Dynamic Inventory Allocation under Imperfect Information and Capacitated Supply

Maher Lahmar • Sylvana S. Saudale

Department of Industrial Engineering, University of Houston, Houston, TX 77204, USA
mlahmar@uh.edu • ssaudale@uh.edu

We consider a distribution system composed of a limited-capacity supplier that serves a set of customers exhibiting stochastic demand. The supplier information is limited to the inventory level observed at the currently visited customer and the previously observed inventory levels at all other customers. The supplier follows a fixed routing and sequentially decides on how much stock to allocate to each customer. We formulate the problem as a Markov Decision Process and discuss optimal replenishment policies. We provide managerial insights on the value of dynamic decision making and information availability in the replenishment of such distribution systems. We show that optimal replenishment policies do not necessarily provide all customers with "fair" service levels. We also show that investing in state-of-the-art information systems may not be justified for systems with very tight or ample supplier capacity. We benchmark the performance of myopic solution approaches to that of the optimal solution procedure. Finally, we develop and test the quality of efficient decomposition-based heuristics that can solve for the homogenous customers case.

Keyword: Inventory allocation, Vendor Managed Inventory (VMI), Markov Decision Process (MDP).

1. Introduction

Vendor Managed Inventory (VMI) has emerged as an effective strategy to coordinate supply chain activities and enhance the system-wide effectiveness. VMI allows the vendor to exploit acquired customer information such as demand forecasts, sales and inventory data to design effective inventory replenishment policies. Stories witnessing the successful implementations of such strategies in a wide range of industries have been reported in the literature ([11], [25], [26]). The Vendor Managed Distribution (VMD) extends the VMI capability by integrating the inventory management and the distribution process. To ensure the success of such coordination strategies, information sharing has to be coupled with the capability to process this information in a timely manner so that distribution routing and stock allocation decisions are efficiently and dynamically undertaken. Equally important, customers have to gain extra benefits that outweigh the risks of sharing their inventory information with a third party and these benefits should be "fairly" distributed among all customers. Since customers often exhibit uncertain demand and have different cost structures, the supplier has to make effective use of available information to allocate its stock efficiently and, if necessary, to ration the supply among customers.

This study is motivated by the need to identify effective policies to manage short-haul distribution systems and determine the factors under which the implementation of such coordination strategies are desirable. More specifically, we address the problem faced by a supplier that manages the distribution operations of a single item to multiple customers through regular replenishments following a fixed routing. The supplier does not necessarily have complete information on demand occurrences and current inventory levels at each customer site. While satisfying the demand of early visited customers can result in the starvation of the remaining unvisited customers, allocating more stock to meet the demand of the late-in-sequence customers can result in excess inventory at the end of the replenishment cycle. In addition, the replenishment decisions that do not account for the system status and the costs incurred in future cycles would only result in myopic policies. Hence the effective deployment of such coordinated policies can reduce the system-wide cost and provide enterprises with key competitive advantages.

Such practice is common in short-haul logistic systems where a single supplier serves a set of customers with relatively low usage, generic (not customer-specific) items or when customers expect frequent replenishments in small quantities from a single supplier. In such systems, the distribution routing is pre-selected, the investment in state-of-the-art communication technology may not be justified and timely information may not be always available. Examples of such distribution systems include the refilling of vending machines, the supply of generic medical supplies to hospitals, the replenishment of convenience stores, and the delivery of water bottles to homes and offices.

In this study, we aim at answering the question of how much stock to allocate to each customer such that the expected system-wide cost is minimized. The answer to this question will help us (i) identify the characteristics of optimal inventory allocation policies, (ii) investigate the impact of such policies on customers' service levels, and (iii) measure the value of inventory information and dynamic decision making on the distribution system performance. We also benchmark the performance of myopic solution approaches against optimal solutions. In addition, we build on these findings to develop more efficient approaches to solve for computationally prohibitive medium and large size problems.

The rest of this paper is organized as follows. In the next section, we provide a brief literature review on related problems. In section 3, we formulate the problem as a finite horizon Markov decision process and present structural properties of the optimal inventory replenishment policies. In section 4, we analyze the impact of optimal replenishment policies on service levels and provide managerial insights on the management of such inventory/distribution systems. In section 5, we test the quality of well-known myopic solutions and that of two decomposition-based heuristics. Finally, we summarize our findings and present potential future research directions.

2. Related Literature

This study is a part of the general effort that aims at enhancing the modeling and practice of VMI coordination strategies. For recent articles addressing a variety of issues related to VMI coordination, the reader may refer to [16], [8], [28], [18], [5], and [6]. Below, we review several research streams that are relevant to our work.

Our research focuses on a special case of the general (IRP), a problem that simultaneously solves for the inventory replenishment decisions as well as the vehicle routing decisions [7]. Given the complexity of such class of problems, most IRP literature focuses on direct shipping policies ([3], [10], [17], [19]). It is until recently that researchers started to show interest in the IRP with stochastic demand, consolidated shipments and general routing policies. For instance, Adelman [1] addresses the problem of inventory replenishment of a set of products for a selected number of locations with limited storage capacity and no stockouts. Kleywegt et al. [20] extend their stock rationing solution procedure presented in [19] to solve for the case of multiple deliveries per trip. Balun [2] models the IRP as a discrete-time Markov Decision Process where the supplier receives updated information on all inventory levels upon arrival at each customer/depot site. At each period, the supplier decides simultaneously on how much to drop at each customer and which customer to visit next. Berman and Larson [4] consider the problem of dynamic delivery of gas to a set of customers such that the total expected discounted cost is minimized. Their objective function incorporates the per-trip fixed cost, the distance-related charges, the earliness and lateness costs, and cost of product shortfall within the objective function. Similar to our work, the amounts of product delivered to each customer are assumed to be unknown until the driver is at the customer location, at which point the customer is either restocked to capacity or left with some residual empty capacity. Under specific settings, they conclude that the fill-fill-dump (FFD) policy is an optimal allocation policy whenever the customers are served according to the descending order of the normalized fixed costs per customer. Due to the complexity of most addressed problems, solving for such problems requires the development of complex models and involved set of approximations.

Since our problem includes dynamic allocation of capacitated resources to several customers, it can also be seen as a multi-period, dynamic, stochastic knapsack problem. The dynamic and stochastic knapsack problem has been addressed in several studies ([21], [22], [23]). Papastavrou et al. [23] address a discrete-time version of the problem where an object possibly with random weight and reward arrive at regular periods. The decision is whether to accept or reject the arriving object given that the weights and rewards associated with each object are only known after arrival. Kleywegt et al. [21] consider the continuous-time version of the problem where objects arrive according to a Poisson process and the admission of objects into the knapsack can be halted at any time. After a preset deadline, the unused capacity is penalized and a waiting time cost is associated before the stopping time. They also analyze the special case with equal size/weight objects for both the infinite and finite horizon problems and identify the optimal admission stopping time. They show that the optimal acceptance policy is given by a simple threshold rule, similar to the result found in [23]. In a following work, Kleywegt et al. [22] extend their study to include objects with random sizes and weights.

The problem we address can also be seen as a version of the capacitated stochastic inventory problem. Under ample capacity, the basic stochastic inventory problem can be optimally solved using a base-stock policy; however, under limited capacity, finding an optimal policy can be a complex task even for small-size problems. Fedregruen and Zipkin ([14], [15]) show that a modified base-stock policy is optimal for the periodic review single-customer capacitated case with backorders. Chen [9] addresses a similar problem with fixed set-up cost and shows that the optimal inventory policy has an X-Y band structure and is not necessarily a modified-base stock policy. Evans [13] also uses a modified base-stock policy to solve for the finite horizon two-customer problem with lost sales and limited supplier capacity. In a more recent work, DeCroix and Arreola Risa [12] extend Evans' results to the infinite horizon backordering case and identify the optimal policy characteristics for the homogenous products problem.

