt

286 Rhodes Hall

School of Operations Research and Industrial Engineering

Cornell University

Ithaca, NY 14853.

December 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.

[1] An earlier version of this paper was titled "Optimality of Multi-tier Base-stock Policies for a Class of Capacitated Serial Systems" and was presented at the Multi-Echelon Inventory Conference in June 2003.

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

Abstract

We study a class of two echelon, serial systems with identical ordering capacities or limits for both echelons. The lead times are initially deterministic. For the case where the lead time to the upstream echelon is one period, the optimality of state-dependent, modified echelon base-stock policies is proved using a decomposition approach. For the case where the upstream lead time is two periods, we introduce a new class of policies called "two-tier, base-stock policies," and prove their optimality. Some insight about the inventory control problem in N echelon, serial systems with identical capacities at all physical stages and stochastic, non-crossing, lead times everywhere is also provided. We argue that a generalization of two-tier, base-stock policies, which we call "multi-tier, base-stock policies" are optimal for these systems. Additional structural results about serial systems with infinite capacities at downstream stages are also presented.

1 Introduction

We consider a periodic review inventory control problem for a three stage supply chain with one supplier, one distribution center and one retailer. The supplier is considered as being external; that is, we are interested in optimally managing inventory at the distribution center and the retailer. Consequently, we refer to this three stage, serial system also as a two-echelon, serial system. The supplier and the distribution center can ship up to C units in any period. The retailer is only an inventory storage stage with unlimited storage capacity. We label the supplier, distribution center and retailer as L_3, L_2, and, L_1, respectively. Inventory at L_1 is used to meet customer demand. Excess demand at L_1 is assumed to be backordered. The costs considered are linear holding costs and linear backorder costs. Customer demands are Markov modulated and lead times are deterministic. Almost all the results proved in this paper hold when lead times are also Markov modulated subject to the restriction that they are non-crossing, that is, units shipped by L_3 (L_2) in this period reach L_2 (L_1) no later than units shipped in a subsequent period.

We prove the following results: (a) the optimal inventory control problem for this system can be decomposed into C problems, each one of which represents a subsystem that consists of a two echelon serial system with unit capacity at each stage, (b) under the additional assumption that the lead time between L_3 and L_2 is one period, the optimal policy is a modified echelon base-stock policy at L_1 and L_2 and (c) when the lead time between L_3 and L_2 is two periods, the optimal policy is a "two-tier, base-stock policy" (we will define this term later) at L_1 and L_2. The decomposition technique used for these two-echelon systems with deterministic lead times also holds for N echelon, serial systems with stochastic, non-crossing lead times and identical capacities everywhere. It can be seen that a straightforward generalization of two-tier, base-stock policies, which we call "multi-tier base-stock policies," are optimal for these systems. In addition, we prove some elementary structural results for more general serial systems consisting of a set of consecutive stages, starting from the most downstream stage, with infinite capacity while the remaining upstream stages are capacitated. We show that the optimal policy is an echelon base-stock policy for the uncapacitated stages of the system. The approach we use is an extension of the "single-unit, single-customer" approach introduced by Katircioglu and Atkins (1998) and Muharremoglu and Tsitsiklis (2003) for uncapacitated systems which was subsequently extended to capacitated, single stage systems by Janakiraman and Muckstadt (2003a).

Next, we briefly review the related literature. We refer the reader to Muharremoglu and Tsitsiklis (2003) and Tayur (1996) for more extensive reviews.

In their seminal paper, Clark and Scarf (1960) showed that echelon base-stock policies are optimal for uncapacitated serial systems with deterministic lead times under the assumption that demands are independent and identically distributed from period to period. The infinite horizon extensions were achieved by Federgruen and Zipkin (1984). A key extension of this result is Chen and Song (2001), where the optimality of state dependent echelon base-stock policies is proved when the demands are driven by a Markov Chain (also known as Markov modulated demand). This result has recently been extended by Muharremoglu and Tsitsiklis (2003) to systems where lead times are stochastic and non-crossing. They allow both lead times and demands to be Markov modulated.

The optimality of modified base-stock policies for a single stage, capacitated system with deterministic lead times and stationary demand was proved by Federgruen and Zipkin (1986a) and Federgruen and Zipkin (1986b). This work was extended to the case of periodic demand processes and Markov modulated demand processes by Aviv and Federgruen (1997) and Kapuscinski and Tayur (1998), respectively. Tayur (1992) uses the "shortfall distribution" using the theory of stochastic storage processes studied by Prabhu (1998) to provide a method to compute the optimal base-stock level for the stationary case. Janakiraman and Muckstadt (2003a) extended these papers further by allowing for stochastic, non-crossing lead times.

The only result that has been proved about the structure of the optimal policy for a serial system with capacities is due to Parker and Kapuscinski (2003). They consider a two echelon serial system of the type we described earlier with the additional assumption that the lead time between L_3 and L_2 is one period and the lead time between L_2 and L_1 is an arbitrary, deterministic integer. There is a capacity of C units per period at L_3 and L_2. (Note: Their model allows for a higher capacity at L_3 than L_2, but the optimal policy is the same as the optimal policy when the capacity at L_3 is replaced by the capacity at L_2.) They show that a modified echelon base-stock policy, specified by two parameters S_1 and S_2, is optimal for this system for both the finite and infinite horizon cases with Markov modulated demands. This policy suggests that L_1 should order up to the level S_1, if possible. L_2 should order as much as possible to raise the echelon inventory position to S_2 or enough to raise the inventory on hand at L_2 to C, whichever is smaller. One of the results we prove in this paper is that the form of the optimal policy is the same even when the lead times between L_2 and L_1 are stochastic and non-crossing. A key difference between their paper and ours is that they use the dynamic programming approach to obtain their results, while we use a decomposition approach to establish ours.

Glasserman and Tayur have made significant contributions to the analysis of multi-echelon inventory systems that have capacities and follow echelon base-stock policies. In Glasserman and Tayur (1994), they study stability conditions and long-run convergence properties. In Glasserman and Tayur (1995), they show how IPA (Infinitesimal Perturbation Analysis) can be used to find near-optimal base-stock levels. They develop simple approximations in Glasserman and Tayur (1996) to find base-stock levels.

Another stream of research that is related to our paper is due to Axsater. In Axsater (1990), he introduces a cost evaluation technique which is based on examining the costs associated with an individual unit and uses this to optimize base-stock levels for two-echelon inventory systems with one-for-one replenishment rules. He extends this technique to systems with batch-ordering in Axsater (1993). It should be noted, however, that Axsater's approach was to evaluate costs and find optimal parameters within the class of one-for-one replenishment policies or re-order point, re-order quantity policies. On the other hand, Katircioglu and Atkins (1998), Muharremoglu and Tsitsiklis (2003), Janakiraman and Muckstadt (2003a) and this paper are concerned with the derivation of the structure of the optimal policies using the single-unit, single-customer approach. Katircioglu and Atkins (1998) study a continuous review, single-stage system with arbitrary inter-arrival distributions with increasing failure rates. Muharremoglu and Tsitsiklis (2003) and Janakiraman and Muckstadt (2003a) study periodic-review systems for uncapacitated, serial systems and a single stage, capacitated system, respectively.

The remainder of the paper is organized as follows. Section 2 describes the inventory systems to be studied in greater detail and the notation used throughout the paper. In section 3, we study the capacitated, two echelon, serial system when the lead time between L_3 and L_2 is either one period or two periods. Specifically, we prove the optimality of modified, echelon base-stock policies (MEBS policies, in short) and two-tier, base-stock policies for these systems, respectively. Our proof methodology is based on a decomposition of a capacitated two-echelon serial system into a collection of two-echelon serial systems with unit capacities at both echelons. In section 4, we consider more general capacitated serial systems and prove elementary structural results for them. A brief discussion on holding and backorder costs that are convex in the waiting time is presented in section 5, and section 6 contains the conclusions.

2 Notation and Preliminaries

The most general system we consider in this paper is a serial system with $N+1$ stages, L_1, L_2, …, L_{N+1}, in series where L_1 is the closest to the customers and L_{N+1} is the farthest from the customer. L_{N+1} is an external supplier with infinite supply. We are interested in determining or characterizing the structure of an optimal inventory policy for stages L_1, L_2, …, L_N. Every stage L_n, $n \in \{2, 3, \ldots, N+1\}$, has a shipping capacity of C_n units per period. L_1 is simply an inventory storage stage that serves the customers and has infinite storage capacity. The amount ordered by L_n, $n \in \{1, 2, \ldots, N\}$, in period t is shipped by L_{n+1} in the same period and this inventory reaches L_n after l_n periods, the lead time for stage L_n. L_n orders q_{nt} units from L_{n+1} in period t only if there are at least q_{nt} units available to be shipped by L_{n+1} in that period and q_{nt} is no larger than the shipping capacity C_{n+1} at L_{n+1}. We refer to L_2, L_3, …, L_{N+1} as "physical stages".

We initially assume the planning horizon consists of T periods, numbered $t = 1$, 2, …, T in that order. In section 3.1.4, we examine the infinite horizon case.

We assume that there is an exogenous, finite-state, ergodic Markov Chain $\{s_t\}$ that governs the demand process. s_t is observed at the beginning of each period t. Ω is the sample space of s_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.

2.1 Customers and Distances

Our analysis is motivated by the "single-unit, single-customer" approach. In this and the following sub-section, we define the concepts of customers and units, and also the associated concepts of distances and locations that are the basis for our analysis. This construction is identical to the construction in Muharremoglu and Tsitsiklis (2003).

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. We index these customers 1, 2, …, v_0 in any order. All subsequent customers are indexed $v_0 + 1$, $v_0 + 2$, … 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, …, 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.

We assume that there is a backorder cost of b associated with every unit backordered at the end of a period, that is, every unit at distance 1.

2.2 Units and Locations

Next, we discuss the concepts of "units" and "locations". Inventory is considered to be discrete throughout this paper and every unit of inventory is referred to as "unit". 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 at L_1, it is said to be in location 1. If it has been shipped by L_{n+1} (in other words, ordered by L_n) t periods ago ($1 \leq t \leq l_n$), it is said to be in location $1+l_1+\ldots+(l_n-t)$. If the unit is waiting at L_{n+1}, it is said to be in location $1+l_1+l_2+\ldots+l_n$. *For compactness, let us denote* $1 + l_1 + l_2 + \ldots + l_{n-1}$ *by* M_n*; that is,* M_n *is the location of stage* L_n. Thus, there are $2 + \sum_{n=1}^{N} l_n$ possible "locations" at which a unit can exist in a N

echelon serial system. (See figure 1 for an example of a two echelon serial system.) Observe that there are $l_{n-1} - 1$ locations in the pipeline between L_n and L_{n-1}; therefore, if $l_2 = 1$, L_3 and L_2 will be adjacent to each other.

