Provably Near-Optimal Balancing Policies for Multi-Echelon Stochastic Inventory Control Models

Retsef Levi*

Robin Roundy[†]

Van Anh Truong [‡]

February 13, 2006

Abstract

We develop the first algorithmic approach to compute provably good ordering policies for a multi-echelon, stochastic inventory system facing correlated, non-stationary and evolving demands over a finite horizon. Our approach is computationally efficient and guaranteed to produce a policy with total expected cost no more than twice the expected cost of an optimal policy. As part of our computational approach, we propose an innovative scheme to account for costs in a multi-echelon, multi-period environment. This scheme, called a *cause-effect cost-accounting scheme*, is significantly different from traditional cost accounting schemes, in that it re-allocates costs with the goal of assigning every unit of cost to the decision that caused the cost to be incurred.

We show that our policy achieves a worst-case expected cost of 2 times the expected cost of the optimal policy under the assumption that holding costs for a stage in the serial system is charged as soon as a unit is ordered by the stage. Under the alternative assumption that holding costs begin to be charged when a unit arrives at a stage, our modified balancing policy achieves a worst-case performance ratio of 3. Additionally, when the backorder cost parameter is greater than the echelon holding cost parameters for all but the final downstream stage, we show that the performance ratio can be reduced to $2\frac{1}{3}$.

^{*}retsef@us.ibm.com. IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598.

[†]robin@orie.cornell.edu. School of ORIE, Cornell University, Ithaca, NY 14853. Research supported by NSF grant DMI-0500263.

[‡]truong@orie.cornell.edu. School of ORIE, Cornell University, Ithaca, NY 14853. Research supported partially by NSF grant DMI-0500263.

1 Introduction

In this paper, we address the challenging task of computing provably good inventory control policies in supply chains with several echelons (locations) and stochastic demands that can be correlated and evolve over time. Specifically, we consider two of the fundamental multi-echelon models in stochastic inventory theory, the single-item, periodic-review, serial system and the single-item periodic-review assembly system. These models have received a lot of attention from numerous authors throughout the years (see, for example, [18] and other references below). The existing literature has mostly been focused on deriving structural results about the form of the optimal policies using a dynamic programming approach. However, computing the optimal policies is tractable only under rather strong assumptions on the demand distributions and there are many important scenarios where the required computations seem unlikely to be tractable. In particular, these models become computationally intractable in the presence of correlated and evolving demands. Such demand structures are typical in scenarios where dynamic forecasts are incorporated into the supply-chain management. In this paper, we provide novel extensions the to the recent papers on single-item, singlelocation models by Levi, Pál, Roundy and Shmoys [11] and Levi, Roundy, Shmoys and Truong [12]. These papers consider single-item, single-location models with very general demand structures and describe a class of policies that are called *Dual-Balancing policies*. These policies are computationally efficient, and are proven to be near-optimal in that they admit a worst-case performance guarantee of 2. That is, the expected cost of the policy is guaranteed to be at most twice the expected cost of an optimal policy, regardless of the input of the problem. In this paper, we establish similar results for the more general multi-echelon and multi-item stochastic inventory models. To the best of our knowledge, these are the first computationally efficient policies for these fundamental models that admit worst-case performance guarantees. As we shall demonstrate, these extensions require several new conceptual and technical ideas that we believe will have additional applications in the future research of multi-echelon and multi-item stochastic inventory models.

The details of the *single-item, periodic-review, serial system* are as follows. A single commodity moves through a supply-chain that consists of a customer, distinct inventory-storage locations called *stages*, and an external supplier. The stages are organized in a *serial system*. The first stage in the series (at the bottom of the series) is facing a sequence of stochastic demands over a planning horizon of finitely many discrete periods. Each stage is supplied by the next stage in the series, and the last stage is supplied by an external supplier with infinite capacity. There are *lead times* between each two consecutive stages in the system that correspond to the number of periods that it takes to ship the commodity from one stage to the next. In each period several types of costs are incurred, per-unit *ordering costs* for ordering inventory at each stage,

per-unit *holding costs* that are incurred for each unit of inventory that is at a stage or in transit between stages at the end of a time period, and a per-unit *backordering penalty cost* for each unit of demand that is not yet satisfied at the end of a period. Unsatisfied demand is fully backordered, i.e., it stays in the system until it is satisfied. In the more general assembly system, the stages in the supply-chain are organized as a directed tree. Each stage produces a different subassembly, and the *root stage* of the tree produces a final product and is facing stochastic demands. Each unit in each stage is assembled from one unit from each of the *parent stages*. There are again lead times between each stage and its *child stage* in the tree. As in the serial system there are per-unit ordering costs and holding costs in each one of the stages, as well as a per-unit backordering penalty cost that is incurred only in the root stage. In both models the goal is to find an ordering policy that respects the system constrains and minimizes the overall expected cost over the entire planning horizon.

As we have already mentioned these fundamental models have attracted a lot of attention from many researchers. The most commonly used paradigm to address these models has been the dynamic programming framework. This approach has been very effective in characterizing the structure of the optimal policies for these models. There are several dynamic-programming-based proofs for the optimality of *echelon state-dependent base-stock policies* in the serial system. The first proof has been established in the seminal paper of Clark and Scarf [4] for a model where the demands in different periods are assumed to be independent and identically distributed. Several subsequent papers (see, for example, [2, 3, 5, 6, 18]) by different authors have established simpler proofs of the initial result of Clark and Scarf and extended it to models with more general assumptions (e.g., more general assumptions on the demand distributions). In particular, the optimality of echelon base-stock policies has been established for models in which demands follow an exogenous Markov-modulated process [2] or an autoregressive process [5].

All of the above-mentioned proofs are based on several concepts regarding the cost accounting scheme. The first major concept is the notion of *echelon inventory levels* and *echelon holding costs*. The echelon inventory level of a stage is the total inventory that is at the stage, that is in transit from the parent stage to the stage, and that is at or in transit to *any* of the downstream stages in the system. The echelon holding cost at a stage is the difference between the cost of holding a unit of inventory at the stage for a time period, and the cost of holding it at the parent stage. Without loss of generality we assume that the echelon holding costs are nonnegative; if this is not the case then no inventory will be stored at the parent stage. Holding costs can be correctly accounted for either in the conventional manner, or using echelon inventories and echelon holding costs - for details see [18] and Section 2 below.

The echelon inventory level provides a compact description of the structure of the optimal policies. In

each period, as a function of the current state of the system (but independent of the inventory control policy used in previous periods), there are target echelon inventory levels for each stage, also called echelon base-stock levels. The optimal policy aims to keep the echelon inventory levels as close as possible to the target base-stock levels. That is, at the beginning of the period, if the echelon inventory level at a stage is below the target, the stage orders enough to bring the echelon inventory up to the target, or it orders all of the inventory that is currently on hand at the parent stage. If the echelon inventory level is above the target, no order is placed.

The research literature on stochastic serial inventory systems has two cost re-allocation schemes, which effectively shift costs from one stage to another. Both of these schemes re-allocate costs between stages for a single time period only, but they do not shift costs from one time period to another. The first published cost re-allocation scheme is the classical, dynamic-programming-based proof of the optimality of an echelon base-stock policy. This proof decomposes the problem into a sequence of single-stage problems, and then proves the optimality of an echelon base-stock policy recursively, stage by stage. Each stage passes an *implicit penalty function* to the parent of the stage. The penalty function captures the impact that decisions made by upper stages will have on the stage in question. This approach was introduced by Clark and Scarf [4], was simplified by Zipkin [18], and has been extended by Dong and Lee [5]. Computationally, a series of single-stage problems are solved to optimality, one for each stage, with each stage assuming that the lower stages have already been optimized. Part of the computation done at each stage is to compute the implicit penalty function that will be passed to the parent stage.

The second cost re-allocation scheme is due to Chen and Zheng [3], for a model in which each stage also has a *fixed set-up cost* that is incurred whenever the stage places an order. In addition to using induced penalty functions, they propose *cost parameter allocation* bounds, in which they create a "component" for each stage in the original system, and a separate multi-stage inventory model for each component. They then allocate the backordering penalty cost parameter and the holding cost parameters to the different components.

Unfortunately, the rather simple state-dependent base-stock form of the optimal policies does not always lead to efficient algorithms for computing the optimal policies. The corresponding dynamic programs are tractable in cases where the demands in different periods are independent and do not evolve over time, although even in such simple scenarios the computations involved can be tedious. The standard dynamic programming approach can still be tractable in models with exogenous Markov-modulated demand but under rather strong assumptions on the size of the state space of the underlying Markov process (see, for example, [17, 2]). However, in many scenarios with more complex demand structure the state space of

the corresponding dynamic programs grows at an exponential rate in the number of periods and explodes very fast. Thus, the straightforward approach to solving the corresponding dynamic programs becomes practically (and often also theoretically) intractable even for single-stage models (see [8, 5] for relevant discussions on the MMFE model). This is especially true in the presence of complex demand structures where demands in different periods are correlated and evolve over time. The difficulty, known as *the curse of dimensionality*, essentially comes from the fact that we need to solve 'too many' subproblems. Because of this phenomenon, it seems unlikely that an efficient algorithm to solve these huge dynamic programs to optimality exists.

Muharremoglu and Tsitsiklis [13] have proposed an alternative approach to the dynamic programming framework. They have observed that this problem can be decoupled into a series of *unit supply-demand subproblems*, where each subproblem corresponds to a single unit of supply and a single unit of demand that are matched together. This novel approach enabled them to substantially simplify some of the dynamic-programming-based proofs on the structure of optimal policies, as well as to prove several important new structural results. In particular, they have established the optimality of echelon state-dependent base-stock policies under general exogenous Markov-modulated demand and stochastic lead times with no order-crossing. Using this unit decomposition, they have also suggested new methods to compute the optimal policies. However, their computational methods are essentially dynamic programming approaches applied to the unit subproblems, and hence they suffer from similar problems in the presence of correlated and evolving demands. Janakiraman and Muckstadt [9] have extended their approach to systems with capacity constraints on the size of the order. Although our approach is very different from theirs, we use their unit-decomposition approach extensively as a technical and descriptive tool.

As a result of this apparent computational intractability, many researchers have attempted to construct computationally efficient (but suboptimal) heuristics for these problems [7, 1, 5, 16, 14]. However, we are aware of no attempts to analyze the worst-case performance of these heuristics. Moreover, we are aware of no computationally efficient policies for which there exists a worst-case analysis that establishes constant performance guarantees.

Rosling [15] has shown that the optimality of echelon base-stock policies holds also in the more general assembly system. Moreover, he has shown that under this condition an assembly system can be transformed into an equivalent serial system, in that policies in one system can be mapped to policies in the other system with the same cost. In particular, echelon base-stock policies in the serial system are mapped to echelon base-stock policies in the assembly system. However, since the assembly system is a generalization of the serial system, it is clear that the above-mentioned computational challenges still exist, and might be even

harder.

This paper extends the recent work by Levi, Pál, Roundy and Shmoys [11] and Levi, Roundy, Shmoys and Truong [12], who have considered, respectively, the uncapacitated and the capacitated single-item, single-stage periodic-review system, with general demand structures that allow correlation and evolution over time. They have proposed a class of computationally efficient policies that they call dual-balancing policies and shown that these policies admit a worst-case performance guarantee of 2. Their work is based on two novel ideas. First, they have introduced a new marginal cost accounting scheme, in which they assign to a decision all present and future costs that, as a result of this decision, become independent of any future decision, and are only a function of future demands. This is in contrast with standard approaches that directly associate with each decision only the costs incurred in that period (or more generally a lead time ahead). Secondly, they have used cost balancing techniques to construct dual-balancing policies. In each period, the dual-balancing policy orders such that the (conditional) expected marginal holding costs incurred by the units ordered is equal to the (conditional) expected marginal backordering cost that will be incurred one lead time into the future. These two costs oppose each other, in that the expected marginal holding cost is an increasing function of the size of the order, while the expected marginal backordering costs is a decreasing function. Thus, it is always possible to order a quantity that will make these two costs equal to each other. The worst-case analysis for the single-item, single-stage models is based on simple, yet powerful, amortization of the costs incurred by the dual-balancing policy and by an optimal policy. In particular, it can be shown that the optimal policy incurs on expectation at least half of the costs incurred by the dual-balancing policy, which implies that the dual-balancing policy has a worst-case performance guarantee of 2.

The models considered in this paper are direct generalizations of the uncapacitated single-stage models in [11]. We use cost balancing techniques, and an innovative, multi-stage version of marginal cost accounting scheme, also called *cause-effect cost accounting*. The goal is to associate each element of cost with the decision that directly causes this cost, where decisions are indexed by the time period and the stage at which the ordering decision arose. Our approach is based on the notion of the *critical chain*. Focus on a single unit of demand and suppose that we know in advance when demand for this unit will occur. It is then clear that we would like to order and deliver this unit of inventory in a just-in-time fashion, so that will arrive at stage 1 exactly on time (incurring no backordering penalty cost) and will not be delayed at any intermediate stage (avoiding unnecessary holding costs). The critical chain is the just-in-time path that the unit would ideally follow through the inventory system, and the *critical time periods* are the time periods when orders for the unit would ideally be placed, at each stage in the system. We use the critical chain and the critical

time periods as points of comparison. There is a critical time period for each unit and each stage.

We categorize all holding costs into three categories, and we assign the holding and backorder costs to ordering decisions, in the following manner. Holding costs that are incurred while units are in transit between stages are called *pipeline costs*. The pipeline costs are inevitable - they will be incurred by every unit that is demanded, regardless of the policy followed. For a given unit, the the echelon-*n* pipeline cost is assigned to stage *n*, and to the time period in which the unit was ordered by stage *n*. Whenever we order a unit at a stage earlier than the critical time we cause *early holding costs* to be incurred. These are extra holding costs, in addition to the pipeline costs. The early holding costs incurred by a given unit are assigned to the stages and the time periods in which the unit was ordered early, relative to the critical time period. Similarly, whenever a unit is ordered at a stage *after* the critical time period, the tardiness of this order indicates that, relative to the critical chain, we will incur both backordering costs and extra holding costs (called *late holding costs*). We assign these costs to carefully selected time periods in which the unit could have been ordered, but was not. For a more detailed discussion, see Subsection 3.3 below.

For a given ordering decision, taken at a given stage and in a given time period, we consider the expectation of all of the costs assigned to this order. We separate these expected costs into two opposing functions. The expected pipeline and early holding costs are increasing in the size of the order placed by the stage in the time period. On the other hand, the expected backordering and late holding costs are decreasing in the size of the order. This leads to a balancing policy that, at each stage in each period, selects an order quantity to balance these two opposing cost functions. This policy is computationally efficient, and can be implemented in an on-line manner, that is, regardless of any future decisions.

The worst-case analysis of the serial case is significantly harder than the single-stage analysis discussed in [11] and [12]. It is again based on amortization of the cost incurred by the balancing policy with the cost incurred by an optimal policy. The presence of multiple stages requires new approaches. Beyond the complexities described above, the main difficulties come from the fact that when the balancing algorithm orders more than the optimal policy, the optimal policy does not necessarily incur the late holding costs that the balancing policy incurs (see Section 5 below).

We show that our policy achieves a worst-case expected cost of 2 times the expected cost of the optimal policy under, the assumption that the holding cost for a stage in the serial system is charged as soon as a unit is ordered by the stage. This cost model will be referred to as *Model 1*. The alternate assumption, referred to as *Model 2*, is that holding costs begin to be charged when a unit arrives at a stage. Under Model 2 our modified balancing policy achieves a worst-case performance ratio of 3. Additionally, when the backorder cost is greater than the echelon holding cost for all but the final downstream stage, we show

that the performance ratio can be reduced to $2\frac{1}{3}$. We also extend our results to assembly systems and discrete-valued demands.

The rest of the paper is organized as follows. Section 2 defines the serial system being studied. In Section 3 we describe in detail our scheme for cost accounting in serial systems and prove that it is consistent and accurate. Section 4 presents the balancing policy and its extensions, and proves the performance bounds. Finally, in Section 6 we use a result of Rosling [15] to extend all of our results to assembly systems.

2 Inventory Control Problem for Serial Systems

In this section, we provide the mathematical formulation of the periodic-review stochastic inventory control problem for serial systems (Model 1), and introduce some of the notation used throughout the paper. We consider a finite planning horizon of T periods numbered $t=1,\ldots,T$. The demands over these periods are random variables, denoted by D_1,\ldots,D_T . We use $D_{[s,t]}$ to denote the accumulated demand over the interval [s,t], i.e., $D_{[s,t]}:=\sum_{j=s}^t D_j$.

There are N stages in the serial system, numbered $1,2,\ldots,N$, with 1 producing the finished product. We use n to refer to a generic stage. Each stage n can order inventory from the on-hand inventory at the preceding stage, stage n+1. The lead time to transport inventory from stage n+1 to n is l_n . Thus, the minimum time that elapses between the time when we order a unit of inventory at stage n, and when we deliver it to the client, is L_n . This is the cumulative time required to transport inventory from stage n+1 to stage 1 with no delays at any intermediate stage. Thus $L_n = l_n + L_{n-1}$, and $L_0 = l_0 = 0$. Both l_n and l_n are integers. We assume that the lead times l_n are strictly positive; otherwise two stages could be merged without loss of generality. By convention, stage n+1 denotes an external supplier with infinite capacity. We assume that it is possible to order a unit of inventory from the external supplier and deliver it to the client before the time horizon ends, i.e., that n+1 we define n+1 we define n+1 and n+1 for n+1 to the client before the time horizon ends, i.e., that n+1 the define n+1 to n+1 the define n+1 to n+1 the product n+1 to n+1 the client n+1 the define n+1 the product n+1 to n+1 the product n+1 the product

As a general convention, we distinguish between a random variable and its realization using capital letters and lower case letters respectively. However the notation T, N and L_n differs from this convention and corresponds to deterministic quantities.

