

A Base Stock Inventory Model for a Remanufacturable Product

Stephen C. Graves

MIT

Abstract—We report on an industrial project in which we developed an inventory model to provide decision support for the design and deployment of the field service support system for a remanufacturable product. The product was a dialysis unit for home use. Each unit that was installed in a home would eventually be removed due to failure, or the need for preventative maintenance, or the termination of the service. Upon removal, each unit was shipped to a central depot for re-manufacturing so that it could be returned to service. We develop a model to determine the inventory requirements at each regional depot, and then describe how to use the model to determine the inventory requirements in the two-echelon system consisting of the central facility and the regional depots.

Index Terms— inventory planning, remanufacturing, reverse logistics, multi-echelon systems, supply chain planning.

I. INTRODUCTION

The supply chains for more and more products must now accommodate a return process so as to facilitate reuse, in one form or another, of the product and its materials. The design of a supply chain with a returns process introduces a number of new challenges. For instance, these return flows can provide a second source of supply if the products can be re-manufactured and returned to use. As such, inventory management must balance procurement of new items from the original supplier with replenishment through re-manufacturing of returned items.

In this paper I report on a model developed for the design of the field service support system for a piece of medical equipment, namely a dialysis unit for home use. This equipment was being developed to allow a person, in need of regular medical treatment, to treat themselves in the comfort of their home, rather than to have to go to a medical facility for the procedure, which is both more expensive and more inconvenient. The equipment was designed to be reused. When a unit was removed from a home, for whatever reason, it would be re-manufactured and then returned to the inventory

for subsequent installation in another home.

The intent of the work was to develop a model that would inform the design of the field service support system. In particular, we wanted to determine the amount of inventory needed and how it should be deployed. We also were asked to determine the workload on the service engineers, as it depends on the volume of transactions. Finally, we wanted to be able to explore how the inventory and workload requirements depend on the key system parameters such as the demand and removal rates for the equipment, as well as various lead times.

In terms of related literature, I would first cite the work on inventory models for repairable items (see Nahmias 1981). The current paper is similar to this literature in that we assume that items can be recovered either through repair or re-manufacturing. But our model differs from this literature in that the demand process (new installations) is composed of an independent process from new patients plus an endogenous process due to equipment failure and preventative maintenance. There is now an emerging literature on models for reverse logistics, within which this paper should fit (see Thierry et al. 1995 and Fleischmann et al. 1997 for an overview). Toktay et al. (2000) develop an inventory model for a remanufacturable product for which key components are recycled. A key difference between their work and this paper is that we have visibility of the returns process, whereas Toktay et al. do not.

II. BACKGROUND

The distribution and service for the equipment were to be done through a service network consisting of a central facility and about 15 regional depots. The depots were geographically distributed across the United States, with each depot having a distinct service area and a set of service representatives. The service representatives were responsible for installing the units in patients' homes, providing service as needed, and disconnecting the units when no longer needed.

The field service system provided a variety of service transactions. When a new patient requests a piece of equipment, a service representative must install it, and will then instruct the patient about the care and use of the unit. The installation time is normally about four hours on site. Once the patient no longer needs the unit, the service representative must disconnect the unit, which requires about two hours on

November 2002. This work was supported in part by the MIT Leaders for Manufacturing Program and the Singapore-MIT Alliance.

Stephen C. Graves is with the Sloan School of Management and the Engineering System Division at MIT, Cambridge MA 02139 USA (email: sgraves@mit.edu).

site. Before using this unit again for a new patient, the unit is returned to a central facility for refurbishment to ensure that it is sufficiently clean and working properly.

Sometimes a unit will not function properly at a patient's home. In such cases the service representative returns to the patient's home to try to fix the unit; if it were not possible to repair the unit on site, then the service representative will swap the failed unit with a good unit. The time to diagnose and swap a unit is three hours. The failure rate for each unit (i.e., the unit cannot be repaired on site by the service rep) is about once every two years. The service representative returns each failed unit to the central facility for repair.