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, ..., location M_{N+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_{N+1}, at the beginning of period 1.

There is an echelon holding cost h_n associated with each unit of inventory downstream of L_{n+1} at the end of a period.

2.3 Sequence of Events

We now define the sequence of events in a period. We will use j and k to denote the indices of both units and customers. We define z_{jt} to be the location of unit j and y_{jt} to be the distance of customer j at the beginning of period t.

Let $\mathcal{S}(\overrightarrow{C})$ refer to the entire system with all the units and all the customers and the capacity constraint of C_n units per period at stage n. The state of the system at the beginning of period t is given by the vector $\mathbf{x}_t = (s_t, (z_{1t}, y_{1t}), (z_{2t}, y_{2t}), \ldots)$. Let Y_{nt} be the amount of inventory on hand at L_n at the beginning of period t. That is, $Y_{nt} = |\{j : z_{jt} = M_n\}|$. The number of backorders at the start of period t is $|\{k : y_{kt} = 1\}|$.

Next, we explain the sequence of events in period t. (Though redundant at this point, we repeat the phrase "in $\mathcal{S}(\overrightarrow{C})$" for the sake of conciseness later in the paper.)
(1) $\mathbf{x}_t$ is observed. **(2)** Next, L_1 places an order for q_{1t} units from L_2, where $0 \leq q_{1t} \leq \min(Y_{2t}, C_2)$, and integer. All units in any of the locations $M_1 + 1$, $M_1 + 2$, ..., $M_2 - 1$ move to the next location. The q_{1t} units move from location M_2 to location $M_2 - 1$. Then, L_2

places an order for q_{2t} units from L_3, where $0 \le q_{2t} \le \min(Y_{3t}, C_3)$, and integer. All units in any of the locations M_2+1, M_2+2, …, M_3-1 move to the next location. The q_{2t} units move from location M_3 to location M_3-1. This process continues sequentially until L_N places an order on the external supplier L_{N+1} for q_{Nt} units, where q_{Nt} is constrained only by C_{N+1}, since L_{N+1} is assumed to carry infinite inventory. The ordering decisions can formally be represented as follows: $u_{jt} \in \{0,1\}$ is decided for all $j \in \mathcal{S}(\overrightarrow{C})$ such that $z_{jt} = M_n$ for some $n \ge 2$. Unit j is ordered (we will use "released" and "ordered" interchangeably) *if and only if* $u_{jt} = 1$. The number of units released from L_{n+1} $(n \ge 1)$ is

$$q_{nt} \;=\; \sum_{j \in \mathcal{S}(\overrightarrow{C}):z_{jt}=M_{n+1}} u_{jt} \; .$$

For *capacity feasibility*, q_{nt} is C_{n+1} or less. **(3)** Demand d_t is realized. That is, customers in $\mathcal{S}(\overrightarrow{C})$ at distances 2, 3, …, $2+d_t-1$ all arrive and are by definition at distance 1. Customers in $\mathcal{S}(\overrightarrow{C})$ currently at distances $2+d_t$, $3+d_t$, … move d_t steps towards distance 1. **(4)** Units on-hand, at stage L_1, in $\mathcal{S}(\overrightarrow{C})$ and waiting customers in $\mathcal{S}(\overrightarrow{C})$ 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 $\mathcal{S}(\overrightarrow{C})$ 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.* Let E_{nt} be the echelon-n inventory position at this point in time. That is,

$$E_{nt} \;=\; |\{j : 1 \le z^{'}_{jt} \le M_{n+1}-1\}| \;-\; |\{k : y^{'}_{kt} = 1\}| \; ,$$

where $z^{'}_{jt}$ and $y^{'}_{jt}$ denote the location of unit j and the distance of customer j at the end of period t, respectively. **(5)** h_n dollars are charged per unit of inventory downstream of stage L_{n+1} in $\mathcal{S}(\overrightarrow{C})$ and b dollars are charged per waiting customer (at distance 1) in $\mathcal{S}(\overrightarrow{C})$. The cost incurred in period t can be written as

$$\sum_{n=1}^{N} h_n \cdot E_{nt} \;+\; (b \;+\; \sum_{n=1}^{N} h_n) \cdot |\{k : y^{'}_{kt} = 1\}| \; .$$

Before proceeding further, we introduce a useful definition.

Definition 1 *For all n, T_n is the smallest positive integer such that when there are at least T_n periods left in the horizon, it is optimal to have a non-negative echelon n inventory position, if possible. In other words, it is optimal to release enough inventory into the pipeline below stage $n+1$ to meet all existing backorders at stage 1, if possible. If the number of periods remaining in the horizon is less than T_n, it is optimal NOT to release any more units from stage $n+1$. Mathematically, T_n is the smallest positive integer such that the discounted cost of backordering a customer for T_n periods exceeds the discounted holding costs accumulated by a unit from the period it is released from stage $n+1$ until the period it is received by stage 1, assuming that the unit is released from every intermediate stage as soon as it is received.*

Note: Though we have not mentioned purchase costs or transportation costs in the model, linear purchase or transportation costs payable at the time of receipt of inventory can easily be accommodated. (See Janakiraman and Muckstadt (2003b) for a general discussion.)

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 constraints $q_{nt} \leq \min(|\{j \in \mathcal{S}(\overrightarrow{C}) : z_{jt} = M_{n+1}\}|, C_{n+1})$ for all n, 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 $\mathbf{x}_t$ such that $z_{jt} = z_{(j+1)t}$. That is, 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, no unit other than j can satisfy customer j's demand since customer demands are satisfied based on the indices. 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*.*

Note: Lemma 1 holds even when the lead times between any two stages are stochastic and non-crossing and the distribution of the lead time of an order shipped from any stage, n, in period t depends only on s_t. The stochastic lead time model is addressed in section 4.

From now on, our attention is restricted to monotone states when analyzing the system $\mathcal{S}(\overrightarrow{C})$ without any loss of generality.

2.4 Basic Properties of the Optimal Policies

We now present two basic properties of the optimal policy in capacitated serial systems with deterministic lead times. Furthermore, assume that L_{n+1} has a higher shipping capacity than L_n, that is, $C_{n+1} > C_n$. It should be intuitive that it will never be optimal to ship more than C_n units from L_{n+1} in any period; in other words, L_n would never order more than C_n units in any period.

Proposition 1 *If* $C_{n+1} > C_n$ *and the lead time between* L_{n+1} *and* L_n *is a constant,* l_n*, then it is never optimal to ship more than* C_n *units from* L_{n+1} *in any period. Consequently, the capacity at stage* L_{n+1} *is effectively* C_n*.*

Proof: Since the proof is straightforward, we only sketch the outline. Consider any policy which suggests that L_{n+1} should ship q_{nt} units in period t and state x_t, where $q_{nt} > C_n$. Consider any sample path of demands for periods $t+1$, $t+2$, …, T. Consider an alternate policy which is identical to the original policy except for delaying the shipment of $q_{nt} - C_n$ units from period t to the earliest possible period after t. The alternate policy incurs the same number of backorders as the original policy in every period and the inventory position in echelon n is smaller in at least one period when the alternate policy is used and all the other echelon inventory positions are the same under both the policies. Consequently, the alternate policy has a lower cost than the original policy. □

A consequence of this proposition is that it is sufficient to consider serial systems such that the capacities are non-increasing in the stage index when all lead times are deterministic.

Note: Proposition 1 does not hold when the lead time between L_{n+1} and L_n is stochastic; it might be optimal to send more than C_n units from L_{n+1} because of the uncertainty in the timing of receipts at stage L_n.

Next, we claim that when stage $j+1$ has more capacity than stage j, it is not optimal to have more than the maximum amount that stage j can ship during a lead time in inventory between the two stages. This result is also intuitive and can be proved in a straightforward manner.

Proposition 2 *Assume that $C_{n+1} \geq C_n$ and that the lead time between L_{n+1} and L_n is a constant l_n and that $E_{n0} - E_{(n-1)0} \leq l_n \cdot C_n$, where E_{n0} is the echelon n inventory position at the start of period 1. Then, any policy that leads to a state where $E_{nt} - E_{(n-1)t} > l_n \cdot C_n$ is not optimal.*

Proof: We sketch the outline of the proof here. Consider a policy that leads to a state where $E_{nt} - E_{(n-1)t} > l_n \cdot C_n$. Let t be the first period where this inequality is satisfied. Consider an alternate policy which delays the shipping of $(E_{nt} - E_{(n-1)t} - l_n \cdot C_n)$ units from stage $n+1$ in period t to the earliest possible period in the future. It is easy to see the superiority of the alternate policy over the original policy. □

This result is useful in deriving the structure of optimal policies in capacitated, serial systems.

Note that Glasserman and Tayur (1994) prove results similar to the two propositions above with the additional assumption that echelon base-stock policies are followed. Parker and Kapuscinski (2003) also prove similar results, which are special cases of the propositions proved above.

3 Two Echelon Serial Systems

In this section, we examine a class of two echelon serial systems (see figure 1) with the feature that the shipping capacities at the two physical stages are identical. That is, C_3 and C_2 are assumed to be equal to C. Furthermore, initially we restrict attention to the case where the lead time between L_3 and L_2 is exactly one period. We study the optimal policy structure for such systems using a decomposition approach. Subsequently, we extend this approach to the case when the lead time between L_3 and L_2 is two periods. The lead time between L_2 and L_1 is deterministic and is l_1 periods. We also comment on the structure of the optimal policy when these lead times are stochastic and non-crossing.

3.1 Case 1: Lead time between L_3 and L_2 is One Period

This is the case where l_2 is one period. Recall that when l_2 is one period, L_3 and L_2 are adjacent to each other. Parker and Kapuscinski (2003) prove the optimality of "modified echelon base-stock policies" for this system. A modified echelon base-stock policy has the following structure. In period t, echelon 1 raises its inventory position to a target level, $S_1(t, s_t)$, if sufficient capacity and inventory are available. If not, the inventory position is raised to the maximum possible level. Furthermore, L_2 should order enough to raise the echelon 2 inventory position to $S_2(t, s_t)$ or enough to raise the inventory on hand at L_2 to C, whichever is smaller. In this section, we will provide an alternate proof of this result. As we will see, our proof of this result holds even when the lead times between L_2 and L_1 are stochastic and non-crossing.

Our proof of this result has the following key steps. We first show that the system can be decomposed into C two-echelon subsystems, each having unit capacity. Subsequently, we prove that each subsystem can be managed optimally by using a "critical distance" policy at each echelon. We also prove that when the same "critical distance" policy is used to manage each subsystem, the system follows a modified echelon base-stock policy.

Note: Throughout this section, we will assume that the number of units at stage 2, that is L_2, is less than or equal to C at the start of period t. Proposition 2 justifies this assumption. This is identical to the assumption of "being within the band" in Parker and Kapuscinski (2003).