Demand can be observed by all the stages in the serial network, but is only satisfied at the end stage, stage 1. As part our demand model, we assume that at the beginning of each period s we are given what we call an *information set*, denoted by f_s . The information set f_s contains all of the information that is available at the beginning of time period s. More specifically, the information set f_s consists of the realized demands (d_1, \ldots, d_{s-1}) over the interval [1, s), and possibly some more (external) information

denoted by (w_1,\ldots,w_s) . The information set f_s in period s is one specific realization in the set of all possible realizations of the random vector $F_s=(D_1,\ldots,D_{s-1},W_1,\ldots,W_s)$. This set is denoted by \mathcal{F}_s . In addition, we assume that in each period s there is a known conditional joint distribution of the future demands (D_s,\ldots,D_T) , denoted by $I_s:=I_s(f_s)$, which is determined by f_s (i.e., knowing f_s we also know $I_s(f_s)$). For ease of notation, D_t will always denote the random demand in period t according to the conditional joint distribution I_s for some $s \leq t$, where it will be clear from the context to which period t we refer. We will use t as the general index for time, and t0 will usually refer to the current period.

The only assumption on the demands is that for each $s=1,\ldots,T$, and each $f_s\in F_s$, the conditional expectation $E[D_t|f_s]$ is well defined and finite for each period $t\geq s$. In particular, we allow non-stationarity and correlation between the demands of different periods. We note again that by allowing correlation we let I_s be dependent on the realization of the demands over the periods $1,\ldots,s-1$ and possibly on some other information that becomes available by time s (i.e., I_s is a function of f_s). Note, however, that the information set f_s as well as the conditional joint distribution I_s are assumed to be independent of the specific inventory control policy being considered. This model accommodates all published mechanisms by which forecasts evolve. (Forecast evolution means that on January 1, 2006, I create a forecast of the demand that will occur during August, 2006. On February 1, 2006, I create a new forecast of the August, 2006 demand. Forecast evolution is a model of the mechanism by which the first of these forecasts evolves into the second one. For more details see [11]).

We assume that every stage has access to complete information, meaning that all of the current inventory levels, and the current information set f_s , are visible to all stages of the inventory system. (This assumption is slightly stronger than what we actually require.)

Cost Minimization

In the periodic-review stochastic inventory control problem for serial systems, our goal is to supply each unit of demand arising at stage 1 while attempting, at each stage in the network, to avoid ordering supply either too early or too late. In period t, t = 1, ..., T, three types of costs are incurred - a per-unit shipping cost c_n per unit ordered at stage n (and shipped from n+1 to n), a unit backordering penalty π that is incurred for each unsatisfied unit of demand at the end of period t at stage 1, and holding costs. Unsatisfied units of demand are usually called *backorders*. Backorders fully accumulate over time until they are satisfied. That is, each unit of unsatisfied demand will stay in the system and will incur a backordering penalty in each period until it is satisfied. The backorder cost π is only charged at the final stage.

In this paper we use the echelon approach to account for holding costs. The "conventional" approach

to accounting for holding costs in serial inventory systems is to charge h'_n dollars at the end of each time period, for each unit currently held in inventory at stage n, or in transit from stage n+1 to stage n. Without loss of generality, $h'_n - h'_{n+1} = h_n \geq 0$. (Otherwise we would merge stage n+1 into stage n. We assume $h_{N+1} = h'_{N+1} = 0$.) The *echelon* approach to holding costs is based on the echelon inventory position which, for stage n, is the total inventory at any stage k or in transit from k+1 to k, for $1 \leq k \leq n$. At the end of every time period, for each stage n, we incur the echelon holding cost h_n for every unit that is part of stage n's echelon inventory position. (To see that the approaches are equivalent, assume that at the end of a given time period the inventory at stage n or in transit to stage n is n, and that the echelon inventory position is n. Then n is n in the equivalent n in the equivalent n in the equivalent n in the equivalent n is n in the equivalent n in the equivalen

Two assumptions about the time at which echelon-n holding costs begin to be charged have appeared in the literature. We begin with Model 1, which assumes that they are charged from the time period in which the unit is ordered by stage n (e.g., [5], [10]). This assumption is reasonable when the cost of capital is larger than the cost of physically storing inventory, and inventory is paid for when orders are placed rather than when they arrive. After deriving results for Model 1, we will extend the analysis to Model 2, which assumes that they are charged from the time period at which the unit arrives at stage n.

The objective of the problem is to find a feasible ordering policy (i.e., one that respects the system constraints) that minimizes the overall expected ordering cost, holding cost and backordering cost. We consider only policies that are *non-anticipatory*, i.e., at time s the information that a feasible policy can use consists only of f_s and the current inventory levels. We use P to denote a generic policy. For a given policy P, conditioning on a specific information set f_s , we know the current on-hand and in-transit inventory levels at all stages deterministically.

System Dynamics

Given a feasible policy P, we describe the dynamics of the system using the following terminology. Let $X_n(t)$ denote the echelon inventory position for stage n at the start of time period t. The echelon inventory position is the echelon inventory, plus the number of units in transit from stage n+1 to stage n, minus the current backorders at stage 1. Stated differently, the echelon inventory position consists of the total number of units at any stage m satisfying $m \le n$, plus the number of units in transit from a stage m+1 to a stage m where again $m \le n$, minus the current backorders at stage 1. Let $Q_n(t)$ be the order quantity for stage n

¹The nouns 'stage' and 'echelon' are usually treated as synonyms. We will use 'stage' to refer to a location where inventory is stored, and 'echelon' when discussing echelon inventory levels or the echelon approach to accounting for holding costs.

at t. Let $Y_n(t)$ denote the echelon inventory position of stage n after ordering at t, but before the demand is satisfied. That is, $Y_n(t) = X_n(t) + Q_n(t)$ and $X_n(t+1) = Y_n(t) - D_t$. We assume that at time t, stage n will not place an order unless there is enough inventory at stage n+1 to fill the order immediately. Thus, $Q_n(t)$ is bounded from above by the inventory physically on hand at stage n+1, after orders have arrived in period t and before orders are placed. In other words, $Y_n(t) = X_n(t) + Q_n(t)$ is bounded from above by the total amount of inventory in the system that reaches stage n+1 by time t, minus the backorders at the start of period t. This quantity is $NI_{n+1}(t)$, the echelon inventory level at stage n+1 at the beginning of time period t. Note that $NI_n(t)$ includes all units in $X_n(t)$ that are not in transit to n after orders have arrived in period t, i.e., $NI_n(t) = X_n(t) - \sum_{j=t-l_n+1}^{t-1} Q_n(j) = Y_n(t) - \sum_{j=t-l_n+1}^{t} Q_n(j)$. Viewed differently, $Y_n(t) - D_{[t,t+l_n)} = NI_n(t+l_n)$.

Since time is discrete, we next specify the sequence of events in each period s. Since s is the current time, the quantities of interest are no longer random, and are written in lower case.

- 1. Period s begins with echelon inventory position $x_n(s)$ and echelon inventory level $ni_n(s-1)-d_{s-1}$. We observe the information set $f_s \in \mathcal{F}_s$, from which we obtain an updated conditional joint distribution I_s for future demands.
- 2. At each stage n, the order placed in period $s l_n$ of $q_n(s l_n)$ units arrives, and the echelon inventory level increases accordingly to $ni_n(s-1) d_{s-1} + q_n(s-l_n) = ni_n(s)$.
- 3. At each stage n the order quantity for period s is selected, i.e., following a given policy $P, q_n(s)$ units are ordered. The order is constrained by $0 \le q_n(s) \le ni_{n+1}(s) x_n(s)$, i.e., the order quantity in period s for echelon n can not exceed the inventory on hand at the preceding stage n+1. Consequently, the echelon inventory position is raised by $q_n(s)$ units, from $x_n(s)$ to $y_n(s)$, where $x_n(s) \le y_n(s) = x_n(s) + q_n(s) \le ni_{n+1}(s)$. This results in a cost of $c_n q_n(s)$.
- 4. We observe the demand d_s in period s, which is realized according to the conditional joint distribution I_s . For each stage n, the echelon inventory level decreases from $ni_n(s)$ to $ni_n(s) d_s$, and the echelon inventory position decreases from $y_n(s)$ to $y_n(s) d_s = x_n(s+1)$.
- 5. Period s ends. Each unit that is backordered at the end of period s results in a backorder cost of π . Each unit that has been ordered by stage n and is in the system results in a holding cost of h_n . In 3 above, recall that node N+1 represents the external supplier. We assume that the external supplier has an infinite supply of inventory, i.e., $ni_{N+1}(s) = \infty$ for all s and all s. (Hence, s infinite.)

At the beginning of the time horizon we inherit an inventory system that is already in operation. Therefore the quantities $x_n(t)$ for $t \le 1$, $y_n(t)$ and $q_n(t)$ for $t \le 0$, and $ni_n(1)$ are pre-determined constants for all stages n. In addition, $NI_n(t)$ for $1 < t \le l_n$ are beyond our control, being functions of past decisions and future demands.

Initially we assume that inventories are measured and managed in continuous, rather than discrete, quantities. Consequently the quantities $x_n(t)$, $y_n(t)$, $q_n(t)$ and $ni_n(t)$ can all assume fractional values. In Section 5.2 below we extend the analysis to accommodate inventory systems that measure and manage inventories using integer-valued quantities, and assumes that the demands are integer-valued.

The manner in which the end-of-horizon costs are defined is most easily described after the concepts in subsection 3.1 have been introduced. Therefore we defer a discussion of end-of-horizon costs until subsection 3.2 below.

3 Cause-Effect Cost Accounting for Serial Systems

In this section, we describe and analyze a new scheme for accounting for costs in serial systems. The phrase "cause-effect cost accounting" reflects the logic that we use to create, and to describe, our cost accounting procedures. The general approach is to take costs incurred in an inventory system, classify them into categories, and use the categories to assign the costs to specific decisions that were made. The dominant logic used to create this assignment is cause and effect - we assign a cost to the specific decision that caused that cost to be incurred. Moreover, the costs assigned to each ordering decision are not affected by any decisions that will be made in the future, and depend only on future demands.

Cause-effect cost accounting differs strongly from most of the literature on inventory systems, which is based on dynamic programming formulations, myopic approaches, or steady-state analysis (see, for example, [18]). In most of that literature costs are accounted for either in the time periods in which they are incurred, or one lead time earlier. In our approach we assign to a decision all costs that were made inevitable by the current decision, whether they are incurred at the present time or in the future. Cause-effect cost accounting for stochastic inventory models originated in [11], where it was called *marginal cost accounting*.

The exposition of our cost accounting scheme, and many of our proofs, are unit-based, i.e., we track the progress of unit k through the system, and assign the costs incurred by unit k to the decisions that caused these costs to be incurred. However we note that this is purely an expositional and an analytical device. From a computational point of view the functions that we manipulate look a lot like functions used in classical inventory theory, and are not unit-based (see Subsections 3.6 and 3.7 below).

First we outline a framework that makes rigorous the notion of a unit-by-unit cost decomposition, and discuss end-of-horizon costs. We then focus on a single unit of inventory and give an overview of our cause-effect approach to cost accounting with reference to this unit. Next we provide a graphical context that we will use in the rest of the paper. In the remainder of the section we give a detailed analysis of our cost accounting mechanisms and establish algebraic expressions for the aggregated costs assigned to each decision.

3.1 Ordering Numbers for Demand and Supply

We will use conventions and techniques similar to the ones used by Muharremoglu and Tsitsiklis [13] (also see [11] for more details). The main idea is that, without loss of generality, we can assume that units of supply are consumed by the demand on a first-ordered-first-consumed basis, and that we can *match* each unit of supply to the specific unit of demand it will be used to satisfy. More rigorously, let \mathcal{L}_D be a half-infinite line segment $[0, \infty)$ that represents the units of demand that might be realized over the planning horizon. If demands are continuous then demand units are of infinitesimal size. The unit that is located a distance of k from the origin is called *unit* k. Without loss of generality, clients purchase and accept delivery of these units in a sequence that is increasing in the distance k. Consequently, the first $\sum_{t=1}^{T} D_t$ of these units correspond to demand that will occur before time period T ends. These units are called *demand units*.

Similarly, let $\mathcal{L}_S = [0, \infty)$ be a half-infinite line segment of supply units, also starting at the origin. It represents all units of inventory that we have obtained, or can obtain, over the planning horizon. The unit of inventory on \mathcal{L}_S that has a Euclidean distance of k from the origin is called *supply unit* k. We assume that supply units are ordered and are delivered to the client in a sequence compatible with this distance, i.e., if $k \leq \sum_{t=1}^T D_t$ then supply unit k will be used to satisfy demand unit k. (This is equivalent to the first-ordered-first-used assumption.) Since the k-th unit on \mathcal{L}_S is matched with the k-th unit on \mathcal{L}_D , we refer to the first $\sum_{t=1}^T D_t$ units on both lines as *demand units*. Units k on both lines that are not required before the time horizon ends (i.e., that satisfy $k > \sum_{t=1}^T D_t$) are called *excess supply units*. Note that we can describe each policy P in terms of the time period in which it orders each unit k, at each stage n.

At the beginning of time period 1 we inherit an inventory system that is already in operation. At this point in time, the existing inventory that is already in the system is located along \mathcal{L}_S according to its proximity to stage 1. This means that the inventory that is currently at stage 1 is located next to the origin on \mathcal{L}_S , followed by the inventory that can or will reach stage 1 in time period 2, etc.

Although we are modeling inventory as a continuous quantity, the exposition of this paper is more natural if we talk about units in discrete terms. In this paper we will mention a number of conditions that define sets

of units. In every case, the sets of units that satisfy the condition are intervals on \mathcal{L}_D (or, equivalently, \mathcal{L}_S). When we speak of the *number of units* that satisfy the condition, we refer to the length of this interval.

We define Δ_k to be the random time period in which supply unit k is demanded. Thus, Δ_k is a random variable that is observed at the end of period δ_k or, equivalently, at the beginning of period $\delta_k + 1$ (where δ_k is again the realization of Δ_k). Note that the demand units are the units k that satisfy $\delta_k \leq T$. If unit k is an excess supply unit then we define $\delta_k = T + 1$. At time $s < \Delta_k$ the random time Δ_k has a conditional distribution that is a function of the current information set f_s , which can be inferred from historical demands and the distribution $I_s(f_s)$ of the future demands (D_s, \ldots, D_T) .

3.2 End-of-Horizon costs

Having discussed unit decomposition, we are ready to specify our end-of-horizon assumptions. After period T no demand occurs. After period $T-L_n$ it is clear that any inventory ordered at stage n will not reach the client before the horizon ends. Consequently we only model ordering decisions at node n taken in time periods s, $1 \le s \le T-L_n$. At stage n after period $T-L_n$, we assume that no orders are placed. That means that no orders will arrive at stage n after time $T-L_{n-1}$. Also, after the demand for unit k has occurred, if it is still possible to deliver the unit to the client before the end of period T, then the optimal policy will not defer ordering unit k at any subsequent stage in the serial system. We limit attention to policies that have these two properties, called *regular policies*.

For regular policies, at the end of period T no units of inventory will be in transit from one stage to another. If an excess supply unit (a unit k with $\delta_k > T$) is at stage n, we assume that the unit is salvaged at the excess end-of-horizon cost of $\sum_{m=n}^{N} (-c_m - L_m h_m)$. This corresponds to refunding the unit purchase costs that were incurred, and the holding costs incurred while unit k was in transit. We make this assumption primarily for ease of exposition, noting that all of the results in this paper hold if we generalize this expression for the end-of-horizon costs to $\sum_{m>n} \xi_m$, where $\xi_m \geq -c_m - L_m h_m$ for $1 \leq m \leq N$.

At the end of period T, if a demand unit (a unit k with $\delta_k \leq T$) is at stage n>1, then we incur an end-of-horizon shortage cost of $\sum_{m=1}^{n-1}(c_m+L_mh_m)$. Note that this is the echelon holding cost and ordering cost that would be incurred for stages downstream from stage n, if we were to advance the unit to stage 1 without any more delays. This assumption is also made for ease of exposition - all of the results in this paper hold if we generalize this expression to $\sum_{m=1}^{n-1}\sigma_m$, where $\sum_{m=1}^{n-1}(\sigma_m-c_m-L_mh_m)\geq 0$ for $2\leq n\leq N+1$. Of course, all end-of-horizon costs are in addition to the typical end-of-period holding and backorder costs incurred at the end of period T.

Our end-of-horizon costs differ from traditional ones in one interesting way - demand units and excess

supply units differ in their end-of-horizon costs. In the past most authors have not done unit matching in the sense of subsection 3.5 and, consequently, have salvaged all end-of-horizon inventory in the same manner. If the inventory system were to cease to operate at the end of time period T that would make sense. However in practice, finite-horizon models are usually used in rolling-horizon mode, in settings where the business will not cease to operate at the end of period T. End-of-horizon costs are used to minimize the end-of-horizon effect. In those settings it makes sense to assign different end-of-horizon costs to demand units and excess supply units. It also makes sense to charge lower end-of-horizon costs for demand units that are closer to the client, because they are likely to reach the client more quickly. The generalizations described above allow end-of-horizon costs for demand units and excess supply units to be either identical or different.

3.3 An Overview of Cause-Effect Cost Accounting

In this subsection, we give an overview of our cause-effect cost accounting scheme. The overview is intended to provide context for the detailed analysis that follows. For ease of exposition, we ignore beginning-of-horizon and end-of-horizon considerations in this subsection.

We use the notation $\langle s,n\rangle$ to refer to the ordering decision taken at stage n at time s. In light of the definition of regular policies in Subsection 3.2 above, we model ordering decisions $\langle s,n\rangle$ for which $1\leq s\leq T-L_n$. We have already defined δ_k as the time at which demand for unit k occurs. We now define u_{kn} as the time at which supply unit k is ordered by stage n, for a given policy P. In other words, for all n, unit k is one of the units of inventory that comprise the order $\langle u_{nk},n\rangle$. Note that whereas Δ_k is an uncontrolled random variable whose value is observed at the end of time period δ_k , u_{kn} is a user-controlled decision variable. For a given policy P the decision U_{kn} is a-priori random. Specifically, in period s, if unit k is ordered by stage s after time s, then the value s0 is probably random, because the timing of the order probably hinges on information that is not yet available. If unit s1 is ordered by stage s2 at or before time s3 then s3 known with certainty, and is less than or equal to s3.

