Traditional Inventory Models in an E-Retailing Setting: A Two-Stage Serial System with Space Constraints

Russell Allgor, Stephen Graves, and Ping Josephine Xu
Amazon.com MIT

Abstract— In an e-retailing setting, the efficient utilization of inventory, storage space, and labor is paramount to achieving high levels of customer service and company profits. To optimize the storage space and labor, a retailer will split the warehouse into two storage regions with different densities. One region is for picking customer orders and the other to hold reserve stock. As a consequence, the inventory system for the warehouse is a multi-item two-stage, serial system. We investigate the problem when demand is stochastic and the objective is to minimize the total expected average cost under some space constraints. We generate an approximate formulation and solution procedure for a periodic review, nested ordering policy, and provide managerial insights on the trade-offs. In addition, we extend the formulation to account for shipping delays and advanced order information.

Index Terms— Inventory, multi-echelon, stochastic demand, periodic review.

I. INTRODUCTION

After a customer orders online at an e-retailer, the order is assigned virtually to one of order fulfillment centers, which are of several hundred thousand square feet. An order fulfillment center is a warehouse consisting of a *picking area*, where items are stored individually in bins, and a *deep-storage area*, where items are stored in bulk on pallets. The customer-ordered items are hand-picked in the picking area and sent for packing afterwards. The deep-storage area receives items from the outsider suppliers and replenishes the picking area.

For e-retailers, which operate with no physical stores, the efficient utilization of inventory, storage space, and labor is paramount to achieving high levels of customer service and company profits. To optimize the storage space and labor, an e-retailer splits the warehouse into two storage regions with different densities. One region is for picking customer orders and the other to hold reserve stock. Consequently, the inventory in the warehouse flows in a serial, two-stage fashion, as illustrated in Figure 1. We investigate the problem of multi-item inventory ordering policy for a two-stage serial system when demand is stochastic and the objective is to minimize

R. Allgor is with Amazon.com at Seattle, WA 98108 USA (email: rallgor@amazon.com).
S. C. Graves is with the Sloan School of Management and the Engineering System Division at MIT, Cambridge MA 02139 USA (email: sgraves@mit.edu).
P. J. Xu is with the Operations Research Center at MIT, Cambridge MA 02139 USA (email: pingx@mit.edu).
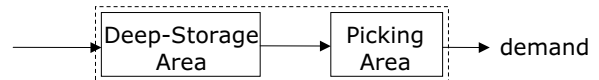
Fig. 1. A serial, two-stage warehouse.

the long run average cost under space constraints. A key assumption of this problem is the existence of economies of scales (in the form of fixed order costs) for replenishing the inventory at both stages.

The problem under consideration is fundamental in warehouse operations, both in a brick-and-mortar or an e-commerce operation. We will, however, consider extension of this problem that are particularly applicable in e-retailing later in the paper. The problem is also well known for its theoretical difficulties. Clark and Scarf [5] concluded that for the unconstrained two-stage, serial inventory system with set-up costs at both stages, the optimal ordering policy, if one exists, must be extremely complex. This observation has driven subsequent research to focus on heuristic policies. There is an extensive literature on various heuristic policies in different multi-echelon, stochastic inventory systems. Here we contribute by generating an approximate formulation and solution procedure for a periodic review, heuristic ordering policy in a multi-item two-stage problem, and provide insights about the intrinsic trade-offs in a constrained warehouse operation.

Using periodic review ordering policy in the inventory model is motivated by several practical considerations. E-retailers often take pride in their wide range of products available for customers. As a result, the number of products an e-retailer orders from a single supplier is often very large. Periodic review policy may reduce fixed replenishment costs by combining order replenishment for different products. A periodic review policy also has the practical benefits of following a regular repeated schedule to coordinate transportation and other logistic considerations [16].