We proceed to discuss how the system under consideration can be decomposed into C subsystems of unit capacity.

3.1.1 Decomposition into Unit Capacity Subsystems

We start by stating two useful definitions.

Definition 2 *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 at stages L_3 and L_2. 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 at stages L_3 and L_2.*

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* at stage $L_3(L_2)$ in any period if unit $j - C$ has still not been released from stage $L_3(L_2)$. This provides a natural connection between unit j and unit $j - C$ for any j.

The sequence of events in $\mathcal{S}$ are steps **(1)**-**(5)** of section 2.3 as applied to a two echelon system with $C_3 = C_2 = C$ and $l_2 = 1$. The sequence of events in $\mathcal{S}_w$ are the same with the additional modifications: $\mathcal{S}(\overrightarrow{C})$ is replaced by $\mathcal{S}_w$ and C_n 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 $j+C$ from stage L_3 or L_2 for any unit j in $\mathcal{S}_w$. A policy for $\mathcal{S}_w$ is *feasible* if it never releases more than one unit from stage L_3 or L_2 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 claim that the subsystems can be optimally managed independently even though the 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 $\mathbf{x}_t^w =_{def} (s_t, (z_{wt}, y_{wt}), (z_{(w+C)t}, y_{(w+C)t}), \ldots)$, that is, the information in $\mathbf{x}_t$ that pertains to $\mathcal{S}_w$. The statement and proof of the following decomposition result are verbatim replicas of the corresponding result for single location capacitated systems presented in Janakiraman and Muckstadt (2003a). They are presented here for the sake of completeness.

Theorem 2 *For any monotone state $\mathbf{x}_t$, the optimal expected discounted (undiscounted) cost in periods t, $t+1$, ... T for system $\mathcal{S}$ equals the optimal expected discounted (undiscounted) cost in periods t, $t+1$, ... 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, $t+1$, ... 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. □

Note: Theorem 2 and the proof hold for serial systems with deterministic lead times and an arbitrary number of stages as long as the capacities are identical. The result also holds for a class of stochastic, non-crossing lead time processes; this is addressed in section 4.

3.1.2 Optimal Policy Structure for a Subsystem

We will now show the existence of an optimal policy with a special structure for every 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 by examining a subsystem $\mathcal{S}_w$. We consider only the class of monotone policies for the subsystems, which contains at least one optimal policy. Let us assume we have used such an optimal policy in periods 1, 2, ..., $t-1$. Therefore, in any period t, the state $\mathbf{x}_t^w$ is monotone. That is, $z_{wt} \leq z_{(w+C)t} \leq \ldots$. Therefore, the units in location $2+l_1$, that is, L_3, are indexed in a serial manner with consecutive indices differing by C. Let j_{wt} be the lowest such index, that is, unit j_{wt} is the only candidate for being released from L_3 in period t in subsystem w. There are two possibilities regarding L_2: either unit $j_{wt} - C$ is present at L_2 or L_2 is empty. L_2 cannot contain more than one unit because both stages have a unit capacity and consequently, it is never optimal to have more than one unit at stage 2 (see Proposition 2).

Recall that in every period, we make the stage 2 decision before the stage 3 decision. We define $U_{2t}^*(\mathbf{x}_t^w) \subseteq \{1, 0\}$, where 1 refers to ordering/releasing the unit and 0 refers to holding the unit, to be the set of optimal stage 2 decisions at time t, if $\mathbf{x}_t^w$ is such that unit $j_{wt} - C$ is present at stage 2, that is, at location $1 + l_1$. If state $\mathbf{x}_t^w$ is such that there is no unit at stage 2, there is no decision to take at stage 2 and consequently, $U_{2t}^*(\mathbf{x}_t^w) = \emptyset$. Similarly, $U_{3t}^*(\tilde{\mathbf{x}}_t^w) \subseteq \{1, 0\}$ is the set of optimal stage 3 decisions for subsystem w in period t if $\tilde{\mathbf{x}}_t^w$ is the state of subsystem w after stage 2 has taken its Release/Hold decision. That is, if the stage 2 decision were to release a unit, then we are examining the subsystem after the unit has been released from stage 2. For example, if $U_{3t}^*(\tilde{\mathbf{x}}_t^w) = \{1\}$ and subsystem w is in state $\tilde{\mathbf{x}}_t^w$ at time t after the stage 2 decision, then it is optimal to release unit j_{wt} from location $2 + l_1$, that is, L_3, and suboptimal to hold it there.

Let us now examine the information that is actually required to manage subsystem w using an optimal, monotone policy. Consider a given t, s_t and j_{wt}. Monotonicity implies that all units indexed below $j_{wt} - C$ in subsystem w have already been released from location $1 + l_1$ (stage 2). 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 subsystem w with indices below $j_{wt} - C$ is unnecessary. Furthermore, we know that the location of all units with indices higher than j_{wt} is $2 + l_1$. It is now clear that s_t, j_{wt}, the location of unit $j_{wt} - C$ and the distances of all customers in w with indices $j_{wt} - C$ 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 $y_{j_{wt}}$, the distance of customer j_{wt} corresponding to the next unit waiting to be released from L_3 in subsystem w. Since unit j_{wt} is still at location $2 + l_1$, y_{wt} cannot be 0. Assume $y_{wt} > 1$, that is, customer j_{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. Also, this means that customer $j_{wt} - C$ is at distance $\max(y_{wt} - C, 1)$ or 0. If unit $j_{wt} - C$ is in location $1 + l_1$ (stage 2), then the distance of this customer cannot be zero and is therefore $\max(y_{wt} - C, 1)$. If unit $j_{wt} - C$ is downstream of stage 2, the cost associated with the unit-customer pair $j_{wt} - C$ is sunk and the distance of customer $j_{wt} - C$ is not required for the decision in this period. That is, if $y_{wt} > 1$, all other information about customer distances necessary to determine the optimal action in this period is already known.

Let us now assume that $y_{wt} = 1$. In this case, customers j_{wt} and $j_{wt} - C$ have arrived and it is not known whether some subsequent customers have also arrived. However, since customer j_{wt} has arrived, it is optimal to release unit j_{wt} from stage 3 if and only if $T - t \geq T_2$ (see Definition 1 for the definition of T_2) and release unit $j_{wt} - C$ from stage 2, if it is located there, if and only if $T - t \geq T_1$. Consequently, any information about other customer distances is unnecessary.

Let us define $i_{wt} \in \{0, 1\}$ to be an indicator of whether unit $j_{wt} - C$ is located at stage 2 or not. In other words, i_{wt} is the indicator of whether stage 2 is empty or not. It is now clear that $(s_t, j_{wt}, y_{wt}, i_{wt})$ is a minimally sufficient information vector to optimally manage subsystem w from period t using a monotone policy. Furthermore, since all subsytems and units are identical, we can use a more compact information vector (s_t, y, i) where $y = y_{wt}$ and $i = i_{wt}$. We define $R^*_{2t}(s_t, y, i) \subseteq \{1, 0\}$ as the set of optimal stage 2 decisions at time t if the state of the exogenous Markov Chain is s_t and if y_{wt} is y and i_{wt} is i. $R^*_{2t}(s_t, y, 0) = \emptyset$ since there is no decision to take at stage 2 if i is zero. Similarly, let $\tilde{i}_{wt} \in \{0, 1\}$ to be an indicator of whether unit $j_{wt} - C$ is located at stage 2 or not after the stage 2 decision. $R^*_{3t}(s_t, y, \tilde{i}) \subseteq$

$\{1,0\}$ is the set of optimal stage 3 decisions at time t if the state of the exogenous Markov Chain is s_t and if y_{wt} is y and $\tilde{i}_{wt}$ is $\tilde{i}$. Proposition 2 implies that $R^*_{3t}(s_t, y, 1)$ is $\{0\}$. That is, if a unit is present at L_2 in a subsystem, it will not be optimal for L_3 to release a unit.

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 subsystem w to release unit $j_{wt} - C$ from L_2 in period t when the system is in the Markovian-state s_t and customer j_{wt} is at a distance $y+1$, then it would be optimal to release it if the customer were any closer. An equivalent claim can be made about releasing unit j_{wt} from L_3.

Lemma 3 *$R^*_{3t}(s_t, y+1, 0) = \{1\}$, for some $y > 1$, implies that $R^*_{3t}(s_t, y, 0) \supseteq \{1\}$. Also, $R^*_{2t}(s_t, y+1, 1) = \{1\}$ for some $y > C+1$, implies that $R^*_{2t}(s_t, y, 1) \supseteq \{1\}$.*

Proof: We start by proving the first part of this lemma. The proof is by contradiction. Assume the statement is not true. That is, there exist t, s_t and y such that $R^*_{3t}(s_t, y+1, 0)$ $= \{1\}$ and $R^*_{3t}(s_t, y, 0) = \{0\}$.

Consider a monotone state of the system, $\tilde{\mathbf{x}}_t$, after stage 2 decisions have been made, with the following characteristics: the exogenous Markov Chain is at state s_t in period t, subsystems w and $w+1$ have no units at L_2, and, $j_{(w+1)t} = j_{wt} + 1$. Assume also that y_{wt} is y for j_{wt} and $y_{(w+1)t}$ is $y+1$ for $j_{(w+1)t}$. Any monotone policy for $\mathcal{S}$ would choose one of the following three pairs of actions for units j_{wt} and $j_{wt}+1$: (a) release both j_{wt} and $j_{wt}+1$, (b) hold both j_{wt} and $j_{wt}+1$ and (c) release j_{wt} and hold $j_{wt}+1$.

Cases (a) and (c) are suboptimal for subsystem w, while cases (b) and (c) are suboptimal for subsystem $w+1$ due to our initial assumption. 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 optimal cost for $\tilde{\mathcal{S}}$ 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 such policy. This completes the proof for the first part of the lemma.

The proof of the second part is identical. □

Also, note that $R^*_{3t}(s,1,0)$ is $\{1\}$ if $T-t \;\geq\; T_2$ because the customer corresponding to the unit under consideration has already arrived. Similarly, $R^*_{2t}(s,y,1)$ is $\{1\}$ if $y \leq C+1$ and $T-t \;\geq\; T_1$ because if customer j_{wt} is within a distance of $C+1$, then customer $j_{wt}-C$ is at distance 1.

Next, we use this lemma to develop the notion of "critical distance" policies. Let us define

$$
\begin{aligned}
y^*_{2t}(s) \;&\stackrel{\text{def}}{=}\; \max\{\, y : \; R^*_{2t}(s,y,1) \;\supseteq\; \{1\}\,\} \text{ and} \\
y^*_{3t}(s) \;&\stackrel{\text{def}}{=}\; \max\{\, y : \; R^*_{3t}(s,y,0) \;\supseteq\; \{1\}\,\} \;.
\end{aligned}
$$

$y^*_{2t}(s)$ is defined in such a way that it is optimal to release unit $j_{wt}-C$ from L_2 if and only if customer j_{wt} is at a distance of $y^*_{2t}(s)$ or closer. This distance $y^*_{2t}(s)$ is a "critical distance" for L_2 at time t and Markovian state s_t for every subsystem. Similarly, $y^*_{3t}(s)$ is a critical distance for stage 3.