A central concept in our cost-accounting scheme is the notion of a *critical time period* associated with each unit k and each stage n. At the beginning of time period 1, focus on some unit k, and momentarily assume that we already know when the demand for unit k will occur, i.e., we know the value of δ_k . It is clear that in order to minimize the total cost incurred, we should order unit k in just-in-time fashion, deferring orders for the unit as long as possible to avoid unnecessary holding costs, while ensuring that it arrives at stage 1 at the beginning of period δ_k to avoid a backordering penalty. In particular, the value of δ_k induces a *critical time period* $\delta_k - L_n$ for each stage n, which is the 'ideal' time period in which stage n should order unit k to minimize the costs incurred. In general, of course, the value of δ_k is not known in advance,

making it difficult to anticipate which time periods are critical, and making the ideal trajectory for unit k impossible to follow.

We now categorize the costs incurred by unit k, and we assign them to different ordering decisions, in the following manner. The rest of this section is mostly written under the assumption that the current time is t = T + 1, so all quantities are deterministic.

Pipeline holding costs are the holding costs incurred by unit k while it is in transit from one stage to another.² After unit k has been ordered at stage n it will spend a total of L_n time periods in transit before reaching stage 1, and will incur a total echelon-n pipeline holding cost of h_nL_n . We assign the echelon-n pipeline holding costs to the order placed for unit k at stage k. The unit shipping costs are charged in like manner. That means that for unit k, a total pipeline cost of k, is assigned to the ordering decision k and k. Note that the total pipeline cost incurred by demand unit k is k and k and independent of k, if we ignore the beginning-of-horizon and end-of-horizon effects (i.e., assuming that unit k is a demand unit that is at stage k at the beginning of time period 1).

Early holding costs are holding costs that are incurred while unit k is held in inventory at some stage of the inventory system, when it is still possible to get the unit to the client by the due date. Early holding costs are incurred in scenarios where unit k is ordered by some stage n prior to the critical period $\delta_k - L_n$, that is, where $\delta_k - L_n - u_{kn} > 0$. As we will argue in subsection 3.6 below, the total number of time periods over which echelon-n early holding costs will be incurred is $(\delta_k - L_n - u_{kn})^+$. This is the total number of time periods that unit k will be held in inventory at some stage between stage n and stage 1, if it is delivered to the client at time δ_k . Therefore the total early holding cost of unit k at stage n is $h_n(\delta_k - L_n - u_{kn})^+$. We assign this cost to the ordering decision $\langle u_{kn}, n \rangle$, adopting the viewpoint that these costs were caused by the decision to order unit k at stage n, taken at time u_{kn} .

Late holding costs are holding costs that are incurred while unit k is held in inventory at some stage of the inventory system, when it is no longer possible to get the unit to the client by the due date. They are incurred in scenarios where, for some stage n, unit k is not ordered by the end of the critical period, i.e., where $\delta_k - L_n - u_{kn} < 0$ for some n. Specifically, assume that unit k is stored at stage n from the some time period s to period s+1, and that after this occurs it is no longer possible to deliver the unit to the client on time (i.e., that $\delta_k \leq s + L_{n-1}$). Every time that this happens, the conventional holding $\cot h'_n = \sum_{m \geq n} h_m$ is incurred for unit k. We adopt the view that this cost was caused by the fact that at time s, stage n-1 did not order unit k. Therefore, whenever a unit k is stored at stage n from some period n0 to period n1, and n2 to period n3 to period n4. Therefore, whenever a unit n5 is assigned to the ordering decision n5.

²This definition will be modified in Subsection 3.6 below, to accommodate end of horizon considerations.

Finally we consider the *backorder costs*. Consider a unit k and a time period t such that $t \geq \delta_k$. If unit k was not delivered to the client before the end of period t then a backorder cost of π is incurred for unit k in period t. This cost would have been avoided if every stage m had ordered unit k by the critical period $t - L_m$. Let n be the largest stage index at which unit k was not ordered by time $t - L_n$. Therefore, at time $t - L_{n+1}$, stage n + 1 had ordered unit k. Hence, at time $t - L_{n+1} + l_{n+1} = t - L_n$, unit k was on-hand at stage n + 1, and could have been ordered by stage n. Our point of view is that the decision to not order unit k at time k - k at stage k - k, caused the backorder that occurred at the end of period k. (Note that time k - k such that unit k - k has not reached the client by the end of period k - k.). Hence, for unit k - k to the ordering decision k - k, where k - k is defined in the manner just described.

Some comments regarding holding costs are in order. First, it is important understand the manner in which we use the terms echelon-n holding cost and holding cost assigned to an order placed at stage n. An echelon-n holding cost is a cost expression that is linear in the coefficient h_n . Such a cost might, or might not, be assigned to an order placed at stage n. Second, note again that at time s the value of Δ_k is often unknown. Therefore, when we make decision $\langle s, n \rangle$ we do not know whether unit k will assign early or late holding costs to $\langle s, n \rangle$, and we do not know the amounts of these costs. Decisions at time s that effect unit k must be made using expected values rather than exact costs. This is made possible by the fact (mentioned in the second paragraph of this subsection) that at the beginning of time period s, subsequent decisions have no impact on the costs assigned to $\langle s, n \rangle$. These costs depend only on the past, on the current ordering decision for stage n, and on future demands. Thirdly, note that if unit k is at stage 1 at time s, and if $\delta_k \leq s$, then we will deliver unit k to the client immediately rather than holding it in inventory to period s+1. Therefore no late holding costs are incurred while unit k is at stage n=1. Fourth, since $h'_{N+1}=0$, no late holding costs are ever assigned to orders placed at stage N. (Backorder costs can be assigned to orders placed at N.) However, to orders placed at all other stages $1 \le n < N$, backordering penalty and late holding costs are assigned together. That is, a decision not to order a unit after the critical period causes one period of backorder costs and one period of late holding costs, both of which are assigned to the order.

The cause-effect assignment of costs usually gives rise to two crucial properties that enable balancing approaches to inventory systems management. These properties will be explicitly stated later on, but briefly, they are the following. First, for each order that is placed, the costs assigned to it are represented by two functions, one being a non-decreasing function of the order quantity, and the other being a non-increasing function of the quantity. Secondly, after the order has been placed, the costs captured by these two functions are now beyond our control, because they are not impacted by any decisions that will be made in the future.

These two properties are fundamental to the analysis of balancing algorithms, and are usually a natural consequence of the process of assigning costs to decisions based on cause-and-effect logic. (In this subsection we have given an overview of the assignment of costs to orders. We will discuss the aggregation of these costs into functions in Subsections 3.6 and 3.7.)

3.4 A Graphical Scheme for Cost Accounting

Before analyzing the properties of our cost accounting scheme, we present a graphic representation for the movement of each unit of supply through the inventory system over the time horizon. The *stage-distance* is defined as the number of time periods required for the unit to arrive to stage 1 if there are no delays at any intermediate stages. For example, a supply unit that is physically located at stage n has a stage-distance of L_{n-1} . A supply unit that was shipped from stage n one time period ago has a stage-distance of $L_{n-1} - 1$.

We can describe the complete history of a supply unit k, from the moment it is ordered by stage N to the moment it is delivered to the client at stage 1, by visualizing its stage-distance as a function of time. Consider a lattice of integer-valued points (t,j) with $1 \le t \le T+1$ and $0 \le j \le L_N$, illustrated in Figure 3.1. A supply unit that has stage-distance j at the beginning of time period t is said to pass through the point (t,j). Thus, in the lattice, time is represented on the horizontal axis and stage-distance on the vertical axis. The time horizon ends at the end of time period t, which is the beginning of time period t and corresponds to the last column in the lattice. The points in the lattice are interconnected by horizontal arcs. In addition, from every point t0 in the lattice there is an arc that descends to the right at a 45-degree angle. All arcs are oriented from left to right, and connect nodes that are one time unit apart.

The point (t, j) is called a *stage point* if $j = L_{n-1}$ for some stage n. When a supply unit is at stage point (s, L_{n-1}) it is at stage n. If unit k is held in inventory at stage n from time t to time t+1 then it traverses the horizontal arc connecting the stage points (t, L_{n-1}) and $(t+1, L_{n-1})$. In this case unit k travels in time only, and does not get closer to the client. The bold horizontal arcs in Figure 3.1 go through the stage points, and are the only horizontal arcs that can carry inventory.

The sequence of points linked by descending arcs is said to constitute a chain. Specifically, *chain t* consists of the set of points $(L_N, t - L_N), (L_N - 1, t - L_N - 1), (L_N - 2, t - L_N - 2), ..., (0,t)$ (i.e., the set of points whose coordinates sum to t), and the arcs that link them. The diagonal paths in Figure 3.1 indicate chains. If a unit is ordered at stage n at time s, then while it is in transit, it travels along diagonal arcs. Thus, if a policy orders a supply unit at time s at stage n, the unit will travel through the points $(s, L_n), (s+1, L_n-1), (s+2, L_n-2), ..., (s+l_n, L_{n-1})$. In other words, it travels both in time and physical distance simultaneously, along part of chain $s+L_n$. In addition, unit s moves from one chain to

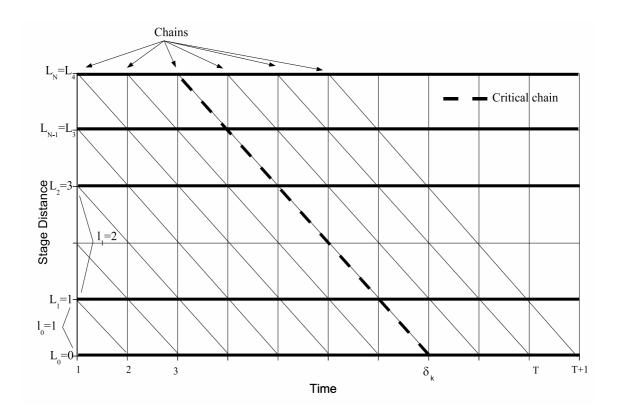


Figure 3.1: The Lattice and the Chains

the next when it is carried in inventory at a certain stage. Specifically, if unit k is held in inventory at stage n from period $t - L_n$ to period $t - L_n + 1$, then unit k moved from chain t to chain t + 1, since now the earliest period by which it can arrive at stage 1 is t + 1. The concept of moving between chains will be used extensively throughout the paper, and especially in the proof of Lemma 5.7 below.

Node N+1 represents the external supplier. We assume that the external supplier has an infinite supply of inventory, i.e., $X_{N+1}(t)=\infty$ for all t. However, inventory that is at the external supplier is not part of our inventory system, and does not incur any costs. If unit k is ordered by stage N at time s, then unit k traverses the arc from point (L_N,s) to point $(L_N-1,s+1)$. At time s, unit k enters the inventory system and begins to incur costs. Unit k leaves the lattice (and the inventory system) when it is delivered to a client, during or after time period δ_k , or at the end of time period T when the time horizon ends. As long as unit k is in the system it traces a monotonic path in the lattice, proceeding to the right and downwards.

Recall that at time 1 we inherit an inventory system that is already in operation. Consequently, at the beginning of time period 1, unit k might already be in the inventory system. In that case the path that unit k traverses through the lattice will start at some point (1, j), $j < L_N$ rather than at point $(1, L_N)$ (the external supplier). Unit k will proceed from its current location in the lattice downwards and to the right, as described

earlier. If unit k is in transit from stage n + 1 to stage n at the beginning of time period 1, then unit k is at a non-stage point (1, j) where $L_{n-1} < j < L_n$. In that case unit k will start out by following the diagonal arcs connecting the lattice points (1, j), (2, j - 1), ..., $(1 + j - L_{n-1}, L_{n-1})$.

As we have mentioned in subsection 3.2 above, unit k will not be ordered at stage n after period $T - L_n$ because it would not reach the client before period T ends. Therefore, to the right of chain T, no inventory will flow along diagonal arcs in the lattice. Inventory will flow horizontally to the right, until it reaches time period T + 1.

Consider a supply unit k, whose demand occurs at time δ_k . We define the *critical chain* to be chain δ_k , which is marked with a dashed line in Figure 3.1. The ideal strategy for unit k would be to keep the unit at the external supplier (node N+1) until the critical time period δ_k-L_N , and then to send it to the client along the critical chain. This is accomplished by ordering unit k at the critical time δ_k-L_n for each stage n. That would incur the pipeline costs, but no early or late holding costs, and no backorder costs. Unfortunately, when the key ordering decisions are made we usually do not know the value of the random variable Δ_k , so this strategy is usually impossible. Whatever strategy is used, the path traversed by a unit as it moves through the lattice contains all of the information required to completely determine the costs that it incurs. We will use the lattice and these paths as a framework to discuss the holding and backorder costs incurred. We start with the holding costs incurred by unit k.

3.5 Cost Decomposition

In the remainder of this subsection we focus on a single policy P and a single supply unit k, with a few exceptions that are specifically noted. We account for the holding costs incurred by unit k by examining the path it traverses in the lattice of Figure 3.1. Unit k incurs holding costs whenever it is in the inventory system at the end of a time period, i.e., whenever it traverses an arc in the lattice, other than the horizontal arcs in the top row (i.e., when it is at stage N+1). For unit k, our analysis splits the arcs in the lattice into three different categories, which correspond to three different types of holding costs that the unit can incur. The types of arcs, and of holding costs, follow.

1. The diagonal arcs in Figure 3.1 are pipeline arcs. The horizontal arcs that lie to the right of chain T+1 are also pipeline arcs. (Recall that to the right of chain T, unit k will only traverse horizontal arcs in the lattice.) *Pipeline holding costs* are incurred when unit k traverses a pipeline arc, i.e., when it traverses a diagonal arc along a chain, and when it traverses a horizontal arc to the right of chain T+1.

The *pipeline costs* consist of the pipeline holding costs just described, plus the per-unit ordering costs. Recall that when unit k is ordered at stage n (i.e., when k leaves the stage point (u_{kn}, L_n) along the diagonal emanating arc) a per-unit ordering cost of c_n is incurred, and echelon-n pipeline holding costs start to be incurred.

- 2. Early arcs are the horizontal arcs to the left of the chain δ_k . Early holding costs are incurred by unit k whenever it traverses an early arc.
- 3. Late arcs are the horizontal arcs that lie between chain δ_k and chain T+1. Late holding costs are incurred whenever unit k traverses a late arc.

In defining early and late arcs, recall that if demand for unit k never occurs we say that $\delta_k = T + 1$. In that event there are no late arcs, and unit k does not incur any late holding costs. Note that the definition of early and late arcs is dependent on the due date for unit k. Thus it differs from one supply unit to another, and for unit k, it is not observed until the end of time period δ_k .

The three types of costs we have described comprise the total holding-plus-ordering cost incurred by a unit. They account for its complete path from upstream to downstream. The ideal path for unit k to follow through the lattice is to follow the critical chain (chain δ_k) from stage N+1 to stage 1. If unit k follows this path, the total cost incurred by the unit is only the pipeline cost.

To obtain a cause-effect cost-accounting scheme we perform a second decomposition of the total holding cost incurred by unit k into increments. An *increment* of holding cost is the echelon cost of holding a unit of inventory for a single time period. Consider a stage n, and consider the set of all arcs [(t,j), (t+1,j')] in the lattice such that $L_n > j'$. Note that $j' \in \{j, j-1\}$. When supply unit k traverses any one of the arcs in this set, it has already been ordered by stage n, so it incurs a holding cost increment h_n . (Note that it simultaneously incurs holding cost increments for other stages as well.) It is obvious that the total holding cost incurred by unit k is the sum of the increments described. A holding cost increment h_n that is incurred when unit k traverses an early arc is called an *increment of echelon-n early holding costs*. We define an *increment of echelon-n late holding costs* and an *increment of echelon-n holding costs* in like manner. Like units, cost increments are of infinitesimal size. The *number of echelon-n holding cost increments* that meet some criterion is the number of units that meet the criterion, which is the length of the interval on \mathcal{L}_D that corresponds to the units that satisfy the criterion.

Similarly, a *increment of backorder cost* π is incurred for each time period during which unit k is backordered. Like the holding cost increments discussed above, backorder cost increments are of infinitesimal

size. When we speak of the *number of backorder cost increments* that meet some criterion, we refer to the length of the interval on \mathcal{L}_D that corresponds to the units that satisfy the criterion.

In the next two subsections we assign these cost increments to ordering decisions taken at specific stages in specific time periods. The way the assignment is done depends on whether the increment is an increment of early holding cost, of late holding cost or of pipeline holding cost. Recall that $\langle s,n\rangle$ is the ordering decision taken at stage n in period s. If unit k passes through the lattice point (t,L_n) (i.e., unit k is at stage n at the beginning of period k), then decision k0 determines whether k1 leaves k1 on the horizontal or the diagonal emanating arc. Therefore the direct impact of decision k2 is felt at the lattice point k3.

3.6 Assignment of Pipeline and Early Holding Costs

We first consider the cost increments that we will assign to a decision to *order* a given unit at a given point in time. (The next subsection deals with cost increments that are assigned to a decision to *not* order a given unit.) These are the pipeline and early holding cost increments. In the remainder of Section 3 we will view the system from time T+1, so all quantities are deterministic. We will use s as the time at which a generic order is placed, and not necessarily as the current time.

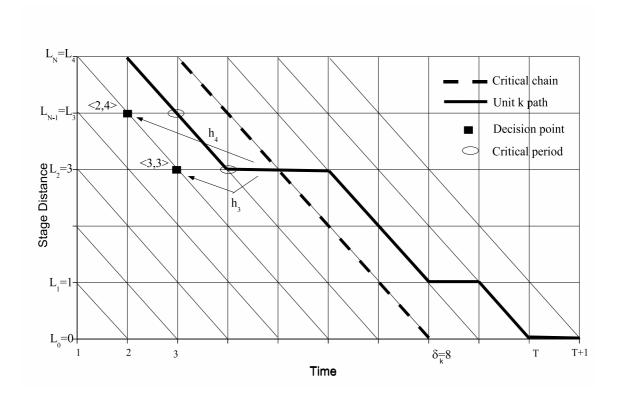


Figure 3.2: Assignment of Early Holding Costs to Orders

The due date $\delta_k = 8$ defines the critical period to be 3 for stage 4, and 4 for stage 3. Stage 4 orders k at time 2, and stage 3 at time 3. The orders at $\langle 2, 4 \rangle$ and $\langle 3, 3 \rangle$ are therefore one period early, and incur h_4 and h_3 in early holding costs, respectively.

