Final Report for MBTC Project 2041:

Integrated Analysis of Transportation and Inventory in Intermodal Distribution Networks

Submitted by:
Scott J. Mason, Ph.D., P.E.
Ghazi M. Magableh
Bashyam Ramasami
Erhan Kutanoglu, Ph.D.

Department of Industrial Engineering
University of Arkansas
April 14, 2004

Table of Contents

Executive Summary

Today's marketplace places a lot of importance on having the right item at the right time in inventory and also being able to provide the item to the customer within the desired timeframe. This requires the simultaneous consideration of both the inventory and transportation functions of the supply chain. Inventory functions drive product fill rate capability, while transportation functions influence the response time windows set for each product. While previous research efforts typically consider these functions in isolation, we present a mixed integer programming model for integrated inventory and transportation decision making for the service parts logistics industry. A number of model formulations are presented that correspond to various static and time-dependent inventory policies. In all cases, customer demand is subject to pre-specified time-based service levels. The model's functionality is evaluated through an extensive experimental design containing multiple transportation modes and product types. Experimental results suggest that employing a dynamic temporal modeling scheme wherein decision makers are allowed to update/change their decisions at the beginning of each time period results in significantly reduced inventory levels, order quantities, and total supply chain costs.

1 Introduction

This research is motivated by the service parts logistics industry where the established network provides after sales support and repair parts to customers who own and operate mission critical equipment. When a customer's part fails due to a part break down, it is the responsibility of the service parts division to be in a position to ship the part to the customer to fix the problem and to get the mission critical parts up and running within the shortest amount of time possible.

The ability to meet contractual service requirement obligations has two dimensions: inventory and transportation logistics. Two key questions are associated with these dimensions:

- Is the part requested in inventory?
- Can the requested item be delivered to the customer within the desired time window?

The first question relates to order fill rate, while the second question speaks to response time. Indeed, these are not two separate questions of inventory and transportation management that should be answered in isolation, but rather two important components of an integrated inventory and transportation decision support system.

Several aspects of this integrated inventory and transportation problem contribute to the problem's complexity. First, the problem involves simultaneous consideration of inventory and transportation; this has not been a common practice in previous research efforts. Further, the problem not only concerns service parts, but includes many different product types characterized by varying service windows and shipping patterns/modes. Third, it is possible that different inventory modeling approaches may be required to handle different product types. Finally, the problem examines multiple, competing modes of transportation, including the possibility of intermodal shipping.

2 Literature Review

The unique nature of this research involving two significant areas of supply chain modeling, inventory and transportation, coupled with the need to consider both fill rate and response time, leads to a wide variety of literature to review. In addition, the existence of multiple transportation modes, multiple product types, and multiple inventory models further increases the breadth of the applicable literature. Therefore, the following sections review key papers in the open literature pertaining to the following general categories:

1. Transportation modeling for service level considerations
2. Inventory modeling involving customer service targets
3. Integrated supply chain modeling that considers inventory and transportation models for fill rate and/or response time evaluation
4. Service parts logistics modeling

2.1 Transportation Modeling For Service Level Considerations

A number of researchers have attempted to model the transportation network problem under the influence of service level considerations. One of the most comprehensive works in this area was the dissertation work of Cole (1996) which includes a comprehensive mathematical model for selecting warehouses, transportation channels, and inventory levels so that cost is minimized while service level constraints are satisfied. Cole (1996) considers production, supply, and warehousing aspects of the supply chain, as well as transportation link issues involving multiple modes of shipping. The model is first developed as a mixed-integer program (MIP) containing nonlinear constraints. The author relaxes a number of problem restrictions, then analyzes the relaxed MIP formulation of the problem.

Lehmusvaara (1998) considers logistics as a strategic level decision tool for organizations. The author considers factors of both transportation and customer service levels that should be incorporated as a part of an organization's logistics strategy. Lehmusvaara (1998) discusses different transportation strategies to minimize transportation cost, as well as the importance of service level considerations as a part of minimizing transport cost. Although the author makes no specific reference to a model, the paper does mention a MIP model and a routing program used to consider two distinct solution phases. The MIP model is used to perform yearly planning studies to determine production volumes and transport prices, while the routing program helps to identify the shortest route(s) between pairs of locations. The author suggests that different transportation policies can be used to evaluate each policy's effect on customer service level.

Hwang (2002) presents a logistics system design that optimizes the performance of logistics systems subject to required service levels. The author considers both the number of warehouses or distribution centers (DCs) in the logistics network, as well as the associated vehicle routing problem. Hwang (2002) defines a probabilistic set covering problem, then solves it using a binary integer programming formulation. The article discusses the complexities of solving the problem as defined, and then suggests metaheuristic approaches such as genetic algorithms be explored to improve the possibility of finding viable solutions.

2.2 Inventory Modeling Involving Customer Service Targets

Evers (1999) discusses the effect of safety stock when forecasts are fairly stable. Safety stock is a function of management-specified customer service levels. After discussing standard safety stock calculation equations, the author suggests that the only major factor that managers and decision makers can influence is lead time, both its average and its variability. Evers (1999) considers different decision rules based on the reduction of either average lead time or the standard deviation of lead time in order to reduce safety stock. This work provides an understanding of the effect of supply lead times on safety stock that eventually affects customer service levels.

