



An Inventory-Location Model: Formulation, Solution Algorithm and Computational Results

MARK S. DASKIN and COLLETTE R. COULLARD

*Department of Industrial Engineering and Management Sciences, Northwestern University,
Evanston, IL 60208, USA*

ZUO-JUN MAX SHEN

*Department of Industrial and Systems Engineering, University of Florida, 303 Weil Hall,
P.O. Box 116595, Gainesville, FL 32611-6595, USA*

Abstract. We introduce a distribution center (DC) location model that incorporates working inventory and safety stock inventory costs at the distribution centers. In addition, the model incorporates transport costs from the suppliers to the DCs that explicitly reflect economies of scale through the use of a fixed cost term. The model is formulated as a non-linear integer-programming problem. Model properties are outlined. A Lagrangian relaxation solution algorithm is proposed. By exploiting the structure of the problem we can find a low-order polynomial algorithm for the non-linear integer programming problem that must be solved in solving the Lagrangian relaxation subproblems. A number of heuristics are outlined for finding good feasible solutions. In addition, we describe two variable forcing rules that prove to be very effective at forcing candidate sites into and out of the solution. The algorithms are tested on problems with 88 and 150 retailers. Computation times are consistently below one minute and compare favorably with those of an earlier proposed set partitioning approach for this model (Shen, 2000; Shen, Coullard and Daskin, 2000). Finally, we discuss the sensitivity of the results to changes in key parameters including the fixed cost of placing orders. Significant reductions in these costs might be expected from e-commerce technologies. The model suggests that as these costs decrease it is optimal to locate additional facilities.

1. Introduction

Managing inventory has become a major challenge for firms as they simultaneously try to reduce costs and improve customer service in today's increasingly competitive business environment. Managing inventory consist of two critical tasks. First, we must determine the *number* of stocking locations or distribution centers to have. Second, we must determine the *amount* of inventory to maintain at each of the centers. Often these tasks are undertaken separately, resulting in a degree of suboptimization. In this paper, we outline an integrated approach to determining the number of distribution centers to establish and the magnitude of inventory to maintain at each center.

We argue that the importance of inventory management as outlined above has, if anything, been heightened by e-commerce. Whether it is business to consumer (B2C) or business to business (B2B), e-commerce end customers expect high levels of service and speedy deliveries. At the same time, many of the inputs to traditional inventory management decision-making are changing rapidly. E-commerce offers the hope of sig-

nificantly reducing order costs thereby allowing smaller more frequent shipments. The model we propose below explicitly includes order shipment costs. As a result it is capable of estimating the impacts of sharply reduced ordering costs on the number of distribution centers, their locations, and the optimal inventory ordering policy. E-commerce can also reduce variability across the supply chain by making end customer orders visible throughout the chain. This visibility should reduce the bullwhip effect [25,26,33]. Our model does not directly account for these effects of e-commerce though the model does explicitly consider safety stock inventories at the distribution centers which are a function of the variability of customer demand.

This work was motivated by a study at a local blood bank conducted by two of the authors. The blood bank supplied roughly 30 hospitals in the greater Chicago area. Our focus was on the production and distribution of platelets, the most expensive and most perishable of all blood products. If a unit of platelets is not used within 5 days of the time it is produced from whole blood, it must be destroyed. The demand for platelets is highly variable as they are needed in only a limited number of medical contexts. When they are used, however, multiple units are often needed. The hospitals supplied by the blood bank collectively owned the blood bank and set prices. As a result they could return a unit of platelets up to the time it outdated and not be charged for it. Thus, there was little incentive to manage inventories in an efficient manner. Many of the larger hospitals ordered almost twice the number of platelet units that they used each year resulting in the need to destroy thousands of units of this expensive blood product. Other hospitals ordered almost all of their needed platelets on a STAT or emergency basis. The blood bank often had to ship the units to these hospitals using a taxi or express courier at significant expense to the system. Clearly an improved system was needed. We proposed a revision to their pricing policies along with a system in which selected hospitals would maintain an inventory of platelets for use in neighboring hospitals. This would allow the system to take advantage of the risk-pooling effect. We selected the hospitals at which inventory would be maintained using a P -median model [7,18,19]. The model did not account directly for the working inventory or safety stock (risk-pooling) effects, for the transport costs from the blood bank to the selected hospitals or the fixed costs of establishing the facilities. Clearly, this was only a first-cut approach. The research outlined below is aimed, in part, at developing a more comprehensive and more accurate model of such a situation.