A. Literature Review

A considerable body of research has evolved in the field of multi-echelon, stochastic inventory systems since the publication of Clark and Scarf in 1960. Axsater [2] and Federgruen [9] provide comprehensive reviews of this literature. In particular, Axsater [2] reviews the literature on continuous review policies for multi-echelon, stochastic systems. Examples include Sherbrooke [19], Graves [12], De Bodt and Graves [6],

Deuermeyer and Schwarz [7], and Svoronos and Zipkin [20]. Federgruen [9] reviews multi-echelon, stochastic models that are centralized and have multiple locations, such as a serial or assembly system. Examples include Clark and Scarf [5], Eppen and Schrage [8], Federgruen and Zipkin [10], [11], and Jackson [14]. However, progress is slow in establishing near-optimal heuristic policies with a guaranteed, worst-case performance. Chen [4] characterizes a continuous review heuristic policy for a two-stage inventory system. The long-run average cost is guaranteed to be within 6% of optimality, where demand is Poisson, leadtime at stage 2 is zero, and both stages incur a fixed order cost.

Closely related to this paper are two publications. De Bodt & Graves [6] develop a similar two-stage serial model for a continuous review (Q, r) policy. They provide approximate performance measures under a nested policy assumption: whenever a stage receives a shipment, a batch must be immediately sent down to its downstream stage. They do not make an assumption about the form of the demand distribution. We, however, consider a periodic review (R, T) policy. Our major assumptions on the policy are different but close in spirit as the model in De Bodt & Graves [6]. Most recently, Rao [16] analyzed the properties of the single-stage (R, T) model, as a counterpart of Roundy [17] and Zheng [21] for a deterministic periodic review model and stochastic (Q, r) model, but with certain demand function restrictions. In the extension, he develops a two-stage serial system which is similar to our model but has different assumptions on the interaction between echelons.

B. Overview

§ II reviews the single-stage periodic review model and its most recent results, and presents the two-stage serial model, along with some space constraints. § III is an extension that accounts for shipping delays and advance demand information.

II. MODEL FORMULATION AND SOLUTION APPROACH

Before proceeding further, we list the following standard definitions:

$I(t)$ —   on-hand inventory or the amount of inventory in the warehouse at time t,

$B(t)$ —   amount of unfulfilled customer demand at t,

$IL(t)$ —   inventory level or net inventory at t, equivalent to $I(t) - B(t)$,

$O(t)$ —   on-order replenishment at t,

$IP(t)$ —   inventory position at t, equivalent to $IL(t) + O(t)$.

Following the literature convention, we denote stage 1 as the downstream stage that faces external demand, and stage 2 as the upstream stage that replenishes stage 1 and is replenished by outsider suppliers. In the e-retailing setting, stage 1 is the picking area and stage 2 is the deep-storage area.

The *basic model* in this paper is an unconstrained single-item two-stage serial model. It extends the single-item single-stage periodic-review (R, T) model ([13], p. 237-245), where

R is the order-up-to level and T is the review period. After every T time units, we order up to R if the current inventory position is below R.

In the next few sub-sections, we first review the single-stage model and then present the basic model.

A. Single-Stage Model Review

We first list the key assumptions of the single-stage model. These assumptions apply to the basic model as well, while additional assumptions will be introduced as we discuss the basic model in detail.

A-1   The inter-arrival times between successive demands are *i.i.d.*. Demand is stationary for the relevant time horizon.

A-2   Each stage has a constant known nonzero lead time.

A-3   When there is no on-hand inventories at stage 1, demand at stage 1 is backlogged, and a penalty cost per backorder is charged.

A-4   Demand backorder quantities are small. We will provide more details on this assumption later in the section.

We need to discuss the validity of the assumptions. We assume stationary demand in A-1. There are usually two distinct demand patterns in the e-retailing setting, namely the off-peak and peak season. Within each season, it is reasonable to assume stationary demand trend. We can treat the two seasons as two separate models. Relaxing A-2 to allow stochastic lead time would not change the formulation much. Unlike in A-3, many models in the literature have the backorder cost as cost charged per item per unit time. Our assumption is more applicable when the fixed cost component of backorder is much larger than the time variable component. We note that the formulation under our backorder cost assumption may be less convenient for theoretical analysis, but it is easier in computation.

We denote:

$C(\cdot)$ —   expected total cost per unit time,

l —   replenishment lead time,

d —   expected demand per unit time,

a —   fixed order, or replenishment, cost,

h —   holding cost per item per unit time,

b —   backorder cost per item,

$f(x|l)$ —   probability density function of lead-time demand, given that the lead time is l.