Erlebacher and Meller (2000) present a location inventory problem with the inventory cost acting as the main objective to be minimized. This objective is particularly pertinent for companies with a large customer base. The problem is first determined to be NP Hard, and then a simplified version of the problem is analyzed to provide a basic set of intuitive results. The authors also present a heuristic technique for addressing the original "full" problem. Although our research is focused on identifying optimal inventory levels that will enable a supplier to meet required response times in an existing network, the work of Erlebacher and Meller (2000) provides good insight into inventory modeling within a supply chain.

Daskin and Coullard (2002) consider a DC location problem that incorporates working inventory and safety stock inventory cost at the DC. The authors formulate the problem as an NLMIP, and then examine the problem using a Lagrangian-relaxation based formulation. The article also discuses a number of heuristic techniques and variable forcing rules for solving the problem in polynomial time. The problem does not consider service level constraints explicitly for the inventory level considerations under study. However, it does incorporate safety stock level modeling to minimize the possibility of stockout.

2.3 Integrated Supply Chain Modeling

Buffa *et al.* (1980) present one of the earliest known articles in the area of customer service level management. This article describes the application of a model used to analyze the relationships between customer service level, inventory investment, and operating expense. The model considers the logistics cost of meeting a set service level and, at the same time, considers the safety stock that has to be maintained to achieve the desired service level. In additional, an attempt is made to develop a model to study the interaction between safety stock and service level. This early research provides insights into understanding the relationship between customer service levels and inventory, as well as to logistics functions.

Perl and Sirisoponsilp (1988) develop an integrated model for the distribution network design problem which explicitly represents the tradeoffs between facility, transportation, and inventory costs. This model differs from a traditional DC location model primarily in its representation of the design objective which attempts to minimize the total distribution cost, warehousing costs, trucking costs, delivery costs, in-transit inventory levels, cycle stock, and safety stock. This model overcomes the shortcomings of traditional DC models in that it does not assume that trucking cost is constant, nor does it assume that it is independent of shipment quantity. The authors provide a complete mathematical formulation to represent the costs associated with multiple transportation options and required levels of customer service.

Jayaraman (1998) simultaneously considers the relationship between inventory management, facility location, and transportation policy in a distribution network design environment. The problem is modeled as a MIP, with the objective being to minimize the total cost of distribution associated with the three decision components being considered. The author also analyzes the interaction between the three factors, as well as how each factor affects the objective function. Although Jayaraman (1998) develops an integrated inventory and transportation model, it is a static (i.e., not time-dependent) formulation. Further, the model does not consider any sort of customer service level requirements.

Qu *et al.* (1999) present an integrated inventory and transportation system with a modified periodic review inventory policy and a traveling salesman problem (TSP) component. Multiple products are considered in this case under a stochastic setting, with decisions being made on inventory and transportation policies simultaneously. The only drawback of the authors' formulation as presented is that it is incapacitated due to service level considerations. However, the research of Qu *et al.* (1999) serves as a good starting point for integrated model development.

A deterministic order up-to level policy is studied for minimizing distribution system costs by Bertazzi *et al.* (2002). The research considers a multi-period inventory and logistics model with a shipping horizon, vehicle capacity and cost, and retailer-determined minimum and maximum inventory levels under deterministic demand. The objective function attempts to minimize the cost of operations. However, when an attempt is made to determine shipping quantities, discrete time units for shipping, and the shipping route to be followed, the problem is identified as NP Hard, as it reduces to the TSP. Hence, this paper also considers a heuristic procedure for solving the deterministic order-up-to inventory problem.

2.4 Service Parts Logistics Modeling

Cohen *et al.* (1997) describe a benchmarking study conducted to study service parts logistics networks. The study was funded by IBM, one of the largest business and consumer electronics providers that also has an extensive service parts book of business. The "after sales" service of technologically sophisticated products is considered, making this previous work very applicable to our research endeavor. The authors discuss the benchmarking methodology adopted and consider the control policies and technology that is in place with respect to the materials management aspect of a service parts logistics network. Cohen *et al.* (1997) examine existing industry practice with respect to logistics networks, support technologies, and parts availability versus response time.

2.5 Literature Review Summary

Based on our review of the four different literature areas presented above, it is apparent that significant work has been done in the area of incorporating service level considerations into both transportation models and inventory models, albeit in isolation. We find no previous research efforts that integrate all three key components of interest in this research: inventory levels, transportation decisions, and service level considerations. Moreover, our review of service parts logistics industry benchmarks and the industry's associated preliminary research efforts suggest that an integrated model capable of optimally determining both replenishment policy parameters and transportation mode/quantity selection to meet desired customer time-based service levels could significantly improve the operational and cost efficiencies of companies involved in this domain.

3 An Integrated Inventory-Transportation Model with Time-Based Service Requirements

Consider a supplier with an establish service parts logistics network containing a known quantity of fixed warehouses, a diversified customer and product base, and multiple transportation mode options with known speeds, quantities, and capacities. Further, this supplier provides its customers with various levels of time-based service, each with an associated cost, with respect to delivering requested products within an agreed upon time window (e.g., within two hours, within eight hours, and so on). The supplier's main objective is to meet its customers' time-based service level requirements with minimum inventory and distribution costs by making appropriate inventory and transportation decisions under varying demand.