The problem addressed in this article is different from most literature in that it combines three main aspects: (i) customers exhibit stochastic demand and are served in a fixed sequence, (ii) the vendor has a limited fixed capacity at the beginning of each replenishment cycle, and (iii) customers share inventory information with the vendor upon contact only. In the rest of this work, we build on the current literature to formulate our problem, design efficient solution procedures, and provide related managerial insights.

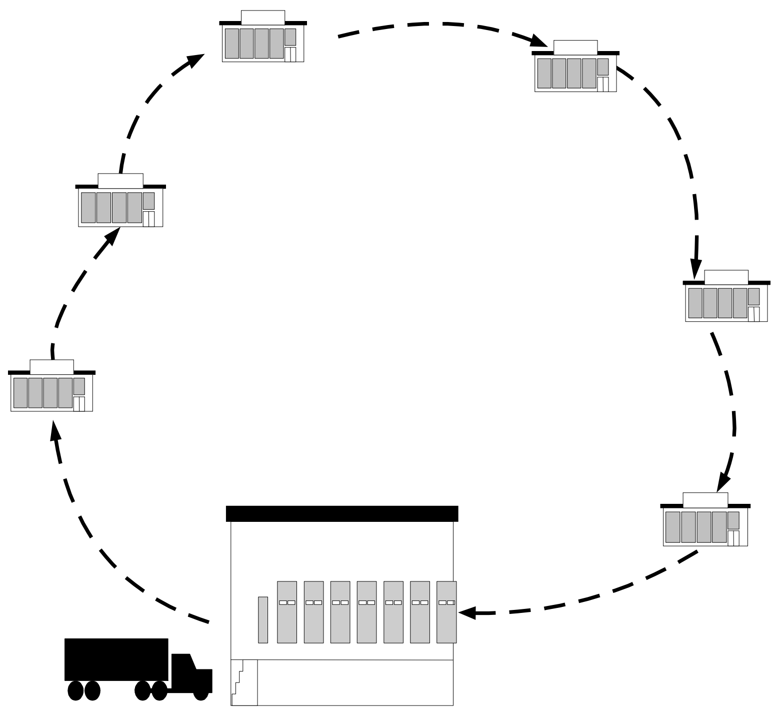

Figure 1: Representation of the inventory/distribution system

3. Problem Description

We consider the problem of a distribution system where a supplier periodically dispatches a truck with capacity C to serve a set of N customers. Customers experience independent stochastic demand with known distributions. The supplier is assumed not to have complete information on the current inventory levels of all customers nor the demand occurrences. Each customer i is assumed to have a limited storage capacity C_i. Upon arrival at each customer location, the supplier observes the current inventory level and decides on how much to drop at the site of that customer. The information regarding the level of inventory observed at each customer during the previous replenishment cycle is carried into the

N:	number of customers.
C:	truck capacity.
C_i:	storage capacity at customer i.
h_i:	unit inventory holding cost at customer i.
b_i:	unit shortage cost at customer i.
p:	unit penalty cost at the end of the route.
x_i:	inventory level at customer i.
c:	actual amount of stock available in the truck.
D_i:	demand random variable at customer i.
λ_i:	average demand of customer i ($E[D_i]$).
$P_i(q)$:	probability that demand of customer i during a cycle is equal to q.
n:	number of replenishment cycle.

Table 1: Notation

next cycle. Since this information will be time-lagged by the time the supplier takes the replenishment decisions in the next cycle, we refer to this information as *imperfect*. Each Customer i incurs an inventory holding cost h_i for each unit kept in stock and a unit shortage cost b_i for stockouts. The supplier follows a fixed routing that includes all N customers (Figure 1). A penalty cost p is charged for each remaining unit in the truck at the end of the route. Unlike the traditional joint-replenishment case, the decision here is based on a trade-off between vendor stock availability and the demand fulfillment for the remaining unvisited customers. Without loss of generality, we assume that it takes one time period to move from one customer to another and that the time to fill in the truck from the depot is negligible. The demand distribution can be easily adjusted to model a system with different deterministic travel times between customers and to/from the depot. A complete list of the adopted notation is presented in Table 1.

3.1 Problem Formulation

To capture the uncertainty in customer demand and the dynamic nature of the replenishment decisions, we formulate our model as a finite horizon Markov Decision Process (MDP) [24]. We define the state of the system at time t when the truck visits customer i at cycle n as $s_t \equiv (c, i, x_{1,t}, \cdots, x_{N,t})$, where c stands for the amount of stock available in the truck, i for the customer currently visited by the supplier such that $t = nN + i$ and $x_{j,t}$ the inventory level last observed at customer j before or at time t. Note that we refer to inventory levels as the stock level last observed at the customer's site; in other words, $x_{i,t}$ stands for the inventory level at customer i seen upon arrival, while $x_{j,t}$ (for $j \neq i$) stands for the inventory level at customer j after the current cycle replenishment when $j < i$, and the inventory level at customer j after the previous cycle replenishment when $j > i$. Thus, the state space can be defined as

$$S = \mathcal{C} \times \mathcal{N} \times \mathcal{X} = \{0, \cdots, C\} \times \{1, \cdots, N\} \times \prod_{i=1}^{N} \{0, \cdots, C_i\}.$$

Upon arrival to customer i, the supplier decides on the amount of stock, denoted a, to be delivered to that customer. We assume that a is a positive integer that is constrained by the available stock and the customer capacity that can be defined on the set of actions

$$a \in A(s) = \{0, 1, \cdots, min(c, C_i - x_i)\}.$$

Given arrival of the supplier to customer i and given the inventory levels x_{jt} at period t, we can express the inventory levels at next period $t+1$ as

$$x_{j,t+1} = \begin{cases} [x_{j,t} - \xi_j]^+ & \text{for } j = mod(i, N) + 1, \\ x_{j,t} + a & \text{for } j = i, \\ x_{j,t} & \text{otherwise,} \end{cases}$$

where $x_{i,t}$ is the inventory level last observed by time t at customer i and ξ_i is the demand at customer i during one replenishment cycle (N time periods). For the sake of simplicity, we drop the time index and define $\mathbf{x} = (x_1, \cdots, x_N)$ to be the vector of inventory levels at each customer where the current system state is $s = (c, i, \mathbf{x}) \in S$.