We assume that discrete units of inventory can be approximated by continuous variables. We follow the inventory literature (e.g., [13], p. 237-245), and the expected total cost per unit time can be approximated as:

$$C(R, T) \cong \frac{a}{T} + h\left(R - d(l + \frac{T}{2})\right) + \frac{b}{T}\int_R^\infty (x - R)f(x|l + T)dx \quad (1)$$

Comparing Equation (1) with the exact model for Poisson demand in Appendix A, we note that the main approximation is the holding-cost term, which is underestimated. Figure 2
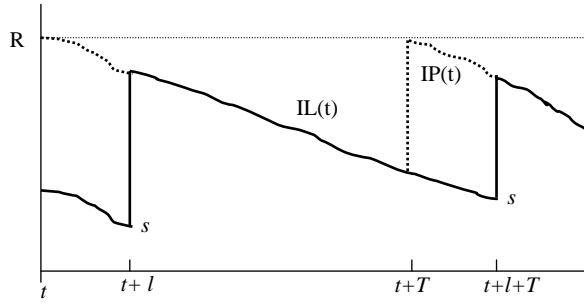
Fig. 2.   Single-stage inventory level and position.

Fig. 3.   A two-stage inventory level and position example, $n = 3$.

is an inventory diagram for the single-stage model. The time between $[t + l, t + l + T]$ is a typical replenishment cycle. We can write the holding-cost term in Equation (1) as $\frac{h}{T} \int_l^{l+T} IL(t)\, dt$, as opposed to $\frac{h}{T} \int_l^{l+T} I(t)\, dt$ in the exact model. That is, we approximate on-hand inventory with net inventory. However, the approximation error is small when backorder is small so that on-hand inventory is close to net inventory. Hence, assumption A-4 guarantees that the approximation is close. The backorder term is, however, slightly overestimated. That is, we assume that we start each replenishment cycle with zero backorders.

For a given value of T, $C(R, T)$ in Equation (1) is convex in R. We can obtain the optimal value of R for a given value of T:

$$\int_R^\infty f(x|l + T) = \frac{hT}{b}. \qquad (2)$$

Moreover, for demand distributions that have $f(x|l + T) > 0, \ \forall x > 0$, $C(R, T)$ is strictly convex in R for a given value of T. Equation 2 then yields a unique value of R. Given a value of R, $C(R, T)$ is not convex in T.

Note that for $T > \frac{b}{h}$, there exists no solution in Equation (2). In this case, we set $R = 0$. We search over values of T in the range $(0, \frac{b}{h})$. It is simple to tabulate over values of T and use Equation (2) to obtain the minimal value of expected total cost.

B. The Basic Model : A Two-Stage Serial Model

Now we consider our basic model, an approximate two-stage serial (R, T) model based on §II-A. We use the concept of *echelon stock*, which is the total inventory in the current stage and all its downstream stages. In our notation, we use subscript *1* to denote echelon 1, that is stage 1, and *2* to denote echelon 2, that is stage 1 and 2.

We first present a main assumption of the basic model.

A-5 Each stage manages its echelon inventory with an (R, T) policy. Furthermore, the ordering policies are nested. That is, stage 1 places a replenishment when stage 2 receives its replenishment. To coordinate the replenishment of both stages, we impose the constraint $T_2 = nT_1$, where n is a positive integer.

Nested policies are applicable when warehouses prefer to move certain shipments from outside suppliers directly to stage 1. Certainly, this assumption on the ordering policy simplifies analysis.
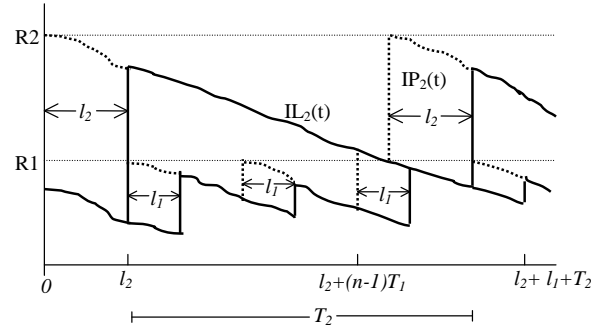
Figure 3 has an example for $n = 3$. We will demonstrate the policy behavior through this example. Echelon 2 receives its replenishment at time l_2 and $l_2 + T_2$. Echelon 1 receives its replenishment at time $l_2 + l_1$, $l_2 + l_1 + T_1$, and $l_2 + l_1 + 2T_1$. Since $n = 3$, there are three inventory replenishment (reviews) of echelon 1 between every consecutive echelon-2 inventory replenishment (reviews). We denote a cycle as the time between consecutive echelon-2 inventory replenishment.