Therefore, the modeling instance can be thought of as *a supplier attempting to minimize total costs by selecting the correct combination of transportation and inventory strategies for their established supply chain network to meet customer specific contractual service requirements.* A model capable of providing decision support assistance for this scenario could serve as an ideal tool for the supplier's logistics and inventory planning divisions. The model could help to identify appropriate inventory and transportation policies for meeting desired customer service requirements with minimum total cost.

3.1 Model Notation

3.1.1 Sets and Parameters

Let I denote the set of warehouses ($i = 1...I$) that the supplier uses to service the demand for a set of customers J ($j = 1...J$). Each warehouse $i \in I$ supply a set of products P ($k = 1...P$) in response to customer j's demand D_{jk} ($j \in J$) for product $k \in P$. Product $k \in P$ has an associated size of z_k units, which can be thought of as either cubic volume or weight. A total of ω_{ik} units of product $k \in P$ are maintained as safety stock at warehouse $i \in I$. It costs O_{ik} dollars to order a standard quantity of π_{ik} units of product $k \in P$ from warehouse $i \in I$.

Assume the period of time over which the integrated model will make inventory and transportation decisions is τ days long ("the decision horizon"). The decision horizon can be decomposed into a set T containing non-overlapping time periods ($t = 1...T$) within which time-based service is offered by the supplier. Let the duration of time period $t \in T$ be denoted by u_t. For example, a decision horizon of $\tau = 1$ day could be made up of $\|T\| = 3$ periods, with

$u_1 = 0.25$ days (i.e., the first six hours), $u_2 = 0.5$ days (i.e., the next 12 hours), and $u_3 = 0.25$ days (i.e., the final six hours of the decision horizon). It is important to note that $\sum_{t=1}^{\|T\|} u_t = \tau$ (i.e., the sum of the individual time periods equals the whole decision horizon). Finally, let the actual time (in days) corresponding to the end of time period $t \in T$ be e_t, which is calculated by $e_t = \sum_{g=1}^{t} u_g$. For example, continuing with the example above, $e_2 = u_1 + u_2 = 0.25 + 0.5 = 0.75$ days or 18 hours.

A set of M different transportation modes ($r = 1...M$) are available for transporting products from each warehouse $i \in I$ to every customer $j \in J$. Transportation mode $r \in M$ has a stated capacity of ϑ_r units. A fixed cost of C_r is incurred for every trip made via transportation mode $r \in M$. Further, a variable cost of c_{ijkr} incurred for every unit of product $k \in P$ transported from warehouse $i \in I$ to customer $j \in J$ via transportation mode $r \in M$. Each mode of transportation $r \in M$ has an associated speed of v_r miles per hour. Assume customer $j \in J$ is located a total of d_{ij} miles from warehouse $i \in I$. Clearly, the time (in hours) required by transportation mode $r \in M$ to travel between warehouse $i \in I$ and customer $j \in J$, denoted by λ_{ijr}, is calculated as $\lambda_{ijr} = {d_{ij}}/{v_r}$. The number of time periods L_{ijr} required to transport goods from warehouse $i \in I$ to customer $j \in J$ by transportation mode $r \in M$ (i.e., lead time) is determined as the minimum value of t such that $\lambda_{ijr} \le e_t$.

Assume the supplier has a contractual obligation to fulfill at least α_{jkt} percent of the total demand from customer $j \in J$ for product $k \in P$ within time period $t \in T$. It follows that at least $\sum_{g=1}^{t} \alpha_{jkg}$ percent of the total demand from customer $j \in J$ for product $k \in P$ must be fulfilled by time $e_t \forall t \in T$. A binary incidence matrix A is used to specify whether or not customer $j \in J$ can be served by warehouse $i \in I$ by transportation mode $r \in M$ in time period $t \in T$. Consider $a_{ijrt} \in A$. If $\lambda_{ijr} \le e_t$, then $a_{ijrt} = 1$; otherwise, $a_{ijrt} = 0$.

At the beginning of the decision horizon, assume that Λ_{ik1} units of initial product $k \in P$ inventory are on hand at warehouse $i \in I$. For all $t > 1$, initial inventory level $\Lambda_{ikt} = 0$. Every unit of product $k \in P$ held in inventory at warehouse $i \in I$ incurs an inventory holding cost of h_{ik} dollars for the decision horizon under study. Finally, let β denote a very large number (i.e., "big M").

3.1.2 Decision Variables

The integrated model will be used to determine values for x_{ijkrt}, the total amount of product $k \in P$ shipped from warehouse $i \in I$ to customer $j \in J$ by transportation mode $r \in M$ in time period $t \in T$. The decision variable y_{ijrt} represents the number of transportation vehicles of type $r \in M$ that are used to ship goods from warehouse $i \in I$ to customer $j \in J$ in time period $t \in T$. A binary indicator variable ψ_{ijrt} is calculated to denote whether transportation mode $r \in M$ is ($\psi_{ijrt} = 1$) used to transport goods between warehouse $i \in I$ to customer $j \in J$ in time period $t \in T$ or not ($\psi_{ijrt} = 0$).