The transition probabilities follow from the demand probabilities. More specifically, we denote $P_i(q)$ to be the demand probability of q units at customer i during a replenishment cycle. We let $s \equiv (c, i, \mathbf{x})$ and $s' \equiv (c', i', \mathbf{x}')$ where $\mathbf{x} = (x_1, \cdots, x_N)$ and $\mathbf{x}' = (x'_1, \cdots, x'_N)$, then the transition probabilities can be depicted as follows:

$$P(s'|s, a) = \begin{cases} P_{i+1}(x_{i+1} - x'_{i+1}) & \text{if} \quad x_{i+1} \geq x'_{i+1} > 0,\ c' = c - a,\ i' = i + 1,\ i \neq N, \\ & \qquad x'_i = x_i + a,\ x'_j = x_j, \quad \text{for } j < i,\ j > i + 1, \\ \sum_{j=x_{i+1}}^{\infty} P_{i+1}(j) & \text{if} \quad x'_{i+1} = 0,\ c' = c - a,\ i' = i + 1,\ i \neq N, \\ & \qquad x'_i = x_i + a,\ x'_j = x_j, \quad \text{for } j < i,\ j > i + 1, \\ P_1(x_1 - x'_1) & \text{if} \quad x_1 \geq x'_1 > 0,\ x'_N = x_N + a,\ c = C,\ i' = 1,\ i = N, \\ & \qquad x'_j = x_j, \quad \text{for } j = 2, \cdots, N - 1, \\ \sum_{j=x_1}^{\infty} P_1(j) & \text{if} \quad x'_1 = 0,\ x'_N = x_N + a,\ c = C,\ c' = 1,\ i = N, \\ & \qquad x'_j = x_j, \quad \text{for } j = 2, \cdots, N - 1, \\ 0 & \qquad \text{otherwise,} \end{cases}$$

The third and fourth set of transition probabilities are specific to the last customer in the sequence (i.e. customer N), while the first and second sets are specific to the rest of customers. The second and fourth sets of probabilities represent the cases where demand is lost while the first and third probabilities represent the cases where a customer is left with a non-zero inventory level.

We define the single period cost function $r_t(s,a)$ at time t as the expected holding and lost sales cost at the visited customer i defined as:

$$r_t(s,a) = \sum_{\xi=0}^{\infty} P_i(D_i = \xi)\Big(h_i[x_i - \xi + a]^+ + b_i[x_i - \xi + a]^-\Big),$$

for $i \neq N$, and

$$r_t(s,a) = \sum_{\xi=0}^{\infty} P_N(D_N = \xi)\Big(h_N[x_N - \xi + a]^+ + b_N[x_N - \xi + a]^-\Big) + p(c-a),$$

for $t = nN$, where $[x]^+ = max(0,x)$ and $[x]^- = max(0,-x)$. The terminal cost at the end of the planning horizon can be expressed as:

$$r_{TN+1}(s) = \sum_{i=1}^{N} d_i x_i,$$

where d_i is a cost incurred due to unsold items (difference between cost of item and its salvage value). We consider the objective of minimizing the finite horizon expected total cost expressed as

$$v^*(s_1) = \min E_{s_1}\Big[\sum_{t=1}^{TN+1} r_t(s_t,a_t)\Big],$$

where s_1 is the initial state at period 1 and π the followed replenishment policy. Since the transition probabilities are Markovian and the set of states and actions are finite, we can restrict our attention to identifying optimal policy π^* among the set of Deterministic Markovian policies (Π) that minimize $v^\pi(s_1)$ such that

$$v^\pi(s_1) = E^\pi_{s_1}\Big[\sum_{t=1}^{TN+1} r_t(s_t, a(s_t))\Big],$$

and therefore

$$v^*(s_1) = \min_{\pi\in\Pi} v^\pi(s_1),\ \forall\ s_1 \in S,\ \pi \in \Pi.$$

3.2 Structural Properties

In this section, we provide partial characterization of the optimal replenishment policies for inventory/distribution systems with identical customers (similar cost parameters and i.i.d.

demand distributions). In the Appendix, we provide the proofs of the structural properties.

Proposition 1: *For each state $s = (c, i, \mathbf{x}) \in S$, the optimal replenishment action $a^*(s) \in A(s)$ at customer i is non-decreasing in the available truck stock c.*

Proposition 1 states that the number of units allocated to each customer upon arrival tends to increase with the amount of stock available in the truck. Moreover, the optimal replenishment policy is of a base-stock nature for fixed inventory levels $\mathbf{x}$, where the stock at customer i is replenished up to $S^*(c^+, i, \mathbf{x}) \geq S^*(c^-, i, \mathbf{x})$ given that $c^+ \geq c^-$.

Proposition 2: *For each state $s = (c, i, \mathbf{x}) \in S$, the optimal replenishment action $a^*(s) \in A(s)$ at customer i is non-increasing in inventory level x_i of customer i.*

Proposition 2 states that the number of units allocated to each customer upon arrival to its site tends to decrease with the amount of stock available at the customer site. This implies that the optimal replenishment policy is of a base-stock nature where the stock at customer i is replenished up to $S^*(c, i, \mathbf{x} + e_i) \leq S^*(c, i, \mathbf{x})$ for fixed values of c and $x_j \ \forall \ j \neq i$ where e_i is a null vector with 1 at the i^th element.

Definition 1 *: The look-ahead inventory level for customer i when the system is at state $s = (c, i, \mathbf{x})$ is defined as the sequence of last observed inventory levels starting at the next customer and is denoted $\mathbf{x^i} = (x^i_{(1)}, \cdots, x^i_{(N)}) = (x_{i+1}, \cdots, x_N, x_1, \cdots, x_i)$.*

Proposition 3: *For each state $s = (c, i, \mathbf{x^i}) \in S$, the optimal replenishment action $a^*(s) \in A(s)$ at customer i is non-decreasing in the sequence of customers for identical look-ahead inventory levels.*

Proposition 3 states that the truck has a tendency to drop more units as it approaches the end of the route when visited customers face similar look-ahead inventory levels. This implies that the optimal replenishment policy is of a base-stock nature where the order-up-to-levels are $S^*(c, i, \mathbf{x^i}) \geq S^*(c, j, \mathbf{x^j})$ for $i > j$, $\mathbf{x^i} = \mathbf{x^j}$, and fixed values of c.

Below, we present an additional property that we state without proof.

Property 4: *For each state $s = (c, i, \mathbf{x}) \in S$, there exists an optimal replenishment action $a^*(s) \in A(s)$ at customer i that is*

a. non-increasing in inventory levels x_j for $j < i$.

b. non-decreasing in inventory levels x_j for $j > i$.

Proposition 4.a and 4.b imply that the replenishment action does not only depend on the location, the currently visited inventory levels, and the available stock units in the truck, but also on the inventory levels at other customer sites. In particular, the optimal action at customer i tends to decrease as the inventory levels of the previously visited adjacent customers $j < i$ increase, and to increase as the inventory levels of the non-visited customers $j > i$ increase. Although, property 4.b may look intuitive, the impact of property 4.a has been mainly observed for problem instances with tight truck capacity relative to total average customers' demand. Properties 1 to 4.a will be used to reduce the action space and the computation times for the modified-policy iteration algorithm used to solve for instances of the dynamic allocation problem in the next section.

4. Numerical Results and Managerial Insights

In this section, we summarize the findings of our numerical experiments and provide managerial insights on the performance of the inventory/distribution system. In particular, we use the modified-policy iteration algorithm [24] to solve for several problem instances and identify the optimal replenishment policies. Through the obtained numerical results, we address the question of how do the implementation of optimal replenishment policies impact the customers' service levels. We also investigate the impact of different cost parameters such as inventory holding cost, lost sales cost, and penalty cost on the service levels. In addition, we explore the value of information and dynamic decision making on the performance of such inventory/distribution systems.