Now, consider the policy

$$
\begin{aligned}
R_{2t}(s,y,1) \;&=\; \{1\} \text{ if and only if } \; y \;\leq\; y^*_{2t}(s) \;, \\
R_{3t}(s,y,0) \;&=\; \{1\} \text{ if and only if } \; y \;\leq\; y^*_{3t}(s) \;, \text{ and} \\
R_{3t}(s,y,1) \;&=\; \{0\} \;\forall\, y \;.
\end{aligned}
$$

It is clear that we have defined the functions R_{2t} and R_{3t} such that they constitute an optimal policy for a subsystem. We prove this result next.

Theorem 4 *Consider subsystem $\mathcal{S}_w$. Assume that there are no units or one unit located at stage 2, that is L_2, at the start of period t. Then, an optimal policy for $\mathcal{S}_w$ is to use R_{2t} at stage 2 and R_{3t} at stage 3 in period t.*

Proof: When i_{wt} is zero, there is no decision to take at L_2. When i_{wt} is one, R_{2t} prescribes an optimal stage 2 decision and this can be seen directly from the definitions of $R_{2t}(y, s, 1)$ and $y^*_{2t}(s)$. Furthermore, if it is optimal to hold unit $j_{wt} - C$ at L_2, Proposition 2 implies that the optimal decision at stage 3 is to hold unit j_{wt} at L_3. For the case where $\tilde{i}_{wt}$ is zero, R_{3t} prescribes an optimal stage 3 decision and this can be seen from the definitions of $R_{3t}(y, s, 0)$ and $y^*_{3t}(s)$. $\square$

3.1.3 Optimality of Modified Echelon Base-Stock Policies

We are ready to prove that when each subsystem follows the policies prescribed by R_{2t} and R_{3t} in every period t, the resulting policy for the original system is of the echelon base-stock type, with the exception that the number of units shipped from either of the two stages, L_3 and L_2, and the inventory at L_2 is never allowed to exceed C. Furthermore, this policy is optimal for the entire system $\mathcal{S}$ according to Theorem 2. Parker and Kapuscinski (2003) introduced the term "Modified Echelon Base-stock Policies" to refer to such policies.

Theorem 5 *Assume that there are $\gamma_2 \leq C$ units at stage 2, that is L_2, at the start of period t. Consider the beginning of period t and some state $s \in \Omega$. Let e_1 and e_2 represent the echelon inventory positions for echelons 1 and 2, respectively. An optimal policy for the system $\mathcal{S}$ from this state is ordering q_1 units at stage 1 and ordering q_2 units at stage 2, that is, shipping q_1 and q_2 units from L_2 and L_3, respectively, where:*

$$
\begin{aligned}
q_1 &= \min(\gamma_2, (y^*_{2t}(s) - (C+1) - e_1)^+) \text{ and} \\
q_2 &= \min((y^*_{3t}(s) - 1 - e_2)^+, C - \gamma_2 + q_1)\ .
\end{aligned}
$$

That is, a state-dependent, modified echelon base-stock policy with base-stock levels $y^*_{2t}(s) - (C+1)$ and $y^*_{3t}(s) - 1$ at echelons 1 and 2, respectively, is optimal for $\mathcal{S}$.

Proof: It is sufficient to prove that the policy stated in the theorem will be followed by $\mathcal{S}$ when each subsystem $\mathcal{S}_w$ follows the policy prescribed by R_{2t} and R_{3t} at stages, 2 and 3, respectively. It is possible to prove this structure algebraically. However, we find that an illustrative example provides better intuition about the policy and we present such an example instead of a more formal algebraic argument.

Let us examine Figure 2. Suppose the state of the system at the beginning of period t is depicted in this figure. The state of the modulating Markov chain is s. Here, units 1 through 11 have already been matched with their corresponding customers. All these units and customers are at location 0 and distance 0, respectively. Customers 12 through 15 are presently backordered and are at distance 1. All customers arriving in a future period occupy the subsequent distances one by one. Also observe that units 12 through 25 are in the pipeline between L_2 and L_1. Units 26, 27 and 28 are at L_2 and all subsequent units are at L_3. Let us assume that $C_3 = C_2 = C = 10$, $y^*_{2t}(s) = 23$ and $y^*_{3t}(s) = 22$. Note that since the capacity is 10, unit j belongs to subsystem j *mod* 10, where 10 *mod* 10 is understood to be 10 instead of 0.

Let us discuss the effect of following R_{2t} at stage 2. Units 26, 27 and 28 are the candidates for release at stage 2. Considering these units to be in their own subsystems, subsystems 6,7 and 8 are the only ones where a decision needs to be made. These three units correspond to $j_{6t} - C$, $j_{7t} - C$, and $j_{8t} - C$ and hence j_{6t}, j_{7t} and j_{8t} are 36, 37 and 38, respectively. It is easy to verify that the corresponding customer distances, y_{36}, y_{37} and y_{38}, are 22, 23 and 24, respectively. Since $y^*_{2t}(s) = 23$, units 26 and 27 are released from stage 2, while 28 is not. Thus, echelon 1's inventory position, that is, inventory on-hand plus inventory in the pipeline minus the backorders (waiting customers), is increased to 12 (27 units have been released from L_2 during periods 1, 2, ..., t and 15 customers have arrived in these periods). Observe that $(y^*_{2t}(s) - 1) - C$ is echelon 1's inventory position at this time.

Next, let us discuss the effect of following R_{3t} at stage 3. The candidates for release in stage 3 are only 29,30, …, 38 since the capacity is 10 units. However, subsystem 8 has unit 28 at stage 2, that is, $\tilde{i}_{8t} = 1$. Therefore, unit 38 cannot be released. The customer distances corresponding to units 29, 30, …, 36 are 15, 16, …, 22. All these units are released from stage 3 because the corresponding customer distances are not larger than $y^*_{3t}(s)$. Units 37 and 38 are held at stage 3. Thus, the echelon 2 inventory position is raised to 21 (36 units have been released from L_3 in periods 1, 2, …, t and 15 customers have arrived during these periods), which is the same as $y^*_{3t}(s) - 1$.

In the particular example discussed above, it was possible to raise the inventory positions of echelons 1 and 2 to the levels $y^*_{2t}(s) - (C+1)$ and $y^*_{3t}(s) - 1$, respectively. We will now change the example and discuss the effects of these changes.

Let us now assume that $y^*_{2t}(s)$ is 18, while $y^*_{3t}(s)$ remains 22. Since the distances for units 36, 37, and 38 are still 22, 23 and 24, respectively, none of these units will be released from stage 2. This is because the echelon 1 inventory position at the start of period t is 10, which exceeds the base-stock level $y^*_{2t}(s) - (C+1)$. It can also be verified that units 29, 30, …, 35 will be released from stage 3. This increases the inventory position of echelon 2 only to 20, which is lower than the base-stock level $y^*_{3t}(s) - 1$. However, the inventory at L_2 reaches C, exceeding which is sub-optimal.

Finally, since there are only C subsystems, the number of units that can be released from either of the stages is bounded from above by C. This limit will be reached when at least C units are required to increase the inventory position of either of the echelons to the corresponding base-stock level. □

3.1.4 Infinite Horizon

Let us briefly discuss two results on the infinite horizon, discounted problem for the two-echelon system studied in this section.

Theorem 6 *Assume that there are C units or less at stage 2, that is L_2, at the start of period t. The class of state-dependent, modified echelon base-stock policies, as described in Theorem 5, is optimal for the system $\mathcal{S}$ when the planning horizon is infinite and the performance measure is the total expected discounted cost. Furthermore, since $\{s_t\}$ is a time-homogeneous Markov Chain, the policy does not depend on the period index, t. That is, for every state $s \in \Omega$, there exist parameters $y_2^*(s)$ and $y_3^*(s)$ such that*

$$\begin{aligned} q_1 &= \min(\gamma_2, (y_2^*(s) - (C+1) - e_1)^+) \text{ and} \\ q_2 &= \min((y_3^*(s) - 1 - e_2)^+, C - \gamma_2 + q_1) \ , \end{aligned}$$

where γ_2, q_1, q_2, e_1 and e_2 represent the same quantities as stated in Theorem 5.

Proof: The existence of optimal stationary policies for the system and the subsystems is a consequence of Theorem 4.1.4 of Sennott (1999). The rest of the proof is the same as the finite horizon proof. □.

Next, we show an additional result for systems where $|\Omega|$ is one, that is, $\{d_t\}$ is a sequence of independent and identically distributed random variables. We show the existence of an optimal policy where the base-stock levels for the two echelons do not differ by more than C. This result is similar to Proposition 1 in Glasserman and Tayur (1994).

Lemma 7 *In addition to the assumptions of Theorem 6, assume $|\Omega|$ is one, that is, the sequence of random vectors $\{(d_t, \rho_{1t})\}$ is independent and identically distributed. In this case, an optimal policy can be defined using two stationary parameters, y_2^* and y_3^*. Also, assume that $e_1 \leq y_2^* - (C+1)$ and $\gamma_2 \leq C$, at the start of period 1, that is, echelon 1's inventory position is lower than its base-stock level and the on-hand inventory at L_2 is not* *more than C. Then, we can find y_3^* such that $y_3^* \leq y_2^*$. In other words, the base-stock level for echelon 2 is at most C in excess of the base-stock level for echelon 1. Using the same notation as Theorem 6, the optimal policy in period t is given by*

$$\begin{aligned} q_1 &= \min(\gamma_2, (y_2^* - (C+1) - e_1)^+) \text{ and} \\ q_2 &= \min[(y_3^* - 1 - e_2)^+, C] \,. \end{aligned}$$

That is, an echelon base-stock policy is optimal.

Proof: The optimal policy is prescribed in Theorem 6, where $y_2^*(s)$ and $y_3^*(s)$ are replaced by y_2^* and y_3^*, respectively. Since e_1 is smaller than the base-stock level at the start of period 1, it will always be smaller than the base-stock level when this policy is followed. That is, $e_1 \leq y_2^* - (C+1)$ at the start of any period t. Since the policy limits the amount of inventory that can be stocked at L_2 to be less than or equal to C, it is clear that the maximum value that echelon 2's inventory position can reach at any time is $y_2^* - 1$. Therefore, we can assume without loss of generality that $y_3^* \leq y_2^*$, because echelon 2's target base-stock level is $y_3^* - 1$. It can be verified that the expression for q_2 stated in this lemma is exactly the same as the corresponding expression in Theorem 6, though the expression is more compact here. □

Let us summarize the main structural results of this section. We showed that a state-dependent, modified echelon base-stock policy is optimal for the system $\mathcal{S}$ for the finite horizon problem and the infinite horizon discounted cost problem. That is, each echelon attempts to increase its inventory position to a target base-stock level if there is sufficient material and capacity with the exception that the inventory at stage 2 will never be allowed to exceed C. When demands are stationary through time, we showed that the optimal policy is an echelon base-stock policy for the infinite horizon, discounted cost problem.