In addition to the transportation-based decision variables, the integrated model will also determine the optimal values for a number of inventory-based variables, such as the reorder point

s_{ikt}, order-up-to level S_{ikt}, total order quantity q_{ikt}, and number of orders placed f_{ikt} for product $k \in P$ at warehouse $i \in I$ in time period $t \in T$. Finally, the average inventory level l_{ik} over the decision horizon is also calculated by the integrated model.

3.2 Model Formulation

The integrated model's objective is to minimize total costs, subject to contractually-specified time-based service requirements for fulfilling customer demand request:

$$\min \sum_i \sum_j \sum_r \sum_t C_r y_{ijrt} + \sum_i \sum_j \sum_k \sum_r \sum_t c_{ijkr} x_{ijkrt} + \sum_i \sum_k h_{ik} l_{ik} + \sum_i \sum_k \sum_t O_{ik} f_{ikt} \qquad (1)$$

In (1), the first term denotes fixed transportation costs, while the second term calculates variable transportation costs per item shipped. The third term in (1) calculates inventory holding costs over the decision horizon under study, while the fourth and final term in (1) tabulates the ordering costs resulting from the inventory policy under study.

As stated above, we assume the supplier has a contractual obligation to fulfill at least α_{jkt} percent of the total demand from customer $j \in J$ for product $k \in P$ within time period $t \in T$:

$$\sum_i \sum_r \sum_{g:g \le t} a_{ijrg} x_{ijkrt} \ge \sum_{w:w \le t} \alpha_{jkw} D_{jk} \quad \forall j \in J, k \in P, t \in T \qquad (2)$$

It is assumed that each customer's total demand should be satisfied over the decision horizon under study (i.e., $\sum_t \alpha_{jkt} = 1 \ \forall j \in J, k \in P$). The model's reorder point s_{ikt} is calculated based on safety stock level and product demand fulfillment estimates based on computed order lead times:

$$s_{ikt} \geq \omega_{ip} + \sum_j \sum_r \alpha_{jkt} D_{jk} L_{ijr} \psi_{ijrt} \quad \forall i \in I, k \in P, t \in T \tag{3}$$

Constraints (3) also guarantee that a minimum of level inventory (i.e., safety stock) is maintained by warehouse for each product type under study. Order up-to-level calculations are based on the reorder point and the quantity of goods shipped out of each warehouse:

$$S_{ikt} \geq s_{ikt} + \sum_j \sum_r x_{ijkrt} \quad \forall i \in I, k \in P, t \in T \tag{4}$$

Both the replenishment policy levels and the initial warehouse inventory are used to calculate the total order quantity q_{ikt} for each product at each warehouse in each time period:

$$q_{ikt} \geq S_{ikt} - s_{ikt} - \Lambda_{ikt} \quad \forall i \in I, k \in P, t \in T \tag{5}$$

The initial inventory term $\Lambda_{ikt} = 0$ in (5) $\forall i \in I, k \in P, t > 1$. The average inventory level l_{ik} over the decision horizon is calculated as the sum of the time-weighted average inventory levels for each product at each warehouse over each time period:

$$l_{ik} \geq \sum_t \left(\frac{u_t}{\tau} \left(q_{ikt} + s_{ikt} - 0.5 \sum_j \sum_r x_{ijkrt} \right) \right) \forall i \in I, k \in P \tag{6}$$

As we assume that demand is uniformly fulfilled over the decision horizon, we employ the common approach of scaling the shipments variable x_{ijkrt} by 0.5 in (6). Further, we assume that the inventory level computed for the beginning of each time period, as determined by the inventory policy under study, is sufficient for filling the demand in that period. The number of orders f_{ikt} that must be placed is calculated using both the total and standard order quantity values for each product at each warehouse in each time period:

$$f_{ikt} \geq \frac{q_{ikt}}{\pi_{ik}} \; \forall i \in I, k \in P, t \in T \tag{7}$$

We assume that the supplier is responsible for all transportation requirements within their supply chain. The required number of transportation vehicles y_{ijrt} of type $r \in M$ used to ship goods from each warehouse to each customer in each time period is calculated as follows:

$$\vartheta_r y_{ijrt} \geq \sum_k x_{ijkrt} z_k \forall i \in I, j \in J, r \in M, t \in T \tag{8}$$

In (8), transportation vehicle capacity ϑ_r is employed to ensure sufficient transportation assets are identified for transporting the necessary goods. Finally, the relationship between transportation mode selection and the quantity of goods shipped is captured in (9) and (10):

$$\psi_{ijrt} \leq y_{ijrt} \; \forall i \in I, j \in J, r \in M, t \in T \tag{9}$$

$$y_{ijrt} \leq \beta \psi_{ijrt} \; \forall i \in I, j \in J, r \in M, t \in T \tag{10}$$

In the model formulation above, all variable are non-negative, with reorder point s_{ikt}, order-up-to level S_{ikt}, total order quantity q_{ikt}, and transportation vehicle count y_{ijrt} being restricted to the set of positive integers.

3.3 Inventory Policy Variations

The formulation presented above in Section 3.2 models a classic (s, S) inventory system. In an (s, S) system, when the on hand inventory for a product under study drops to some level σ that is below the reorder point s (i.e., $\sigma \leq s$), an order for ($S - \sigma$) units is placed to replenish inventory levels back up to the order-up-to level S. In addition to the (s, S) inventory system modeled above, we investigate two other inventory policies in this research: a base stock policy and a classic (r, Q) policy.