Early Holding Costs

If policy P orders supply unit k at stage n at time s, this decision causes unit k to travel diagonally along chain $s+L_n$, from point (s,L_n) to point $(s+l_n,L_{n-1})$ (i.e., unit k will reach stage n at time $s+l_n$). If this chain falls to the left of the critical chain (i.e., if $\delta_k > L_n + s$) then because of this ordering decision, unit k will incur $\delta_k - (L_n + s)$ increments of echelon-n, early holding cost. We adopt the view that the ordering decision at time s at stage s caused all of these cost increments to be incurred. Consequently, we assign the total echelon-s0 early holding cost s1 incurred by unit s2 to the ordering decision at time s3 at stage s3. Note that s3. Note that s4 to the ordering decision at time s5 at stage s5. See Figure 3.2. Note that s5 to the ordering decision 3.2).

We need to formally establish that if unit k is included in the order $\langle s, n \rangle$ then unit k will incur a total of $[\delta_k - (L_n + s)]^+$ increments of echelon-n early holding costs. The proof of the claim is simple. Recall

that $\delta_k \leq T+1$, and that unit k cannot leave the inventory system before time δ_k . The unit first reaches stage n when it reaches the lattice point $(s+l_n,L_{n-1})$, on chain $s+L_n$. From that point onwards, unit k will incur one increment of echelon-n early holding cost whenever it traverses an early arc, i.e., whenever it transitions from one chain to the next, until it reaches chain δ_k . This happens a total of $[\delta_k - (L_n + s)]^+$ times.

In executing the assignment procedure described above we consider every unit k that is in the system at any time during the time horizon [1,T], and the costs that are incurred during time periods [1,T]. The essential point is that, when applied in this manner, the procedure defined above assigns every increment of early holding cost to a unique ordering decision. Note, however, that some of the ordering decisions we assign costs to took place in the past (i.e., in periods t, $t \le 0$) rather than in the future. For example, suppose that $l_1 = 3$, $\delta_k = 5$, and unit k traverses the diagonal arc connecting lattice points (1,1) and (2,0), followed by the horizontal arc connecting points (2,0) and (3,0). Then unit k was ordered by stage 1 at time $t = 2 - l_1 = -1$, and the increment of echelon-1 early holding cost that unit k incurs when it traverses the arc from (2,0) to (3,0) is assigned to the order $\langle -1,1\rangle$, placed at time -1. We define H_{\ll} to be the total early holding cost of all increments assigned to all orders $\langle s,n\rangle$ that occurred in the past (i.e., having $s \le 0$). The algebraic form of H_{\ll} does not concern us, but two facts regarding H_{\ll} are important. First, the costs included in H_{\ll} are uncontrollable, i.e., they are independent of the policy used. In addition, $H_{\ll} \ge 0$. Consequently the value of H_{\ll} does not effect our approximation results.

We define \widehat{H}^P_{sn} to be the number of echelon-n early holding cost increments assigned to order $\langle s,n\rangle$ (placed at time s stage n), aggregated over all supply units k, assuming that policy P is followed. In other words, $h_n \widehat{H}^P_{sn}$ is the total early holding cost assigned to the ordering decision $\langle s,n\rangle$, where $1 \leq s \leq T - L_n$. We now derive a closed-form algebraic expression for \widehat{H}^P_{sn} . As we have just seen, the number of increments of early holding cost that supply unit k contributes to order $\langle s,n\rangle$ is equal to the number of chains k satisfying k and k are k and k are k are k and k are k and k are k are k are k and k are k are k and k are k are k are k are k are k and k are k are k are k are k are k and k are k and k are k a

- 1. $t < \delta_k$, i.e., demand for unit k occurs after time t,
- 2. Demand unit k has been ordered at stage n by the end of time period s, and
- 3. At the beginning of time period s, demand unit k had not been ordered at stage n.

For each fixed value of t we want to find the total number of units that meet these requirements. We will then sum these numbers over $t \geq s + L_n$ to obtain \widehat{H}_{sn}^P . Given t, from the definition of $y_n^P(s)$ we see that $(y_n^P(s) - d_{[s,t]})^+$ is the number of supply units k that meet conditions 1 and 2. (Recall that $d_{[s,t]}$ is

the realized demand over the interval [s,t], and corresponds to the random variable $D_{[s,t]}$.) Similarly, the definition of $x_n^P(s)$ implies that $(x_n^P(s)-d_{[s,t]})^+$ is the number of supply units k that meet condition 1 and fail condition 3. Finally, note that any supply unit that fails condition 3 must necessarily satisfy condition 2. Therefore, the number of supply units that satisfy conditions 1-3 is $(y_n^P(s)-d_{[s,t]})^+-(x_n^P(s)-d_{[s,t]})^+$. Summing over all chains $t \geq L_n + s$, and using the fact that $y_n^P(s) = x_n^P(s) + q_n^P(s) \geq x_n^P(s)$, we obtain

Lemma 3.1 The total early holding cost assigned to order $\langle s, n \rangle$ is $h_n \widehat{H}_{sn}^P$, where

$$\hat{H}_{sn}^{P} = \sum_{t=s+L_{n}}^{T} \left[\left(Y_{n}^{P}(s) - D_{[s,t]} \right)^{+} - \left(X_{n}^{P}(s) - D_{[s,t]} \right)^{+} \right]
= \sum_{t=s+L_{n}}^{T} \left[Q_{n}^{P}(s) - \left(D_{[s,t]} - X_{n}^{P}(s) \right)^{+} \right]^{+}.$$
(1)

Furthermore, $H_{\ll} + \sum_{1 \leq n \leq N} \sum_{1 \leq s \leq T - L_n} h_n \widehat{H}_{sn}^P$ is the sum of all early holding costs incurred by policy P.

The proof of the lemma is in the preceding paragraphs. We note that the expression in (1) is similar to the expression developed for the marginal holding cost in the single-stage model (see [11]).

Pipeline Costs and Excess End-of-Horizon Costs

We now consider the pipeline and excess end-of-horizon costs incurred by policy P. As we have mentioned in Subsection 3.5, the pipeline costs are the sum of the pipeline holding costs and the per-unit ordering costs.

Assume that unit k is ordered by echelon n before period $T-L_n+1$. The definitions of pipeline arcs and holding cost increments in subsection 3.5 indicate that an increment of echelon-n pipeline holding costs is incurred whenever unit k traverses an arc that enters into a stage point (s,j) satisfying $1 \leq j < L_n$, if the arc is either a diagonal arc or a horizontal arc to the right of chain T+1. Each echelon-n increment results in an (echelon) holding cost of h_n . Considering Figure 3.1, and recalling that no unit traverses any diagonal arcs to the right of chain T+1 in the lattice, it is not hard to see that for unit k, an increment of echelon-n pipeline holding costs is incurred exactly L_n times, resulting in a total echelon-n pipeline holding cost of $h_n L_n$. We assign this cost to the ordering decision $\langle u_{kn}, n \rangle$, as Figure 3.3 indicates. (Recall that u_{kn} is the time when unit k was ordered at stage n, and that u_{kn} might be negative.) We also assign the unit shipping cost c_n for stage n to decision $\langle u_{kn}, n \rangle$. Consequently the total pipeline cost assigned to $\langle u_{kn}, n \rangle$ is $c_n + h_n L_n$ for each unit k ordered at stage n time u_{kn} . Aggregated over all units, the total pipeline cost assigned to order $\langle s, n \rangle$ is $\langle c_n + h_n L_n \rangle q_n^p(s)$, if $s \geq 1$.

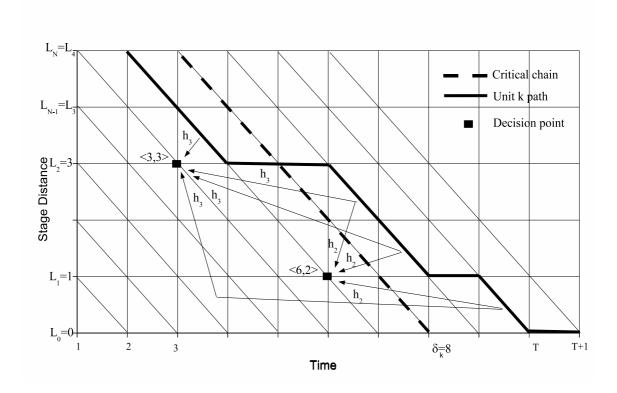


Figure 3.3: Assignment of Pipeline Holding Costs to Orders (order $\langle 2, 4 \rangle$ omitted)

The pipeline cost h_3 is incurred when unit k traverses a diagonal arc with stage distance L_3 or less. Similarly, h_2 is incurred when k traverses a diagonal arc with stage distance L_2 or less.

However, we assumed in Subsection 3.2 that at the end of the time horizon, each excess supply unit will be salvaged by refunding the pipeline costs that the unit incurred. Of the $q_n^P(s)$ units ordered at stage n at time s, $[q_n^P(s) - (d_{[s,T]} - x_n^P(s))^+]^+$ units will receive the refund because they will not be demanded before the end of period T. Consequently³, the sum of the pipeline and excess end-of-horizon costs allocated to $\langle s,n\rangle$ is $(c_n+h_nL_n)\{q_n^P(s)-[q_n^P(s)-(d_{[s,T]}-x_n^P(s))^+]^+\}$.

The sum of the total pipeline and excess end-of-horizon costs that are assigned to all orders $\langle s, n \rangle$ such that $s \leq 0$ is \bar{H}_{\ll} which, like H_{\ll} , is non-negative and policy-independent. We have proven the following lemma.

Lemma 3.2 The sum of the total pipeline and excess end-of-horizon costs assigned to order $\langle s, n \rangle$ with $s \geq 1$ is $(c_n + h_n L_n) \{Q_n^P(s) - [Q_n^P(s) - (D_{[s,T]} - X_n^P(s))^+]^+\}$. The sum of all pipeline and excess

³If we use the general expression for excess end-of-horizon costs described in Subsection 3.2 then this cost expression becomes $(c_n + h_n L_n)q_n^P(s) + \xi_n[q_n^P(s) - (d_{[s,T]} - x_n^P(s))^+]^+$.

end-of-horizon costs incurred by policy P is $\bar{H}_{\ll} + \sum_{1 \leq n \leq N} \sum_{1 \leq s \leq T-L_n} (c_n + h_n L_n) \{Q_n^P(s) - [Q_n^P(s) - (D_{[s,T]} - X_n^P(s))^+]^+\}$.

Summary

We summarize this subsection as follows. We have discussed the early holding, pipeline and excess end-of-horizon costs, which are the cost increments that we assign to a decision to order a unit at a given point in time. For an arbitrary policy P, the sum of these costs is

$$H_{\ll} + \bar{H}_{\ll} + \sum_{1 \le n \le N} \sum_{1 \le t \le T - L_n} H_{tn}^P$$
, where (2)

$$H_{tn}^{P} = h_n \hat{H}_{tn}^{P} + (c_n + h_n L_n) \{ Q_n^{P}(t) - [Q_n^{P}(t) - (D_{[t,T]} - X_n^{P}(t))^+]^+ \}$$

$$= h_n \sum_{s=t+L_n}^T [Q_n^P(t) - (D_{[t,s]} - X_n^P(t))^+]^+ + (c_n + h_n L_n) \{Q_n^P(t) - [Q_n^P(t) - (D_{[t,T]} - X_n^P(t))^+]^+\}.$$

3.7 Assignment of Backorder Costs and Late Holding Costs

Now we will consider cost increments that are caused because policy P does *not* order a given unit at a given stage and time period. These are the backorder costs, late holding costs and end-of-horizon shortage costs.

As in the previous subsection, we can account for the total backorder cost incurred by supply unit k over the horizon by examining the path it follows through the lattice of points. Recall that an increment of backorder cost is the cost π of backordering a unit of inventory for a single time period. Unit k incurs an increment of backorder cost in each period, starting in period δ_k (when unit k is demanded), and ending with either the arrival of unit k at stage 1 or the beginning of period k (when unit k incurs is equal to the number of late arcs (horizontal arcs in the lattice between chain k and chain k incurs is equal to the number of late arcs (horizontal arcs in the lattice between chain k and chain k incurs no backorder cost and no late holding cost, because it arrives at stage 1 before it is demanded.

We now describe the assignment of backorder costs and late holding costs to ordering decisions $\langle s, n \rangle$. Specifically, let $[(s, L_n), (s+1, L_n)]$ be a late arc in the lattice that the path for unit k traverses; hence $s + L_n \geq \delta_k$. Then unit k was at stage n+1 at time s and could have been included in the order $\langle s, n \rangle$ that was placed at stage n at time s, but it was not included. This decision caused unit k to traverse the late arc $[(s, L_n), (s+1, L_n)]$ and, in so doing, to incur the (echelon) late holding cost incre-

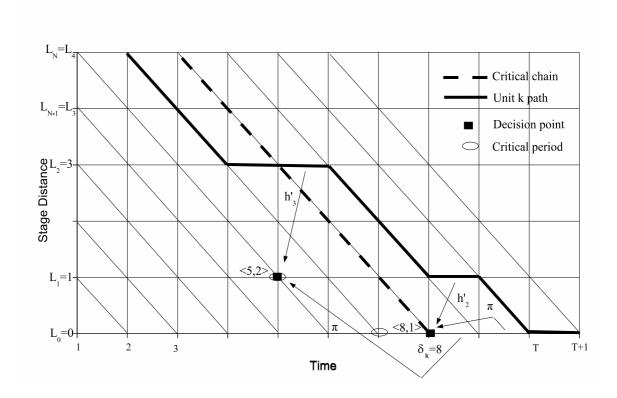


Figure 3.4: Assignment of Late Holding and Backorder Costs to Orders

The critical period for stage 2 is 5. Unit k was available for stage 2 to order in period 5 (i.e., was on-hand at stage 3), but was not included in order $\langle 5, 2 \rangle$. Therefore, $\langle 5, 2 \rangle$ is assigned both π and $h_3' = h_3 + h_4$. $\langle 8, 1 \rangle$ is similar.

ments $h_{n+1}, h_{n+2}, ..., h_N$, whose sum is h'_{n+1} . All of these late holding cost increments are assigned to decision $\langle s, n \rangle$. The decision also shifted unit k from chain $s + L_n$ to chain $s + L_n + 1$, where $\delta_k + 1 \le s + L_n + 1 \le T + 1$. In so doing, it caused one more increment of backorder cost for unit k to be incurred - a cost of π . We assign this backorder cost increment to order $\langle s, n \rangle$ as well. Therefore, the total backorder-plus-late-holding cost assigned to $\langle s, n \rangle$ by unit k is $h'_{n+1} + \pi$. This is illustrated in Figure 3.4. It is not difficult to verify that this procedure assigns each increment of late holding costs and of backorder costs to a unique ordering decision $\langle s, n \rangle$ satisfying $s \le T - L_n$.

We define Π_{\ll} to be the backorder and late holding cost assigned to orders $\langle s, n \rangle$ placed in the past (meaning that $s \leq 0$), aggregated over all units k. As before, Π_{\ll} is policy-independent and non-negative.

Let $(h'_{n+1} + \pi) \widehat{\Pi}^P_{sn}$ be the total backorder-plus-late-holding cost incurred by policy P. As before, we seek a closed-form expression for the number $\widehat{\Pi}^P_{sn}$ of backorder increments assigned to decision $\langle s, n \rangle$. We

have seen that $h'_{n+1} + \pi$ in backorder and late holding costs are assigned to decision $\langle s, n \rangle$ by unit k if and only if the following conditions hold:

- 4. Unit k has reached stage n+1 by the beginning of time period s.
- 5. At the end of time period s, unit k has still not been ordered by stage n.
- 6. Demand for unit k occurs at time δ_k , where $1 \le \delta_k \le s + L_n$.

For each ordering decision $\langle s,n\rangle$, we are interested in the total number of units k that meet these conditions. Since stage n cannot order a unit that has not yet reached stage n+1, the set of units that fail condition 4 is a subset of the set of units that satisfy condition 5. Following the argument that was used for early holding costs and the definition of $y_n^P(s)$, the number of units k that meet conditions 5 and 6 is $(d_{[s,s+L_n]}-y_n^P(s))^+$. Similarly, the number of units k that meet condition 6 and fail condition 4 is $(d_{[s,s+L_n]}-ni_{n+1}^P(s))^+$. Consequently, the number of units that satisfy conditions 4-6 is $(d_{[s,s+L_n]}-y_n^P(s))^+-(d_{[s,s+L_n]}-ni_{n+1}^P(s))^+$. For each of these units, a cost of $h'_{n+1}+\pi$ is assigned to the ordering decision $\langle s,n\rangle$, comprised of late holding and backorder costs. The arguments in the preceding paragraphs prove the following lemma. (The second equality in (3) follows from $ni_{n+1}^P(s) \geq y_n^P(s) = q_n^P(s) + x_n^P(s)$).

Lemma 3.3 The total backorder and late-holding cost assigned to $\langle s, n \rangle$ is $(h'_{n+1} + \pi) \widehat{\Pi}^P_{sn}$, where

$$\widehat{\Pi}_{sn}^{P} = \left[\left(D_{[s,s+L_n]} - Y_n^P(s) \right)^+ - \left(D_{[s,s+L_n]} - NI_{n+1}^P(s) \right)^+ \right]$$

$$= \min \left[\left(D_{[s,s+L_n]} - \left(Q_n^P(s) + X_n^P(s) \right) \right)^+, NI_{n+1}^P(s) - \left(Q_n^P(s) + X_n^P(s) \right) \right].$$
(3)

Furthermore, the total backorder-plus-late-holding cost of policy P is

$$\Pi_{\ll} + \sum_{1 \leq n \leq N} \sum_{1 \leq s \leq T - L_n} (h'_{n+1} + \pi) \, \widehat{\Pi}^P_{sn}$$

End-of-Horizon Shortage Costs We now consider the end-of-horizon shortage costs incurred by policy P. These costs are only incurred if unit k is a demand unit (i.e., if $\delta_k \leq T$) that is at some stage n+1 where $1 \leq n \leq N$, at the beginning of time period T+1. In that case, an end-of-horizon shortage cost of $\sum_{m=1}^{n} (c_m + h_m L_m)$ is incurred. This cost was caused by the fact that unit k was at stage k at time k we assign a cost of k did not include it in the order k placed at stage k. Consequently, for unit k we assign a cost of k did not include it in k to k did not include it in the order k did not include it in the order k placed at stage k.