To simplify the formulation, we have the following additional assumptions.

A-6 $IP_1(l_2 + (n - 1)T_1) = IL_2(l_2 + (n - 1)T_1)$, given stage 2 orders at time $t = 0$. That is, during the last replenishment in a cycle, stage 1 orders *all* of the remaining on-hand inventory from stage 2.

We call the last echelon-1 replenishment in a cycle an *exhaustive replenishment*, and the other $n - 1$ echelon-1 replenishment *normal replenishment*. We claim that we have $n-1$ normal replenishment for every exhaustive replenishment in this ordering policy. At the exhaustive replenishment, our order-up-level could be less than, equal to, or greater than R_1. This assumption allows us to simplify the formulation without having to track which cycle has extra inventory in echelon 2 at the end of the cycle. On average, the amount of such extra stock is small. Otherwise, we could decrease the value of R_2. Also, there is no value to leave the extra inventory in echelon 2, since a replenishment for echelon 2 will arrive next. Therefore, since the stock is already in the warehouse at the exhaustive replenishment, it may be more cost effective to move the extra stock to stage 1.

Given stage 2 orders at time $t = 0$, we assume that

A-7 $IL_2(l_2 + (n - 2)T_1) > R_1$.

This assumption ensures that echelon 2 has sufficient inventory to raise the stage 1 inventory position to R_1 for every normal replenishment. Since $IL_2(\tau)$ is nondecreasing for $l_2 \leq \tau \leq l_2 + T_2$, this assumption also implies that $IL_2(l_2 + mT_1) > R_1, \ \forall m < n - 1$. However, because of assumption A-6, inventory level at the exhaustive replenishment is not restricted, $IL_2(l_2 + (n - 1)T_1) \geq R_1$ or $< R_1$.

A-8 $IP_1^-(l_2) < R_1$.

We denote $IP^-(t)$ as the inventory before the event takes place at time t. We assume the inventory level in echelon 2 (equivalent to inventory level in echelon 1 due to A-6) shortly before the shipment from the outside supplier

arrives is less than R_1. Otherwise, stage 2 doesn't need to send any shipment to stage 1 at time l_2.

We denote $D(t, t+\tau)$ as the total demand placed from time t to $t + \tau$. If echelon 2 orders up to R_2 at time t, then for A-7 to be valid, we must have that

$$D(t, t + l_2 + (n-2)T_1) \le R_2 - R_1. \tag{3}$$

For A-8 to be valid, we must have that

$$D(t, t + l_2 + T_2) \ge R_2 - R_1. \tag{4}$$

We expected that the accuracy of our cost expressions will depend on the probability of the above two equations.

To develope the cost expressions, we derive the cost elements separately. The expected set-up cost per unit time is

$$\frac{a_1}{T_1} + \frac{a_2}{T_2}. \tag{5}$$

As a result of the assumptions, we order after every review period.

To derive the holding cost element, we examine echelon 1 and 2 separately. It is easy to see that the holding cost of echelon 2 can be approximated as in the single-stage model, $\frac{h_2}{T_2} \int_{l_2}^{l_2 + T_2} IL_2(t)\, dt$, that is

$$h_2 \left(R_2 - d \left(l_2 + T_2/2 \right) \right). \tag{6}$$

The holding cost for echelon 1 needs a little more discussion. For the $n-1$ normal replenishment cycles, we can derive the holding cost for each cycle just as in the single-stage model. However, the expected inventory level for a exhaustive replenishment cycle needs to be re-derived. Referring to Figure 3, the time between $(l_2 + l_1 + (n-1)T_1, l_2 + l_1 + T_2)$ is an exhaustive replenishment cycle for stage 1. Let $D(0, t)$ be the demand up to time t. The inventory level at the start of the cycle is $R_2 - D\left(0, l_2 + l_1 + (n-1)T_1\right)$. The inventory level at the end of the cycle is $R_2 - D\left(0, l_2 + l_1 + T_2\right)$. The average net inventory in the cycle is, therefore,

$$\frac{1}{T_1} \int_{l_1+l_2+(n-1)T_1}^{l_1+l_2+T_2} IL_1(t)\, dt = R_2 - d\left(l_1 + l_2 + T_2 - \frac{T_1}{2} \right)$$