This concludes our discussion of the two echelon system with a one period lead time between L_3 and L_2. Next, we present some results for the case where this lead time is two periods.

3.2 Case 2: Lead time between L_3 and L_2 is Two Periods

In this section, we consider two echelon serial systems with identical capacities at stages 2 and 3 and a two period lead time between these stages. Clearly, the important question is whether the class of MEBS policies is optimal for these problems. If not, what can we say about the optimal policy? How complicated can the structure of the optimal policy be? We answer these questions here.

Let us first extend the definition of MEBS policies to these systems. When the lead time between L_3 and L_2 was one period, we knew that it was never optimal to stock L_2 in excess of C. Now, since the lead time is two periods, it is never optimal for the stock on hand at L_2 plus the stock in the pipeline from L_3 to L_2 to exceed $2C$. This should be the $modification$ to echelon base-stock policies. However, this does not appear to be the form of the optimal policy according to an example presented in Speck and van der Wal (1991). They show that in the optimal policy, the number of units released from stage 2 depends non-trivially on the number of units in stock at L_2 and in the pipeline between L_3 and L_2. In particular, as this total amount of inventory increases, the amount released from stage 2 may increase and this is observed to happen even when the quantity released from stage 2 is initially strictly less than the stock there. (This unintutive phenomenon occurs because the marginal value of a unit of available capacity at stage 2 in the next period depends on the number of units in stock at L_2 at the beginning of the next period.) This clearly violates the conditions of an MEBS policy. As Parker and Kapuscinski (2003) comment, it is still possible that the structure of the policy is a modification of echelon base-stock policies in some other way. From a more abstract perspective, it might be possible to find an optimal policy that depends only on two parameters, one for each echelon for a given set of problem data, in every period and state $s \in \Omega$. We show that an optimal policy will depend on a maximum of four parameters, rather than two. The rest of this section is devoted to the development of this result. We omit the proofs of the results leading to the main result because they are identical to the corresponding results for the one period lead time case.

First, we know from Lemma 1 that monotone policies are optimal for this system. We also know from Theorem 2 that this system can be decomposed into C subsystems, each with unit capacities at L_2 and L_3. Optimally managing each of these subsystems is an optimal policy for the entire system.

Let us now examine the subsystem w's decision problem in period t. We start by finding a sufficient information vector for the subsystem.

As before, let j_{wt} be the lowest index of the units in subsystem w located at stage 3, that is, location $l_1 + 3$. In other words, unit j_{wt} is the only candidate for being released from L_3 in period t. Note that location $l_1 + 2$ represents the *pipeline* between L_3 and L_2. Location $l_1 + 1$ is L_2. There are five possibilities regarding units being present at locations $l_1 + 2$ and $l_1 + 1$. Let $(i_1, i_2) \in \mathcal{I}_2$, where $\mathcal{I}_2$ is $\{(0, 2), (1, 1), (1, 0), (0, 1), (0, 0)\}$, represent the number of units at these two locations. That is, at the time of stage 2's release decision, there are i_1 units in the pipeline between stages 2 and 3, and, i_2 units at stage 2. Note that $\mathcal{I}_2$ represents all the possible realizations of (i_1, i_2) if an optimal policy has been followed in periods 1, 2, $\ldots$ $t - 1$ and there were 2 units or less, in total, in locations $2 + l_1$ and $1 + l_1$ at the start of period 1. This can be seen from the following facts: (i) i_1 is the number of units shipped by L_3 in the previous period and is constrained to be 0 or 1 because of the unit capacity restriction and (ii) $i_1 + i_2 \leq 2$ when an optimal policy is followed, as proved in Proposition 2. Let y be the distance of customer j_{wt}. Using exactly the same arguments as in Section 3.1, we can see that (t, s, i_1, i_2, y) is a sufficient information vector required by subsystem w at the start of period t. Similarly, let i_2' be the number of units at L_2 after the stage 2 shipments are sent out and the i_1 units are moved from the pipeline to stage 2 inventory, but before the stage 3 decision. Clearly $i_2' \leq 2$ and (t, s, i_2', y) is a sufficient information vector required at the time of stage 3's decision.

Let us now examine stage 2's decision problem closely. Clearly, the only states where a decision needs to be taken at stage 2 are states such that $i_2 \geq 1$, that is, $(i_1, i_2) \in \{(0,1),(1,1),(0,2)\}$. Observe that the optimal release decision for stage 2 is the same when (i_1, i_2) is either $(1,1)$ or $(0,2)$ because any stage 3 decision in $\{1,0\}$ leads to exactly the same state of the system after this decision has been made, in either case.

Let $R^*_{2t}(s, y, i_1, i_2)$ and $R^*_{3t}(s, y, i_2^{'}) \subseteq \{1,0\}$ be the set of optimal stage 2 and stage 3 decisions given the information vectors. Using arguments similar to those in the proof of Lemma 3, we can prove the following lemma.

Lemma 8

$$
\begin{aligned}
&\textit{(0)} \quad R^*_{2t}(s_t, y+1, 1, 1) \;=\; R^*_{2t}(s_t, y+1, 0, 2) \\
&\textit{(i)} \quad R^*_{3t}(s_t, y+1, 0) \;=\; \{1\}\;,\; \textit{for some}\; y \;>\; 1 \;\Rightarrow\; R^*_{3t}(s_t, y, 0) \;\supseteq\; \{1\}\;, \\
&\textit{(ii)} \quad R^*_{3t}(s_t, y+1, 1) \;=\; \{1\}\;,\; \textit{for some}\; y \;>\; 1 \;\Rightarrow\; R^*_{3t}(s_t, y, 1) \;\supseteq\; \{1\}\;, \\
&\textit{(iii)} \quad R^*_{2t}(s_t, y+1, 0, 1) \;=\; \{1\}\;,\; \textit{for some}\; y \;>\; C+1 \\
&\qquad \Rightarrow\; R^*_{2t}(s_t, y, 0, 1) \;\supseteq\; \{1\}\; \textit{and} \\
&\textit{(iv)} \quad R^*_{2t}(s_t, y+1, 1, 1) \;=\; R^*_{2t}(s_t, y+1, 0, 2) \;=\; \{1\}\;,\; \textit{for some}\; y \;>\; 2C+1 \\
&\qquad \Rightarrow\; R^*_{2t}(s_t, y, 1, 1) \;=\; R^*_{2t}(s_t, y, 0, 2) \;\supseteq\; \{1\}\;,
\end{aligned}
$$

It is also clear that $R^*_{3t}(s_t, y, 2) = \{0\}$ due to Proposition 2. Furthermore, we know that $R^*_{3t}(s_t, 1, 0)$ is $\{1\}$ if $T - t \;\geq\; T_2$ because the customer corresponding to the unit under consideration has arrived already. Similarly, when $T - t \;\geq\; T_1$, $R^*_{2t}(s_t, y, 0, 1)$ is $\{1\}$ if $y \leq C + 1$, and, $R^*_{2t}(s_t, y, 1, 1)$ and $R^*_{2t}(s_t, y, 0, 2)$ are both $\{1\}$ if $y \leq 2C + 1$. This is because if $i_1 + i_2$ is two, then the unit waiting to be released at stage 2 is $j_{wt} - 2C$. So, if customer j_{wt} is within a distance of $2C + 1$, customer $j_{wt} - 2C$ has arrived.

Let us now proceed in exactly the same fashion as the one period lead time case and develop the notion of "critical distance" policies. Let us define

$$
\begin{aligned}
y^*_{2t}(s,2) &\stackrel{\text{def}}{=} \max\{\, y : R^*_{2t}(s,y,1,1) = R^*_{2t}(s,y,0,2) \supseteq \{1\} \,\}\,, \\
y^*_{2t}(s,1) &\stackrel{\text{def}}{=} \max\{\, y : R^*_{2t}(s,y,0,1) = \{1\} \,\}\,, \\
y^*_{3t}(s,0) &\stackrel{\text{def}}{=} \max\{\, y : R^*_{3t}(s,y,0) \supseteq \{1\} \,\}\,, \text{ and} \\
y^*_{3t}(s,1) &\stackrel{\text{def}}{=} \max\{\, y : R^*_{3t}(s,y,1) = \{1\} \,\}\,.
\end{aligned}
$$

Observe that the definitions of $y^*_{2t}(s,1)$ and $y^*_{3t}(s,1)$ are different in form from $y^*_{2t}(s,2)$ and $y^*_{3t}(s,0)$, respectively. This definition turns out to be useful in the proof of Lemma 10 to be stated later.

Now, consider the policy

$$
\begin{aligned}
R_{2t}(s,y,1,1) &= R_{2t}(s,y,0,2) = \{1\} \text{ if and only if } y \leq y^*_{2t}(s,2)\,, \\
R_{2t}(s,y,0,1) &= \{1\} \text{ if and only if } y \leq y^*_{2t}(s,1)\,, \\
R_{3t}(s,y,0) &= \{1\} \text{ if and only if } y \leq y^*_{3t}(s,0)\,, \\
R_{3t}(s,y,1) &= \{1\} \text{ if and only if } y \leq y^*_{3t}(s,1) \text{ and} \\
R_{3t}(s,y,2) &= \{0\}\ \forall\, y\,.
\end{aligned}
$$

This is clearly an optimal policy for the subsystem. We state this result formally.

Lemma 9 *Assume that the total amount of inventory at stage* 2 *and in the pipeline to stage* 2 *is not more than two units at the start of period t. Then, an optimal policy for any subsystem is to use R_{2t} at stage* 2 *and R_{3t} at stage* 3 *in period t.*

We are now ready to state a lemma that relates the critical distances for each echelon through inequalities.

Lemma 10 *(i)* $y^*_{2t}(s,2) \geq y^*_{2t}(s,1) + C - 1$ *and (ii)* $y^*_{3t}(s,0) + 1 \geq y^*_{3t}(s,1)$.