Clearly, the end-of-horizon shortage costs are all allocated to ordering decisions in a proper manner. Note that all end-of-horizon shortage costs are assigned to orders $\langle s, n \rangle$ such that the stage point (s, L_n) is on chain T (i.e., $s = T - L_n$). Recall that $T \geq L_N + 1$, which implies that all points (t, T - t) on chain T

satisfy $t \ge 1$. We conclude that all end-of-horizon shortage costs are assigned to orders $\langle s, n \rangle$ placed after the beginning of the planning horizon, i.e., at or after time s = 1.

For a given ordering decision $\langle s,n\rangle$, we seek a closed-form algebraic expression for the total end-of-horizon shortage cost incurred by P and assigned to $\langle s,n\rangle$, aggregated over all supply units k. This number is zero if $s \neq T - L_n$. If $s = T - L_n$ then unit k contributes an end-of-horizon shortage cost of $\sum_{m=1}^{n} (c_m + h_m L_m)$ to decision $\langle T - L_n, n\rangle$, if conditions 4-6 above hold (with $s = T - L_n$)⁴. By the argument given above, the number of units k that satisfy conditions 4-6 is $\widehat{\Pi}_{T-L_n,n}^P$.

Let $\mathbb{1}(S) = 1$ if S is a true statement, and $\mathbb{1}(S) = 0$ otherwise. We have proven the following lemma.

Lemma 3.4 The total end-of-horizon shortage cost assigned to order $\langle s, n \rangle$ is

$$1 (s = T - L_n) \left(\sum_{m=1}^n (c_m + h_m L_m) \right) \widehat{\Pi}_{sn}^P$$
. The sum of all end-of-horizon shortage costs incurred by policy P is $\sum_{1 \leq n \leq N} \sum_{1 \leq s \leq T - L_n} \left[1 (s = T - L_n) \left(\sum_{m=1}^n (c_m + h_m L_m) \right) \widehat{\Pi}_{s,n}^P \right]$.

Summary

The total backorder, late holding and end-of-horizon shortage cost incurred by an arbitrary policy P is

$$\Pi_{\ll} + \sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} \Pi_{sn}^P, \text{ where}$$

$$\Pi_{sn}^P = \left[(h'_{n+1} + \pi) + \mathbb{1}(s = T - L_n) \left(\sum_{m=1}^{n-1} (c_m + h_m L_m) \right) \right] \widehat{\Pi}_{sn}^P$$

$$= \left[(h'_{n+1} + \pi) + \mathbb{1}(s = T - L_n) \left(\sum_{m=1}^{n-1} (c_m + h_m L_m) \right) \right] \times$$

$$\left[\left(D_{[s,s+L_n]} - (Q_n^P(s) + X_n^P(s)) \right)^+ - \left(D_{[s,s+L_n]} - NI_{n+1}^P(s) \right)^+ \right].$$
(4)

4 Balancing Policy

In this section, we describe a policy for the periodic-review serial inventory system with stochastic demands. Following [11, 12], we call it a *balancing policy*. We shall show that this policy, when applied to Model 1, has a worst-case performance guarantee of 2, i.e., for each instance of the problem, the expected cost of the policy is at most twice the expected cost of an optimal policy. We next describe the policy and its worst-case

⁴If we use the general expression for end-of-horizon shortage costs described in Subsection 3.2, then in this subsection we replace all instances of $(c_m + h_m L_m)$ with σ_m .

analysis under the assumptions described in Section 2, and then discuss the worst-case performance of the policy under different assumptions on the model.

The balancing policy presented in this paper is based on the cause-effect cost accounting scheme described in Section 3 above and extends the cost-balancing techniques used in [11, 12] for single-stage models. Without capacity constraints, the cause-effect (marginal) cost-accounting scheme in the single-stage models associates with each decision two costs, the marginal holding costs incurred by the units ordered in the current period over the entire horizon, and the marginal backordering penalty cost incurred a lead time ahead. These costs are opposing in that the holding costs are increasing in the size of the order and the backordering penalty cost is decreasing in the size of the order. In each period s, conditioned on the observed information set f_s , the dual-balancing policy balances the expected marginal holding cost of the units ordered in the period against the (conditional) expected marginal backordering cost incurred a lead time into the future (see [11, 12] for details). Next we shall show how to generalize this approach to the multi-stage serial system.

Using the cause-effect cost accounting scheme described in Section 3, the effective expected cost of each policy P can be written as

$$E[C(P)] = E\left[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} H_{sn}^P + \Pi_{sn}^P\right],\tag{5}$$

where again

$$H_{sn}^{P} = h_n \widehat{H}_{sn}^{P} + (c_n + h_n L_n) Q_n^{P}(s) - (c_n + h_n L_n) \left(Q_n^{P}(s) - (D_{[s,T]} - X_n^{P}(s))^+ \right)^+$$

denotes the respective early holding, pipeline and excess end-of-horizon costs, and Π^P_{sn} denotes the backordering penalty, late holding and end-of-horizon shortage costs, all assigned to decision $\langle s,n\rangle$ (see Lemmas 3.1, 3.2, 3.3 and 3.4). Equation (5) above ignores the beginning-of-horizon expected costs (the terms H_{\ll} , \bar{H}_{\ll} and Π_{\ll} in Section 3 above) because they are incurred by any feasible policy. In addition, Equation (5) naturally decomposes into single-stage cost functions, where in each period s, the cost function of stage s is s in s in

This leads to the following natural balancing policy. In each period s, conditioned on the observed information set f_s , consider each of the functions $E[H_{sn}^P(q_n^P(s)) + \Pi_{sn}^P(q_n^P(s))|f_s]$ (for $n=1,\ldots,N$) separately, and apply cost-balancing similar to the single-stage cost-balancing used in [11, 12]. That is, the size of the order placed by the balancing policy at stage n in period s, denoted by $q_n'(s)$, is such that $E[H_{sn}^P(q_n'(s))|f_s] = E[\Pi_{sn}^P(q_n'(s))|f_s]$. However, before we show that the above policy is well-defined we first discuss a more refined view of the ordering decisions of regular policies.

Consider some unit k, and let s be the first period such that $s > \delta_k$ and unit k is at, say, stage n + 1, at the beginning of period s. That is, at the beginning of s, the demand for unit k has already occurred, and unit k is part of the on-hand inventory at stage n+1. Assume that demand unit k can be satisfied before time T+1, i.e., that $s \leq T-L_n$. Because we restrict attention to regular policies, we assume, without loss of generality, that in the future any policy P will not defer ordering unit k at any stage in the serial system. Moreover, it is clear that the chain on which unit k will be sent (specifically, chain $s + L_n$) is to the right of the critical chain δ_k . Thus, by definition, unit k will not incur any early holding costs for stages $n, \ldots, 1$ or be in excess inventory at the end of the horizon (i.e., it will not incur any salvage costs at the end of the horizon). Furthermore, since unit k will not be delayed at any stage, it is clear that it incurs no backordering penalty or late holding costs that will be assigned to decisions taken after period δ_k . Consequently, the only costs that unit k assigns to any of the decisions $\langle s', n' \rangle$ with $s' > \delta_k$ are the pipeline costs $c_m + h_m L_m$, assigned to the orders $\langle s + L_n - L_m, m \rangle$ for which $1 \leq m \leq n$. We use the term *immediate orders* to refer to orders of units that have already occurred (like unit k above). Denote by $\bar{Q}_n^P(s) \geq 0$ the size of the immediate order placed by policy P at stage n in time period s. Let $\tilde{Q}_n^P(s) = Q_n^P(s) - \bar{Q}_n^P(s)$ be the size of the order for units whose demand has not yet occurred, and call it a regular order. Note that at the beginning of period s all the immediate orders in that period are known. We will adopt the view that these orders are not part of the decisions taken by the policy P at the beginning of period s. Following this approach let $\tilde{X}_n^P(s) = X_n^P(s) + \bar{Q}_n^P(s)$ be the respective echelon inventory position of policy P in stage n at the beginning of period s after immediate orders are placed (by all stages) and before regular orders are placed, that is, $X_n^P(s) \le \tilde{X}_n^P(s) \le Y_n^P(s)$.

We claim that, under the new definitions, $\tilde{Q}_n^P(s)$ can be positive only if $\tilde{X}_n^P(s)$ is nonnegative. Note that if $\tilde{X}_n^P(s)$ is strictly negative then we have not been able to place immediate orders for all of the units whose demand has already occurred, because stage n+1 does not have enough on-hand inventory. We conclude that $NI_{n+1}^P(s) = \tilde{X}_n^P(s) < 0$, i.e., after placing immediate orders the on-hand inventory at stage n+1 is 0, and so no regular order can be placed. Because the regular order is also constrained by the availability of inventory at stage n+1, $\tilde{Q}_n^P(s) \leq NI_{n+1}^P(s) - \tilde{X}_n^P(s) = 0$.

We have already seen that units included in immediate orders do not incur any early holding costs. Thus, it is readily verified that, given the definitions of $\tilde{Q}_n^P(s)$ and $\tilde{X}_n^P(s)$, the expression in Lemma 3.1 for early holding costs is still valid for the regular order placed at stage n at the beginning of period s. In particular, the early holding costs assigned to decision $\langle s,n\rangle$ can be expressed as

$$h_n \hat{H}_{sn}^P = h_n \sum_{t=s+L_n}^T \left[\tilde{Q}_n^P(s) - \left(D_{[s,t]} - \tilde{X}_n^P(s) \right)^+ \right]^+.$$
 (6)

Similarly, we claim that the expression in Lemma 3.2 for the pipeline costs with $\tilde{Q}_n^P(s)$ and $\tilde{X}_n^P(s)$ instead of $Q_n^P(s)$ and $X_n^P(s)$, respectively, is still valid in that it accurately accounts for the pipeline costs incurred by regular orders placed in period s at stage n. (The proof of Lemma 3.2 applies, when applied to the regular order only, after the immediate order has been placed.) That is, the pipeline costs incurred by the regular order placed in period s at stage s can be expressed as

$$(c_n + h_n L_n) \left[\tilde{Q}_n^P(s) - \left(\tilde{Q}_n^P(s) - (D_{[s,T]} - \tilde{X}_n^P(s))^+ \right)^+ \right]. \tag{7}$$

Finally, because $\tilde{Q}_n^P(s) + \tilde{X}_n^P(s) = Q_n^P(s) + X_n^P(s) = Y_n^P(s)$, the expression in Lemma 3.3 for the backordering penalty and late holding costs is unchanged (again with $\tilde{Q}_n^P(s)$ and $\tilde{X}_n^P(s)$ instead of $Q_n^P(s)$ and $X_n^P(s)$), i.e.,

$$\Pi_{ns}^{P} = \left(h'_{n+1} + \pi + \mathbb{1}(s = T - L_n) \sum_{m=1}^{n} (c_m + h_m L_m)\right) \times$$

$$\min \left[\left(D_{[s,s+L_n]} - (\tilde{Q}_n^P(s) + \tilde{X}_n^P(s)) \right)^+, NI_{n+1}^P(s) - (\tilde{Q}_n^P(s) + \tilde{X}_n^P(s)) \right].$$
(8)

Let \tilde{H}_{sn}^P be the early holding costs and pipeline costs incurred by the regular order placed in period s at stage n (i.e., the sum of (6) and (7)). Recalling that immediate orders incur no salvage costs, we now conclude that Equation (5) can be equivalently expressed as

$$E[C(P)] = E\left[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} (\tilde{H}_{sn}^P + \Pi_{sn}^P + (c_n + h_n L_n) \bar{Q}_n^P(s))\right]. \tag{9}$$

Consider now the following policy, called again the *balancing policy* B. At the beginning of time period s, conditioned on the observed information set f_s , we first place immediate orders, at each stage $n=1,\ldots,N$, as described above. Next we place regular orders. The regular order at stage n is set to balance the conditional expected early holding and pipeline costs incurred by the units included in the regular order in that period, against the conditional expected backordering penalty and late holding costs assigned to the decision $\langle s,n\rangle$. That is, we set $\tilde{q}_n^B(s)=q_n'(s)$ such that

$$E[\tilde{H}_{sn}^{B}(q_n'(s))|f_s] = E[\Pi_{sn}^{B}(q_n'(s))|f_s]. \tag{10}$$

The random variables \tilde{H}_{sn}^B and Π_{sn}^B are functions of many other quantities. However, at the beginning of period s, all of these quantities have been observed except for future demands (which we cannot control), and the decision variable $\tilde{q}_n^B(s)$ (which is no longer random). Consequently the expectations in (10), which are conditioned on the current information set f_s , are deterministic functions of the size of the regular order $\tilde{q}_n^B(s)$ placed at the beginning of period s. Next we show that the above policy is well-defined, i.e., that

indeed in each period s, the quantity $q'_n(s)$ is well-defined, for each $n=1,\ldots,N$. First, we claim that $E[\tilde{H}^B_{sn}(\tilde{q}^B_n(s))|f_s]$ is an increasing function that starts at 0 when $\tilde{q}^B_n(s)=0$, and diverges to infinity as $\tilde{q}^B_n(s)$ grows.

Lemma 4.1 For each stage n, period $1 \le s \le T - L_n$ and information set $f_s \in \mathcal{F}_s$, the function $E[\tilde{H}_{sn}^B(\tilde{q}_n^B(s))|f_s]$ is increasing in $\tilde{q}_n^B(s)$, is equal to 0 for the value $q_n^B(s) = 0$, and goes to infinity as $\tilde{q}_n^B(s)$ approaches infinity.

The proof of Lemma 4.1 follows immediately from the fact that \tilde{H}^B_{sn} is the sum of (6) and (7) above. Next we claim that the $E[\Pi^B_{sn}(\tilde{q}^B_n(s))|f_s]$ is a convex decreasing function function of $\tilde{q}^B_n(s)$ that is nonnegative for the value $\tilde{q}^B_n(s)=0$ and is equal 0 for the value $\tilde{q}^B_n(s)=ni^B_{n+1}(s)-\tilde{x}^B_n(s)$ (except for the case n=N).

Lemma 4.2 For each stage n = 1, ..., N, period s between l and $T - L_n$, and information set $f_s \in \mathcal{F}_s$, the function $E[\Pi_{sn}^B(\tilde{q}_n^B(s))|f_s]$ is nonnegative, and is convex and decreasing in $\tilde{q}_n^B(s) \geq 0$. As \tilde{q}_n^B approaches infinity, $E[\Pi_{sn}^B(\tilde{q}_n^B(s))|f_s]$ approaches 0. Furthermore, if n < N, then the function $E[\Pi_{sn}^B(\tilde{q}_n^B(s))|f_s]$ is equal to 0 when $q_n^B(s) = ni_{n+1}^B - \tilde{x}_n^B(s)$ (i.e., when all of the available inventory at stage n+1 has been ordered).

Proof: Note again that conditioning on the information set f_s , the value $\tilde{x}_n^B(s)$ is already known deterministically. It follows from equations (3) and (8) that $E[\Pi_{sn}^B(\tilde{q}_n^B(s))|f_s]$ is indeed a function only of $\tilde{q}_n^B(s)$ with the claimed nonnegativity, monotonicity, convexity and convergence properties. Moreover, Equation (8) implies that for n < N, $E[\Pi_{sn}^B(\tilde{q}_n^B(s))|f_s] = 0$ for the value $\tilde{q}_n^B(s) = ni_{n+1}^B - \tilde{x}_n^B(s)$.

The fact that fractional orders are allowed implies that the functions $E[\tilde{H}^B_{sn}(\tilde{q}^B_n(s))|f_s]$ and $E[\Pi^B_{sn}(\tilde{q}^B_n(s))|f_s]$ are continuous. Thus, Lemmas 4.1 and 4.2 imply that $\tilde{q}^B_n(s) = q'_n(s)$ is indeed well-defined for each $n=1,\ldots,N$ and $s=1,\ldots,T-L_n$.

As argued in [11, 12], the balancing policy can be implemented efficiently as long as there exist efficient procedures to evaluate the functions $E[\tilde{H}^B_{sn}(\tilde{q}^B_n(s))|f_s]$ and $E[\Pi^B_{sn}(\tilde{q}^B_n(s))|f_s]$. We note that in most common scenarios there are efficient and straightforward ways to evaluate these functions (for a more detailed discussion see [11, 12]). In fact, the computations that we perform at each stage and in each time period, are computationally very similar to the ones performed in each time period in the single-stage models [11, 12]. Since $E[\tilde{H}^B_{sn}(\tilde{q}^B_n(s))|f_s]$ and $E[\Pi^B_{sn}(\tilde{q}^B_n(s))|f_s]$ are functions only of $\tilde{q}^B_n(s)$ the balancing policy can be implemented in an on-line manner, i.e., the decision made in the current period is independent of any future decision. It is readily verified that the computational effort is modest, in contrast to the standard dynamic programming approach, which requires a state space that specifies both the inventory levels and all possible

realizations of the information set f_s . In the next subsection we shall describe the worst-case analysis of the balancing policy described above.

5 Analysis

In this section, we shall show that the balancing policy described in Section 4 above has a worst-case performance guarantee. The worst-case analysis is based on an amortization of the costs incurred by the balancing policy (which is again denoted by the superscript B) against the costs incurred by an optimal policy (denoted by the superscript OPT). We first analyze the worst-case performance guarantee under the assumptions described in Section 2, and then discuss several important variants and extensions.

First, we express the expected cost of the balancing policy. For each decision $\langle s, n \rangle$ where $n = 1, \dots, N$ and $s = 1, \dots, T - L_n$, let Z_{sn} be the following random variable:

$$Z_{sn} := E[\tilde{H}_{sn}^B | F_s] = E[\Pi_{sn}^B | F_s].$$

Recall that F_s is the random information set that will be realized at the beginning of period s (or equivalently at the end of period s-1). Thus, Z_{sn} is a random variable that is realized at the beginning of period s as a function of the observed information set f_s . The second equality follows from the construction of the balancing policy (see Section 4). Using equation (9) and standard arguments of conditional expectation, we establish the following lemma.