We can then write the holding cost at stage 1 as

$$h_1 \left(\frac{n-1}{n} \left(R_1 - d\left(l_1 + T_1/2 \right) \right) \right.$$
$$\left. + \frac{1}{n} \left(R_2 - d\left(l_1 + l_2 + T_2 - T_1/2 \right) \right) \right). \tag{7}$$

We again derive the backorder costs for normal and exhaustive replenishment separately. The expected number of backorders during a normal replenishment cycle is $\int_{R_1}^{\infty} (x - R_1) f(x|T_1 + l_1) dx$. The expected number of backorders during an exhaustive replenishment cycle is $\int_{R_2}^{\infty} (x - R_2) f(x|T_2 + l_1 + l_2) dx$. We can express the expected backorder cost per unit time as

$$\frac{b}{T_1} \left(\frac{n-1}{n} \int_{R_1}^{\infty} (x - R_1) f(x|T_1 + l_1) dx \right.$$
$$\left. + \frac{1}{n} \int_{R_2}^{\infty} (x - R_2) f(x|T_2 + l_1 + l_2) dx \right). \tag{8}$$

By summing up Equations (5) to (8), we have the expected average total cost $C(R_1, R_2, T_1, T_2, n)$. Substituting the constraint nT_1 for T_2, we have the cost function $C(R_1, R_2, n, T_1)$. The optimization problem, P, can be written as:

$$\min \quad C(R_1, R_2, T_1, n)$$
$$n \in \mathbb{Z}^+$$
$$R_1, R_2, T_1 \ge 0$$

For given values of (T_1, n), the cost function $C(R_1, R_2, n, T_1)$ is a convex function in R_1, R_2. We can find solutions of R_1, R_2 according to the following equations:

$$\int_{R_1}^{\infty} f(x|T_1 + l_1)dx = \frac{h_1 T_1}{b}, \tag{9}$$

$$\int_{R_2}^{\infty} f(x|T_2 + l_1 + l_2)dx = \frac{h_1 T_1}{b} + \frac{h_2 T_2}{b}. \tag{10}$$

Equations (9) and (10) are a result of setting $\frac{\partial C}{\partial R_1}, \frac{\partial C}{\partial R_2}$ to be zero. For Equation (9) and (10) to have unique solutions, we need to have $\frac{\partial^2 C}{\partial R_1^2} = \frac{b}{T_1} \frac{n-1}{n} f(R_1|T_1 + l_1) > 0$ and $\frac{\partial^2 C}{\partial R_2^2} = \frac{b}{T_2} f(R_2|T_2 + l_1 + l_2) > 0$. As in the single-stage model, for demand distributions that have $f(x|t) > 0, \ \forall x > 0, \ t > 0$, Equations (9) and (10) have unique solutions. However, the cost function $C(R_1, R_2, n, T_1)$ is not convex in T_1 or n.

We can use a simple search method, where we search over given values of T_1 and n. The value of n is a positive integer. Note that for large value of T_1 or n in Equation (10), we set $R_2 = 0$. Therefore, we search over the range of values of (T_1, n) such that $(h_1 + nh_2)T_1 < b$. If the value of T_1 is restricted to be a multiple of some minimal review period (e.g., a day), it is simple just to tabulate over the values of T_1 and n. For problems that have a large range of (T_1, n), we consider using simple gradient methods like Newton's method or Steepest Descent method where the step sized is determined by Amijo's rule. We can use the starting value of $T_1^D = \frac{EOQ}{d} = \sqrt{\frac{2a_1}{h_1 d}}$ from the single-stage deterministic problem. We can use the starting value of $n^D \approx \sqrt{\frac{a_2}{a_1} \frac{h_1}{h_2}}$ from the deterministic demand two-stage problem. The starting values of R_1 and R_2 can be determined accordingly given T_1^D and n^D.