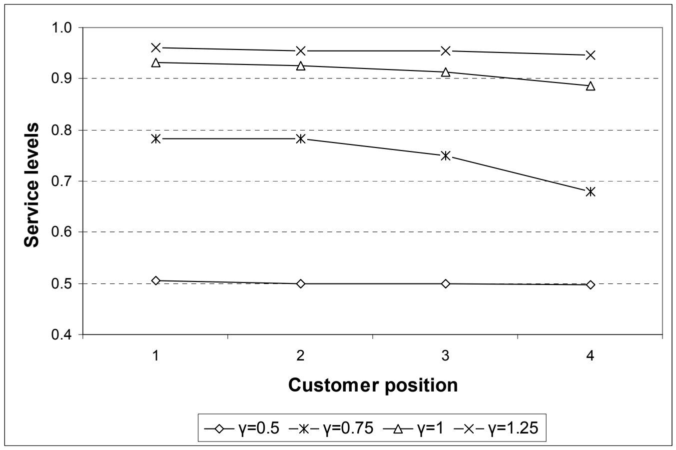

Figure 2: Impact of capacity ratio on the service levels for a system with identical customers ($N = 4, p = 4, h = 1, b = 4, \lambda_i = 2\ \forall\ i$).

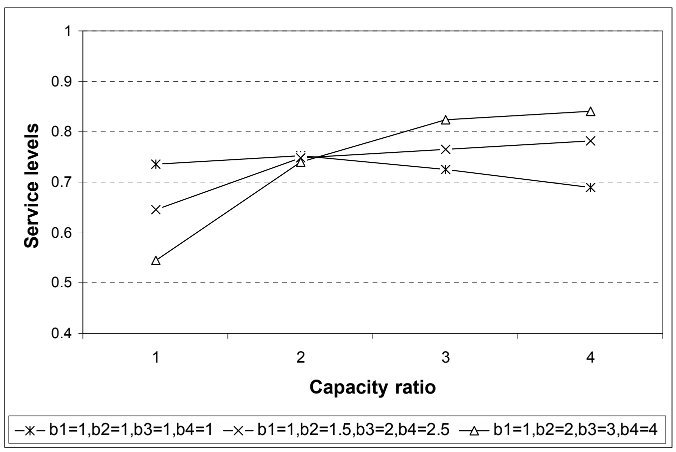

Figure 3: Impact of shortage cost on service levels for systems with customers ordered in increasing shortage cost ($N = 4, p = 1, \gamma = 0.75, h = 1, \lambda_i = 2\ \forall\ i$).

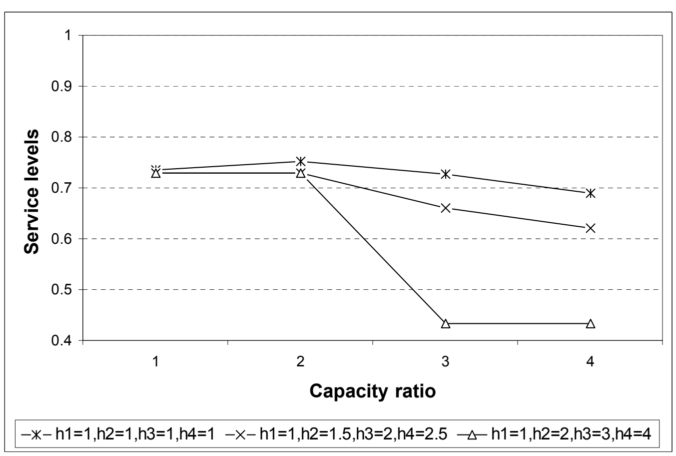

Figure 4: Impact of inventory holding cost on service levels for systems with customers ordered in increasing inventory holding cost ($N = 4, p = 1, \gamma = 0.75, b = 1, \lambda_i = 2\ \forall\ i$.)

4.1 Service Levels

The obtained optimal replenishment policies are initially designed to minimize the total expected cost for the inventory/distribution systems. However, it is not clear whether these policies provide equal benefits to all served customers. To answer this question, we first examine the impact of capacity on the incurred service levels (the ratio of satisfied demand relative to total demand) due to optimal replenishment policies. We refer to the capacity level (γ) as the ratio of the truck capacity to the average total customers' demand (i.e. $\gamma = C/\sum_i \lambda_i$). In our experiments, we consider several capacity levels that range between 0.5 (very tight capacity) and 1.25 (ample capacity). Figure 2 illustrates the service levels for different capacity ratios for problem instances with four identical customers.

Our results show that, in general, the customers visited towards the end of the route tend to have lower service levels than the earlier visited customers. In other words, an optimal replenishment policy for a distribution system with identical customers prioritizes service to the early visited customers. This has been mainly observed for cases with tight ($\gamma = 0.75$) and medium ($\gamma = 1$) capacity ratios. However, for ample ($\gamma = 1.25$) and very tight ($\gamma = 0.5$) cases, the service levels are insensitive to the customer position. In fact, when an abundant stock is available, all customers are able to fill their stock up to their needs, whereas when only a limited stock is available, all customers are served poorly regardless of their respective positions in the route.

To understand the impact of the cost parameters on the service levels, we run two sets of experiments where customers are ordered according to their shortage costs and inventory holding costs respectively. In each set, we consider three scenarios where customers are sequenced in the increasing order of their costs with different increments. Figures 3 and 4 illustrate a sample of the obtained numerical results for both sets. In particular, Figures 3 (4) show that the increasing patterns of shortage costs (inventory holding costs) increase (decrease) the service levels of customers visited at later stages of the route.

We also examine the combined impact of the unit inventory holding cost (h), unit shortage cost (b) and penalty cost (p) on the service levels in systems with identical customers. In particular, we solve several problem instances for different values of b and h where $\mu = h/b$ takes values of 0.25, 1 and 4. Figures 5, 6, and 7 show that with the increase in p and b, the supplier tends to drop more items at the customers' sites, whereas it tends to drop less items with the increase in h. We also observe that although the service levels increase with

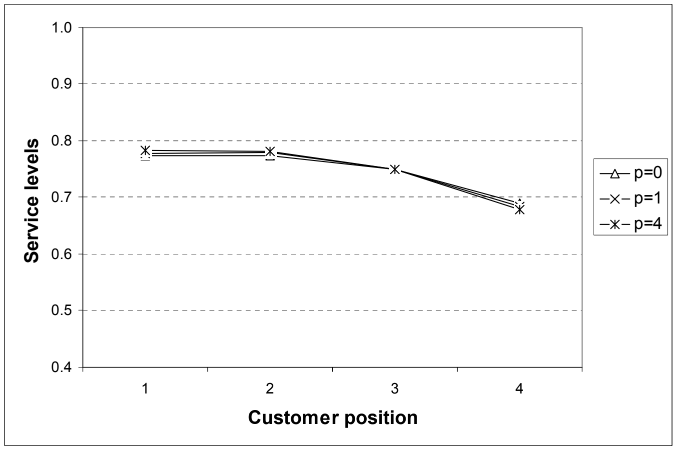

Figure 5: Impact of remaining stock penalty on the service levels for $\mu = 0.25$ ($N = 4, \gamma = 0.75, h = 1, b = 4, \lambda_i = 2\ \forall\ i$).

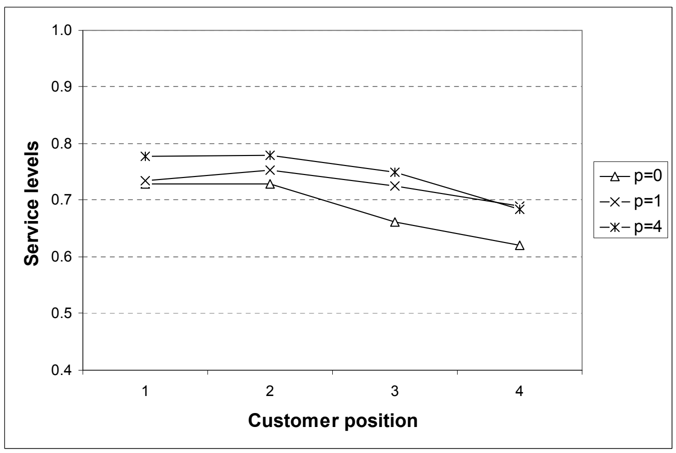

Figure 6: Impact of remaining stock penalty on the service levels for $\mu = 1$ ($N = 4, \gamma = 0.75, h = 1, b = 1, \lambda_i = 2\ \forall\ i$).