The remainder of this paper is organized as follows. In section 2 we review relevant related literature. The model we propose is formulated in section 3. The formulation we obtain is a mixed integer non-linear programming problem which can be viewed as an extension of the traditional uncapacitated fixed charge facility location problem. In addition to the standard facility location and local distribution costs, the model includes cost components representing working and safety stock inventories at the distribution centers as well as transport costs from the supplier(s) to the distribution centers. These inventory and supplier-to-DC transport costs introduce significant non-linearities into the model. Section 4 outlines a number of key properties of the model we propose and introduces an additional assumption on the demand distributions that we use to develop

a solution algorithm for the problem. Section 5 outlines our solution procedure for the problem. Computational results are presented in section 6. In section 7, we present our conclusions and outline directions for future work.

2. Literature review

Inventory theory literature tends to focus on finding optimal inventory replenishment strategies at the DCs and the retailer outlets. This work usually assumes that the number and locations of the DCs are given. See, for example, Graves et al. [17], Nahmias [29] and Zipkin [38]. On the other hand, location theory tends to focus on developing models for determining the number of DCs and their locations, as well as the DC-retailer assignments. This work usually includes fixed facility setup costs and transportation costs, but the operational inventory and shortage costs are typically ignored. Daskin and Owen [8] provide an overview of facility location modeling as do the recent texts by Daskin [7] and Drezner [9].

Eppen [10] studied the so-called “risk pooling effects”, namely the effects of inventory-cost savings achieved by grouping retailers. Assume customer demands are normally distributed with a mean μ_i and a standard deviation σ_i for customer i . Then the expected total inventory cost under the decentralized mode for n retailers is $K \sum_{i=1}^n \sigma_i$. If the demands of the n retailers are independent, the optimal cost under a centralized mode can be expressed by $K \sqrt{\sum_{i=1}^n \sigma_i^2}$, which is less than $K \sum_{i=1}^n \sigma_i$, where K is a constant depending on the holding and penalty costs and the standard normal loss function.

Recently, there are several new studies that combine inventory management and routing decisions. For example, Federgruen and Zipkin [14], Federgruen and Simchi-Levi [13], Viswanathan and Mathur [36], Chan, Federgruen, and Simchi-Levi [5] and Kleywegt, Nori, and Savelsbergh [21].

Also, several models combine location and routing decisions; for instance, Laporte [23], Min, Jayaraman and Srivastava [27], Laporte and Dejax [24], Berman, Jaillet and Simchi-Levi [4], and Berger, Coullard and Daskin [3].

Shen, Coullard and Daskin [32] studied model presented below, but they use a set partitioning approach. We compare our computational results with theirs. Teo, Ou and Goh [35] studied the impact on inventory costs with consolidation of distribution centers. They also propose an algorithm that solves for a distribution system with the total fixed facility location cost and inventory costs within $\sqrt{2}$ of the optimal. But they ignore the costs to ship from the supplier to the DCs and from the DCs to the retailers in their model. Finally, Erlebacher and Meller [11] formulate a highly non-linear integer location-inventory model. They attack the problem by using a continuous approximation as well as a number of construction and bounding heuristics. For problems with 16 customers, they obtained solutions that were between 3.78% and nearly 36% of a lower bound. An exchange heuristic improved the solution considerably. Computation times

on a 600 node problem using the exchange heuristic averaged 117 hours on a Sun Ultra Sparcstation.

3. Model formulation

We consider a three-tiered system consisting of one or more suppliers, distribution centers and retailers. We assume that the locations of the suppliers and the retailers are known and that the suppliers have infinite capacity at least from the perspective of the system being modeled. The problem is to determine the optimal number of distribution centers, their locations, the retailers assigned to each distribution center, and the optimal ordering policy at the distribution centers. We do not explicitly model the inventory maintained by the retailers themselves. A key problem is that the demand that is seen by each distribution center is a function of the demands at the retailers assigned to the distribution center. Thus, the inventory policy – the reorder interval, reorder size, and safety stock – at the distribution center is a function of the assignment of retailers to the distribution center. Since these assignments are not known a priori, the inventory policy must also be endogenously determined.

The formulation allows for multiple production plants but assumes that the uncapacitated distribution centers receive the product from the (uncapacitated) plant with the smallest total shipping cost to the DC. Since we will be further assuming that the shipping costs to the DC depend only on the DC and possibly on the distance between the DC and the plant and not on the plant per se, the identity of the best plant for each DC can readily be identified a priori. We also will assume that the plant to DC lead time is the same for all plant/DC combinations and so the choice of the plant to supply each DC will not affect the required level of safety stock at the DCs. With these assumptions, the model is mathematically equivalent to assuming that there is only a single plant.