For $n = 1$, the problem can be solved as a single-stage problem whose cost parameters are $h = h_1 + h_2, a = a_1 + a_2$, and $l = l_1 + l_2$. However, the cost of the $n = 1$ two-stage problem is not equivalent to such a single-stage problem due to a minor accounting difference in holding cost.

C. Multi-Item Two-Stage Model with Space Constraints

A multi-item two-stage problem based on the basic model can be solved separately for each item as described in the previous section. A space constraint couples all items together. We consider two different space constraints: *i)* on the total space in echelon 2, which corresponds to the picking and deep-storage area, and *ii)* on the space in echelon 1 (i.e. stage 1) only, which is the picking area. We need to introduce

additional notations:

M – number of items in storage,

γ_{ik} – space taken by an item i, (e.g., cubic in. per item), in stage k. Typically, $\gamma_{i1} > \gamma_{i2}$.

A_{ij} – average inventory per unit time of item i in echelon j,

S_j – available space in echelon j.

1) Space Constraint on Echelon 2: In the context of an e-retailer, stage 1 and stage 2 are in the same warehouse, and therefore, share the total space in the warehouse. Imposing a constraint on the total space seems natural. Denote C_i as the total expected cost per unit time of item i, then the problem can be formulated as:

$$\min \quad \sum_{i=1}^{M} C_i(R_{i1}, R_{i2}, T_{i1}, n_i)$$

$$\text{s.t.} \quad \sum_{i=1}^{M} \gamma_{i1} A_{i1} + \gamma_{i2}(A_{i2} - A_{i1}) \leq S_2 \quad (11)$$

$$n_i \in \mathbb{Z}^+, \quad \forall i$$

$$R_{i1}, R_{i2}, T_{i1} \geq 0, \quad \forall i,$$

where for each item i, A_1 can be found in Equation (7). Equation (7) is equivalent to $h_1 A_1$. The term A_2 can be found in Equation (6). Equation (6) is $h_2 A_2$. We use the average inventory in the space constraint as an approximation. This approximation is adequate as the number of items M becomes large.

We solve the problem by solving the dual problem. Denote θ to be the Lagrangian Multiplier. Given θ, the Lagrangian function is:

$$L(\bar{R}_1, \bar{R}_2, \bar{n}, \bar{T}_1, \theta) = \sum_{i=1}^{M} C_i(R_{i1}, R_{i2}, n_i, T_{i1}) \quad (12)$$

$$+ \theta \left(\sum_{i=1}^{M} \gamma_{i1} A_{i1} + \gamma_{i2}(A_{i2} - A_{i1}) \right) - \theta S_2$$

$$= \sum_{i=1}^{M} \tilde{C}_i(R_{i1}, R_{i2}, n_i, T_{i1}, \theta) - \theta S_2, \quad (13)$$

where $\bar{R}_1, \bar{R}_2, \bar{n}, \bar{T}_1$ are vectors whose ith component is for item i, and the cost function $\tilde{C}_i$ has the same cost structure as C_i but with modified holding costs. Specifically, the holding costs in $\tilde{C}_i$, denoted as $\tilde{h}_{ij}$, can be set as:

$$\tilde{h}_{i1} \leftarrow h_{i1} + \theta(\gamma_{i1} - \gamma_{i2})$$
$$\tilde{h}_{i2} \leftarrow h_{i2} + \theta\gamma_{i2}.$$

The dual function q can be written as:

$$q(\theta) = \min_{R_{i1}, R_{i2}, T_{i1} \geq 0, \ n_i \in \mathbb{Z}^+} \sum_{i=1}^{M} \tilde{C}_i(R_{i1}, R_{i2}, n_i, T_{i1}, \theta) - \theta S_2.$$

$$(14)$$

Equation (14) can be solved as M separable problems, and each can be solved as a single-item problem with modified

holding costs. The value θS_2 is merely a constant. We can then solve the dual problem:

$$\max \quad q(\theta)$$
$$\text{s.t.} \quad \theta \geq 0.$$

We can search over the values of θ to solve the dual problem. This separable problem structure is helpful in solving the dual problem. The cost function $\sum_{i=1}^{M} C_i(R_{i1}, Ri2, T_{i1}, n_i)$ is not convex and the feasible solution set is not convex since n_i is discrete. The separable structure is additionally helpful here because it is known to have relatively a small duality gap, and the duality gap can be shown to diminish to zero relative to the optimal value as M increases [3].