Lemma 5.1 The expected cost of the balancing policy is equal to twice the expected sum of the Z_{sn} variables plus the expected pipeline costs incurred by immediate orders. That is,

$$E[C(B)] = 2E\left[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} Z_{sn}\right] + E\left[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} (c_n + h_n L_n) \bar{Q}_n^B(s)\right].$$

Proof: Using standard arguments of conditional expectations and the definition of the Z_{sn} variables we get

$$E[\tilde{H}_{sn}^B + \Pi_{sn}^B] = E[E[\tilde{H}_{sn}^B + \Pi_{sn}^B | F_s]] = 2E[Z_{sn}].$$

Using this and Equation (9) applied to the balancing policy, we have

$$E[C(B)] = \sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} E[\tilde{H}_{sn}^B + \Pi_{sn}^B] + E[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} (c_n + h_n L_n) \bar{Q}_n^B(s)] =$$

$$2 \sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} E[Z_{sn}] + E[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} (c_n + h_n L_n) \bar{Q}_n^B(s)].$$
(11)

This concludes the proof of the lemma.

Next we show that the expected cost C(OPT) incurred by OPT is at least

$$\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} E[Z_{sn}] + E[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} (c_n + h_n L_n) \bar{Q}_n^B(s)].$$
(12)

This and Lemma 5.1 above imply that the balancing policy has a worst-case performance guarantee of 2.

For the optimal policy OPT we do not use the cause-effect cost accounting scheme of Section 3 to assign costs to ordering decisions. The costs incurred by OPT are not assigned to decisions at all. Recall that a cost increment is an *echelon-n* cost if it is linear in the coefficient h_n . We describe and analyze the costs incurred by OPT as echelon-n early holding costs, echelon-n late holding costs, echelon-n excess end of horizon costs, echelon-n end of horizon shortage costs, and backorder costs.

Partition all of the ordering decisions $\langle s, n \rangle$ (for n = 1, ..., N and $s = 1, ..., T - L_n$) into two sets according to the echelon inventory positions of OPT and the balancing policy in each period at each stage. The set \mathcal{T}_{Π} contains all ordering decisions $\langle s, n \rangle$ such that the echelon inventory position of the balancing policy at stage n after ordering in period s is (strictly) higher than the respective echelon inventory position of OPT. That is, at the end of period s the balancing policy has more units than OPT at or in transit to stages n, ..., 1. Similarly, the complement set \mathcal{T}_H contains the ordering decisions $\langle s, n \rangle$ such that the echelon inventory position of the balancing policy at stage n after ordering in period s is not higher than the respective echelon inventory position of OPT. Specifically,

$$\mathcal{T}_{\Pi} := \{ \langle s, n \rangle : Y_n^B(s) > Y_n^{OPT}(s) \}$$

and

$$\mathcal{T}_H := \{ \langle s, n \rangle : Y_n^B(s) \le Y_n^{OPT}(s) \}.$$

Note that T_H and T_{Π} are random sets that are not fully realized until period $T - L_1$. Thus, they induce a random partition of the ordering decisions $\langle s, n \rangle$.

As we have already mentioned, the analysis is based on an amortization of the costs incurred by the balancing policy against costs incurred by OPT. Specifically, we shall show that the costs incurred by OPT are always higher than a portion of the costs incurred by the balancing policy. The expectation of this portion is at least half of the total costs incurred by the balancing policy. Central to the strategy just described is the following inequality that, as we shall show below, holds with probability 1:

$$C(OPT) \ge \sum_{\langle s,n\rangle \in \mathcal{T}_H} \tilde{H}_{sn}^B + \sum_{\langle s,n\rangle \in \mathcal{T}_\Pi} \Pi_{sn}^B + \sum_{\langle s,n\rangle} (c_n + h_n L_n) \bar{Q}_n^B(s). \tag{13}$$

Taking expectations and again using properties of conditional expectations, this implies that

$$E[\mathcal{C}(OPT)] \geq \sum_{\langle s,n\rangle} E\left[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B} + (c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)\right] =$$

$$\sum_{\langle s,n\rangle} E\left[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B} + (c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)\right] =$$

$$\sum_{\langle s,n\rangle} E\left[E[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B}|F_{s}]\right] + \sum_{\langle s,n\rangle} E[(c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)] =$$

$$\sum_{\langle s,n\rangle} E\left[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})E[\tilde{H}_{sn}^{B}|F_{s}] + \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})E[\Pi_{sn}^{B}|F_{s}]\right] + \sum_{\langle s,n\rangle} E[(c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)] =$$

$$\sum_{\langle s,n\rangle} E[Z_{sn}] + \sum_{\langle s,n\rangle} E[(c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)].$$

The third equality follows from the fact that conditioning on the information set at the beginning of period s, we already know deterministically the echelon inventory positions of OPT and the balancing policy (i.e., we know $y_n^{OPT}(s)$ and $y_n^B(s)$). Thus, it is already known whether $\langle s, n \rangle \in \mathcal{T}_H$ or $\langle s, n \rangle \in \mathcal{T}_\Pi$. Lemma 5.1 and (14) imply that indeed the balancing policy has a worst-case performance guarantee of 2.

The proof of the key inequality, (13), is based on a unit by unit analysis. We consider each unit k and compare the respective costs it incurs under the balancing policy and OPT. Consider any information set $f_{T+1} \in \mathcal{F}_{T+1}$ (i.e., a complete evolution of the system until the end of the horizon). Focus on on some unit k and let \mathcal{I}^k be the set of all stages at which the balancing policy ordered unit k. In particular, $\mathcal{I}^k = [m_1, m_2)$ for some $1 \le m_1 \le m_2 \le N+1$. Note that m_1 is the realization of M_1 which is a-priori a random variable. Similarly, the set \mathcal{I}^k is a-priori random. However, m_2 denotes the first stage to which unit k arrived at or after the beginning of period 1. This corresponds either to stage N+1, or to a stage with a lower index if unit k was ordered by stage N prior to time period 1. Thus, m_2 was already known deterministically at the beginning of period 1.

In the remainder of this section we use $u_{kn}=u_{kn}^B$ to indicate the period in which the balancing policy ordered unit k at stage n. In particular, for each $n\in\mathcal{I}^k$ we have $u_{kn}\geq 1$. Similarly, $u_{km_2}\leq 0$ denotes the period in which unit k was ordered at stage m_2 (defined above), that is, unit k reached stage m_2 at the beginning of time period $u_{km_2}+l_{m_2}\geq 1$. For completeness, if $m_2=N+1$, then $u_{k,N+1}=0$ and $l_{N+1}=1$.

Partition the stages in \mathcal{I}^k into three (disjoint) sets. Let $\mathcal{I}^k_E \subseteq \mathcal{I}^k$ be all of the stages in which unit k was included in an immediate order placed by the balancing policy. Since immediate orders occur only when demand for unit k has already occurred and can be satisfied before the end of the horizon, it is readily

verified that if $\mathcal{I}_E^k \neq \emptyset$, then $m_1 = 1$ and $\mathcal{I}_E^k = [1, m]$ for some $1 \leq m < m_2$. Let $\mathcal{I}_H^k \subseteq \mathcal{I}^k \setminus \mathcal{I}_E^k$ be the set of all stages n such that unit k was included in a regular order placed by the balancing policy at stage n in period u_{kn} , where the corresponding echelon inventory position of the balancing policy was not higher than the echelon inventory position of OPT. That is, $\mathcal{I}_H^k = \{n \in \mathcal{I}^k \setminus \mathcal{I}_E^k : \langle u_{kn}, n \rangle \in \mathcal{I}_H \}$. (Note that being sets, \mathcal{I}_H , \mathcal{I}_Π , \mathcal{I}^k , \mathcal{I}_H^k , \mathcal{I}_E^k are written in upper case, even after they have been realized and have become deterministic.) Finally, let $\mathcal{I}_\Pi^k = \mathcal{I}^k \setminus (\mathcal{I}_E^k \cup \mathcal{I}_H^k)$ be the set of all stages such that unit k has been included in a regular order placed by the balancing policy at stage n in period u_{kn} , where the corresponding echelon inventory position of the balancing policy was (strictly) higher than the respective echelon inventory position of OPT, i.e. $\mathcal{I}_\Pi^k = \{n \in \mathcal{I}^k \setminus (\mathcal{I}_E^k \cup \mathcal{I}_H^k)\}$. Observe that the set $\{\langle u_{kn}, n \rangle : n \in \mathcal{I}_\Pi^k\}$ is contained in \mathcal{I}_Π .

Next we state and prove several lemmas, from which the proof of inequality (13) above follows. In all of the subsequent lemmas we focus on a specific unit k, and a specific information set $f_{T+1} \in \mathcal{F}_{T+1}$ that specifies the complete evolution of the system. In particular we know δ_k , i.e., if and when unit k was demanded. (Recall that if unit k was not demanded then $\delta_k = T + 1$, and k is an excess supply unit rather than a demand unit). We will prove the following lemmas for each unit k and each information set k0, establishing that inequality (13) holds with probability 1.

The first of these lemmas amortizes the pipeline costs incurred due to immediate orders placed by the balancing policy, i.e., $\sum_{\langle s,n\rangle}(c_n+h_nL_n)\bar{Q}_n^B(s)$. Specifically, we focus on the pipeline costs unit k has assigned under the balancing policy to stages $n\in\mathcal{I}_E^k$. Under the balancing policy, we say that a specific cost increment is assigned by unit k to stage n if unit k assigns the cost increment to any one of the ordering decisions $\langle s,n\rangle,\ 1\leq s\leq T-L_n$ placed at stage n. These costs can be amortized against the echelon-n pipeline costs and end of horizon shortage costs incurred by unit k under OPT, for stages $n\in\mathcal{I}_E^k$.

Lemma 5.2 Consider a stage n at which unit k was included in an immediate order placed by the balancing policy, i.e., a stage $n \in \mathcal{I}_E^k$. Under the balancing policy, all of the costs that unit k assigned to n are pipeline costs, in the amount $c_n + h_n L_n$. Moreover, under OPT, unit k incurred echelon-n pipeline and end of horizon shortage costs, which are greater than or equal to the pipeline costs assigned to n under the balancing policy. That is, for each stage $n \in \mathcal{I}_E^k$, unit k incurred under OPT echelon-n pipeline and end of horizon shortage costs of at least $c_n + h_n L_n$.

Proof: We have already shown in Section 3 that all units k included in an immediate order at a stage $n \in I_E^k$ are demand units (i.e., $\delta_k \leq T$), and that to stage n they assign only the echelon pipeline costs $c_n + h_n L_n$. Thus, if $\mathcal{I}_E^k \neq \emptyset$, then there are two possible cases. Either unit k was satisfied by OPT before

the end of the horizon, i.e., OPT ordered unit k at stages $n=N,\ldots,1$ and incurred echelon-n pipeline costs $c_n+h_nL_n$ for each n, or OPT has not satisfied demand unit k at the end of the horizon. In the latter case, unit k is at some stage n'>1 at the end of the horizon. In particular, it has incurred echelon-n pipeline costs $c_n+h_nL_n$ for each stage $n\in[n',N]$. In addition, unit k has incurred echelon-n end of horizon shortage costs $c_n+h_nL_n$ for each $n\in[1,n')$. Applying this argument to stages $n\in I_E^k$ (and only to those stages), the lemma follows.

In the next lemma we show how to amortize the echelon pipeline costs, early holding costs and excess end-of-horizon costs incurred when the balancing policy places a regular order for unit k, in periods where the echelon inventory position of the balancing policy is no higher than that of OPT. That is, we amortize the pipeline, early holding and excess end-of-horizon costs incurred by unit k under the balancing policy that are assigned to ordering decisions in the set $\mathcal{T}_H^k = \{\langle u_{kn}, n \rangle : n \in \mathcal{I}_H^k \}$ (this set is contained in \mathcal{T}_H). These costs are amortized against the echelon-n pipeline, early holding and excess end-of-horizon costs incurred by unit k under OPT, for stages $n \in \mathcal{I}_H^k$. More generally, this lemma shows how to amortize the costs $\sum_{\langle s,n \rangle \in \mathcal{T}_H} \tilde{H}_{sn}^B$.

Lemma 5.3 Consider the pipeline, early holding and excess end-of-horizon costs incurred by unit k under the balancing policy that are assigned to ordering decisions in the set $\mathcal{T}_H^k = \{\langle u_{kn}, n \rangle : n \in \mathcal{I}_H^k \}$. To each such ordering decision $\langle u_{kn}, n \rangle$, unit k assigns pipeline, early holding and excess end-of-horizon costs that are equal $h_n(\delta_k - L_n - u_{kn})^+ + \mathbb{1}(\delta_k \leq T)(c_n + h_nL_n)$ (where again u_{kn} is the unique period in which the balancing policy has ordered unit k at stage n). Moreover, these costs are no greater than the corresponding costs unit k has incurred under OPT. In particular, for each $n \in \mathcal{I}_H^k$, unit k incurred under OPT echelon-n pipeline, early holding and excess end-of-horizon costs of at least $h_n(\delta_k - L_n - u_{kn})^+ + \mathbb{1}(\delta_k \leq T)(c_n + h_nL_n)$.

Proof: Consider an ordering decision $\langle u_{kn}, n \rangle$ with $n \in \mathcal{I}_H^k$. Recall the pipeline, early holding and excess end-of-horizon cost definitions in Section 3. Specifically, the pipeline, early holding and excess end-of-horizon costs assigned to the ordering decision $\langle u_{kn}, n \rangle$ are $h_n(\delta_k - L_n - u_{kn})^+ + \mathbbm{1}(\delta_k \leq T)(c_n + h_n L_n)$ (the pipeline costs are salvaged if unit k is an excess supply unit, i.e., if $\delta_k = T + 1$). Our assumptions imply that the ordering decision $\langle u_{kn}, n \rangle$ belongs to the set \mathcal{T}_H , i.e., that $y_n^B(u_{kn}) \leq y_n^{OPT}(u_{kn})$. However, this implies that unit k was ordered by OPT at stage n at the same time period in which it was ordered by the balancing policy at that stage, or earlier. That is, $u_{kn}^{OPT} \leq u_{kn}$ (where u_{kn}^{OPT} is the time period in which unit k was ordered by OPT at stage n). The echelon-n pipeline, early holding and excess end-of-horizon costs incurred by unit k under OPT are equal $h_n(\delta_k - L_n - u_{kn}^{OPT})^+ + \mathbbm{1}(\delta_k \leq T)(c_n + h_n L_n)$ (see again

Section 3). As we have already observed, $u_{kn}^{OPT} \leq u_{kn}$, from which the proof of lemma follows.

Finally, the next lemmas show how to amortize the late holding and backordering penalty costs incurred by unit k under the balancing policy that are assigned to ordering decisions in \mathcal{T}_{Π} . In Section 3, we have already seen that late holding and backordering penalty costs are assigned to decisions not to order a unit at a given stage, in a time period when the unit was on hand at the parent stage. Recall that unit k was ordered by the balancing policy at stages $n \in \mathcal{I}^k$, where $\mathcal{I}^k = [m_1, m_2)$. Moreover, the balancing policy placed these orders for unit k in the time periods $\{u_{kn}: n \in \mathcal{I}^k\}$. Recall that $m_2 \leq N+1$ is the first stage to which unit k arrives (at or after time period 1) due to decisions made prior to time period 1 (i.e., $u_{km_2} \leq 0$). It is readily verified that under the balancing policy, all of the late holding and backordering penalty costs incurred by unit k, are assigned to decisions $\langle s, n \rangle$ made in time periods s that fall in the intervals $\{[u_{k,n+1}+l_{n+1},u_{kn}): n+1\in\mathcal{I}^k\cup\{N+1\}\}$ where, for completeness, we define $u_{k,N+1}=0$, $l_{N+1}=1$ and $u_{k,m_1-1}=T+1$. Specifically, over the interval $[u_{k,n+1}+l_{n+1},u_{kn})$, unit k was on hand at stage n+1 and was not ordered by stage n. In particular, we amortize the part of the late holding and backordering penalty costs that were incurred by unit k under the balancing policy, which were assigned to ordering decisions in the set

$$\mathcal{T}_{\Pi}^{k} = \left\{ \langle s, n \rangle \in \mathcal{T}_{\Pi} : n+1 \in \mathcal{I}^{k} \cup \{N+1\}, \ s \in [u_{k,n+1} + l_{n+1}, u_{k,n}) \right\}.$$

More generally, the next lemmas show how to amortize the costs $\sum_{\langle s,n\rangle\in\mathcal{T}_\Pi}\Pi^B_{sn}$. In the next lemma we amortize the backordering and end of horizon shortage costs incurred by unit k under the balancing policy that are assigned to ordering decisions in \mathcal{T}^k_Π above.

Lemma 5.4 Consider the part of the backordering costs incurred by unit k under the balancing policy which is assigned to ordering decisions in the set T_{Π}^k defined above. These costs are no higher than the total backordering costs incurred by unit k in OPT. Similarly, the end of horizon shortage costs incurred by unit k under the balancing policy and assigned to decisions in T_{Π}^k , are no higher than the echelon-n end of horizon shortage costs incurred by k under OPT, summed over $n \notin \mathcal{I}^k$.

Proof : First observe that unless unit k is a demand unit it incurs no backordering or end of horizon shortage costs. In other words, we must have $\delta_k \leq T$. Consider any period $t' \geq \delta_k$ in which unit k has been backordered under the balancing policy. Further assume that the respective backordering penalty cost increment π is assigned to some ordering decision $\langle s, n \rangle \in \mathcal{T}_{\Pi}^k$. From the cause-effect cost accounting scheme we know that $s = t' - L_n$, and that under the balancing policy, throughout time period s, unit k was on hand at stage n + 1, i.e., at stage-distance L_n . By the assumption that $\langle s, n \rangle \in \mathcal{T}_{\Pi}^k$ we conclude

that $y_n^B(s) > y_n^{OPT}(s)$. This implies that at the end of period $s = t' - L_n$, under OPT, unit k was at stage-distance L_n or higher. However, this implies that in period t' unit k was backordered under OPT and incurred a backordering penalty π . It is now readily verified that the overall backordering cost incurred by unit k under OPT is greater than or equal to the backordering penalty costs it has incurred under the balancing policy that are assigned to orders $\langle s, n \rangle \in \mathcal{T}_{\Pi}^k$.