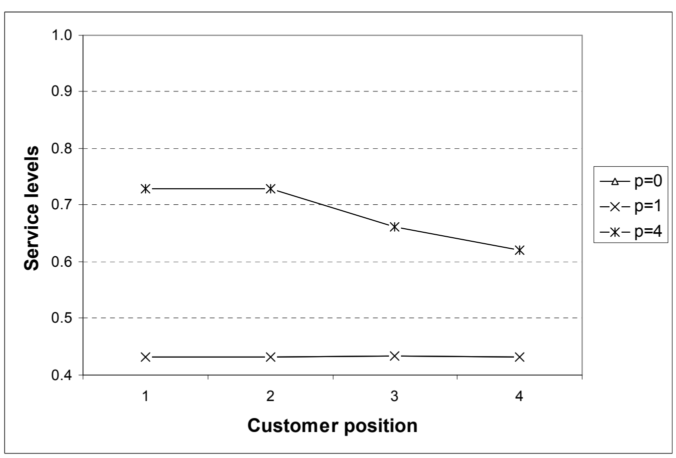

Figure 7: Impact of remaining stock penalty on the service levels for $\mu = 4$ ($N = 4, \gamma = 0.75, h = 4, b = 1, \lambda_i = 2\ \forall\ i$).

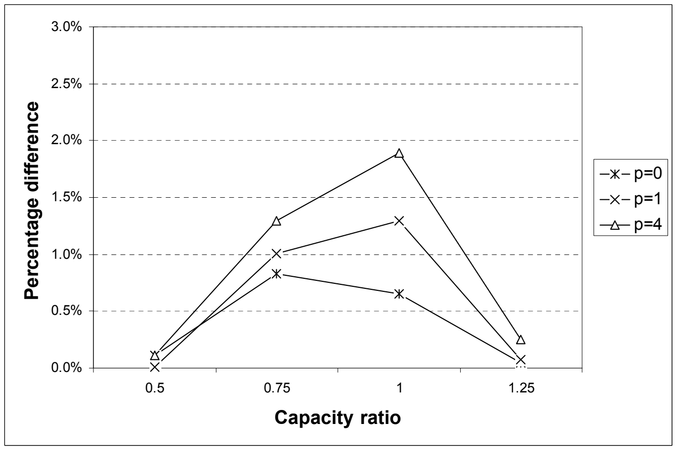

Figure 8: Comparison of average expected cost for DAUI and DAII models relative to penalty costs ($N = 4, h = 1, b = 4, \lambda_i = 2 \; \forall \; i$).

the value of penalty cost p (Figure 6 and 7), the impact of p becomes less significant for systems with low values of μ (Figure 5).

The above results imply that, unless the supplier has ample or very tight capacity relative to total average customers' demand, the vendor has to create a financial mechanism that redistributes marginal revenues among customers to guarantee service fairness. Other possible remedies can involve the cyclical rerouting of the distribution process to allow a balanced long-run average service levels. The impact of cost parameters on the service levels can also provide guidelines to how these incentives and balancing mechanisms can be implemented.

4.2 Value of Information

In this section, we aim at answering the question of how much does updated information availability and dynamic decision making contribute to the reduction of average expected cost. To measure the combined value of information and dynamic decision making on the inventory/distribution system performance, we compare the results of our dynamic inventory allocation procedure under imperfect information (DAII) with those of:

- The dynamic inventory allocation problem using updated information (DAUI), and
- The static joint inventory allocation problem under complete information (SACI).

In the DAUI problem, upon arrival to the customer site, the supplier receives updated information about the inventory levels of all customers. The supplier then uses this information to decide about the amount of stock to drop at the visited retailer. This is similar

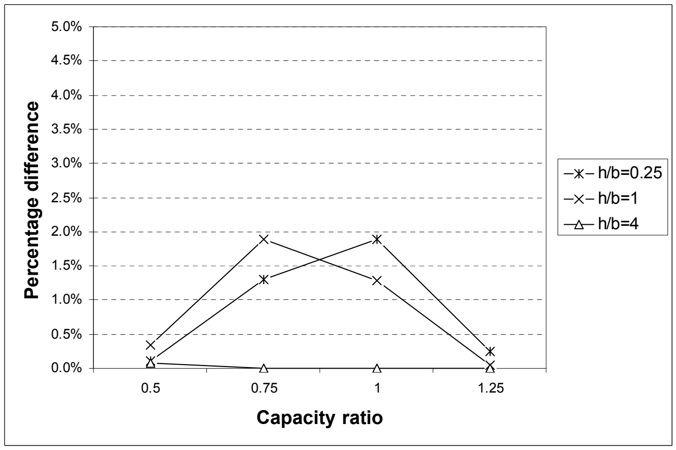

Figure 9: Comparison of average expected cost for DAUI and DAII models relative to inventory holding and shortage costs ($N = 4, p = 4, \lambda_i = 2\ \forall\ i$).

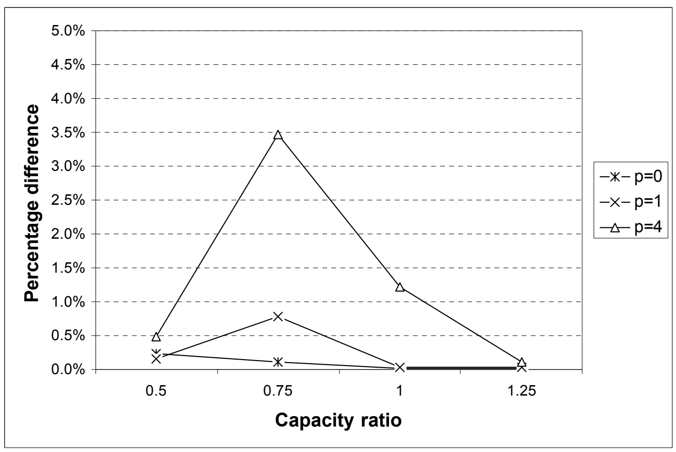

Figure 10: Comparison of average expected cost for SACI and DAII models relative to penalty costs ($N = 4, h = 1, b = 4, \lambda_i = 2\ \forall\ i$).

to the problem addressed in [2] in the absence of customer routing decisions. To allow an adequate comparison with our approach, we consider a problem with inventory holding costs equal to h_i/N and Poisson distributed customer's demand with mean λ_i/N. We solve for the resulting problem using a modified policy iteration algorithm. For the SACI problem, the supplier decides about the allocation of the stock among all customers before leaving the depot, based on complete updated inventory information. To solve for the SACI model, we develop an MDP model where at each time period, the decisions of how much to drop at each customer site are taken simultaneously.