2) Space Constraint on Echelon 1: In the context of e-retailing, echelon 1 is the picking area. The larger the picking area, the more costly it would be to pick items efficiently. For example, a worker picks items from a list of customer orders. The larger the picking area, the longer the route he or she may have to walk to complete the task. Therefore, labor costs are higher per customer order when the picking area is larger. A space constraint on echelon 1 can ensure efficient picking or efficient utilization of labor. Suppose the total warehouse space may be augmented easily, such as adding some trailers in the yard or finding a close storage building. Then imposing a constraint only on echelon 1 is reasonable. The problem with only echelon 1 is constrained can be formulated as the following:

$$\min \quad \sum_{i=1}^{M} C_i(R_{i1}, R_{i2}, T_{i1}, n_i)$$

$$\text{s.t.} \quad \sum_{i=1}^{M} \gamma_{i1} A_{i1} \leq S_1, \quad (15)$$

$$n_i \in \mathbb{Z}^+, \quad \forall i$$

$$R_{i1}, R_{i2}, T_{i1} \geq 0, \quad \forall i,$$

Similar to the procedures in the previous section, we again solve for the dual problem. Given the value of θ, the dual function can be solved by solving M separable minimization single-item problems, where the holding costs can be set as $\hat{h}_{ij}$:

$$\hat{h}_{i1} \leftarrow h_{i1} + \theta\gamma_{i1}$$
$$\hat{h}_{i2} \leftarrow h_{i2}$$

III. EXTENSION - ALLOCATING SPACE FOR WIP

In most of the inventory models in the literature, inventory units disappear from the warehouse as soon as they meet demand. This modeling assumption is reasonable in most applications. However, it is not so realistic in the e-retailing setting. Demanded items often may be in a multi-item order, where some of the items may be out of stock and the entire order stays in the warehouse until all items are available. In this sense, the order fulfillment process is more like a make-to-assemble process: products can be any subset of all items, and customer orders are assembled after their replenishment are received.

We call items in the warehouse that were assigned to customer orders but are waiting to be shipped as Work In Process (WIP). One way to incorporate WIP into the current inventory model is to allocate some space for WIP. Once an item is ordered, we virtually direct the item to a WIP $M/G/\infty$ queue. The item leaves the virtual queue when all items in the same order are assembled. Note that we do not assume the form of the demand distribution in the basic model, but we assume Poisson demand in this extension. More formally, for single item, we denote

λ $-$ Poisson demand arrival rate,

Y $-$ time from the item is assigned to a warehouse

      to the time until it is assembled with all other items

      in the order and sent to the shipping department.

Let $N(t)$ be the number of demand arrivals for the item in $(0, t)$ still in service at τ, then $N(t)$ has non-homogenous Poisson rate, $\lambda P(Y > \tau - t)$. In steady state,

$$\lim_{\tau \to \infty} E[N(\tau)] = \lambda \int_0^\infty P(Y > t) \, dt$$
$$= \lambda E[Y]$$

The distribution of Y may depend on the ordering policy, multi-item demand patterns, and assembly priority policy. As an approximation, we can eliminate the complexity by determining Y from historical data. Then we can incorporate the term into our models by leaving the WIP queue in stage 1,

$$min \quad \sum_{i=1}^M C_i(R_{i1}, R_{i2}, T_{i1}, n_i)$$
$$s.t. \quad \sum_{i=1}^M \gamma_{i1}\left(A_{i1} + d_i E[Y_i]\right) \leq S_1, \qquad (16)$$
$$n_i \in \mathbb{Z}^+, \qquad \forall i$$
$$R_{i1}, R_{i2}, T_{i1} \geq 0, \qquad \forall i,$$

where to be consistent with our previous definition, $\lambda = d$.