In what follows, we will model the inventory problem faced by the DC using a (Q, r) inventory model with type I service [20,29]. Axsater [1] notes that it is common to approximate the (Q, r) model using two steps with the order quantity determined by an EOQ model in which the mean demand is used to represent the stochastic demand process and the reorder point is determined in a second step. Zheng [38] showed that the maximum relative error introduced by using the EOQ quantity instead of the optimal (Q, r) quantity is 0.125. Axsater [1] later showed that this bound can be tightened slightly to 0.118. Zheng notes (p. 99) that, in most cases, the relative increase in total costs will be much less than these worst-case bounds. Thus, this is the approach we adopt. Specifically, we will approximate the (Q, r) model by assuming that the DC orders inventory from the plant using an EOQ model. The reorder point, and hence the safety stock, will be determined to ensure that the probability of a stockout at the DC is less than or equal to some specified value.

To begin modeling the problem, let us assume for the moment that we know which customers are to be assigned to a specific distribution center. Assume that the demand at each retailer is Normally distributed with a daily mean of μ_i and a daily variance of σ_i^2 and let S be the set of customers assigned to the distribution center. Let L be the lead

time in days for deliveries from the supplier to the distribution center. Assuming that the daily demands at each retailer are uncorrelated over time and across retailers, the lead time demand at the distribution center is Normally distributed with a mean of $L \sum_{i \in S} \mu_i$ and a variance of $L \sum_{i \in S} \sigma_i^2$. The safety stock required to ensure that stockouts occur with a probability of α or less is

$$z_\alpha \sqrt{L \sum_{i \in S} \sigma_i^2} \quad (1)$$

where z_α is a standard Normal deviate such that $P(z \leq z_\alpha) = \alpha$.

For the moment, let D be the expected annual demand (i.e., $D = \chi \sum_{i \in S} \mu_i$), where χ is a constant used to convert daily demand into annual demand (e.g., 365 if demands occur every day of the year), let h be the holding cost per item per year and let F be the fixed cost of placing an order from the distribution center to the supplier. Then the annual cost of ordering inventory from the supplier at the distribution center is approximated by

$$Fn + \beta v\left(\frac{D}{n}\right)n + \theta \frac{hD}{2n} \quad (2)$$

where n is the (unknown) number of orders per year, $v(x)$ is the cost of shipping an order of size x from the supplier, and β and θ are weights that we assign to transportation and inventory costs respectively so that we can later test the effects of varying the importance of these costs relative to the fixed facility costs.

The first term of (2) represents the total fixed cost of placing n orders per year. The second term represents the shipment cost $v(D/n)$ per shipment multiplied by the number of shipments per year and the weight, β , associated with transport. D/n is the expected shipment size per shipment. The third term is cost of the average working inventory. On average, there will be $D/(2n)$ items of working inventory on hand at a cost of h per item per year. Taking the derivative of this expression with respect to n , the number of orders per year, and setting the derivative to zero, we obtain

$$F + \beta v\left(\frac{D}{n}\right) - \beta n v'\left(\frac{D}{n}\right) \frac{D}{n^2} - \theta \frac{hD}{2n^2} = F + \beta v\left(\frac{D}{n}\right) - \beta v'\left(\frac{D}{n}\right) \frac{D}{n} - \theta \frac{hD}{2n^2} = 0. \quad (3)$$

If $v(x)$ is linear in x (e.g., if $v(x) = g + ax$) then $v'(x)$ is a constant (e.g., a) and the expression above becomes

$$F + \beta g + \beta a \frac{D}{n} - \beta a \frac{D}{n} - \theta \frac{hD}{2n^2} = F + \beta g - \theta \frac{hD}{2n^2} = 0. \quad (4)$$

In this case, g is the fixed cost of a shipment from the plant to the DC and a is the volume dependent cost of a shipment from the plant to the DC. Solving for n , we obtain

$n = \sqrt{\theta h D / (2(F + \beta g))}$. Substituting this into the cost function (2) we obtain an annual working inventory cost of

$$\begin{aligned} & F \sqrt{\frac{\theta h D}{2(F + \beta g)}} + \beta g \sqrt{\frac{\theta h D}{2(F + \beta g)}} + \beta a D + \theta \frac{h D}{2} \sqrt{\frac{2(F + \beta g)}{\theta h D}} \\ &= \sqrt{\frac{\theta h D (F + \beta g)}{2}} + \beta a D + \sqrt{\frac{\theta h D (F + \beta g)}{2}} = \sqrt{2 \theta h D (F + \beta g)} + \beta a D. \end{aligned} \quad (5)$$

Recall that this cost includes the costs of placing orders at the distribution center, transporting goods from the supplier to the distribution center and holding the working inventory at the DC.