After one year of continuous use, a unit must undergo a mandatory preventative maintenance (PM); this service is performed at the central facility, and is similar to the refurbishment that is done on each disconnected unit. In other words, if a unit has been at a patient's home for a full year, and has not failed during that year, then the service rep swaps the unit with a new or refurbished unit, and returns the removed unit to the central facility for preventative maintenance.

Each regional depot batch ships its units in need of repair, refurbishment and preventative maintenance. That is, the depot will accumulate all of the returned units and send them in one load to the central facility to save transportation costs. The policy is expressed in terms of the accumulation time, e.g., an accumulation time of 14 days means that shipments are made once every 14 days to the central facility.

Each regional depot controls a base stock of units in order to be able to meet customer demand both for new service as well as for the replacements (swaps) described above. The central facility replenishes the inventory at each regional depot on a one-for-one basis. Whenever a unit is returned to the central facility from the regional depot, the central facility is notified and a replacement unit is shipped from the central facility to the regional depot. These replacement units are typically not new units, but units that have been refurbished or repaired.

In order to provide service to the regional depots, the central facility maintains an inventory of spares. Shipments to the regional depots are made out of this inventory. Furthermore, units that have been refurbished or repaired by the central facility are put into this inventory and held until needed by one of the regional depots.

The primary issue at hand is how much inventory is needed in the system to provide a high level of service to the customers at the least cost. A secondary issue is to model the workload for the service representatives at each depot, so as to plan staff levels. Costs include inventory holding, transportation and the cost for the service representatives.

III. ASSUMPTIONS

In order to develop a model to examine the inventory requirements and service transactions for this system, we state a series of assumptions.

1. We assume that the size of the system is stationary; that

is, we assume on average, the number of new demands equals the number of disconnects or terminations. Nevertheless, there is randomness in the arrival of new customers, and in the occurrences of failures and disconnects.

2. We assume that new demands arrive to the system according to a Poisson process.

3. We assume that both the length of time until a failure and the usage duration (time until disconnect) have exponential distributions. One consequence of this assumption is that the remaining lifetime (either time until failure or time until disconnect) does not depend on the current age of the unit, namely the time since last failure or time since installation.

4. Preventative maintenance is done on a regular basis, once a year.

5. The central facility does all repair, refurbishment and preventative maintenance. For the initial model development we assume that the central facility does not maintain any safety stock of spare units; hence, each depot is returned the exact same unit that it had sent for repair, refurbishment or preventative maintenance.

Later in the paper we discuss the extension in which the central facility holds inventory. In this case, the replenishment time seen by each depot will be reduced.

6. We assume the transit time to (from) the central facility from (to) each depot to be deterministic; the time to the central facility from a depot need not be the same as the time from the central facility to the depot. This is a simplification, especially for the transit time from the depot to the central facility. The units to be returned to the central facility will experience a random delay at the depot, while a batch is being accumulated to ship to the central facility. We discuss later in the paper how to extend the analysis to stochastic transit times.

7. We assume the service time at the central facility to be deterministic and the same for all units and for all types of service. That is, the time to repair or refurbish a unit or to perform a preventative maintenance is the same for all units.

Discussion of Assumptions: The first two assumptions were viewed as reasonable within the context. The demand process would be for single units to be installed in individual homes, and the assumption of a Poisson process was most natural. As this work examined the deployment of a new product, clearly there would be a transitional period in which the system was growing; that is, more units would be installed than would be removed for some period of time. However, we decided to ignore this growth phase and focus on the equilibrium phase when the removal rate matched the installation rate for new patients. This was acceptable given our purpose to determine the inventory requirements for the system and to show how these requirements depend on system design and operating parameters.

The third assumption of exponential times is certainly driven by a desire for analytical tractability. But at the time of the study there were no histories on the failure rates for the units, or on the time until a patient disconnects and returns the

unit. In the absence of data, we opted to start with this common assumption.