Before proving this lemma, let us discuss the intuition behind the second and third statements. First, consider the inequality relating the stage 2 critical distances. Let us consider two scenarios, (A) and (B), at the start of period t in state $s \in \Omega$. In (A), subsystem w has units j and $j + C$ located at stage 2 and unit $j + 2C$ waiting at stage 3. In (B), unit j is located at stage 2 while unit $j+C$ is waiting at stage 3. Let us assume the vector of customer distances are the same in both the scenarios. In (A), observe that unit j is constraining unit $j + C$, in the sense that unit $j + C$ cannot be released prior to unit j and is forced to wait at stage 2 until unit j is released. Consequently, if we $Hold$ unit j in this period, we are constraining unit $j + C$ for at least one more period. In (B), if we $Hold$ unit j, we will not constrain unit $j + C$ in the next period because it cannot reach stage 2 in the next period even if it is released from stage 3 in the current period. This is an intuitive reason to believe that a unit is more likely to be released from stage 2 if one more unit belonging to the same subsystem is located at stage 2 or is in transit to stage 2 than otherwise. It is easy to see that this notion can be formalized by the inequality $y^*_{2t}(s, 2) \geq y^*_{2t}(s, 1) + C$. Statement (i) of the Lemma is a slightly weaker representation of this idea. Using identical reasoning, we can argue intuitively that a unit is less likely to be released from stage 3 if a unit belonging to the same subsystem is located at stage 2 at the time of the stage 3 decision than otherwise. This notion could be expressed by the inequality $y^*_{3t}(s, 0) \geq y^*_{3t}(s, 1)$. Statement (ii) of the Lemma is a weaker version of this idea.

Proof of Lemma 10: Let us first prove the inequality relationship between the echelon 2 critical distances.

Let A_1 and A_2 be the number of units in the system in transit to stage 2 and at stage 2, respectively, at the beginning of period t. Without loss of generality, we can assume that the A_2 units at stage 2 are numbered 1, 2, …, A_2 and the A_1 units in transit to 2 are numbered $A_2 + 1$, $A_2 + 2$, …, $A_1 + A_2$. Notice that this numbering determines the number of units at stage 2, i_{2w}, and number of units in transit to stage 2, i_{1w}, in each subsystem w.

Let us now classify the C subsystems into 5 mutually exclusive categories. Let $N_{ab} = \{w : i_{1w} = a, i_{2w} = b\}$ for $(a, b) \in \{(0, 0), (0, 1), (1, 0), (1, 1), (0, 2)\}$. Let $n_{ab} = |N_{ab}|$. Clearly, $n_{02} = (A_2 - C)^+$, because subsystems 1, 2, ..., $(A_2 - C)^+$ have two units present at stage 2 and none in transit. That is, $N_{02} = \{1, 2, \ldots, (A_2 - C)^+\}$. Similarly, $N_{11} = \{(A_2 - C)^+ + 1, (A_2 - C)^+ + 2, \ldots, (A_1 + A_2 - C)^+\}$ and consequently, $n_{11} = (A_1 + A_2 - C)^+$ - $(A_2 - C)^+$. (Recall that $A_1 + A_2 \leq 2C$ by assumption and $A_1 \leq C$ because C is the maximum number of units that could have been released by stage 3 in the previous period. These inequalities are useful in verifying the composition of these 5 sets.) Furthermore, we have $n_{11} + n_{01} + 2n_{02} = A_2$ and consequently, $n_{01} = \min(A_2, C) - (A_1 + A_2 - C)^+$. Also, $N_{01} = \{(A_1 + A_2 - C)^+ + 1, (A_1 + A_2 - C)^+ + 2, \ldots, \min(A_2, C)\}$. Sets N_{00} and N_{10} are not important now because there is no unit that can be released from stage 2 in these subsystems in this period.

Consider a state such that $A_1 \leq C - 1$, $A_2 \leq C$ and $A_1 + A_2 > C$. So, there are $n_{02} + n_{11}$ (this is strictly positive) subsystems with two units in stage 2 or in transit to stage 2 and n_{01} (this is strictly positive) subsystems with exactly one unit in stage 2 and no unit in transit to 2. Now, consider unit $(A_1 + A_2 - C)^+ + 1$. This is the least indexed element of N_{01}. By definition of $y^*_{2t}(s, 1)$, $\{1\}$ is the only optimal decision for this unit if and only if the distance of customer $(A_1 + A_2 - C)^+ + C + 1$, the customer waiting at stage 3 in the same subsystem, is $y^*_{2t}(s, 1)$ or less. Next, consider unit $(A_1 + A_2 - C)^+$, the highest indexed element in $N_{11} \cup N_{02}$. By definition of $y^*_{2t}(s, 2)$, $\{0\}$ is the only optimal decision for this unit if and only if the distance of customer $(A_1 + A_2 - C)^+ + 2C$, the customer waiting at stage 3 in this subsystem, is greater than $y^*_{2t}(s, 2)$. Using arguments similar to those presented in the proof of Lemma 3, we can show that it can never be strictly optimal to *Release* unit $(A_1 + A_2 - C)^+ + 1$ and strictly optimal to *Hold* unit $(A_1 + A_2 - C)^+$ at stage 2 in the same period. Observe that customer $(A_1 + A_2 - C)^+ + 2C$ is always at a distance of $C - 1$ higher than customer $(A_1+A_2-C)^+ + 1$, unless customer $(A_1+A_2-C)^+ + 2C$ is within a distance of $C-1$. Combining these ideas, we can see that $y^*_{2t}(s,2) \geq y^*_{2t}(s,1) + C - 1$. The proof of the statement about the stage 3 critical distances is almost identical and is omitted. □

In the next section, we deduce the structure of the optimal policy in $\mathcal{S}$ based on the structure of the optimal policy in every subsystem.

3.2.1 Optimality of Two-tier Base-stock Policies

Let us first discuss the optimal ordering policy for echelon 1, that is, the optimal release policy for stage 2. The following theorem characterizes the structure of an optimal policy, which is derived by using policy R_{2t} at stage 2 in every subsystem in period t. The proof is similar to the proof of Theorem 5 and is omitted.

Theorem 11 *Assume that there are* $2C$ *units or less at stage* 2 *and in transit to stage 2 at the start of period* t. *Consider the beginning of period* t *and some state* $s \in \Omega$. *Let* N_{02}, N_{01}, N_{11}, n_{02}, n_{01} *and* n_{11} *denote the sets and quantities defined in the proof of Lemma 10. Let* E_1 *be the echelon 1 inventory position at this time. An optimal ordering policy for echelon 1, or equivalently, an optimal release policy for stage 2, dictates the release of* q_1 *units from stage 2, where* $q_1 = q_1^{(1)} + q_1^{(2)}$, *where,* $q_1^{(1)}$ *and* $q_1^{(2)}$ *are computed as follows.*

$$q_1^{(1)} \;=\; \min\Big((y^*_{2t}(s,2) - (2C+1) - E_1)^+, n_{02} + n_{11}\Big)\,.$$

If $q_1^{(1)} < n_{02} + n_{11}$, *then* $q_1^{(2)} = 0$. *Otherwise,*

$$q_1^{(2)} \;=\; \min\Big((y^*_{2t}(s,1) - (C+1) - (E_1 + q_1^{(1)}))^+, n_{01}\Big)\;.$$

In words, the optimal policy says the following: first, echelon 1 should order-up-to $y^*_{2t}(s,2) - (2C+1)$, *if possible, using only the units that are elements of* $N_{02} \cup N_{11}$. *If all the elements of* $N_{02} \cup N_{11}$ *have now been released, then echelon 1 should order up to* $y^*_{2t}(s,1) - (C+1)$, *if possible, using the elements of* N_{01}. *(We refer to this policy as a "two-tier, echelon base-stock policy".)*

Let us now proceed to analyze the structure of an optimal policy for echelon 2. That is, let us examine the stage 3 release decision after the stage 2 release decision has been taken, and the units released from stage 2 have moved to the subsequent location and the units in transit to stage 2 have moved to stage 2. So there are no units in transit to stage 2 at this point in time.

Let A be the number of units at stage 2. Without loss of generality, we can assume that the A units at stage 2 are numbered 1, 2, …, A. Notice that this numbering determines $\tilde{i}_w$, the number of units at stage 2 in each subsystem w. Let us now classify the C subsystems into 3 mutually exclusive categories. Let $N_a = \{w : \tilde{i}_w = a\}$, $a \in \{0, 1, 2\}$ and let n_a denote $|N_a|$. It is easy to verify that $n_0 = (C - A)^+$ and $n_1 = \min(A, C) - (A - C)^+$. Since there are C subsystems in total, n_2 is $C - (n_0 + n_1)$, that is, $(A - C)^+$.

The following theorem characterizes the structure of an optimal policy, which is derived by using policy R_{3t} at stage 3 in every subsystem in period t. The proof is similar to the proof of Theorem 5 and is omitted.

Theorem 12 *Assume there are* $2C$ *units or less at stage* 2 *plus in transit to stage 2 at the start of period* t. *Consider the point in time when stage 3's release decision has to be made in period* t *and some state* $s \in \Omega$. *Let* E_2 *be the echelon 2 inventory position at this time. An optimal ordering policy for echelon 2, or equivalently, an optimal release policy for stage 3, dictates the release of* q_2 *units from stage 3, where* $q_2 = q_2^{(1)} + q_2^{(2)}$, *where,* $q_2^{(1)}$ *and* $q_2^{(2)}$ *are computed as follows.*

$$q_2^{(1)} = \min\left((y_{3t}^*(s,0) - 1 - E_2)^+, n_0\right).$$

If $q_2^{(1)} < n_0$, *then* $q_2^{(2)} = 0$. *Otherwise,*

$$q_2^{(2)} = \min\left((y_{3t}^*(s,1) - 1 - (E_2 + q_2^{(1)}))^+, n_1\right) .$$

In words, the optimal policy says the following: first, echelon 2 should order-up-to $y_{3t}^*(s,0) - 1$, *if possible, using only the units that are elements of* N_0. *If all the elements of* N_0 *have now* *been released, then echelon 2 should order up to* $(y^*_{3t}(s,1)-1)$*, if possible, using the elements of* N_1.

Theorem 12 indicates that this optimal policy for echelon 2 is also a "two-tier" echelon base-stock policy.

Note that Theorems 11 and 12 can also be shown to hold in the infinite horizon, discounted cost version of the problem. See the discussion in section 3.1.4.

One final comment about the "two-tier" echelon base-stock policies: note that the policy for echelon 1 requires the knowledge of A_1 and A_2, the inventory in transit to stage 2 and on hand at stage 2 explicitly. An echelon base-stock policy would have required the knowledge of A_2 only. Similarly, the "two-tier" policy at echelon 2 requires the knowledge of A, the number of units at stage 2, which an echelon base-stock policy would not.

Next, we comment briefly about the optimal policy when the lead time between stages 3 and 2 is an arbitrary integer and/or when there are more than three stages in the serial system.

3.3 Optimality of Multi-Tier Base-stock Policies

In this section, we briefly examine serial systems with an arbitrary number of stages with identical capacities and stochastic, non-crossing lead times everywhere. The following is the lead time model we consider. This model was first introduced by Kaplan (1970) and subsequently redefined and simplified by Nahmias (1979), both in the context of single stage systems.