Unfortunately, the derivation of the safety stock (1) and the working inventory cost (5) assumes that we know the assignment of retailers to the distribution center as denoted by the set S . The identity of this set is *not* known a priori and must be determined endogenously.

To simultaneously determine the locations of the DCs, assignments of the retailers to the DCs, and working and safety stock inventory costs, let us define the following additional inputs and sets:

- I set of retailers indexed by i ,
- J set of candidate DC sites indexed by j ,
- f_j fixed (annual) cost of locating a distribution center at candidate site j , for each $j \in J$,
- d_{ij} cost per unit to ship between retailer i and candidate DC site j , for each $i \in I$ and $j \in J$,
- χ days per year (used to convert daily demand and variance values to annual values).

In addition, we define the following decision variables:

$$\begin{aligned} X_j &= \begin{cases} 1 & \text{if we locate at candidate site } j, \\ 0 & \text{if not,} \end{cases} \\ Y_{ij} &= \begin{cases} 1 & \text{if demands at retailer } i \text{ are assigned to a DC at candidate site } j, \\ 0 & \text{if not.} \end{cases} \end{aligned}$$

With this notation, we can formulate the problem as follows:

$$\begin{aligned} \text{Minimize } & \sum_{j \in J} f_j X_j + \left(\beta \sum_{j \in J} \sum_{i \in I} \chi d_{ij} \mu_i Y_{ij} \right) \\ & + \left(\sum_{j \in J} \sqrt{2 \theta h (F_j + \beta g_j)} \sum_{i \in I} \chi \mu_i Y_{ij} + \beta \sum_{j \in J} a_j \sum_{i \in I} \chi \mu_i Y_{ij} \right) \\ & + \theta h z_\alpha \sum_{j \in J} \sqrt{\sum_{i \in I} L \sigma_i^2 Y_{ij}} \end{aligned} \quad (6)$$

$$\text{subject to } \sum_{j \in \mathbf{J}} Y_{ij} = 1, \quad \forall i \in \mathbf{I}, \quad (7)$$

$$Y_{ij} \leq X_j, \quad \forall i \in \mathbf{I}, \forall j \in \mathbf{J}, \quad (8)$$

$$X_j \in \{0, 1\}, \quad \forall j \in \mathbf{J}, \quad (9)$$

$$Y_{ij} \in \{0, 1\}, \quad \forall i \in \mathbf{I}, \forall j \in \mathbf{J}. \quad (10)$$

The first term of the objective function (6) is the fixed cost of locating facilities. The second term represents the local delivery cost. Note that $\sum_{i \in \mathbf{I}} \chi \mu_i Y_{ij}$ represents the total annual demand assigned to distribution center j . Thus, the third term represents the total working inventory cost (5) where we have added the obvious subscripts j to the fixed and unit shipping costs g and a respectively obtaining g_j and a_j . In addition, the fixed cost of placing an order, F , has been made DC-specific, obtaining F_j . The fourth term represents the safety stock inventory cost. Constraint (7) states that each demand node must be assigned to a DC. Constraint (8) stipulates that the assignments can only be made to open DCs. Finally, constraints (9) and (10) are standard integrality constraints, with (10) representing single-sourcing constraints, meaning that all of the demand at a retailer must be assigned to the same DC.

The objective function can be rearranged as follows:

$$\begin{aligned} \text{Minimize } & \sum_{j \in \mathbf{J}} f_j X_j + \left(\beta \sum_{j \in \mathbf{J}} \sum_{i \in \mathbf{I}} \chi \mu_i (d_{ij} + a_j) Y_{ij} \right) \\ & + \left(\sum_{j \in \mathbf{J}} \sqrt{2\theta h (F_j + \beta g_j) \sum_{i \in \mathbf{I}} \chi \mu_i Y_{ij}} \right) \\ & + \theta h z_\alpha \sum_{j \in \mathbf{J}} \sqrt{\sum_{i \in \mathbf{I}} L \sigma_i^2 Y_{ij}} \\ = & \sum_{j \in \mathbf{J}} \left\{ f_j X_j + \sum_{i \in \mathbf{I}} \hat{d}_{ij} Y_{ij} + K_j \sqrt{\sum_{i \in \mathbf{I}} \mu_i Y_{ij}} + \Theta \sqrt{\sum_{i \in \mathbf{I}} \hat{\sigma}_i^2 Y_{ij}} \right\} \quad (11) \end{aligned}$$

where

$$\begin{aligned} \hat{d}_{ij} &= \beta \chi \mu_i (d_{ij} + a_j), \\ K_j &= \sqrt{2\theta h \chi (F_j + \beta g_j)}, \\ \Theta &= \theta h z_\alpha, \\ \hat{\sigma}_i^2 &= L \sigma_i^2. \end{aligned}$$