The fourth and fifth assumptions were policy decisions. One purpose of our work was to explore the implications from these decisions.

We can relax the sixth and seventh assumptions to permit random times for transit and service; the critical assumption is that units in the replenishment pipeline do not cross (i.e., the order in which units are sent to the central facility is preserved.). The assumption that the service times for repair, refurbishment and PM are the same is a simplifying assumption. We expect there would be some variation depending on the type of service, although the service times are all of the same order of magnitude; to distinguish the service times in the model requires keeping track of the different type of service classes, which adds another level of complexity. We decided to start simple.

IV. DEPOT SERVICE TRANSACTIONS

To build the system model, we first examine a single depot, and model the field service transactions. We base the analysis on the observation that every installation of a unit will eventually result in a removal of the unit; that is, there is a balance relationship for the service transactions. There are three types of removal transactions: removal due to service termination; removal due to unit failure; and removal due to preventative maintenance.

Each unit that is installed is eventually removed by one of these transactions. We can determine the probability for each type of removal transactions. Let μ denote the termination or disconnect rate for each unit, and ρ denote the failure rate for each unit.

A unit is removed for maintenance if it operates for one year. This happens if the unit survives a year without a failure and without a service termination; the probability of this is:

$$e^{-(\mu+\rho)} \quad (1)$$

A unit that fails is removed for repair. The probability that an installation results in a failure is the probability that the unit doesn't survive a year, times the conditional probability that the unit fails given that it either has failed or is terminated:

$$\frac{\rho}{\mu+\rho} \left(1 - e^{-(\mu+\rho)} \right) \quad (2)$$

When a household terminates a unit, the unit is removed. For each installation, the probability that it is removed as the result of a termination is the probability that the unit does not survive a year, times the conditional probability that the unit terminates given that it either failed or is terminated:

$$\frac{\mu}{\mu+\rho} \left(1 - e^{-(\mu+\rho)} \right) \quad (3)$$

We next characterize the number of transactions that a depot performs. We define the following variables:

- I number of installations per year
- P number of preventative maintenance requests served per year
- D number of terminations or disconnects performed per year
- Rp number of repairs performed per year

Every installation results in a repair if it fails within a year, or a disconnect if it is terminated within a year, or a preventative maintenance if it survives one year of usage. Thus, under the assumption that the system is in steady state, we have that the expected number of installations equates to the expected number of service transactions:

$$E[I] = E[P] + E[D] + [Rp]. \quad (4)$$

From the probabilities for the removal transactions we have

$$E[P] = E[I] e^{-(\mu+\rho)} \quad (5)$$

$$E[D] = E[I] \frac{\mu}{\mu+\rho} \left(1 - e^{-(\mu+\rho)} \right) \quad (6)$$

$$E[Rp] = E[I] \frac{\rho}{\mu+\rho} \left(1 - e^{-(\mu+\rho)} \right) \quad (7)$$

If we assume that the field service system, as a whole, is neither expanding nor contracting, then we should have that the new installations should match the number of disconnects in expectation:

$$\lambda = E[D] \quad (8)$$

where λ denotes the annual rate of new installations at the depot. From (5) - (8), we then get the steady state values for the key service events:

$$E[I] = \lambda \left[\frac{\mu}{\mu+\rho} \left(1 - e^{-(\mu+\rho)} \right) \right]^{-1} \quad (9)$$

$$E[P] = \lambda e^{-(\mu+\rho)} \left[\frac{\mu}{\mu+\rho} \left(1 - e^{-(\mu+\rho)} \right) \right]^{-1} \quad (10)$$

$$E[D] = \lambda \quad (11)$$

$$E[Rp] = \lambda \frac{\rho}{\mu} \quad (12)$$

From these relations, we can estimate the workload for the service representatives for each depot, and determine the transportation costs between the central facility and each depot.

V. DEPOT INVENTORY REQUIREMENTS

The prior analysis characterizes the number of transactions at each depot. In this section we will determine how large of a stock of units each depot needs to carry.