We solve for several problem instances with different cost combinations. Obviously, due to the updated information availability, both models outperform the DAII model. Figures 8

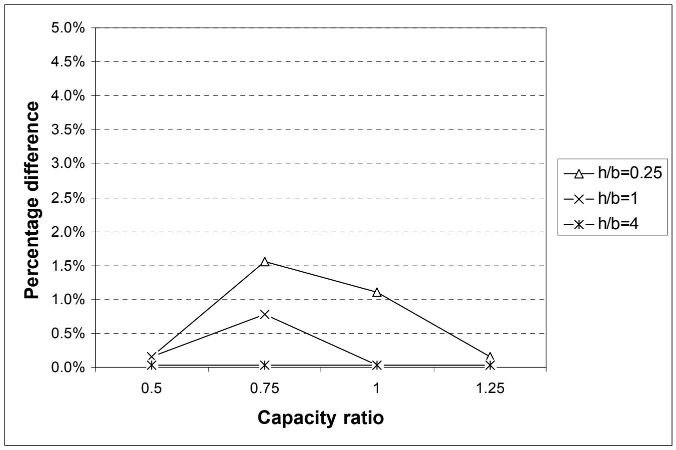

Figure 11: Comparison of average expected cost for SACI and DAII models relative to inventory holding and shortage costs ($N = 4, p = 1, \lambda_i = 2 \; \forall \; i$).

and 9 illustrate the percentage differences between the optimal average cost for the DAUI and DAII models. Our results show that the performance gap between both models is minimum for systems with ample ($\gamma = 1.25$) or very tight ($\gamma = 0.5$) capacity ratios. In addition, the gap increases with the penalty cost and decreases with the h/b ratio. Figures 10 and 11 illustrate the percentage differences between the optimal average cost for the SACI and DAII models. Similar to the DAUI model, the gap between the SACI and DAII models is largest for systems with tight and medium capacity ratios, high penalty costs, and low h/b ratios.

These results suggest that the value of information depends mainly on the characteristics of the distribution system in terms of capacity and cost. In particular, the availability of updated information has little to no impact on the performance of such systems where either ample or little stock is available for allocation. The value of updated information is further dampened by high inventory holding cost. More importantly, for all tested problem instances, the percentage average cost difference between the DAII model and the DAUI and SACI models did not exceed 5% and %4 respectively. Therefore, investing in state-of-the-art information technology to acquire more updated information may not be always justified when appropriate dynamic inventory allocation decisions are taken in such inventory/distribution systems.

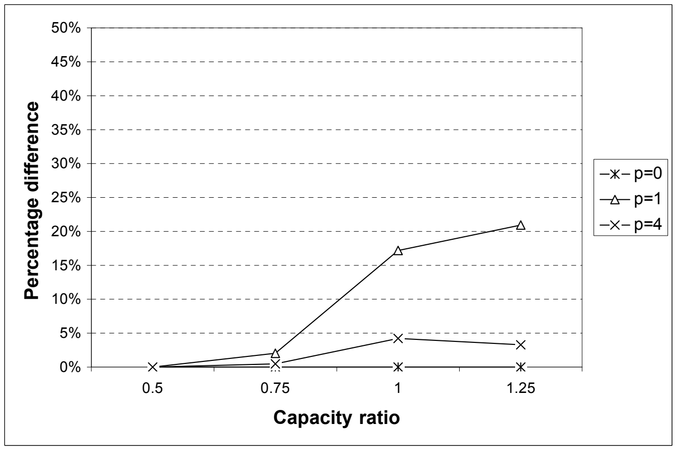

Figure 12: Comparison of average expected cost for the single-cycle dynamic inventory allocation policy (SCDA) and the dynamic allocation under imperfect information policy (DAII) relative to p ($N = 4, \gamma = 0.75, h = 1, b = 1, \lambda_i = 2 \ \forall \ i$).

5. Heuristics Performances

Our numerical experiments show that the computation time necessary to solve for the DAII model is relatively prohibitive even for small-size problems. In particular, the number of states increases exponentially in the number of customers and each customer's storage capacity. In this section, we evaluate the efficiency of alternative myopic policies that are easy to find and implement. We also develop and test the quality of iterative decomposition-based approaches.

5.1 Myopic Heuristics

We investigate the efficiency of two myopic policies with respect to the optimal replenishment policy; the single cycle dynamic allocation policy (SCDA) and the independent base stock policy (IBS). The SCDA policy is a finite-horizon replenishment myopic policy that minimizes the total cost of serving all customers during a single cycle. In each cycle, the supplier starts at the depot with full load and dynamically replenishes the stock of each customer upon contact. At the end of each cycle, the supplier is penalized for the remaining stock.

The SCDA policy falls short from accounting for the impact of current actions on the costs to be incurred in future cycles. Our numerical results show that the percentage difference between the SCDA and DAII models can ba as high as 35% for some problem instances. However, for problem instances with low penalty cost ($p = 0$), the myopic procedure can result in near-optimal solutions (Figure 12).

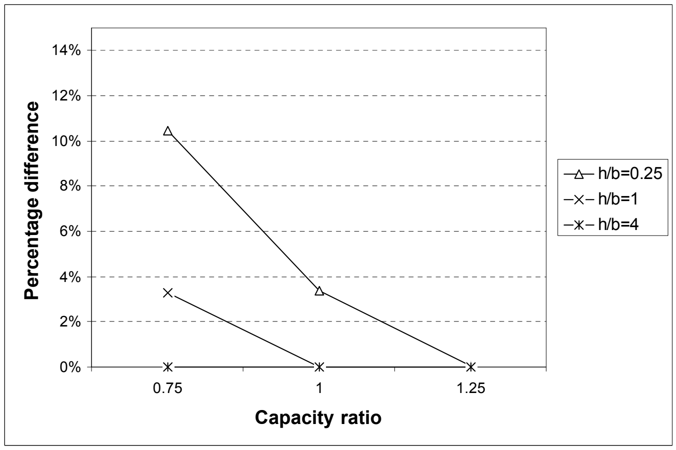

Figure 13: Comparison of average expected cost for the independent base-stock policy (IBS) and the dynamic allocation under imperfect information policy (DAII) relative to μ ($N = 4, \gamma = 0.75, p = 1, \lambda_i = 2 \ \forall \ i$).

The IBS is a myopic policy where the supplier replenishes each customer inventory up to the identified independent base-stock level subject to stock availability. Due to the myopic nature of the IBS heuristic, the supplier capacity, the penalty cost, and the inventory level information of the customers are not fully incorporated in the decision making process and therefore fail to ration the available stock among different customers. Our numerical results show that the percentage difference between the DAII and the IBS heuristic results can be as high as 27%. Figures 13 and 14 show that the performance of the IBS policy deteriorates with the increase in inventory holding cost and penalty cost respectively. Both figures also show that the performance of the IBS policy improves with the increase in capacity ratio. Therefore a base-stock policy can be effectively deployed to manage such distribution/inventory systems with low penalty cost and ample supplier capacity. In addition, the IBS will be more effective for systems with high inventory holding costs.

Since the SCDA and IBS policies are relatively easier to identify and implement than the DAII solution procedure, such policies can be recommended in lieu of the dynamic inventory allocation whenever the aforementioned conditions exist.