In (11) we have grouped the linear term representing the marginal transport costs from the supplier to the distribution center (represented by terms in a_j) with the local delivery costs (represented by terms in d_{ij}). Thus, \hat{d}_{ij} captures local delivery costs from the DCs to the retailers as well as the marginal cost of shipping a unit from a supplier to a DC; K_j captures the working inventory effects due to the fixed ordering costs at the DC as

well as the fixed transport costs from a supplier to a DC. Finally, Θ captures the safety stock costs at the distribution centers.

The constraints (7)–(10) are identical to those of the traditional uncapacitated fixed charge facility location problem [2,12,22]. Thus, the solution approach that we propose below mirrors some of the solution approaches used for that problem. However, these approaches must be modified significantly to account for the final two non-linear terms in the objective function (11).

4. Model properties

Before outlining the solution approach, it is worth noting a number of properties of the model. First, unlike the traditional uncapacitated fixed charge location model in which it is always optimal to assign demands to the facility that can serve the demands at least cost (i.e., the facility j with the smallest value of d_{ij} for retailer i), in this problem it may be optimal to assign retailers to a more remote distribution center. Doing so may reduce the working inventory, safety stock inventory and transport costs from the supplier to the DC sufficiently to offset the increased local delivery costs. Figure 1 illustrates a small example in which this occurs. Table 1 provides the input data and parameters for the problem while table 2 compares the total costs for different DC locations and different customer assignments for the problem. The first column gives the DC locations; the next three give the retailer assignments to DCs; the next three give the facility, local delivery and inventory cost terms; and the last two columns give the total cost and the difference between the optimal total cost and the cost of the solution indicated in that row. In this example, it is optimal to assign demands at retailer B to a facility at A with a cost per unit shipped of 101 despite the fact that there is a facility at C with a smaller unit shipping cost for the B to C channel.

Not only does this phenomenon occur in small (contrived) cases, but we have found that it occurs in our computational results using realistic national demand data, particularly when the inventory-related costs are large relative to the other costs (i.e., when θ is large relative to β). Typically we find this behavior associated with small retail nodes that are almost equidistant to two different DCs. We also can construct examples

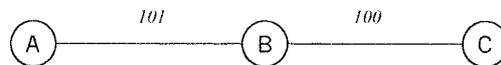


Figure 1. Sample figure for non-closest assignment.

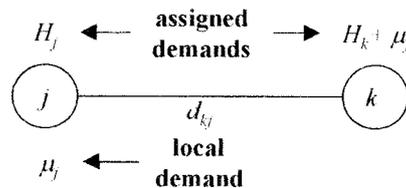


Figure 2. Small network for proof of theorem 1.

Table 1a
Data for network in figure 1.

	Node		
	A	B	C
Fixed cost	10	1000	10
Mean	100	1	1

Table 1b
Parameters for network in figure 1.

F	Fixed order cost	10
$g_j, \forall j \in J$	Fixed shipping cost from supplier to DC	1
$a_j, \forall j \in J$	Unit shipping cost from supplier to DC	1
β	Transport weight	1
z_α	Service level parameter	1.96
θ	Inventory weight	1
L	Lead time	1
h	Unit holding cost	2
χ	Days per year	1

Table 2
Costs of different DC locations and retailer assignments.