Model $\mathcal{L}$: There is a random variable ρ_{nt}, whose distribution is determined completely by s_t, that specifies the least "age" of units that will be delivered in period t at L_n. This means all units shipped by L_{n+1}, or ordered by L_n, in period $t-\rho_{nt}$ or earlier are delivered at L_n

no later than period t. We assume that the sample space of the random variable ρ_{nt} is finite and that the probability mass function of ρ_{nt} is known for every possible value of s_t.

Theorem 2 can be used to decompose N echelon, serial systems with identical capacities at all physical stages and stochastic, non-crossing, lead times everywhere into serial systems with unit capacities at all physical stages. Now, we can use the proof technique we used for the two echelon system with a two-period upstream lead time. For every echelon n, we can classify the subsystems into several categories based on the positioning of inventories within each subsystem. For each of these categories, the optimality of monotone policies leads to the existence of a critical distance, which in turn, leads to a base-stock level. This is a multi-tier, base-stock policy, in the sense that there is a base-stock level corresponding to every category of subsystems. Thus, for each echelon, a "multi-tier base-stock policy" is optimal. The number of "tiers" grows exponentially in the total leadtime of the system between stages 2 and $N + 1$. In fact, it is easy to show that the number of tiers at each echelon is less than $2^{M_{N+1}-M_2}$. In particular, for a two echelon serial system where l_2 is three, there are 4 tiers for echelon 1 and 4 tiers for echelon 2.

In the next section, we present additional structural results for capacitated, N echelon, serial systems.

4 More General Serial Systems

The assumption that the capacities at all stages of the serial system are identical is crucial to the analysis presented in section 3. In this section, we address the question: "What can we say about systems with non-identical capacities using our approach?". We also allow stochastic lead times according to Model $\mathcal{L}$.

In this section, we also consider a set-up cost of K_n at stage n $(1 < n \leq N+1)$ in period t if any positive quantity is shipped out of stage n in that period.

The first result we present in this section is a decomposition result for serial systems where a set of consecutive downstream stages starting from stage 1 are uncapacitated and the upstream locations are capacitated.

Lemma 13 *Consider an N echelon serial system, that is, an $N+1$ stage serial system, such that $C_2 = C_3 = \ldots C_n = \infty$ and $K_2 = K_3 = \ldots K_n = 0$ for some $2 \leq n \leq N$. The lead times between any pair of consecutive stages are stochastic and non-crossing as explained in Model $\mathcal{L}$. At the beginning of period t, let us define $J_1(t)$ to be $\{j$: unit j is located at stage $n+1$ or upstream of stage $n+1\}$ and $J_2(t)$ to be $\{j$: unit j is downstream of stage $n+1\}$. The discounted cost, optimal control problem for a finite or an infinite horizon can be decomposed into two independent control problems, one for $J_1(t)$ and another for $J_2(t)$. Furthermore, the optimal control problem for $J_2(t)$ can be decomposed into several single-unit problems, one for each unit.*

Proof: Let us consider the class of monotone policies, which we know is optimal. Any unit in $J_1(t)$ is operationally independent of any unit in $J_2(t)$ in the sense that the release decision corresponding to some unit $j_1 \in J_1(t)$ in any period $t^{'} \geq t$ is not constrained by the location of any unit $j_2 \in J_2(t)$ in that period $t^{'}$. This is because all stages downstream of $n+1$ are uncapacitated. Similarly, this availability of infinite shipping capacity at stages n $\ldots$ 1 also implies that every unit in $J_2(t)$ is independent of any other unit in $J_2(t)$. Thus we have established the absence of any constraints linking $J_1(t)$ and $J_2(t)$ and any constraints between any two units in $J_2(t)$. Furthermore, since all costs are linear, the objective function in the optimal control problem is decomposable into several functions, one for every unit in $J_1(t) \cup J_2(t)$. $\square$

Next, we state another important result about the structure of the optimal policy in the serial system examined in Lemma 13.

Theorem 14 *Consider the serial system examined in Lemma 13. (a) There exists an optimal policy such that the ordering decisions of echelons n, $n+1$, $\ldots$ N, or equivalently, the release decisions of stages $n+1$, $n+2$, $\ldots$, $N+1$, depend on the vector of inventories downstream of stage $n+1$ only through the echelon n inventory position. (b) There exists an optimal policy such that the ordering policy for echelons 1, 2, $\ldots$, $n-1$ is an echelon base-stock policy. Furthermore, at stages 1, 2, $\ldots$, $n-1$, there exist optimal base-stock levels which are independent of the policies in the upstream stages. In particular, the optimal base-stock levels for stages 1, 2, $\ldots$, $n-1$ of an uncapacitated, N echelon, serial system with the same holding and backorder cost parameters, zero set-up costs everywhere, and the same lead time and demand processes, are optimal for these stages in the original system.*

Proof: Both the statements are direct consequences of Lemma 13. □

Next, we extend the decomposition result shown in Theorem 2 to serial systems that have a set of uncapacitated stages downstream.

Lemma 15 *Consider the serial system examined in Lemma 13. Assume that the set-up costs are zero and that $C_{n+1} = C_{n+2} = \ldots C_{N+1} = C$. The decomposition result stated in Theorem 2 holds for this system in the following sense. This system can be decomposed into C subsystems, each having unit capacity at stages $n+1$, $n+2$, $\ldots$, $N+1$ and infinite capacities downstream.*

Proof: This statement is a direct consequence of Lemma 13 and the proof of Theorem 2. □

The next result is an extension of the main results in section 3.

Theorem 16 *Consider an N echelon serial system with no set-up costs such that $C_{N+1} = C_N$, and $C_2 = C_3 = \ldots C_{N-1} = \infty$. If the lead time between stages $N+1$ and N is one period , then an Echelon Base-stock Policy is optimal for echelon $N-1$ and a Modified Echelon Base-stock Policy is optimal for echelon N. Similarly, when the lead time between* *stages N and $N-1$ is two periods, "Two-tier" Base-stock Policies are optimal for echelons N and $N-1$.*

Proof: This theorem is a direct consequence of Lemma 13, the sequence of results developed in section 3 and their proofs. □

5 Convex Penalty Costs

One of the main reasons why we are able to decompose capacitated systems into C subsystems is the fact that the cost incurred by the system in a period is nothing but the sum of the costs associated with each unit-customer pair. Clearly, this fact holds even when the holding and backorder cost associated with a unit-customer pair is non-linear in the length of time the unit or the customer waits at a stage. This leads to the question of whether all the results contained in this paper hold when such non-linear costs are used. We address this issue in this section.

Let $h_n(t)$ denote the holding cost associated with a unit if it stays at stage n for exactly t time periods. $h_{N+1}(t)$ is zero. Similarly, $b(t)$ denotes the backorder cost associated with a unit of demand that is backordered for exactly t time periods. $h_n(0)$ is zero for all n and $b(0)$ is zero.

Let us now see if monotone policies are optimal in the serial system with this new cost structure. Consider an example where $h_n(t)$ is concave. Then, an optimal policy will always release units with the highest indices from stage n because "younger" units have a greater marginal holding cost with respect to time. This violates monotonicity. So, the optimality of monotone policies clearly does not hold for arbitrary non-linear holding cost structures. However, if $h_n(t)$ is convex for all n, then it is easy to observe that monotone policies are indeed optimal. Similarly, if $b(t)$ is convex, customer demands are met in the order of their arrivals. It can easily be verified that all the results in sections 3 and 4 hold under the assumptions that (a) $h_n(t)$ is convex and non-decreasing for all n, and, (b) $b(t)$ is convex and non-decreasing. This is related to Derman and Klein (1958); they present conditions under which FIFO (first in, first out) or LIFO (last in, first out) release policies are optimal in an environment where units of different ages are available to meet future demands.

One of the drawbacks of the decomposition approach is that convex holding and backorder costs as considered in traditional inventory theory (the holding cost incurred in a period is convex in the number of units held and the backorder cost is convex in the number of backorders) cannot be accomodated. Interestingly, however, a model in which the cost of backordering a unit of demand is convex in the waiting time is easily accomodated using our decomposition approach whereas it seems less easy using traditional dynamic programming methods. Furthermore, the assumption that the backorder cost is convex in the waiting time is, perhaps, more appropriate in several situations. An example is a supplier who has an advertised delivery-time promise which is backed-up by discounts to the customer when the promise cannot be met.

6 Conclusions

We have extended the "single-unit single-customer" approach introduced in the context of uncapacitated systems by Katircioglu and Atkins (1998) and Muharremoglu and Tsitsiklis (2003) to a class of capacitated serial systems. In particular, two echelon serial systems with identical capacities at both echelons are studied in detail. When the lead time at the upstream echelon is one period, and the lead times to the downstream echelon are stochastic and non-crossing, we show that modified echelon base-stock policies are optimal using a decomposition scheme introduced in Janakiraman and Muckstadt (2003a). This is a generalization of the optimality result shown in Parker and Kapuscinski (2003) for such systems with a deterministic lead time at the downstream echelon. When the lead time to the upstream echelon is two periods, the class of "two-tier base-stock policies" are shown to be optimal. We also argue that a generalization of these policies, which we call "multi-tier base-stock policies", are optimal for multi-echelon serial systems with identical capacities at all physical stages and stochastic, non-crossing lead times. Elementary structural results on decomposing serial systems with a set of uncapacitated downstream stages and capacitated upstream stages are also given. Finally, all these results are shown to hold when inventory holding costs and backorder costs are convex and non-decreasing in the waiting times of units and customers, respectively.

7 Acknowledgement

We are thankful to Ton de Kok and Michal Tzur for several useful comments on an earlier version of this paper.