Suppose each depot maintains a base stock of S units. Each unit in the base stock can be in one of three places: installed at a household; as a spare at the depot; or in the service cycle. We say that a unit is in the service cycle if it is in transit between the depot and the central facility, or in service at the central facility, undergoing repair, refurbishment or preventative maintenance.

We let $N(t)$ denote the number of units installed at households at time t , and let $R(t)$ be the number of units in the service cycle at time t . Both $N(t)$ and $R(t)$ are random processes.

By assumption, if a depot has a base stock of S units, then there are $S - N(t) - R(t)$ units in spares at the depot. We desire to set S so that there is a very low probability of there being no spare units at the depot when a unit is needed for installation. Thus, with high probability a service representative should be able to complete a transaction requiring an installation, whether it is for a new service or as a replacement due to a failure or PM.

At a depot, for a given value of S , both $N(t)$ and $R(t)$ depend on S , since S effectively constrains how many units can be installed or in the service cycle at any point in time, i.e., $N(t) + R(t) \leq S$. However, to determine S , we will ignore this dependence and will model $N(t)$ and $R(t)$ as if they were not constrained by S . In effect we will model these two variables, assuming that there is always ample stock available at the depot. Then we will propose to set S , so that the assumption of ample stock is approximately true. For instance, we will set S so that there is a high fill rate; that is, the probability that $N(t) + R(t) \leq S$ is high, say 0.98. [An alternative interpretation is to assume that we can delay or backorder requests for installations.]

Since we assume that new installations occur as a Poisson process at rate λ , and that the mean time until a household disconnects its service is $1/\mu$, then we can model $N(t)$ as the occupancy level of an $M/G/\infty$ system. By the assumption of ample stock, a new installation never waits for a unit to become available. In steady state, then, $N(t)$ is a Poisson random variable with mean λ/μ . [Note that we do not require the assumption of an exponential distribution of time until termination here, but we did use this assumption for (1) – (3) and the development thereafter.]

To model $R(t)$, we need to describe the process for repair or refurbishment or preventative maintenance. A unit that has been removed from the field is shipped from the depot back to the central service facility where it undergoes servicing. Upon completion of its service, a unit is returned to the depot. We assume that the total time to service a unit, including transportation time, is a constant τ . We term τ to be the length of the service cycle for a unit. We will consider later the cases when the service cycle is stochastic, such as would occur when the central facility holds a safety stock.

We will use the following key property for determining S :

$$N(t + \tau) + R(t + \tau) = N(t) + I(t, t + \tau) \quad (13)$$

where $I(t, t + \tau)$ denotes the number of installations performed over the time interval $(t, t + \tau]$.

To show (13), consider each unit that is in use at time t , i.e., part of $N(t)$. At time $t + \tau$, each unit will either still be in use or have been removed for maintenance, repair, or a disconnection. If it has been removed, it will enter the service cycle. Since the service cycle is a fixed length τ , a unit removed in the time interval $(t, t + \tau]$ is still in the service cycle at time $t + \tau$. Thus, we have that each installed unit at time t , $N(t)$, is either part of $N(t + \tau)$ or $R(t + \tau)$ at time $t + \tau$. And each unit that is installed during the time interval $(t, t + \tau]$ either is in use at time $t + \tau$ or has entered the service cycle by time $t + \tau$ and is still there. Thus we have the inequality:

$$N(t + \tau) + R(t + \tau) \geq N(t) + I(t, t + \tau).$$

Furthermore, each unit in use at time $t + \tau$ was either in use at time t and has not been removed in the time interval $(t, t + \tau]$ or was installed in the time interval $(t, t + \tau]$ and has not been removed by time $t + \tau$. Similarly, each unit in the repair or refurbishment cycle at time $t + \tau$ was either in use at time t and has been removed in the time interval $(t, t + \tau]$ or was installed in the time interval $(t, t + \tau]$ and was subsequently removed prior to time $t + \tau$. Thus, we have the inequality:

$$N(t + \tau) + R(t + \tau) \leq N(t) + I(t, t + \tau),$$

which completes the argument for (13). Thus, at time $t + \tau$, the number of units in use or in the service cycle equals the number of units in use at time t , plus all installations within the time window $(t, t + \tau]$.