Suppose that the balancing policy assigned end of horizon shortage costs to an order $\langle s,n\rangle\in\mathcal{T}_\Pi^k$; hence $s=T-L_n$. (Note that for each n there is at most one such order). As we explained in Subsection 3.7, the costs assigned to $\langle s,n\rangle$ are $\sum_{m=1}^n (c_m+h_mL_m)$. We also know that unit k was at stage n+1 at time $T-L_n$ (on chain T), and k was not included in the order $\langle T-L_n,n\rangle$. In other words, under the balancing policy unit k first reached chain T+1 at stage n+1, and it stayed at n+1 from period $T-L_n$ to period T+1. Consequently $\{m:1\leq m\leq n\}\cap\mathcal{I}^k=\emptyset$ (i.e., the balancing policy did not order unit k at these stages), and $n+1\in\mathcal{I}^k\cup\{N+1\}$. Because $\langle T-L_n,n\rangle\in\mathcal{T}_\Pi^k$, under OPT unit k was at stage-distance L_n or higher at the end of period $T-L_n+1$, which implies that at period T+1 it was at some stage $n'\geq n+1$. Therefore, the total end of horizon shortage cost incurred by the balancing policy for unit k is $\sum_{m=1}^n (c_m+h_mL_m)=\sum_{m\notin\mathcal{I}^k} (c_m+h_mL_m)$, and for each stage $m\notin\mathcal{I}^k$, OPT incurred an echelon-m end of horizon shortage cost of $c_m+h_mL_m$.

The fact that Lemma 5.4 uses only echelon-n end of horizon shortage costs incurred by k under OPT for stages that are not in \mathcal{I}^k , is crucial. This fact insures that none of OPT's end of horizon costs that are used in Lemma 5.4, are also used in Lemmas 5.2 and 5.7. The next lemma bounds the number of increments of late holding cost incurred by unit k that are assigned to ordering decisions at a given stage, under the balancing policy. Recall that in the lattice, when unit k traversed the late arc $[(s, L_n), (s, L_n + 1)]$ at stage n+1, it generated a late holding cost increment of h_m for each m, $n < m \le N$, all of which were assigned to the order (s, n) (see Figure 3.4).

Lemma 5.5 Consider some stage n. The late holding cost that unit k assigns to n is at most $h'_{n+1}(L_n+1)$. In other words, the overall late holding cost incurred by unit k under the balancing policy and assigned to any of the ordering decisions made at stage n is at most $h'_{n+1}(L_n+1) = (\sum_{m \geq n+1} h_m)(L_n+1)$. Moreover, for each stage m > 1, the total echelon-m late holding cost incurred by unit k under the balancing policy is at most $h_m(L_{m-1}+1)$ (i.e., there are at most $L_{m-1}+1$ echelon-m late holding cost increments).

Proof: If $\delta_k = T + 1$ then unit k incurs no late holding costs and the lemma holds, so we assume that $\delta_k \leq T$. Because the balancing policy is a regular policy, starting at time period $\delta_k + 1$, unit k incurred no late holding costs. (Either end of horizon assumptions were in effect because unit k had already reached

chain T+1, meaning that at time δ_k+1 it was already impossible for unit k to reach stage 1 before the end of the time horizon, or from period δ_k+1 onwards unit k was included in immediate orders until it was satisfied).

Let $n'\geq 2$ be the minimum index stage which unit k had reached by the end of period δ_k . In other words, at the beginning of period δ_k+1 unit k was at stage-distance j (on chain δ_k+1+j), such that $L_{n'-2}< j\leq L_{n'-1}$. Then unit k incurred no echelon-m late holding costs for any $m\leq n'-1$, and assigned no late holding costs to orders placed at any stage $n, n\leq n'-2$. Furthermore, because we have a regular policy, the last time when unit k was backordered was at the end of period $\min(T,\delta_k+j)$. Thus, the number of backorder cost increments incurred by k, which is also the total number of late arcs traversed by k, is $\min(T,\delta_k+j)-\delta_k+1\leq j+1\leq L_{n'-1}+1$. (Recall that for each time period in which unit k incurs late holding costs, there is a uniquely corresponding time period in which unit k is backordered, and that the correspondence is based on chains (see Figure 3.4).) Hence $L_{n'-1}+1$ is an upper bound on the number of echelon-m late holding cost increments incurred by unit k for any m, $m\geq n'$. In addition, to any stage $n\geq n'-1$, a total late holding cost of at most $h'_{n+1}(L_{n'-1}+1)\leq h'_{n+1}(L_n+1)$ was assigned.

The above lemma has an immediate corollary.

Corollary 5.6 There can be at most one stage n < N, for which the number of late holding increments that unit k incurred under the balancing policy, that are assigned to ordering decisions in the set $\{\langle s, n \rangle : s \in [u_{k,n+1} + l_{n+1}, u_{kn})\}$, is equal $L_n + 1$.

Proof: Focus on some stage n < N+1. From the proof of Lemma 5.5 above, it is clear that the upper bound L_n+1 is achieved only in one scenario – if unit k was on hand at stage n+1 from the beginning of time period $\delta_k - L_n$ to the beginning of time period $\delta_k + 1$. Clearly, this can happen for at most one stage n.

Lemma 5.7 Consider the part of the late holding costs incurred by unit k under the balancing policy, which is assigned to the set of ordering decisions in \mathcal{T}_{Π}^k defined above. These costs are no higher than the echelon-n pipeline and end of horizon shortage costs incurred by unit k under OPT for stages $n \in \mathcal{I}_{\Pi}^k$, plus the sum of all late holding costs incurred by k under OPT.

Proof: If unit k did not incur any late holding costs under the balancing policy then the lemma holds. Assume that unit k did incur late holding costs and, consequently, that unit k is a demand unit (i.e., $\delta_k \leq T$). Let \mathcal{J} be the set of stages (possibly empty) at which unit k was ordered before the time horizon began, i.e., if $\mathcal{I}^k = [m_1, m_2)$ then $\mathcal{J} = \{m: m_2 \leq m \leq N\}$. Consider the total late holding cost assigned by unit k to the ordering decisions in \mathcal{T}_Π^k under the balancing policy. We divide this cost into three quantities – the total echelon-n late holding cost summed over all $n \in \mathcal{I}_E^k$, the total echelon-n cost for $n \in \mathcal{I}_\Pi^k$, and the total echelon-n cost for $n \in \mathcal{I}_H^k \cup \mathcal{J}$. With respect to the first quantity, observe that $\mathcal{I}_E^k = [1, n']$ for some n' (we allow n' = 0, meaning that $\mathcal{I}_E^k = \emptyset$), and that unit k was included in immediate orders placed at each stage $n \in \mathcal{I}_E^k$. Hence under the balancing policy, for each $n \in [1, n']$, unit k is not delayed at stage n and does not incur any echelon-n late holding costs. Consequently, the first quantity is zero.

As for the second quantity, recall that there are no echelon-1 late holding costs, because we will never backorder inventory at stage 1. For any given $n \geq 2$, Lemma 5.5 above implies that the total echelon-n late holding cost incurred by unit k under the balancing policy is at most $h_n(L_{n-1}+1) \leq h_n L_n$ (recall that $l_n \geq 1$ and that $L_n = L_{n-1} + l_n$). This is bounded by the echelon-n pipeline and end of horizon shortage costs incurred by unit k under OPT for stage n. To be specific, by the arguments used in the proof of Lemma 5.2, we know that under OPT, either unit k incurred an echelon-n pipeline cost of $h_n L_n$, or it incurred an echelon-n end of horizon shortage cost of $h_n L_n$. We apply this logic to the stages $n \in \mathcal{I}_{\Pi}^k$, proving that the second of the three quantities described above, is bounded by the echelon-n pipeline and end of horizon shortage costs for stages $n \in \mathcal{I}_{\Pi}^k$, incurred for unit k under OPT. Observe that these costs are distinct from the echelon-n end of horizon shortage costs incurred under OPT for stages $n \in \mathcal{I}_E^k$ and $n \notin \mathcal{I}^k$, which were discussed in Lemmas 5.2 and 5.4, respectively.

We will complete the proof by proving that the third quantity described above is bounded by the late holding costs incurred by unit k under OPT. The third quantity is the total echelon-m late holding cost assigned by unit k to ordering decisions in \mathcal{T}_{Π}^k under the balancing policy, summed over all $m \in \mathcal{I}_H^k \cup \mathcal{J}$. Consider a given increment of echelon-m late holding costs incurred by unit k under the balancing policy at some stage $m \in \mathcal{I}_H^k \cup \mathcal{J}$, which is assigned under the cause-effect cost accounting scheme to some ordering decision $\langle s,n\rangle \in \mathcal{T}_{\Pi}^k$ (see Figure 5.5). Since the echelon-m late holding cost increment is assigned to $\langle s,n\rangle$, it is clear that n < m and that under the balancing policy unit k was on hand at stage n+1 (where $n+1 \leq m$) at both the beginning and the end of period s. (Recall that in the lattice, stage n+1 corresponds to row L_n .) In other words, the echelon-m late holding cost increment was incurred when the balancing policy sent unit k from chain $t' = s + L_n$ to chain t'+1 (see Section 3 above), at stage n+1. (Because it is a late holding cost increment, $\delta_k \leq t' = s + L_n$. Recall that $\delta_k - L_n \leq s$ is the critical period for stage n induced by unit k, so k is backordered in period t'.) We will amortize this echelon-t late holding cost increment using the echelon-t late holding cost increment that was incurred when, under t' to chain t' to

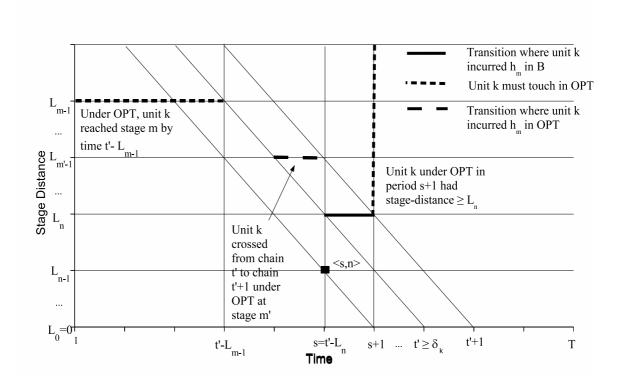


Figure 5.5: Proof of Lemma 5.7, the Third Quantity

most one time, this amortization does not use any increment of echelon-m late holding cost incurred by unit k under OPT more than once. To complete the proof we need to show that the echelon-m late holding cost increment in question was actually incurred under OPT, i.e., that under OPT unit k moved from chain t' to chain t'+1 at some stage m' satisfying $1 < m' \le m$, after the time horizon began. (Recall that k moved from chain t' to chain t'+1 at m' if k was held at m' throughout time period $t'-L_{m'-1}$.)

Because $\langle s,n\rangle\in \mathcal{T}_\Pi^k$, and because under the balancing policy unit k was on hand at stage n+1 at both the beginning and the end of period s, we know that under OPT at time s+1, unit k was either at stage n+1 or had not yet reached stage n+1 (i.e., k was at stage-distance $\geq L_n$; see Figure 5.5). Either way, under OPT unit k crossed from chain $t'=s+L_n$ to chain t'+1 at some stage $m'\geq n+1$. In other words, under OPT there exists some maximum index $m'\geq n+1$ such that unit k was not ordered by OPT at stage m'-1 in period $t'-L_{m'-1}$, although it was on hand at stage m' at that time. This implies that m' does exist. It now suffices to show that $m'\leq m$ and $t'-L_{m'-1}\geq 1$ (i.e., that under OPT unit k crossed from chain t' to chain t'+1 at or below stage m after the time horizon began).

Observe that at time $t' - L_n$, under the balancing policy unit k was at stage n + 1, on chain t'. Hence, for each stage n' > n, unit k was ordered at stage n' at or before time $t' - L_{n'}$. We first show that under

OPT unit k was ordered at stage m by period $t'-L_m$, which implies that $m' \leq m$. It is sufficient to show that at stage m unit k was ordered by OPT at the same period it was ordered by the balancing policy or even earlier. The claim is trivially true if $m \in \mathcal{J}$ since unit k was ordered at these stages prior to period 1, and these decisions are common to all feasible policies. If $m \in \mathcal{I}_H^k$, then the claim follows because at period u_{km} when unit k was ordered at stage m by the balancing policy, OPT had higher echelon-m inventory position. We conclude that indeed $m' \leq m$. Recall that $\mathcal{I}^k = [m_1, m_2)$ and m_2 is defined to be the first period at which unit k arrives after the beginning of the horizon. We finish the proof by showing that $m' \leq m_2$, which implies that $t' - L_{m'-1} \geq 1$. However, this again follows from the fact that at all stages $n' \geq m_2$ unit k was ordered prior to period 1 in the same period by all feasible policies, and that the balancing policy 'ordered' unit k at these stages by period $t' - L_{n'}$.

In Lemmas 5.2, 5.3, 5.4 and 5.7, we have amortized the following costs incurred by unit k under the balancing policy: pipeline costs assigned to immediate orders (i.e., pipeline costs unit k assigned to stages $n \in \mathcal{I}_E^k$; Lemma 5.2), pipeline, early holding and excess end of horizon costs assigned to orders in \mathcal{I}_H (Lemma 5.3), and backordering, end of horizon shortage and late holding costs assigned to orders in \mathcal{I}_{Π} (Lemmas 5.4 and 5.7). Summing over all units k, we have shown how to amortize the costs

$$\sum_{\langle s,n\rangle\in\mathcal{T}_H}\tilde{H}^B_{sn} + \sum_{\langle s,n\rangle\in\mathcal{T}_\Pi}\Pi^B_{sn} + \sum_{\langle s,n\rangle}(c_n+h_nL_n)\bar{Q}^B_n(s).$$

Observe that the amortization uses the following costs incurred by unit k under OPT: echelon-n pipeline and end of horizon shortage costs for stages $n \in \mathcal{I}_E^k$ (Lemma 5.2), echelon-n pipeline, early holding and excess end of horizon costs for stages $n \in \mathcal{I}_H^k$ (Lemma 5.3), echelon-n end of horizon shortage costs for stages $n \notin \mathcal{I}^k$, backorder costs (Lemma 5.4), echelon-n pipeline and end of horizon shortage costs for stages $n \in \mathcal{I}_\Pi^k$, and late holding costs at all stages (Lemma 5.7). It is readily verified that this cost amortization does not use any increment of cost incurred by OPT more than once. Thus, inequality (13) indeed holds with probability 1, which implies that (14) holds. We have established the following theorem.

Theorem 5.8 The balancing policy has a worst-case performance guarantee of 2 for the periodic-review serial inventory system with stochastic demands, under the assumptions of Model 1.

We note that this analysis is tight. That is, there exists a sequence of problem instances such that the ratio between the expected cost of the balancing policy and the expected cost of the optimal policy is asymptotically 2. We obtain the problems by modifying the sequence of problem instances discussed in [11] for the single stage so that they hold under the assumption that holding costs are incurred from the moment a

unit is ordered. (In other words, we make the problems conform to Model 1, whereas they currently conform to Model 2.)

5.1 Model 2

The only difference between Model 2 and Model 1 is the fact that in Model 2, echelon holding costs begin to be charged when a unit arrives at the stage, and not when it was ordered by the stage. Specifically, if a unit was ordered at stage n in period s, the echelon-n holding cost for this unit will be charged from period $s + l_n$ until it leaves the system (instead of charging the echelon-n holding cost starting in period s). We make corresponding changes to the end of horizon assumptions, incurring an excess end of horizon cost of $\sum_{m=n}^{N} (c_m + h_m L_{m-1})$ for each excess supply unit that is at stage n in period T + 1, and an end of horizon shortage cost of $\sum_{m=1}^{n-1} (c_m + h_m L_{m-1})$ for each demand unit that is at stage n at time T + 1.

The different assumptions above do not conceptually change the cause-effect cost accounting scheme described in Section 3. In fact the only change is in the pipeline costs that are now equal to $c_n + h_n L_{n-1}$ (instead of $c_n + h_n L_n$ before). In equations (7) and (9) above the corresponding changes occur $(c_n + h_n L_{n-1})$ is used instead of $c_n + h_n L_n$). Similarly, the excess end of horizon cost assigned to the order $\langle s, n \rangle$ is now $(c_n + h_n L_{n-1})Q_n^P$, and the end of horizon shortage cost assigned to $\langle T - L_n, n \rangle$ is now $\sum_{m=1}^{n-1} (c_m + h_m L_{m-1})$ $\widehat{\Pi}_{T-L_n,n}^P$. Moreover, we can still apply the same balancing policy described in Section 4. That is, in each period, conditioned on the observed information set f_s and after immediate orders are placed, at each stage n we order a quantity $q'_n(s)$ such that $E[\widehat{H}_{sn}^B(q'_n(s))|f_s] = E[\Pi_{sn}^B(q'_n(s))|f_s]$.

Next we shall show that in model 2 the balancing policy has a worst-case performance guarantee of 4. Then we discuss certain modifications in the policy that improve the guarantee to 3. Observe that if we replace all instances of $c_n + h_n L_n$ with $c_n + h_n L_{n-1}$, the new assumption on the echelon holding costs does not affect the proofs of Lemmas 5.1, 5.2, 5.3 and 5.4, in which we amortize the pipeline costs incurred by immediate orders, the pipeline, early holding and excess end of horizon costs assigned to ordering decisions in \mathcal{T}_H , and the backordering, end of horizon shortage and late holding costs assigned to ordering decisions in \mathcal{T}_{Π} . In addition, this does not change the proofs of Lemma 5.5 and Corollary 5.6, in which we bound the number of echelon late holding cost increments assigned to ordering decisions at stage n by $L_{n-1} + 1$.