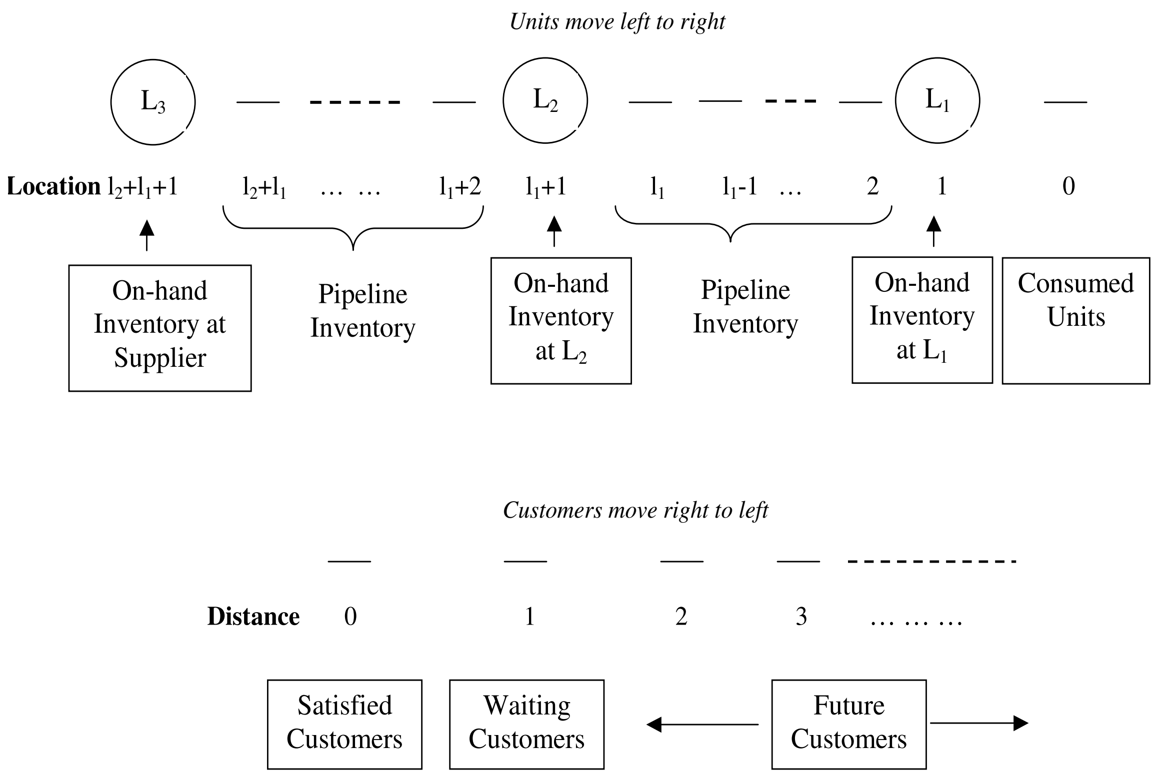

Figure 1: Locations of Units and Distances of Customers

References

Aviv, Y., and A. Federgruen. 1997. Stochastic Inventory Models with Limited Production Capacity and Periodically Varying Parameters. *Probability in the Engineering and Informational Sciences* 11:107–135.

Axsater, S. 1990. Simple Solution Procedures For A Class of Two-Echelon Inventory Problems. *Operations Research* 38 (1): 64–69.

Axsater, S. 1993. Exact and Approximate Evaluation of Batch-Ordering Policies for Two-Level Inventory Systems. *Operations Research* 41 (4): 777–785.

Chen, F., and J. Song. 2001. Optimal Policies for Multi-echelon Inventory Problems with Markov-modulated demand. *Operations Research* 49 (2): 226–234.

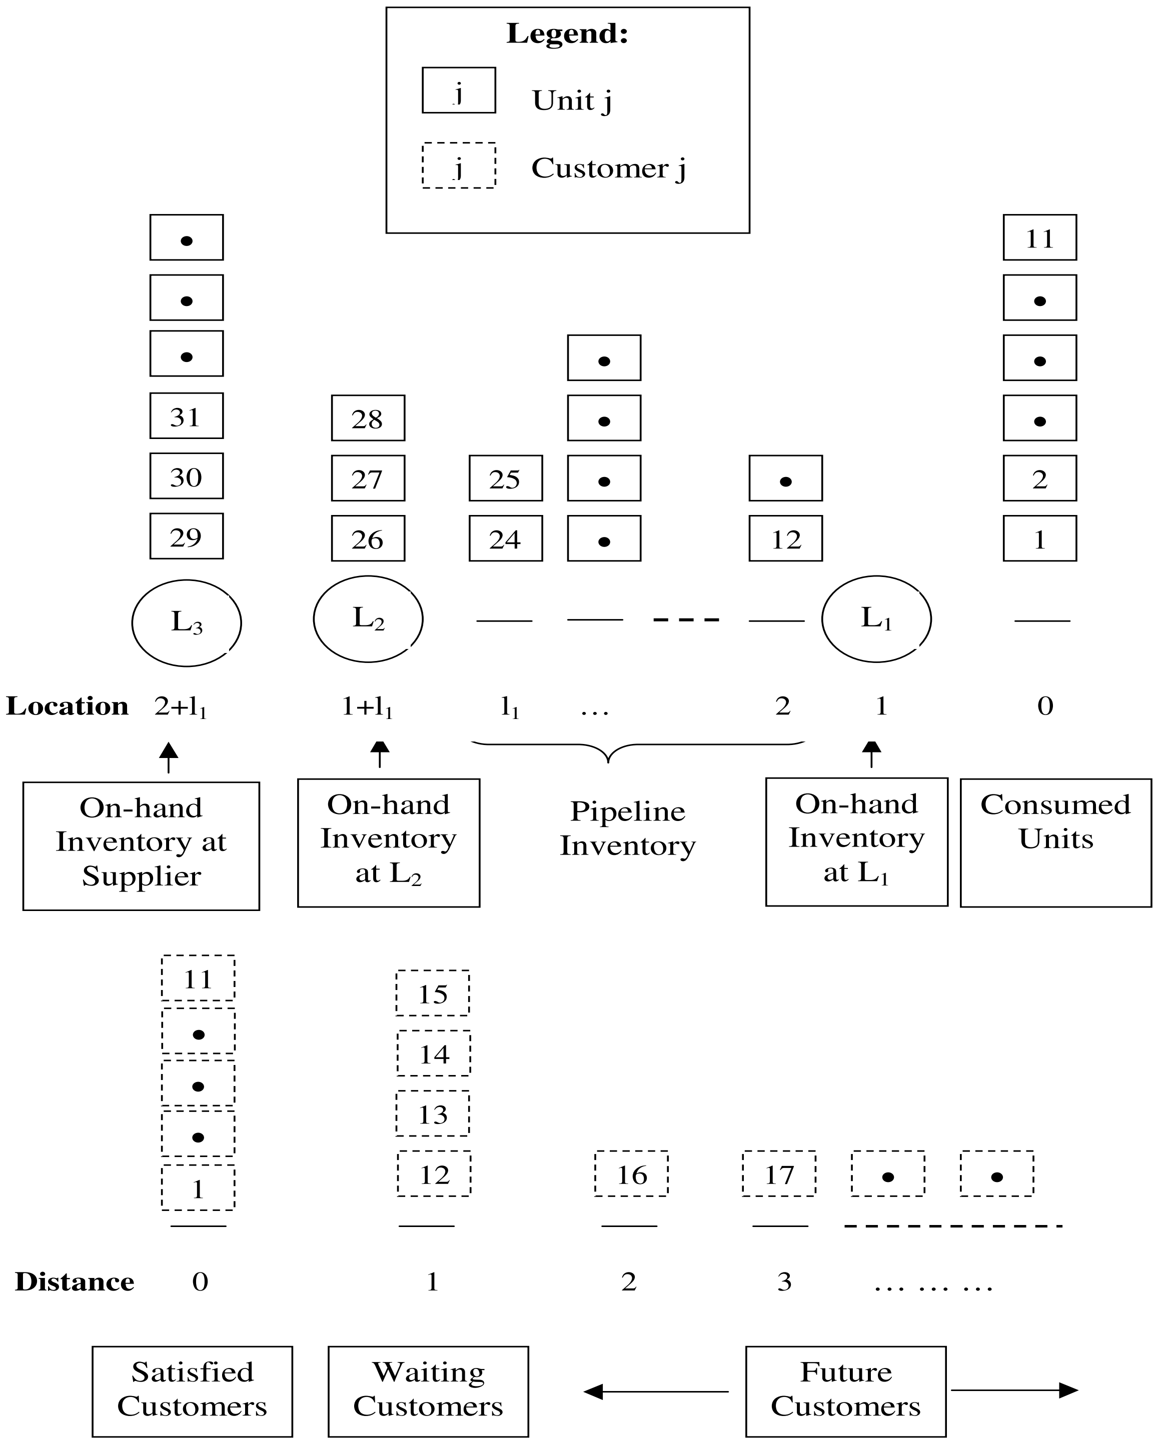

Figure 2: Proof of Theorem 5

Clark, A., and H. Scarf. 1960. Optimal Policies for a Multiechelon Inventory Problem. *Management Science* 6:475–490.

Derman, C., and M. Klein. 1958. Inventory Depletion Management. *Management Science* 4 (4): 450–456.

Federgruen, A., and P. Zipkin. 1984. Computational Issues in an Infinite-Horizon, Multiechelon Inventory Model. *Operations Research* 32 (4): 818–836.

Federgruen, A., and P. Zipkin. 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 P. Zipkin. 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., and S. Tayur. 1994. The Stability of a Capacitated, Multi-Echelon Production-Inventory System Under a Base-Stock Policy. *Operations Research* 42 (5): 913–925.

Glasserman, P., and S. Tayur. 1995. Sensitivity Analysis for Base-stock Levels in Multiechelon Production-inventory Systems. *Management Science* 41 (2): 263–281.

Glasserman, P., and S. Tayur. 1996. A Simple Approximation for a Multistage Capacitated Production-Inventory System. *Naval Research Logistics* 43:41–58.

Janakiraman, G., and J. Muckstadt. 2003a. A Decomposition Approach for a Capacitated, Single-stage, Production-Inventory System. Technical Report 1360, School of O.R.I.E., Cornell University.

Janakiraman, G., and J. Muckstadt. 2003b. Inventory Control in Directed Networks: A Note on Linear Costs. *Operations Research (forthcoming)*.

Kaplan, R. 1970, March. A Dynamic Inventory Model With Stochastic Lead Times. *Management Science* 16 (7): 491–507.

Kapuscinski, R., and S. Tayur. 1998, Nov.-Dec.. A Capacitated Production-Inventory Model with Periodic Demand. *Operations Research* 46 (6): 899–911.

Katircioglu, K., and D. Atkins. 1998. New Optimal Policies for A Unit Demand Inventory Problem. Working Paper.

Muharremoglu, A., and J. N. Tsitsiklis. 2003. A Single-Unit Decomposition Approach to Multi-Echelon Inventory Systems. http://web.mit.edu/jnt/www/publ.html.

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 R. Kapuscinski. 2003. Optimal Inventory Policies for a Capacitated Two-Echelon System. *Operations Research (forthcoming)*.

Prabhu, U. 1998. *Stochastic Storage Systems: Queues, Insurance Risk, and Dams*. New York: Springer.

Sennott, L. 1999. *Stochastic Dynamic Programming and the Control of Queueing Systems*. John Wiley and Sons, Inc.

Speck, C., and J. van der Wal. 1991. The Capacitated Multi-Echelon Inventory System with Serial Structure: 1. The 'Push-Ahead'- effect. Memorandum COSOR 91-39, Eindhoven University of Technology, The Netherlands.

Tayur, S. 1992. Computing the Optimal Policy for Capacitated Inventory Models. *Communications in Statistics: Stochastic Models* 9:585–598.

Tayur, S. 1996. Recent Developments in Single Product, Discrete-time, Capacitated Production-Inventory Systems. citeseer.nj.nec.com/tayur96recent.html.