3.3.1 *Base Stock Inventory Policy*

Base stock inventory policies involve a single critical parameter S, the base stock level. Under a base stock policy, an order is placed as soon as the number of units in stock is smaller than S. Further, customer demand is assumed to arrive in single units increments. This is often the case for slow moving items such as refrigerators and engines.

A base stock policy can be thought of as an (s, S) inventory policy with $s = S - 1$, as items are ordered one at a time. Four modifications are required to the formulation above to accurately model a base stock policy. First, constraints (4) are simplified to denote the fact that only one unit is ordered at a time:

$$S_{ikt} \geq s_{ikt} + 1 \; \forall i \in I, k \in P, t \in T \qquad (11)$$

The total order quantity q_{ikt} is simply calculated as the difference between the demand shipped and the initial inventory level for each product at each warehouse in each time period. This order quantity is required to bring the inventory level up to the base stock level S_{ikt} by the end of the time period:

$$q_{ikt} \geq \sum_j \sum_r x_{ijkrt} - \Lambda_{ikt} \quad \forall i \in I, k \in P, t \in T \qquad (12)$$

The average inventory level l_{ik} constraints (6) are modified also, as the total order quantity q_{ikt} is replaced by the average order quantity (i.e., $0.5\,q_{ikt}$) because items are ordered one at a time under a base stock policy:

$$l_{ik} \geq \sum_{t} \left(\frac{u_t}{\tau} \left(0.5 q_{ikt} + s_{ikt} - 0.5 \sum_{j} \sum_{r} x_{ijkrt} \right) \right) \forall i \in I, k \in P \tag{13}$$

In other words, if a total of q_{ikt} units are needed during time period t, then the average number of units needed in inventory during time period t is 0.5 q_{ikt} under a base stock policy with single unit orders. Finally, the number of orders f_{ikt} that must be placed is simply equal to the total order quantity q_{ikt}, as we are ordering one unit at a time under a base stock policy (i.e., $\pi_{ik} = 1$):

$$f_{ikt} \geq q_{ikt} \ \forall i \in I, k \in P, t \in T \tag{14}$$

3.3.2 (r,Q) Inventory Policy

Under a classic (r,Q) inventory policy, when the on hand inventory for a product under study drops to some level r or below, an order for Q units is placed to replenish inventory levels. This approach uses a fixed order quantity of Q, rather than the potentially variable order quantity of $S - \sigma$ as discussed above for the classic (s,S) inventory system.

We now introduce two new decision variables for modeling an (r,Q) inventory policy. First, let Q_{ikt} denote the order quantity for product $k \in P$ in warehouse $i \in I$ in time period $t \in T$. This new variable Q_{ikt} replaces the total order quantity variable q_{ikt} used in both the (s,S) and base stock policy formulations above. Further, let r_{ikt} replace s_{ikt} as the reorder point under an (r,Q) inventory policy. Both Q_{ikt} and r_{ikt} are restricted to be non-negative integers $\forall i \in I, k \in P, t \in T$.

With the addition of these new decision variables, four modifications must again be made to the original (s, S) formulation presented above in Section 3.2. First, constraints (3) are modified by updating the notation for the reorder point decision variable:

$$r_{ikt} \geq \omega_{ip} + \sum_{j}\sum_{r}\alpha_{jkt} D_{jk} L_{ijr} \psi_{ijrt} \quad \forall i \in I, k \in P, t \in T \tag{15}$$

The order quantity Q_{ikt} is simply calculated as the difference between the demand shipped and the initial inventory level for each product at each warehouse in each time period:

$$Q_{ikt} \geq \sum_{j}\sum_{r} x_{ijkrt} - \Lambda_{ikt} \quad \forall i \in I, k \in P, t \in T \tag{16}$$

To calculate the average inventory level l_{ik} over the decision horizon, constraints (6) are updated with the new decision variable notation for the (r, Q) system:

$$l_{ik} \geq \sum_{t}\left(\frac{u_t}{\tau}\left(Q_{ikt} + r_{ikt} - 0.5\sum_{j}\sum_{r} x_{ijkrt}\right)\right) \forall i \in I, k \in P \tag{17}$$

Finally, constraints (7) are also updated with the new decision variable notation to calculate the number of orders f_{ikt} that must be placed for each product at each warehouse in each time period:

$$f_{ikt} \geq \frac{Q_{ikt}}{\pi_{ik}} \; \forall i \in I, k \in P, t \in T \tag{18}$$

4 Experimental Study

The integrated inventory-transportation models presented in Section 3 are MIPs containing linear objective functions, linear constraints, and non-negative decision variables, some of which are restricted to the set of positive integers. A full-factorial experimental design was conducted to examine the integrated model's performance under a wide range of operating conditions and inventory policies. Each model was coded in AMPL (Fourer *et al.* 1993), and then solved using the MIP solver in CPLEX (ILOG 1997).

In our experimentation, the key performance measures of interest include the total cost objective function and the inventory positions resulting at each warehouse from the inventory policy under study. A surrogate performance measure of interest was model computation time, measured as seconds required for CPLEX to produce the optimal solution for the problem instance under study.