Facilities	Retailer Assignments			Facility cost	Local delivery cost	Inventory cost	Total cost	Cost – min cost
	A	B	C					
A	A	A	A	10	304	106.58	420.58	181.97
B	B	B	B	1,000	10,301	106.58	11,407.58	11,168.97
C	C	C	C	10	20,301	106.58	20,417.58	20,178.97
A,C	A	C	C	20	101	120.46	241.46	2.84
A,C	A	A	C	20	102	116.61	238.61	0.00
A,B	A	B	B	1,010	101	120.46	1,231.46	992.84
A,B,C	A	B	C	1,020	0	126.64	1,146.64	908.03

in which it is optimal to *locate* a DC at a particular node, but for demands from that node to be assigned to a different DC. In other words, it is conceivable that the optimal solution would locate a facility in Chicago, for example, but would assign demands from Chicago to a facility in Minneapolis. Again, the intuitive reason for this is that the reduction in inventory and supplier-to-DC transport costs (all of which entail risk-pooling terms) more than offset the increased local delivery costs from the DC to the retailers. The reader is referred to Shen [31] or Shen, Coullard and Daskin [32] for a simple example in which this occurs. In practice, it is unlikely that a supply chain manager would organize her distribution system in this way even if doing so would result in small cost savings. Fortunately, we will be able to prove that this does not occur once we make one additional assumption.

To simplify the model somewhat further, we assume that the variance-to-mean ratio at each retailer is identical for all retailers. In other words, we assume that $\sigma_i^2/\mu_i = \gamma$, $\forall i \in \mathbf{I}$. While this may seem like a restrictive assumption, if demands arise from a Poisson process, then this assumption is exact and we are merely approximating a Poisson demand process by Normally distributed demands, which is a good approximation for sufficiently large demand values [28]. Also, in many other contexts (e.g., transportation planning with random travel times) it is common to assume a single variance-to-mean ratio [30]. With this assumption, the objective function can be rewritten as

$$\begin{aligned}
\text{Minimize } & \sum_{j \in \mathbf{J}} \left\{ f_j X_j + \sum_{i \in \mathbf{I}} \hat{d}_{ij} Y_{ij} + K_j \sqrt{\sum_{i \in \mathbf{I}} \mu_i Y_{ij}} + \Theta \sqrt{\sum_{i \in \mathbf{I}} \hat{\sigma}_i^2 Y_{ij}} \right\} \\
& = \sum_{j \in \mathbf{J}} \left\{ f_j X_j + \sum_{i \in \mathbf{I}} \hat{d}_{ij} Y_{ij} + K_j \sqrt{\sum_{i \in \mathbf{I}} \mu_i Y_{ij}} + \Theta \sqrt{\sum_{i \in \mathbf{I}} L \gamma \mu_i Y_{ij}} \right\} \\
& = \sum_{j \in \mathbf{J}} \left\{ f_j X_j + \sum_{i \in \mathbf{I}} \hat{d}_{ij} Y_{ij} + \hat{K}_j \sqrt{\sum_{i \in \mathbf{I}} \mu_i Y_{ij}} \right\} \tag{12}
\end{aligned}$$

where $\hat{K}_j = K_j + \Theta \sqrt{L \gamma}$.

This assumption does two things. First, it reduces the number of non-linear terms from two to one. This will facilitate the solution approach outlined below in section 5. Second, as shown in the following theorem, under this assumption it is never optimal to open a DC at a node and to then serve the demands from that node from another DC.

Before stating the theorem, we note that if customer c is optimally assigned to a distribution center at k , then for any other open DC j ,

$$\mu_c \tilde{d}_{cj} + \hat{K}_j [(H_j + \mu_c)^{1/2} - H_j^{1/2}] \geq \mu_c \tilde{d}_{ck} + \hat{K}_k [H_k^{1/2} - (H_k - \mu_c)^{1/2}], \quad \forall j \neq k,$$

where $\tilde{d}_{cj} = \beta \chi (d_{cj} + a_j) = \hat{d}_{cj} / \mu_c$ and H_j is the total demand assigned to a facility at j . In other words, the incremental cost of assigning customer c to DC k is less than or equal to the incremental cost of assigning c to any other DC j . We now have the following theorem.

Theorem. If $(d_{ck} + a_k) - (d_{cj} + a_j) \geq (d_{ek} + a_k) - (d_{ej} + a_j)$ (or equivalently, if $d_{ck} - d_{cj} \geq d_{ek} - d_{ej}$) for all DCs $j (j \neq k)$, $\sigma_i^2/\mu_i = \gamma$, $\forall i \in \mathbf{I}$, all mean demands are positive, and if customer c is optimally assigned to distribution center k , then customer e is also optimally assigned to DC k .¹

¹ The authors gratefully acknowledge Ann Chan who suggested this theorem as a replacement for a weaker theorem in an earlier draft of the paper that simply showed that if the variance to mean ratio is the same for all customers then demands at a DC node must be assigned to that DC.