To use (13) we need to model $I(t, t + \tau)$. The installations over the interval $(t, t + \tau]$ consist of new installations plus replacements due to unit failures or PM's.

The number of new installations over a time window of length τ is a Poisson random variable with mean $\lambda \tau$.

The number of replacements over $(t, t + \tau]$ is a superposition of the failure processes for the installed units, truncated by the annual PM refurbishments, and does not have a simple characterization, as it depends on the value of $N(t)$ and new installations within the time window. *As an approximation*, we first ignore the possibility that a new installation in $(t, t + \tau]$ could also result in a subsequent replacement in this interval; this is reasonable as the length of the transit time is on the

order of a few weeks, and the time between removals is on the order of several months. Second, we assume that the expected number of replacements over $(t, t + \tau]$ is proportional to $N(t)$. Thus, we approximate the expected number of replacements over the interval $(t, t + \tau]$ by $p \tau N(t)$, where p is an unknown constant denoting the expected number of replacements per installed unit per year.

By combining the two types of installations, we have that the expected number of installations over the interval $(t, t + \tau]$ is given by:

$$\lambda \tau + p \tau N(t).$$

Finally we *approximate* $I(t, t + \tau)$ to be a Poisson random variable with the above mean.

For some justification, we note that the new installation process is Poisson, and the replacement process is a superposition of the replacement processes for the installed base of $N(t)$ units.

To determine p , we note that the above expression implies that the expected installation rate is

$$E[I] = \lambda + p E[N(t)] = \lambda + p \frac{\lambda}{\mu}.$$

Then we can equate this to the expected installation rate, given by (9), to find:

$$p = \frac{\rho + \mu e^{-(\mu + \rho)}}{1 - e^{-(\mu + \rho)}}. \quad (14)$$

For a given value of $N(t) = n$, we thus assume that $I(t, t + \tau)$ is Poisson with mean

$$\tau \left[\lambda + n \frac{\rho + \mu e^{-(\mu + \rho)}}{1 - e^{-(\mu + \rho)}} \right].$$

We now use this result to characterize $N(t) + I(t, t + \tau)$. Conditioned on $N(t) = n$, we find the mean and variance:

$$E[N(t) + I(t, t + \tau) | N(t) = n] = n + \lambda \tau + n p \tau$$

$$\text{Var}[N(t) + I(t, t + \tau) | N(t) = n] = \lambda \tau + n p \tau$$

where p is given by (14). Now using the fact that $N(t)$ is Poisson with mean λ/μ , we find the unconditioned mean and variance:

$$E[N(t) + I(t, t + \tau)] = \frac{\lambda}{\mu} + \lambda \tau + \frac{\lambda p \tau}{\mu} \quad (15)$$

$$\text{Var}[N(t) + I(t, t + \tau)] = \frac{\lambda}{\mu} (1 + p \tau)^2 + \lambda \tau + \frac{\lambda p \tau}{\mu}. \quad (16)$$

But, from (13), we observe that (15) and (16) provides the

steady-state mean and variance for $N(t + \tau) + R(t + \tau)$, or equivalently $N(t) + R(t)$, since the choice of t is arbitrary. We observe that the variance is larger than the mean, where the difference depends upon the value of $p \tau$, the expected number of replacements per installed unit over the service cycle. For the case at hand, $p \tau$ will be fairly small (always less than 0.2) and hence the variance will be modestly larger than mean. Indeed, we have found that the distribution of $N(t) + R(t)$ is very Poisson-like, albeit more variable.

We now propose to set S for each depot such that $S > N(t) + R(t)$ with high probability. Provided that the expectation of $N(t) + R(t)$ is of modest size, e. g., at least 30, the distribution of $N(t) + R(t)$ can be reasonably approximated by a normal distribution. From the mean and variance of $N(t) + R(t)$ given by (15)-(16), we can then easily set S to achieve a desired coverage level or fill rate.