4.1 Experimental Parameter Settings

Each product's associated size z_k is sampled from a uniform distribution over the interval [1, 200] cubic feet. Safety stock levels ω_{ik} are set at 20% of the maximum realized demand for product $k \in P$ at warehouse $i \in I$, rounded up to the next whole integer. Ordering costs $O_{ik} = 0.50$ dollars, as the Internet has dramatically driven down the cost to place an order. The standard order quantity of π_{ik} units of product $k \in P$ ordered from warehouse $i \in I$ is set to 25% of the maximum realized demand, again rounded up to the next whole integer. The initial supply of Λ_{ik1} units of product $k \in P$ inventory is set to 25% of the maximum realized demand for this product, rounded up to the next whole integer. The holding cost h_{ik} associated with carrying this product at warehouse $i \in I$ is sampled from a uniform distribution over the interval [0.5, 15] dollars.

The decision horizon of interest is $\tau = 3$ days (i.e., 72 hours). Four different periods of time-based service will be modeled: delivery of customer orders within two hours ($u_1 = 2/24$ days, $e_1 = u_1$); within eight hours ($u_2 = \frac{8-2}{24} = 6/24$ days, $e_2 = \sum_{t=1}^{2} u_t = 8/24$ days); within 24 hours ($u_3 = \frac{24-8}{24} = 16/24$ days, $e_3 = 1$ day); and within 72 hours ($u_4 = \frac{72-24}{24} = 2$ days, $e_4 = \tau = 3$ days). In terms of the supplier's contractual obligation to fulfill at least α_{jkt} percent of customer j's total demand for product $k \in P$ in time period $t \in T$, $\alpha_{jkt} \sim U[0.05, 0.05 + 0.1t]$ $\forall t = 1...3$. Finally, $\alpha_{jk4} = 1 - \sum_{t=1}^{3} \alpha_{jkt}$.

Two different temporal schemes will be analyzed during our experiments. First, the *static* case of $\|T\| = 1$ will be examined. Under this scenario, inventory and transportation decisions are made only once, at the beginning of the decision horizon. The second case to be considered is the *dynamic* scenario. The dynamic scenario will examine the potential benefits gained by being able to make and/or update model decisions at the beginning of each time period $t = 1...4$, where t maps to the time-based service periods described in the preceding paragraph.

Three different modes of transportation are assumed to be available for goods shipment. Table 1 describes each mode, along with its associated speed (in miles per hour), capacity (in cubic feet), and both fixed and variable costs which are expressed in dollars. The variable cost values describe the bounds of the uniform distribution from which each value is sampled.

Table 1. Transportation Mode Information

Transportation Mode	**Example**	**Speed (v_r)**	**Capacity (ϑ_r)**	**Fixed Cost (C_r)**	**Variable Cost (c_{ijkr})**
Low Speed	Train	25	9,000	25	U[0.01,0.5]
Medium Speed	Truck	50	3,600	50	U[0.1,1]
High Speed	Plane	500	13,200	500	U[5,50]

The parameter d_{ij} representing the distance between supplier warehouse $i \in I$ and customer $j \in J$ is sampled with equal probability from two different discrete uniform distributions: $d_{ij} = DU[10,250]$ (to represent "nearby" customers) and $d_{ij} = DU[250,1000]$ (to represent "distant" customers). Both order lead times L_{ijr} and the individual elements of the binary incidence matrix A are computed using the realized value of each warehouse to customer distance, time-based service period lengths u_t, and transportation mode speed v_r.

4.2 Experimental Design

Table 2 presents the list of experimental factors, along with the corresponding levels that are investigated in our full-factorial experiment. Demand is assumed to be discrete, as it involves service parts. Therefore, the demand is modeled as a Poisson random variable. The realization of customer j's demand for product $k \in P$ is produced by sampling from a Poisson distribution via uniform distribution transformation. The sampled Poisson demand value is then rounded up to the next whole integer value.

Table 2. Experimental Design Factors and their Corresponding Levels

Factor	Symbol	Levels
Warehouses	I	1, 2
Customers	J	5, 10
Products	P	1, 5
Mean of Poisson Demand	λ_{jk}	1, 2

As many of the experimental parameters described in the previous section are sampled from a probability distribution, 10 replications will be performed for each experimental design factor combination. This leads to a total of $2^4 \times 10 = 160$ model runs. The output for a given problem instance will be compared both in terms of static and dynamic model results, as well as across the inventory policies under study to assess whether any temporal and/or inventory policy approach is significantly better than its competitors.

4.3 Experimental Results

During the experimentation, it is assumed that $(q+s)$ are ordered each period and that these units are available for demand consumption at the beginning of the period. As mentioned earlier, an initial inventory of Λ_{ik1} units of product $k \in P$ inventory are on hand at warehouse $i \in I$. Product demand is assumed to be distributed/realized evenly over the decision period.

Table 3 presents the average values across all 10 replications of the inventory performance measures of interest for different experimental design factor combinations. It is clear that for a given inventory model, the dynamic model wherein decision makers are allowed to update/change their decisions at the beginning of each time period always outperforms the static case, as expected. Average inventory levels, the number of orders placed each period (order frequency), and order quantity all are substantially reduced under a dynamic model paradigm. Further, for the (s, S) and base stock inventory models, both the order up to level (S) and the reorder point (s) associated with the dynamic cases result in significantly lower inventory on hand on average.