APPENDIX A: EXACT MODEL OF THE SINGLE-STAGED (R, T) MODEL FOR POISSON DEMAND

Following the literature (e.g., [13]), we write the exact expected total cost per unit time for Poisson demand as:

$$C(R, T) = \frac{a}{T}P(1|T) + \frac{b}{T}\sum_{x=R}^{\infty}(x - R)p(x|l + T)$$
$$- \frac{b}{T}\sum_{x=R}^{\infty}(x - R)p(x|l) + h\left(R - d(l + \frac{T}{2})\right)$$
$$+ \frac{h}{T}\int_l^{l+T}\sum_{x=R}^{\infty}(x - R)p(x|t)dt,$$

where $p(x|\tau)$ is the PMF of demand during time τ, and $P(x|\tau)$ is the right-hand CDF of demand during time τ.

REFERENCES

[1] Axsater, S. 1990. Simple solution procedures for a class of two-echelon inventory problems. *Oper. Res.* 38, 64-69.
[2] Axsater, S. 1993. Continuous review policies for multi-level inventory systems with stochastic demand. S. Graves, A. Rinnooy, Kan and P. Zipkin, eds. *Handbook in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory.* North Holland.
[3] Bertsekas, D. 1999. *Nonlinear Programming.* Athena Scientific, Belmont, Massachusetts.
[4] Chen, F.. 1999. 94%-Effective policies for a two-stage serial inventory system with stochastic demand. *Management Science.* 45(12), 1679-1696.
[5] Clark, A., H. Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6, 475-490.
[6] DeBodt, M. and S.C. Graves. 1985. Continuous-review policies for a multi-Echelon inventory problem with stochastic demand. *Management Science.* 31(10), 1286-1299.
[7] Deuermeyer, B., L. Schwarz. 1981. A model for the analysis of system service level in warehouse/retailer distribution systems: The identical retailer case. L. Schwarz, ed. *Studies in the Management Sciences: The Multi-Level Production Inventory Control Systems,* Vol. 16. North-Holland, Amsterdan. 163-193.
[8] Eppen, G. and L. Schrage. 1981. Centralized ordering policies in a multiwarehouse system with leadtimes and random demand. L. Schwarz, ed. *Multi-Level Production Inventory Control Systems: Theory and Practice.* North Holland, 51-69.
[9] Federgruen, A. 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. S. Graves, A. Rinnooy Kan and P. Zipkin, eds. *Handbook in Operations Research and Management Science, Vol. 4, Logistics of Production and Inventory.* North Holland, Amsterdam.
[10] Federgruen, A., P. Zipkin. 1984a. Approximation of dynamic, multi-location production and inventory problems. *Management Science.* 30, 69-84.
[11] Federgruen, A., P. Zipkin. 1984b. Allocation policies and cost approximation for multi-location inventory systems. *Naval Res. Logist. Quart.* 31, 97-131.
[12] Graves, S.C. 1985. A multi-echelon inventory model for a reparable item with one-for-one replenishment. *Management Science.* 31, 1247-1256.
[13] Hadley, G. and T.M. Whitin. 1963. *Analysis of Inventory System.* Prentice-Hall, Englewood Cliffs, N.J..
[14] Jackson, P. 1988. Stock allocation in a two-echelon distribution system of "What to do until your ship comes in". *Management Science.* 34, 880-895.
[15] Naddor, E.. 1982. *Inventory Systems.* Robert E. Krieger Publishing Co., Malabar, Florida.
[16] Rao, U.S.. 2003. Properties of the periodic review (R, T) inventory control policy for stationary, stochastic demand. *Manufacturing & Service Operations Management.* 5(1), 37-53.
[17] Roundy, R. 1986. A 98% effective lot sizing rule for multi-product, multi-stage production inventory systems. *Math. OR.* 11, 699-727.
[18] Schwarz, L. B.. 1973. A simple continuous review deterministic one-warehouse N-retailer inventory problem. *Management Science.* 19(5), 555-566.
[19] Sherbrooke, C. 1968. METRIC: A multi-echelon technique for recoverable item control. *Operations Research.* 16, 122-141.
[20] Svoronos, A., P. Zipkin. 1988. Estimating the performance of multi-level inventory systems. *Operations Research.* 36, 57-72.
[21] Zheng, Y.S.. 1992. On properties of stochastic inventory systems. *Management Science.* 38(1), 87-103.