Proof. Assume that customer e is optimally assigned to some DC j other than k . Then

$$\begin{aligned} & \mu_e \tilde{d}_{ej} + \hat{K}_j [H_j^{0.5} - (H_j - \mu_e)^{0.5}] \\ &= \mu_e \tilde{d}_{ej} + \hat{K}_j [H_j^{0.5} - (H_j - \mu_e)^{0.5}] \\ & \quad - \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{cj} + \hat{K}_j [(H_j + \mu_c)^{0.5} - H_j^{0.5}] \} \\ & \quad + \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{cj} + \hat{K}_j [(H_j + \mu_c)^{0.5} - H_j^{0.5}] \} \end{aligned} \quad (13)$$

$$\begin{aligned} &= \mu_e [\tilde{d}_{ej} - \tilde{d}_{cj}] + \hat{K}_j \left\{ [H_j^{0.5} - (H_j - \mu_e)^{0.5}] \right. \\ & \quad \left. - \frac{\mu_e}{\mu_c} [(H_j + \mu_c)^{0.5} - H_j^{0.5}] \right\} \\ & \quad + \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{cj} + \hat{K}_j [(H_j + \mu_c)^{0.5} - H_j^{0.5}] \} \end{aligned} \quad (14)$$

$$\begin{aligned} &> \mu_e [\tilde{d}_{ek} - \tilde{d}_{ck}] \\ & \quad + \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{ck} + \hat{K}_k [H_k^{0.5} - (H_k - \mu_c)^{0.5}] \} \end{aligned} \quad (15)$$

$$\begin{aligned} &> \mu_e [\tilde{d}_{ek} - \tilde{d}_{ck}] + \hat{K}_k \left\{ [(H_k + \mu_e)^{0.5} - H_k^{0.5}] \right. \\ & \quad \left. - \frac{\mu_e}{\mu_c} [H_k^{0.5} - (H_k - \mu_c)^{0.5}] \right\} \\ & \quad + \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{ck} + \hat{K}_k [H_k^{0.5} - (H_k - \mu_c)^{0.5}] \} \end{aligned} \quad (16)$$

$$\begin{aligned} &= \mu_e (\tilde{d}_{ek}) + \hat{K}_k [(H_k + \mu_e)^{0.5} - H_k^{0.5}] \\ & \quad - \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{ck} + \hat{K}_k [H_k^{0.5} - (H_k - \mu_c)^{0.5}] \} \\ & \quad + \frac{\mu_e}{\mu_c} \{ \mu_c \tilde{d}_{ck} + \hat{K}_k [H_k^{0.5} - (H_k - \mu_c)^{0.5}] \} \end{aligned} \quad (17)$$

$$= \mu_e \tilde{d}_{ek} + \hat{K}_k [(H_k + \mu_e)^{0.5} - H_k^{0.5}]. \quad (18)$$

Equality (13) holds by simple addition and subtraction of the second and third (identical) terms. Equality (14) holds by regrouping of terms. Inequality (15) holds because

$$(a) \quad \tilde{d}_{ck} - \tilde{d}_{cj} \geq \tilde{d}_{ek} - \tilde{d}_{ej} \text{ or } \tilde{d}_{ej} - \tilde{d}_{cj} \geq \tilde{d}_{ek} - \tilde{d}_{ck},$$

(b) the assumed optimality of assigning c to k means that

$$\mu_c \tilde{d}_{cj} + \hat{K}_j [(H_j + \mu_c)^{1/2} - H_j^{1/2}] \geq \mu_c \tilde{d}_{ck} + \hat{K}_k [H_k^{1/2} - (H_k - \mu_c)^{1/2}]$$

and

(c) concavity of the square root term, meaning

$$\left\{ [H_j^{0.5} - (H_j - \mu_e)^{0.5}] - \frac{\mu_e}{\mu_c} [(H_j + \mu_c)^{0.5} - H_j^{0.5}] \right\}$$

can be rewritten as follows after multiplying both sides by μ_c :

$$\begin{aligned} & \mu_c [H_j^{0.5} - (H_j - \mu_e)^{0.5}] - \mu_e [(H_j + \mu_c)^{0.5} - H_j^{0.5}] \\ &= \mu_c H_j^{0.5} + \mu_e H_j^{0.5} - \mu_c (H_j - \mu_e)^{0.5} - \mu_e (H_j + \mu_c)^{0.5} \\ &= (\mu_c + \mu_e) H_j^{0.5} - \mu_c (H_j - \mu_e)^{0.5} - \mu_e (H_j + \mu_c)^{0.5} \\ &= (\mu_c + \mu_e) \{ H_j^{0.5} - \alpha (H_j - \mu_e)^{0.5} - (1 - \alpha) (H_j + \mu_c)^{0.5} \} \\ &> 0 \end{aligned}$$

where $\alpha = \mu_c / (\mu_c + \mu_e)$ and the inequality follows from concavity of the square root operator since

$$H_j = \frac{\mu_c}{(\mu_c + \mu_e)} (H_j - \mu_e) + \frac{\mu_e}{(\mu_c + \mu_e)} (H_j + \mu_c).$$