Table 3. Experimental Results for Low Demand Cases ($\lambda_{jk} = 1$)

Warehouses	Customers	Inventory Model	Model Type	Inventory Level	Order Frequency	Order Quantity	Order Up To Level	Reorder Point
1	5	(s, S)	Dynamic	11.6	3.1	16.8	26.1	12.2
			Static	55.7	7.8	48.8	90.3	34.7
		Base Stock	Dynamic	14.8	15.4	15.4	15.6	14.5
			Static	32.5	48.8	48.8	36.9	35.9
		(r, Q)	Dynamic	11.6	2.7	16.8		
			Static	53.0	7.8	48.8		
	10	(s, S)	Dynamic	18.6	4.9	29.5	42.4	18.5
			Static	97.9	13.4	88.5	152.9	57.2
		Base Stock	Dynamic	19.3	27.1	27.1	21.8	20.7
			Static	53.3	88.5	88.5	57.9	56.9
		(r, Q)	Dynamic	18.6	4.5	29.5		
			Static	94.8	13.4	88.5		
2	5	(s, S)	Dynamic	6.1	2.7	12.4	14.5	7.5
			Static	19.3	3.7	18.7	35.9	14.5
		Base Stock	Dynamic	6.9	11.5	11.5	9.1	8.0
			Static	13.2	19.6	19.6	17.1	16.1
		(r, Q)	Dynamic	6.0	2.2	12.2		
			Static	19.3	3.7	18.7		
	10	(s, S)	Dynamic	8.1	3.6	20.7	24.4	11.2
			Static	31.8	6.1	41.5	60.3	17.1
		Base Stock	Dynamic	6.8	20.6	20.6	12.4	11.4
			Static	14.5	45.3	45.3	19.0	18.0
		(r, Q)	Dynamic	8.3	3.2	20.5		
			Static	31.8	6.1	41.5		

Similar results are evident when product demand levels are increased (Table 4). Again, all static model instances are outperformed by their dynamic model counterparts. At higher levels of product demand, this performance difference between the static and dynamic cases is magnified. As both the number of warehouses and number of customers per warehouse increase, the three competing inventory models trend towards common average inventory levels and order quantities in the dynamic cases. The base stock policy, which models one-for-one replenishment, is characterized by a significantly higher order frequency, as expected. Clearly, given the resulting levels of ordering frequency, the products modeled in our experiments would not be considered "slow moving" items. The further explains base stock's poor performance in our experimentation, as this inventory approach is best suited for application on slow moving items.

Table 4. Experimental Results for High Demand Cases ($\lambda_{jk} = 2$)

Warehouses	Customers	Inventory Model	Model Type	Inventory Level	Order Frequency	Order Quantity	Order Up To Level	Reorder Point
1	5	(s, S)	Dynamic	50.0	3.8	78.7	114.0	51.9
			Static	230.1	9.2	221.9	379.1	131.8
		Base Stock	Dynamic	38.7	78.5	78.5	55.2	54.2
			Static	119.1	221.8	221.8	132.8	131.8
		(r, Q)	Dynamic	50.0	3.3	78.7		
			Static	230.0	9.2	221.8		
	10	(s, S)	Dynamic	110.4	5.1	160.5	235.6	105.1
			Static	535.4	14.2	486.9	831.2	309.4
		Base Stock	Dynamic	94.2	152.2	152.2	113.4	112.4
			Static	287.4	486.9	486.9	305.9	304.9
		(r, Q)	Dynamic	110.4	4.7	160.5		
			Static	532.7	14.2	486.9		
2	5	(s, S)	Dynamic	24.8	3.1	62.2	67.1	35.5
			Static	63.8	3.6	102.9	128.4	49.2
		Base Stock	Dynamic	23.5	59.0	59.0	37.9	36.8
			Static	36.9	111.4	111.4	50.1	49.1
		(r, Q)	Dynamic	24.8	2.5	62.2		
			Static	63.8	3.6	102.9		
	10	(s, S)	Dynamic	74.2	3.2	92.8	131.5	65.5
			Static	339.1	7.5	265.4	529.4	225.6
		Base Stock	Dynamic	70.2	86.0	86.0	70.8	69.7
			Static	147.0	239.9	239.9	167.2	166.2
		(r, Q)	Dynamic	74.2	2.6	92.7		
			Static	339.1	7.5	265.4		

With this knowledge, we now consider the total costs associated with each inventory model and temporal paradigm. Table 5 presents the average ratio across all 10 problem instances of static model total cost to dynamic model total cost for each inventory model type for both levels of product demand. As all ratios of static to dynamic total cost are greater than one, this demonstrates that the dynamic model paradigm produces superior performance in terms of total costs, in addition to providing the "best" inventory solution. Table 5's results are not surprising, as the primary difference in objective function values between static and dynamic model instances is attributable to inventory costs, primarily ordering and holding costs.