The new assumption on the echelon holding costs does affect, however, Lemma 5.7, which needs to restated. In particular, the proof of Lemma 5.7 is divided into three parts, with each part bounding a separate quantity. The discussion of the first and third quantities still applies, but not that of the second one. The second quantity we bound is the echelon-n late holding cost assigned to ordering decisions in \mathcal{T}_{Π}^k under the balancing policy, summed over $n \in \mathcal{T}_{\Pi}^k$. We bound (i.e., amortize) this cost by the echelon-n pipeline and

end of horizon shortage costs for stages $n \in \mathcal{I}_{\Pi}^k$, incurred under OPT. However, under the new assumption, OPT's echelon-n pipeline holding and end-of-horizon shortage costs are only h_nL_{n-1} . This might not be sufficient to cover the balancing policy's echelon-n late holding costs, which in the worst case can be equal to $h_n(L_{n-1}+1)$. Specifically, it is possible that the balancing policy might have an increment of echelon-n holding cost h_n , for a single stage n, that can not be amortized against the costs of OPT, i.e., we might be able to amortize only h_nL_{n-1} out of the overall cost $h_n(L_{n-1}+1)$.

This implies that we can only prove a weaker version of Inequality (13), that is,

$$C(OPT) \ge \sum_{\langle s,n\rangle \in \mathcal{T}_H} \tilde{H}_{sn}^B + \frac{L_n}{L_n+1} \sum_{\langle s,n\rangle \in \mathcal{T}_\Pi} \Pi_{sn}^B + \sum_{\langle s,n\rangle} (c_n + h_n L_n) \bar{Q}_n^B(s). \tag{15}$$

Observe that the expression $\frac{L_n}{L_n+1}$ is increasing in L_n , i.e., with longer lead times we can amortize a higher fraction of the late holding costs of OPT. Moreover, since $L_n \geq 1$, we conclude that $\frac{L_n}{L_n+1} \geq \frac{1}{2}$. Similar to (14), we can write

$$E[\mathcal{C}(OPT)] \geq \tag{16}$$

$$\sum_{\langle s,n\rangle} E\left[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \frac{1}{2} \cdot \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B} + (c_{n} + h_{n}L_{n-1})\bar{Q}_{n}^{B}(s)\right] =$$

$$\sum_{\langle s,n\rangle} E\left[E[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \frac{1}{2} \cdot \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B}|F_{s}]\right] + \sum_{\langle s,n\rangle} E[(c_{n} + h_{n}L_{n-1})\bar{Q}_{n}^{B}(s)] =$$

$$\sum_{\langle s,n\rangle} E\left[\mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{H})E[\tilde{H}_{sn}^{B}|F_{s}] + \frac{1}{2} \cdot \mathbb{1}(\langle s,n\rangle \in \mathcal{T}_{\Pi})E[\Pi_{sn}^{B}|F_{s}]\right] + \sum_{\langle s,n\rangle} E[(c_{n} + h_{n}L_{n-1})\bar{Q}_{n}^{B}(s)] \geq$$

$$\frac{1}{2}\sum_{\langle s,n\rangle} E[Z_{sn}] + \sum_{\langle s,n\rangle} E[(c_{n} + h_{n}L_{n-1})\bar{Q}_{n}^{B}(s)].$$

Lemma 5.1 and (16) imply the balancing policy has a worst-case performance guarantee of 4. We have proven the following theorem.

Theorem 5.9 The balancing policy has a worst-case performance guarantee of 4 for model 2 of the periodic-review serial inventory system with stochastic demands.

Observe that second line of (16) is not balanced, in that the first term has weight 1 and the second term has weight $\frac{1}{2}$. Next we shall show that by modifying the balancing policy above, we can exploit this fact to improve the worst-case performance guarantee to 3. However, before we explain the details of the modifications, we note that the above problem does not exist (i.e., the balancing policy still has a worst-case performance guarantee of 2) under either one of the following two assumptions:

Assumption 1. Consider model 2 but under continuous review. It is readily verified that in a continuous review system, we can strengthen Lemma 5.5 above and bound the number increments of n-echelon late holding costs ($n \ge 2$) by L_{n-1} (rather than $L_{n-1} + 1$). This implies that Lemma 5.7 is again valid, since we can again amortize the full echelon-n late holding costs with the echelon-n pipeline costs.

Assumption 2. Suppose that, for each $n \ge 2$, we have $c_n \ge h_n$. It is then readily verified that the echelon-n late holding costs can be amortized with the echelon-n pipeline costs; thus, Lemma 5.7 is again valid.

Consider again model 2 under no further assumptions. We shall describe a modified policy that is called a β -balancing policy. This policy takes into account the fact that in model 2, the echelon late holding costs may not be fully amortized. As we have already mentioned this policy has an improved worst-case performance guarantee. Specifically, in each period s, conditioned on the observed information set f_s , the β -balancing policy first places immediate orders as before, and then places a regular order of $q'_n(s)$ units at each stage n, such that $E[\tilde{H}^B_{sn}(q'_n(s))|f_s] = \beta E[\Pi^B_{sn}(q'_n(s))|f_s]$. That is, instead of equating the quantities $E[\tilde{H}^B_{sn}(q^B_n(s))|f_s]$ and $E[\Pi^B_{sn}(q'_n(s))|f_s]$, it balances them in a ratio of β :1. Note that this induces a family of policies one of which is the balancing policy described in Section 4 above (this policy corresponds to the value $\beta = 1$).

The next lemma is a generalization of Lemma 5.1, and is proven using similar arguments. For each $n=1,\ldots,N$ and $s=1,\ldots,T-L_n$, let $Z_{sn}:=E[\Pi_s^B|F_s]$. Observe that by the construction of the β -balancing policy, $E[\tilde{H}_{sn}^B]=\beta Z_{sn}$, with probability 1.

Lemma 5.10 The expected cost of the balancing policy is equal to $(1 + \beta)$ times the expected sum of the Z_{sn} variables plus the expected pipeline costs incurred by immediate orders placed the β -balancing policy. That is,

$$E[\mathcal{C}(B)] = (1+\beta)E[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} Z_{sn}] + E[\sum_{1 \le n \le N} \sum_{1 \le s \le T - L_n} (c_n + h_n L_{n-1}) \bar{Q}_n^B(s)].$$

The rest of the analysis is identical with the exception that (16) is now written as

$$E[\mathcal{C}(OPT)] \geq \sum_{\langle s,n \rangle} E\left[\mathbb{1}(\langle s,n \rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \frac{1}{2} \cdot \mathbb{1}(\langle s,n \rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B} + (c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)\right] = (17)$$

$$\sum_{\langle s,n \rangle} E\left[E[\mathbb{1}(\langle s,n \rangle \in \mathcal{T}_{H})\tilde{H}_{sn}^{B} + \frac{1}{2} \cdot \mathbb{1}(\langle s,n \rangle \in \mathcal{T}_{\Pi})\Pi_{sn}^{B}|F_{s}]\right] + \sum_{\langle s,n \rangle} E[(c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)] =$$

$$\sum_{\langle s,n \rangle} E\left[\mathbb{1}(\langle s,n \rangle \in \mathcal{T}_{H})E[\tilde{H}_{sn}^{B}|F_{s}] + \frac{1}{2} \cdot \mathbb{1}(\langle s,n \rangle \in \mathcal{T}_{\Pi})E[\Pi_{sn}^{B}|F_{s}]\right] + \sum_{\langle s,n \rangle} E[(c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)] \geq$$

$$\min(\frac{1}{2},\beta) \sum_{\langle s,n \rangle} E[Z_{sn}] + \sum_{\langle s,n \rangle} E[(c_{n} + h_{n}L_{n})\bar{Q}_{n}^{B}(s)].$$

Thus the worst-case performance guarantee of the β -balancing policy is $(1+\beta)/\min(\frac{1}{2},\beta)$. It is readily verified that this expression is minimized by the value $\beta = \frac{1}{2}$ and its minimal value is equal to 3.

Theorem 5.11 The $\frac{1}{2}$ -balancing policy has a worst-case performance guarantee of 3 for model 2 of the periodic-review serial inventory system with stochastic demands.

Finally, we consider model 2 under an additional assumption that $\pi \geq h_n$, for each $n \geq 2$. Focus again on unit k. We have already observed that each late holding cost increment assigned to an ordering decision $\langle s, n \rangle$ corresponds to one increment of backordering penalty cost incurred in period $s + L_n$. Thus, for each increment of late holding cost due to unit k that is assigned to $\langle s, n \rangle$, there is an increment of backordering penalty cost due to unit k that is also assigned to $\langle s, n \rangle$.

Now consider some stage n', the ordering decisions $\langle s',n'\rangle \in \mathcal{T}_\Pi$, and the total late holding cost assigned to these decisions by unit k (under the balancing policy). Lemma 5.5 and Corollary 5.6 imply that there is only one scenario in which we can not necessarily fully amortize these costs against respective costs of OPT. Specifically, that scenario is when unit k was on hand at stage n+1 from the beginning of period $\delta_k - L_n$ to the beginning of period $\delta_k + 1$, under the balancing policy. In that case, the total late holding and backorder cost assigned to ordering decisions in $\mathcal{D}_{n'}$ by unit k are $(h'_{n+1} + \pi)(L_n + 1)$. Lemmas 5.4 and 5.7 imply that we can amortize all of these costs except possibly one increment of echelon-n+1 holding cost h_{n+1} . Since $\pi \geq h_{n+1}$, it is readily verified that $[h'_{n+1}L_n+\pi(L_n+1)]/[(h'_{n+1}+\pi)(L_n+1)] \geq 1-\frac{1}{2(L_n+1)}$ of the cost can be amortized. Recall that $L_n \geq 1$, which implies that at least $1-\frac{1}{2(L_n+1)} \geq \frac{3}{4}$ of the cost can be amortized.

Consider the $\frac{3}{4}$ -balancing policy as described above. Using similar arguments as in Lemma 5.10, the expected cost of this policy can be expressed as

$$E[\mathcal{C}(B)] = \frac{7}{4} \sum_{\langle s,n \rangle} E[Z_{sn}] + \sum_{\langle s,n \rangle} E[(c_n + h_n L_{n-1}) \bar{Q}_n^B(s)].$$

In addition, similar to (18) above, we conclude that

$$E[\mathcal{C}(OPT)] \ge \frac{3}{4} \sum_{\langle s,n \rangle} E[Z_{sn}] + \sum_{\langle s,n \rangle} E[(c_n + h_n L_n) \bar{Q}_n^B(s)]. \tag{18}$$

This implies that the worst-case performance guarantee of the $\frac{3}{4}$ -balancing policy is $\frac{7}{4}/\frac{3}{4}=2\frac{1}{3}$.

Theorem 5.12 The $\frac{3}{4}$ -balancing policy has a worst-case performance guarantee of $2\frac{1}{3}$ for model 2 of the periodic-review serial inventory system with stochastic demands, under the assumption that $\pi \geq h_n$, for each n = 2, ..., N.

We note that this analysis for Model 2 is not tight. The sequence of problem instances discussed in [11] shows a lower bound of 2 on the worst-case guarantee, while the analysis that we present establishes guarantees higher than 2.

5.2 Integer-Valued Demand

In this section, we briefly discuss a randomized version of the β -balancing policy that can be applied in scenarios, where orders are restricted to be integers, and the demands are integer-valued random variables. The randomized policy achieves the same worst-case performance guarantees described above for the case where fractional orders are allowed. The modifications to the policy and the analysis are almost identical to what is described in [11, 12] for the single-stage models. Thus, we discuss only the high level ideas and refer the reader to these references for more details.

For ease of exposition we consider the case where $\beta=1$. In each period, conditioned on the observed information set f_s , we first place immediate orders as before. In each stage n, we again consider the functions $E[\tilde{H}^B(q_n^B(s))|f_s]$ and $E[\Pi^B(q_n^B(s))|f_s]$. These functions are originally defined only for integer values of $q_n^B(s)$. We define these functions for any value of $q_n^B(s)$ by interpolating piecewise linear extensions of the integer values. It is clear that these extended functions preserve the properties of convexity and monotonicity discussed in the previous (continuous) case. However, it is still possible (and even likely) that the value $q_n'(s)$ that balances these functions is not an integer. Instead we consider the two consecutive integers $q_n^1(s)$ and $q_n^2(s) := q_n^1(s) + 1$ such that $q_n^1(s) < q_n'(s) < q_n^2(s)$. In particular, $q_n'(s) := \lambda q_n^1(s) + (1 - \lambda)q_n^2(s)$ for some $0 < \lambda < 1$. The randomized balancing policy orders $q_n^1(s)$ units with probability λ and λ and

In [11, 12] it has been shown that by slightly changing the definition of the sets T_H and T_{Π} above, a similar worst-case analysis goes through and theorems analogous to Theorems 5.8, 5.9, 5.11 and 5.12 can be established.

6 Assembly Systems

In assembly systems, each stage may have more than one predecessor. Each predecessor stage has its own non-zero lead time. We assume that each unit produced at a stage is assembled using one unit from each of its predecessors. For convenience, we use stage numbers to denote unit types. Thus, the end unit is of type 1.

Let M_n denote the total lead time for stage n and all its successors. Thus, M_n is the minimum number of time periods that can possibly elapse, starting when we order a unit at stage n, and ending when that unit reaches stage 1. (M_n is a generalization of L_n). We assume that the stages are numbered so that M_n increases with n. This implies that whenever stage k succeeds stage n, k < n.

We continue to use the notion of stage-distance to mean the number of periods required for a unit to pass from its present position in the assembly system to stage 1 if there are no delays at intermediate stages. For $j \leq M_n$, let $X_n^j(t)$ be the number of supply units of type n in the system at stage-distances less than or equal to j, in period t. We say a system is in *long-run balance* in period t if and only if

$$X_n^j(t) \le X_{n+1}^j(t) \tag{19}$$

for all $j \leq M_n$ and all n.

Assuming that a system is in long-run balance in the first period, Rosling shows that the assembly system can be mapped to a corresponding series system. In the series system stage n immediately succeeds stage n-1, and the lead time to make a shipment from stage n+1 to stage n is M_n-M_{n-1} . All feasible policies for the series system naturally give rise to feasible policies for the assembly system, which have the same cost, and which keep the assembly system in balance. (Rosling's results are formally stated for stationary, discounted, infinite-horizon systems, but the relevant proofs clearly apply to the systems that we are studying as well.) Therefore, any policy for the serial system can be used to manage an in-balance assembly system. The important fact is that in so doing there is no loss of optimality.

Theorem 6.1 (Rosling) If the echelon holding costs h_n and backorder cost π are non-negative, and the assembly system starts out in long-run balance, then an optimal policy for the assembly system naturally gives rise to an optimal policy for the corresponding pure series system, and these policies have the same cost.

The long-run balance condition is a natural requirement for the following reason:

Theorem 6.2 (Rosling) If the echelon holding costs h_n and backorder cost π are non-negative then the optimal policy eventually leads the assembly system into long-run balance and keeps it in balance.

These theorems of Rosling are stated for Model 2 with continuous order quantities, infinite time horizons and base stock policies. However the logic of the proofs applies directly to general policies, to finite time horizons, to Model 1, and to systems with integer-valued demands as well. Thus, we establish similar results for the assembly system.

Theorem 6.3 The balancing policy has a worst-case performance guarantee of 2 for the assembly system under model 1; a performance guarantee of 3 for the assembly system under model 2; and a performance guarantee of $2\frac{1}{3}$ under model 2 with the assumption that the backorder parameter is greater than the echelon holding cost parameters of all stages but stage 1.

References

- [1] E. W. M. Chan. *Markov Chain Models for multi-echelon supply chains*. PhD thesis, School of OR&IE, Cornell University, Ithaca, NY, January 1999.
- [2] F. Chen and J. Song. Optimal policies for multi-echelon inventory problems with Markov-modulated demand. *Operations Research*, 49:226–234, 2001.
- [3] F. Chen and Y. Zheng. Lower bounds for multi-echelon stochastic inventory systems. *Management Science*, 40:1426–1443, 1994.
- [4] A. Clark and H. Scarf. Optimal policies for a multiechelon inventory problem. *Management Science*, 6:475–490, 1960.
- [5] L. Dong and H. L. Lee. Optimal policies and approximations for a serial multiechelon inventory system with time-correlated demand. *Operations Research*, 51, 2003.
- [6] A. Federgruen and P. Zipkin. Computational issues in an infinite-horizon multi-echelon inventory model. *Operations Research*, 32:818–836, 1984.
- [7] T. Iida. The infinite horizon non-stationary stochastic multi-echelon inventory problem and near myopic policies. Working paper, 1998.
- [8] T. Iida and P. Zipkin. Approximate solutions of a dynamic forecast-inventory model. Working paper, 2001.
- [9] G. Janakiraman and J. A. Muckstadt. A decomposition approach for a capacitated singel-stage inventory system. Technical Report TR1360, ORIE Department, Cornell University, 2003.

- [10] G. Janakiraman and R. O. Roundy. Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Operations Research*, 52:785–803, 2004.
- [11] R. Levi, M. Pal, R. O. Roundy, and D. B. Shmoys. Approximation algorithms for stochastic inventory control models. Technical Report TR1412, ORIE Department, Cornell University, 2004. Submitted.
- [12] R. Levi, R. O. Roundy, D. B. Shmoys, and V. A. Truong. Approximation algorithms for capacitated stochastic inventory control models. submitted, 2004.
- [13] A. Muharremoglu and J. N. Tsitsiklis. A single-unit decomposition approach to multi-echelon inventory systems. Working paper, 2001.
- [14] Ö. Özer and G. Gallego. A new algorithm and a new heuristic for serial supply systems. To appear in Operations Research Letters, 2005.
- [15] K. Rosling. Optimal inventory policies for assembly systems under random demands. *Operations Research*, 37:565–579, 1989.
- [16] H. K. Shang and J. Song. Newsvendor bounds and heuristics for optimal policies in a serial suplly-chains. *Management Science*, 49:618–638, 2003.
- [17] J. Song and P. Zipkin. Inventory control in a fluctuating demand environment. *Operations Research*, 41:351–370, 1993.
- [18] P. H. Zipkin. Foundations of inventory management. The McGraw-Hill Companies, Inc, 2000.