5.2 Decomposition-based Heuristics

In this section, we propose an iterative dynamic decomposition-based approach that exploits the fact that the MDP solution procedures with a small number of customers (two or three) can be efficiently solved in a reasonable amount of time. The solutions of such problems can

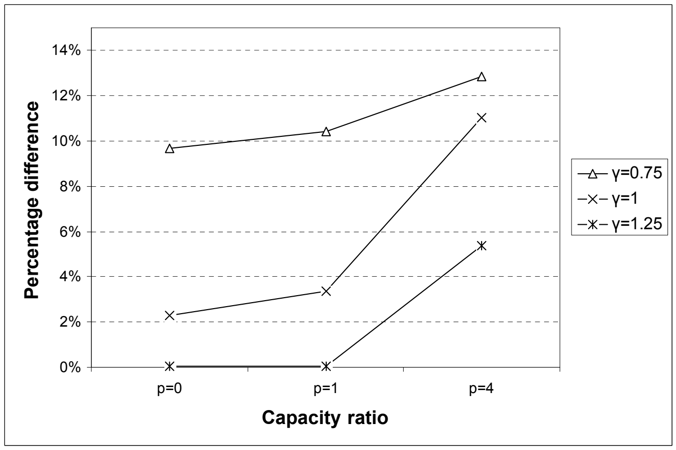

Figure 14: Comparison of average expected cost for the independent base-stock policy (IBS) and the dynamic allocation under imperfect information policy (DAII) relative to p ($N = 4, \gamma = 0.75, h = 1, b = 4, \lambda_i = 2 \; \forall \; i$).

then be employed in solving for the original problem. For each state of the MDP, the replenishment problem for multiple-customers is reduced to a two-customer or three-customer sub-problem. The reduction of the original problem into a set of smaller sub-problems is achieved through the pooling of the inventory levels and the aggregation of demands for the visited customers and the remaining customers. While the first heuristic considers problems where inventory and demand is pooled for the remaining non-visited customers, the second heuristic considers problems where inventory and demand is pooled for the previously visited (within the cycle) as well as the remaining customers. For some demand distributions such as the Poisson and Normal distributions, demand pooling is straightforward, however, for problems with other demand distributions, the calculation of several distribution convolutions may be required. The cost function for the grouped customers reflects the total expected cost for the individual customers. Both heuristics result in an approximate action to each state in the original problem. Below, we provide a detailed description of both heuristics.

5.2.1 Heuristic 1

Step 0: Consider the original problem composed of a supplier with capacity C and N customers. Each customer i experiences a Poisson demand of mean λ_i.

Step 1: Consider the N two-customer sub-problems where in sub-problem i, the first customer has demand with mean λ_i and the second customer has a demand with mean $\sum_{j=i+1}^{N} \lambda_j$ with states $s' \equiv (c, i, x_1, x_2)$. Notice that due to the pooling effect,

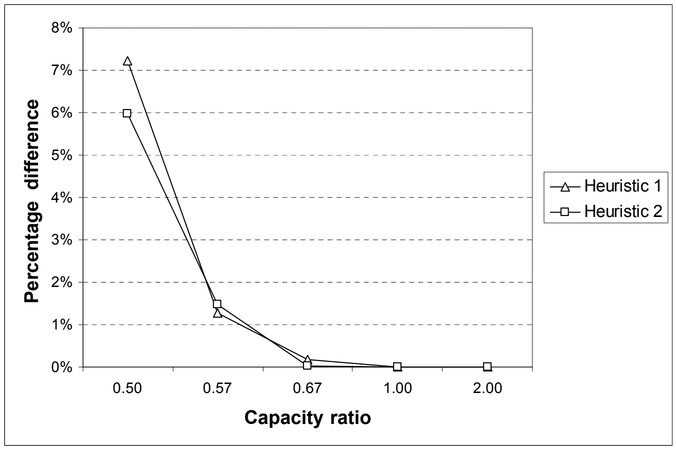

Figure 15: Comparison of decomposition heuristics average expected cost relative to optimal results ($N = 5, p = 1, h = 1, b = 1$).

sub-problem N is only composed of customer N with states $s^{'} \equiv (c, x_N)$. Use the modified-policy iteration algorithm to solve optimally for all N sub-problems. Identify all optimal actions $a^*(s^{'})\ \forall\ s^{'} \in S^{'}$.

Step 2: Consider state $s \equiv (c, 1, x_1, \cdots, x_N) \in S$ for the original problem.

Step 3: Transform current state s to state $s^{'}$ such that:

$$s \equiv (c, i, x_1, \cdots, x_N) \rightarrow\ s^{'} \equiv \begin{cases} (c, i, x_i, \sum_{j=i+1}^{N} x_j) & \text{for } i = 1, \cdots, N-1, \\ (c, x_N) & \text{for } i = N. \end{cases}$$

Step 4: Set optimal actions $a^*(s) = a^*(s^{'})$.

Step 5: Repeat steps 3 and 4 for all states $s \in S$.

Heuristic 2 can be described as follows:

5.2.2 Heuristic 2

Step 0: Consider the original problem composed of a supplier with capacity C and N customers. Each customer i experiences a Poisson demand of mean λ_i.

Step 1: Consider the N three-customer sub-problems where in sub-problem i, the first customer has demand with mean $\sum_{j=1}^{i-1} \lambda_j$, the second customer has a demand with mean λ_i, and the third customer has a demand with mean $\sum_{j=i+1}^{N} \lambda_j$ with states $s^{'} \equiv (c, i, x_1, x_2, x_3) \in S^{'}$. Notice that due to the pooling effect, sub-problem 1 is composed of two customers where the first customer has demand with mean λ_1 and the second customer has a demand with mean $\sum_{j=2}^{N} \lambda_j$ with states $s^{'} \equiv (c, i, x_1, x_2)$. Similar to Heuristic 1, sub-problem N is composed only of customer N with states $s^{'} \equiv (c, x_N)$. Use the modified-policy iteration algorithm to solve optimally for all N sub-problems. Identify all optimal actions $a^*(s^{'})\ \forall\ s^{'} \in S^{'}$.

Step 2: Consider state $s \equiv (c, 1, x_1, \cdots, x_N) \in S$ for the original problem.

Step 3: Transform current state s to state $s^{'}$:

$$s \equiv (c, i, x_1, \cdots, x_N) \rightarrow\ s^{'} \equiv \begin{cases} (c, 1, x_1, \sum_{j=2}^{N} x_j) & \text{for } i = 1, \\ (c, i, \sum_{j=1}^{i-1} x_j, x_i, \sum_{j=i+1}^{n} x_j) & \text{for } i = 2, \cdots, N-1, \\ (c, x_N) & \text{for } i = N. \end{cases}$$

Step 4: Set optimal actions $a^*(s) = a^*(s^{'})$.

Step 5: Repeat steps 3 and 4 for all for all states $s \in S$.

In addition to the number of customers in the resulting sub-problems, the main difference between Heuristics 1 and 2 lies in the transformation of the states at Step 3. Note also that in Step 1 of heuristics 1 and 2, sub-problem N is composed only of customer N and therefore can be solved as a single customer stochastic inventory problem with limited capacity.

5.2.3 Results

We run several experiments to test the quality of a variety of heuristics and develop managerial insights on the desirability of replenishment policies for such distribution systems. In our experiments, we simulate Heuristics 1 and 2 for a large set of problems. To ensure the quality of our solutions, heuristic results for small-size problems were benchmarked against the optimal expected average cost function obtained by the modified-policy algorithm.

Figure 15 illustrates the average percentage difference between optimal results and both heuristics results for problem instances with 5 homogeneous customers. To understand the impact of the capacity on the performance of the heuristics, we run problem instances for different capacity to demand ratios that range between 0.5 and 2.

In general, Heuristic 2 performs better than Heuristic 1 since it accounts for the pooled information of visited customers as well as that of remaining customers in each cycle. In most tested cases, the gap between the results of both heuristics did not exceed %8 of optimality.