Table 5. Average Ratio of Static Model Total Cost to Dynamic Model Total Cost

Warehouses	Customers	Inventory Model	Demand Level $\lambda_{jk} = 1$	$\lambda_{jk} = 2$
1	5	(s, S)	1.298	1.571
		Base Stock	1.116	1.274
		(r, Q)	1.280	1.571
	10	(s, S)	1.268	1.610
		Base Stock	1.110	1.280
		(r, Q)	1.256	1.599
2	5	(s, S)	1.270	1.591
		Base Stock	1.095	1.306
		(r, Q)	1.271	1.592
	10	(s, S)	1.265	1.702
		Base Stock	1.109	1.269
		(r, Q)	1.266	1.703

It should be noted that even though the ratios in Table 5 for the base stock inventory model are often smaller in value (i.e., approximately 1.11 vs. 1.25+ on average for the $\lambda_{jk} = 1$ case), both (s, S) and (r, Q) inventory approaches produced inventory solutions with less total cost on average (Table 6). The larger ratio values in Table 5 simply suggest that less of a difference exists between static and dynamic total cost values for the base stock inventory approach as compared to the other two competing inventory approaches. Even though base stock produces superior total costs under static decision-making conditions, switch to a dynamic temporal paradigm also results in a change in the "best" inventory modeling policy in terms of reducing average total costs.

Table 6. Average Total Cost Comparison for Two Levels of Demand

Inventory Model	**Model Type**	**Demand Level** $\lambda_{jk} = 1$	$\lambda_{jk} = 2$
(s, S)	Dynamic	$ 4,574	$ 9,783
	Static	$ 5,934	$ 16,383
Base Stock	Dynamic	$ 4,650	$ 10,024
	Static	$ 5,220	$ 13,119
(r, Q)	Dynamic	$ 4,571	$ 9,781
	Static	$ 5,890	$ 16,359

5 Conclusions and Future Research

In this research, we present a mixed integer programming model for integrated inventory and transportation decision making for the service parts logistics industry. A number of model formulations are presented that correspond to various static and time-dependent inventory policies. In all cases, customer demand is subject to pre-specified time-based service levels. Experimental results suggest that employing a dynamic temporal modeling scheme wherein decision makers are allowed to update/change their decisions at the beginning of each time period results in significantly reduced inventory levels, order quantities, and total supply chain costs.

On average, total costs were 21.3% (52.7%) higher for the static decision-making instances as compared to the dynamic decision-making instances when the supply chain was subject to low (high) product demand. The experimental results suggest only a very small difference in performance exists between the (s,S) and (r,Q) inventory modeling approaches. This difference is statistically insignificant at the $\alpha = 0.05$ level. The supply chain modeled in this research did not lend itself to a base stock inventory policy, as the demanded items are not easily classified as "slow moving" due to their demand levels.

Future research should consider a number of important extensions to this research effort for added insights and realism. For example, period review inventory policies could be investigated and compared with the continuous review policies modeled in this research. Customer service level could be modeled as an optimization model output variable that depends on both time period and demand requirements. In addition, different time period configurations could be examined to further understand the dynamic behavior of the system. The availability of less-than-truckload (LTL) could be incorporated into the model for additional testing to represent this important segment of the current transportation logistics landscape. Finally, backorder costs were not considered as part of this research effort. The inclusion of these costs into the objective function could provide additional useful insights into system performance.

6 References

Bertazzi, L., Paletta, G., Speranza, M. G. 2002, Deterministic Order-Up-To level policies in an inventory routing problem, *Transportation Science,* 36(1) 119-132.

Buffa, F. P., Bryant, T. R., 1980. Reflecting logistics costs in customer service level targets, *Production and Inventory Management*, 21(1) 45-55.

Cohen, M. A., Zheng, Y., Agarwal, A., 1997. Service parts logistics: A benchmark analysis, *IIE Transactions*, 29 627-639.

Cole, M. H. 1996. Service considerations and the design of strategic distribution systems, Doctoral Dissertation, Department of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia.

Daskin M. S., Coullard, C. R., 2002. An inventory-location model: Formulation, solution algorithm and computational results, *Annals of Operations Research* 110 83-106

Erlebacher, S. J., Meller, R. D., 2000. The interaction of location and inventory in designing distribution systems, *IIE Transactions*, 32 155-166.

Evers, P. T., 1999. The effect of lead times on safety stock. *Production and Inventory Management Journal,* 40(1) 6-10.

Fourer, R., Gay, D.M., Kernighan, B.W., 1993. *AMPL: A Modeling Language for Mathematical Programming*, Scientific Press, Massachusetts.

Hwang, H. 2002. Design of supply chain logistics system considering service level. *Computers & Industrial Engineering,* 43(1-2) 283-297.

ILOG, Incorporated, 1997. *Using the CPLEX Callable Library*, ILOG, Nevada.

Jayaraman, V., 1998. Transportation, facility location and inventory issues in distribution network design – An investigation. *International Journal of Operations & Production Management,* 18(5) 471-494.

Lehmusvaara, A., 1998. Transport time policy and service level as components in logistics strategy: A case study. *Production Economics,* 56-57 379-387.

Perl, J., Sirisoponsilp, S. 1988. Distribution Networks: Facility location, transportation and inventory. *International Journal of Physical Distribution and Logistics Management,* 18(6) 18-26.

Qu, W. W., Bookbinder, J. H., Iyogun, P., 1999. An integrated inventory-transportation system with modified periodic policy for multiple products, *European Journal of Operational Research*, 115 254-269.