VI. EXAMPLE

To illustrate the model, we provide an example with representative numbers from the application. In Table 1, we provide the inputs. We illustrate the model with three possible rates for new installations, as might be seen by a regional depot. The firm estimated that there was a probability of 0.6 that a patient would disconnect within one year; from this and the assumption of exponentially-distributed termination times, we find the disconnect rate. Similarly, the firm estimated that there was a probability of 0.5 that a unit would fail in any given year; with the assumption of exponentially-distributed failures, we find the failure rate. For the example we assumed a service cycle of twenty-one business days (five days at central facility, plus sixteen days for round-trip transit plus accumulation time at the depot), with 250 days per year.

Annual new installation rate λ	Disconnect rate μ	Failure rate ρ	Service cycle time τ
75, 90, 135	0.916	0.693	0.084

Table 1: Input parameters for example

In Table 2, we present the outputs. The top four rows of the table give the expected number of transactions per year, as given by (9) – (12); one can translate these outputs into the workload for service representatives, and hence find the number required for the depot. We then compute the expected number of units in use, and the expectation and variance of $N(t) + I(t, t + \tau)$, as given by (15) – (16). By the equivalence (13), we then have the first two moments for $N(t) + R(t)$, the number of units in use or in the service cycle. We can then use these two moments to prescribe a base stock level to satisfy a service target. In the table we report the base stock needed for a 98% fill rate ($z = 2.05$), under the assumption of the normal approximation.

	$\lambda = 75$	$\lambda = 90$	$\lambda = 135$
$E[I]$	164.7	197.6	296.4
$E[P]$	32.9	39.5	59.3

E[D]	75.0	90.0	135.0
E[Rp]	56.7	68.1	102.1
E[N(t)]	81.9	98.2	147.3
E[N(t) + I(t, t+τ)]	95.7	114.8	172.2
Var[N(t) + I(t, t+τ)]	111.4	133.7	200.6
Base Stock S	117.3	138.5	201.3

Table 2: Outputs for example

VII. STOCHASTIC LEAD-TIMES

We now assume that the time to replenish the regional depot from the central facility is stochastic. In particular, we continue to assume that the central facility does not maintain any safety stock of spare units; hence, each unit that is sent for repair, refurbishment or preventative maintenance returns to the depot from which its service cycle originated. But now we will assume that the service cycle, which is the sum of the accumulation time at the depot, the round-trip transit time from the depot to the central facility and back, plus the service time at the central facility, is a random variable.

We model the lead time as being generated by an exogenous, sequential supply system, in the terminology of Zipkin (2000) In particular, we assume there is a stochastic process $L(t)$, corresponding to the virtual lead time at time t . That is, $L(t)$ is the lead time for any service cycle initiated at time t . We assume that $L(t)$ is independent of $N(t)$, the number of units in use; consequently, $L(t)$ is independent of the number of units in service at time t . We also assume that $L(t)$ is independent of the process for new installations. Finally, we assume that $t + L(t)$ is non-decreasing, and thus, that there is no order-crossing.

The development of the inventory model for the case of a stochastic lead-time parallels that for the case of a deterministic lead-time. We can rewrite the equivalence (13) as:

$$N(t + L(t)) + R(t + L(t)) = N(t) + I(t, t + L(t)). \quad (17)$$

To characterize the first two moments for the right-hand-side of (17), we first condition on both $N(t) = n$ and $L(t) = \tau$:

$$\begin{aligned} E[N(t) + I(t, t + \tau) | N(t) = n, L(t) = \tau] \\ = n + \lambda\tau + np\tau \end{aligned}$$

$$Var[N(t) + I(t, t + \tau) | N(t) = n, L(t) = \tau] = \lambda\tau + np\tau$$

where p is given by (14). We now use the fact that $N(t)$ is Poisson with mean λ/μ , and that $L(t)$ and $N(t)$ are independent to find the unconditioned mean and variance:

$$E[N(t) + I(t, t + L(t))] = \frac{\lambda}{\mu} + E[L] \left(\lambda + \frac{\lambda p}{\mu} \right) \quad (18)$$

$$\begin{aligned} Var[N(t) + I(t, t + L(t))] = \\ \frac{\lambda}{\mu} (1 + pE[L])^2 + E[L] \left(\lambda + \frac{\lambda p}{\mu} \right) \\ + Var[L] \left(\left(\lambda + \frac{\lambda p}{\mu} \right)^2 + \frac{\lambda p^2}{\mu} \right). \end{aligned} \quad (19)$$

As before, we note from (17) that (18) and (19) provide the steady-state mean and variance for $N(t+L(t)) + R(t+L(t))$. We observe that the variance is larger than the mean, where the difference depends directly upon the variance of the lead-time. We use (18)-(19) to set the base stock S for each depot. For instance, we can approximate the distribution of $N(t) + R(t)$ by a normal distribution with mean and variance given by (18) and (19). We then set the base stock S such that the probability that $S > N(t) + R(t)$ meets the desired coverage level.

VIII. SAFETY STOCK AT CENTRAL FACILITY

We now describe how to extend the analysis to permit the central facility to hold a safety stock. As such, we develop an approximate multi-echelon model in the spirit of the extensive literature in this area (see Axsater 1993, and Federgruen 1993). The central facility maintains a base stock of spare units, call it S_0 , in order to provide service to the regional depots. The central facility replenishes the inventory at each regional depot on a one-for-one basis. Whenever a regional depot returns a unit to the central facility, it will notify the central facility and the central facility will ship a replacement to the regional depot. When the returned unit reaches the central facility, the unit enters the service process, for repair, refurbishment, or preventative maintenance. Upon completion of the service, the unit becomes part of the available stock for the central facility.

The inventory in the base stock S_0 can be in one of three places – in-transit from a regional depot back to the central facility, in service at the central facility, or in stock at the central facility. When there is a replenishment request, the central facility will ship a unit immediately, provided it has one in stock. When it does not have a unit in stock, the central facility delays shipping (i.e., backorders) the replenishment request until a unit becomes available.

When the central facility holds an inventory, the replenishment lead-time seen by a regional depot is uncertain, as it depends on whether or not there is stock on hand when a request is made. The actual distribution of the depot lead-time is a function of the base stock at the central facility. Once we characterize the lead-time, we can find the inventory required at the depot, as shown in the prior section, to meet a specified service target. Thus, for a given base-stock level at the central facility, we can find the base-stock levels at the regional depots. We then search over possible values for the central facility's base stock to find the setting that minimizes the total

inventory in the system.

Let $L_i(t)$ denotes the virtual lead-time at time t for depot i . For notational convenience, we'll drop the time index t . We can express the lead-time as follows:

$$L_i = \tau_i + \Delta_i(S_0)$$

where τ_i is the one-way transit time from the central facility to depot i , and $\Delta_i(S_0)$ is the delay at the central facility in filling a replenishment request, which is a function of the base stock. We will assume for ease of presentation that the transit time is a known constant; however, the delay is a random variable. Thus,

$$E[L_i] = \tau_i + E[\Delta_i(S_0)] \quad (20)$$

$$Var[L_i] = Var[\Delta_i(S_0)]. \quad (21)$$

Thus, if we can model the first two moments for the delay at the central facility, then we use (20) – (21) in (18) – (19) to find the base stock required at each depot.

There are various ways to model the delay at the central facility, depending upon what is assumed about the demand (return) process and about the service processes. (see Axsater 1993 and Federgruen 1993) We illustrate here how we approached this in the context of the application.