Our results show that for the ample and medium capacity cases ($\gamma \in [1, 2]$), Heuristics 1 and 2 result in near-optimal results; however, Heuristic 2 tends to outperform Heuristic 1 for problems with tighter capacity ($\gamma > 1$). The optimality gap for both heuristics was also found to increase with the capacity ratio. For instance, at tight capacity ($\gamma = 0.5$), the average difference between the optimal solution and Heuristic 1 and 2 results are 7.2% and 6% respectively. In fact, at tight capacity, the value of stock allocation rationing becomes more crucial and the optimality gap widens. Since, Heuristic 1 is computationally more efficient than Heuristic 2, it is sufficient to use Heuristic 1 to solve effectively for problems with medium to ample capacity.

6. Conclusion

In this research, we consider the problem where a supplier sequentially replenishes a set of customers experiencing stochastic demand at regular intervals. Inventory information is shared with the supplier only upon contact. We model the problem as a finite-horizon MDP such that the total expected cost is minimized. We also develop structural properties characterizing the optimal replenishment policies for each customer.

We investigate the impact of the cost parameters on the performance of the inventory/distribution system. In particular, we show that for systems with identical customers, the optimal replenishment policies tends to provide higher service levels for customers that are earlier visited in the route. However, for systems with ample or very tight capacity, all customers tend to receive similar service levels. Our results imply that an incentive mechanism or a rerouting solution may be necessary to guarantee "fair" service to all customers. Our results also demonstrate that the value of inventory information may not be that critical to the system performance and that an investment in information technology may not be always justified especially for systems with ample or very tight capacity.

Given the complexity of the problem, we investigate the efficiency of alternative myopic solution approaches such as the single-cycle dynamic allocation policy and the state-independent base-stock policy. We also develop a decomposition-based solution approach that can solve efficiently for the distribution-inventory problem. Through our numerical results, we show that our solution approach provides a near-optimal solution in a reasonable computation time and outperforms myopic solution approaches.

Since the success of coordination strategies necessitates that all involved customers fairly share their benefits, it is important to identify mechanisms that can optimize the performance of the inventory/distribution system while balancing customers' service levels. So far, we have tested for dynamic transshipments as a possible mechanism that can provide an additional coordination mechanism to minimize the expected total system-wide cost while balancing service levels among customers. We have also expanded our model to incorporate the possibility of transshipments between different customers to allow the supplier to simultaneously determine the pick-up and drop-off amounts at each customer. Our preliminary results show that dynamic transshipments provide marginal improvement to the system performance, and therefore do not represent an appropriate mechanism to balance service levels. It would be beneficial to investigate whether other coordination mechanisms can provide "fair" service levels to all customers. In addition, there is a need to identify a heuristic that can solve for complex problems with customers exhibiting different demand patterns and having cost characteristics.

More importantly, the problem of identifying optimal inventory policies for two-echelon inventory systems with dynamic decision making is still wide open. Solving for such problem requires the simultaneous optimization of the supplier replenishment policy and inventory allocation to customers such that the system-wide cost is minimized.

Acknowledgement: This work was supported by the Department of Research at the University of Houston through a Grant to Enhance and Advance Research (GEAR) grant.

References

[1] D. Adelman. Price-directed replenishment of subsets: Methodology and its application to inventory routing. *Manufacturing & Service Operations Management*, 5(4):348–371, 2003.

[2] P. Balun. *A Stochastic Vendor Managed Inventory Problem and its Variations*. PhD thesis, School of Industrial and Systems Engineering, Georgia Institute of Technology, 2004.

[3] D. Barnes-Schuster and Y. Bassok. Direct shipping and the dynamic single-depot/multi-retailer inventory system. *European Journal of Operational Research*, 101:509–518, 1997.

[4] O. Berman and R. C. Larson. Deliveries in an inventory/routing problem using stochastic dynamic programming. *Transportation Science*, 35(2):192–213, 2001.

[5] F. Bernstein and A. Fedregruen. Pricing and replenishment strategies in a distribution system with competing retailers. *Operations Research*, 51(3):409–426, 2003.

[6] G. Cachon and M. Fisher. Campbell soup's continuous production replenishment program: Evaluation and enhanced decision rules. *Production and Operations Management*, 6:266–275, 1997.

[7] A. Campbell and M. Savelsbergh. *The Vehicle Routing Problem*, chapter Inventory routing in practice, pages 309 – 330. SIAM Monographs on Discrete Mathematics and Applications. SIAM, 2002.

[8] S. Çetinkaya and C. Y. Lee. Stock replenishment and shipment scheduling for vendor-managed inventory systems. *Management Science*, 46(2):217–232, 2000.

[9] S. Chen. The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy. *Operations Research*, 52(3):409–421, 2004.

[10] T. W. Chien. Determining profit-maximizing production/shipping policies in a one-to-one direct shipping, stochastic demand environment. *European Journal of Operational Research*, 64(1):83–102, 1993.

[11] T. Clark. Campbell soup company: A leader in continuous replenishment innovations. Case 9-195-124, Harvard Business School, 1994.

[12] G. DeCroix and A. Arreola-Risa. Optimal production and inventory policy for multiple products under resource constraints. *Management Science*, 44:950–961, 1998.

[13] R. Evans. Inventory control of a multi-product system with a limited production resource. *Naval Research Logistics Quarterly*, 14:173–184, 1967.

[14] A. Fedregruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands (i) the average-cost criterion. *Mathematics of Operations Research*, 11(2):193–207, 1986.

[15] A. Fedregruen and P. Zipkin. An inventory model with limited production capacity and uncertain demands (ii) the discounted-cost criterion. *Mathematics of Operations Research*, 11(2):208–215, 1986.

[16] M. Fry, R. Kapuscinski, and T. Olsen. Coordinating production and delivery under a (z,Z) type vendor-managed inventory contract. *Manufacturing and Service Operations Management*, 3(2):151–173, 2001.

[17] G. Gallego and D. Simchi-Levi. On the effectiveness of direct shipping strategy for the one-warehouse multi-retailer systems. *Management Science*, 36(2):240–243, 1990.

[18] J. Hammond and T. Clark. Reengineering channel reordering process to improve total supply chain performance. *Production and Operations Management*, 6(3):248–265, 1997.

[19] A. J. Kleywegt, V. S. Nori, and M. W. P. Savelsbergh. The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36(1):94–118, 2002.

[20] A. J. Kleywegt, V. S. Nori, and M. W. P. Savelsbergh. Dynamic programming approximations for a stochastic inventory routing problem. *Transportation Science*, 38(1):42–70, 2004.

[21] A. J. Kleywegt and J. D. Papastavrou. The dynamic and stochastic knapsack problem. *Operations Research*, 46(1):17–35, 1998.

[22] A. J. Kleywegt and J. D. Papastavrou. The dynamic and stochastic knapsack problem with random sized items. *Operations Research*, 49(1):26–41, 2001.

[23] J. D. Papastavrou, S. Rajagopalan, and A.J. Kleywegt. The dynamic and stochastic knapsack problem with deadlines. *Management Science*, 42(12):1707–1718, 1996.

[24] M. L. Puterman. *Markov Decision Processes*. John Wiley & Sons, New-York, NY, 1994.

[25] J. B. Shah. St, hp vmi program hitting its stride. *Electronics Supply and Manufacturing*, April 2002.

[26] D. Simchi-Levi, P. Kaminsky, and E. Simchi-Levi. Designing and managing the supply chain: Concepts, strategies, and case studies. *Mc-Graw Hill-Irwin*, 2003.

[27] D. M. Topkis. Minimizing a submodular function on a lattice. *Operations Research*, 26:305–321, 1978.

[28] A. Toptal, S. Çetinkaya, and C. Lee. The buyer-vendor coordination problem: Modeling inbound and outbound cargo capacity and costs. *IIE Transactions*, 35(11):987–1002, 2003.