We assume that the aggregate return process seen by the central facility could be modeled as a Poisson process with rate Λ given by:

$$\Lambda = \sum_i \left(\lambda_i + \frac{\lambda_i}{\mu} p \right)$$

where p is given by (14), and we assume that each depot experiences the same disconnect rate μ and failure rate ρ , and has a rate λ_i for new installations.

We assume that the time to return a unit from the depot, plus the time to service it is the same for all depots and all types of service, and is given by a constant θ . As justification for this assumption, we found that the depots had comparable return times, since much of this was the time to accumulate a shipping load. And the time to service a unit (including queuing) once it arrived would be roughly the same for all units. However, the total time for return and service was stochastic, whereas we assume it is a constant for the purposes of modeling the inventory requirements.

Suppose we number the return events according to the order in which they are initiated. Then, under a first-in, first-out dispatch rule, the central facility will use the n^{th} unit that is returned to fill the replenishment request associated with the $S_0 + n^{\text{th}}$ return. There is a delay whenever the $S_0 + n^{\text{th}}$ return occurs before the n^{th} unit can be returned and serviced; that is, there is a delay if the time between the occurrence of n^{th} return and that for the $S_0 + n^{\text{th}}$ return is less than the time to return and service the n^{th} unit, namely θ . Indeed we can express the delay as

$$\Delta_i(S_0) = [\theta - T_0]^+ \quad (22)$$

where T_0 is a random variable equal to the interarrival time for the occurrence of S_0 returns to the central facility.

With the assumption that the return process is a Poisson process with rate Λ , then T_0 has a gamma distribution with parameters S_0 and Λ ; thus, T_0 has mean $E[T_0] = S_0/\Lambda$ and variance $Var[T_0] = S_0/\Lambda^2$. For a candidate base stock S_0 , we can use (22) to find the moments for the delay, which we input to (20) – (21) to find the moments for the depot lead time. We can then use (18) – (19) to determine the base stock at each depot to satisfy a service target. We then search over possible values for the base stock S_0 for the central facility to find the solution that minimizes the total system inventory.

ACKNOWLEDGMENT

S. C. Graves thanks Curt Coopridier for his assistance in implementing the model, conducting computational experiments, and in transferring the results to the industrial client.

REFERENCES

- [1] Axsater, S. (1993), "Continuous Review Policies for Multi-Level Inventory Systems with Stochastic Demand," Chapter 4 in Graves, S., A. Rinnooy Kan and P. Zipkin (eds.) *Logistics of Production and Inventory*, Handbooks in Operations Research and Management Science, Vol. 4, Elsevier (North-Holland), Amsterdam.
- [2] Federgruen, A. (1993), "Centralized Planning Models for Multi-Echelon Inventory Systems under Uncertainty," Chapter 3 in Graves, S., A. Rinnooy Kan and P. Zipkin (eds.) *Logistics of Production and Inventory*, Handbooks in Operations Research and Management Science, Vol. 4, Elsevier (North-Holland), Amsterdam.
- [3] Fleischmann, M., J. M. Bloemhof-Ruwaard, R. Dekker, E. van der Laan, J. van Nunen, and L. Van Wassenhove (1997), "Quantitative Models for Reverse Logistics: A Review," *European Journal of Operational Research*, Vol. 103, pp. 1 – 17.
- [4] Nahmias, S. (1981), "Managing Repairable Item Inventory Systems: A Review," Chapter 13 in Schwarz, L. (ed.), *Multi-Level Production/Inventory Control Systems: Theory and Practice*, North-Holland, Amsterdam, pp. 253 – 277.
- [5] Thierry, M., M. Salomon, J. van Nunen, and L. Van Wassenhove (1995), "Strategic Issues in Product Recovery Management," *California Management Review*, Vol. 37, No. 2, pp. 114 – 135.
- [6] Toktay, B. L., L. M. Wein, and S. A. Zenios (2000), "Inventory Management of Remanufacturable Products," *Management Science*, Vol. 46, No. 11, pp. 1412 – 1426.
- [7] Zipkin, P. H., (2000), *Foundations of Inventory Management*, McGraw Hill, Boston.