Inequality (16) holds because of concavity of square root term meaning

$$\left[(H_k + \mu_e)^{0.5} - H_k^{0.5} \right] - \frac{\mu_e}{\mu_c} [H_k^{0.5} - (H_k - \mu_c)^{0.5}]$$

can be rewritten as follows after multiplying both sides by μ_c :

$$\begin{aligned} & \mu_c (H_k + \mu_e)^{0.5} - \mu_c H_k^{0.5} - \mu_e H_k^{0.5} + \mu_e (H_k - \mu_c)^{0.5} \\ &= \mu_c (H_k + \mu_e)^{0.5} + \mu_e (H_k - \mu_c)^{0.5} - (\mu_c + \mu_e) H_k^{0.5} \\ &= (\mu_c + \mu_e) \{ \alpha (H_k + \mu_e)^{0.5} + (1 - \alpha) (H_k - \mu_c)^{0.5} - H_k^{0.5} \} \\ &< 0. \end{aligned}$$

Finally, equality (17) holds by regrouping of terms and equality (18) holds by the elimination of the second and third terms of (17) which are identical.

But, $\mu_e \tilde{d}_{ej} + \hat{K}_j [H_j^{0.5} - (H_j - \mu_e)^{0.5}] > \mu_e \tilde{d}_{ek} + \hat{K}_k [(H_k + \mu_e)^{0.5} - H_k^{0.5}]$ violates the assumed optimality of assigning customer e to a DC at j . Thus, we have that if $(d_{ck} - a_k) - (d_{cj} + a_j) \geq (d_{ek} + a_k) - (d_{ej} + a_j)$, then assigning customer c to DC k implies that customer e has to be assigned to DC k as well. \square

This theorem has two important implications. First, if we let $e = k$ in the distance condition above ($(d_{ck} + a_k) - (d_{cj} + a_j) \geq (d_{kk} + a_k) - (d_{kj} + a_j)$) then we get $d_{ck} - d_{cj} \geq -d_{kj}$ or $d_{ck} + d_{kj} \geq d_{cj}$ which is true by the triangle inequality. Thus, for distance metrics for which the triangle inequality holds, it is always optimal to assign demands at a node that has a distribution center to itself. Second, each distribution center has a region of service. In other words, if we have network distances, then if customer c is optimally assigned to distribution center k then it is optimal to assign all customers e which are on the path from c to k to the distribution center at k .

5. Solution approach

5.1. Finding a lower bound

The model formulated in section 3, which accounts for the fixed DC location costs, local distribution costs from the DCs to the retailers, working and safety stock inventory at the DCs and transport costs from the supplier(s) to the DCs, is a variant of the uncapacitated fixed charge location problem. To solve this problem, we will use Lagrangian relaxation [15,16] embedded in branch and bound. In particular, we will relax constraint (7) to obtain the following Lagrangian Dual problem.

$$\begin{aligned} \text{Max}_{\lambda} \text{Min}_{X,Y} \sum_{j \in J} & \left\{ f_j X_j + \sum_{i \in I} \hat{d}_{ij} Y_{ij} + \hat{K}_j \sqrt{\sum_{i \in I} \mu_i Y_{ij}} \right\} + \sum_{i \in I} \lambda_i \left(1 - \sum_{j \in J} Y_{ij} \right) \\ & = \sum_{j \in J} \left\{ f_j X_j + \sum_{i \in I} (\hat{d}_{ij} - \lambda_i) Y_{ij} + \hat{K}_j \sqrt{\sum_{i \in I} \mu_i Y_{ij}} \right\} + \sum_{i \in I} \lambda_i \quad (19) \\ \text{subject to } & Y_{ij} \leq X_j, \quad \forall i \in I, \forall j \in J, \quad (8) \\ & X_j \in \{0, 1\}, \quad \forall j \in J, \quad (9) \\ & Y_{ij} \in \{0, 1\}, \quad \forall i \in I, \forall j \in J. \quad (10) \end{aligned}$$

For fixed values of the Lagrange multipliers, λ_i , we want to minimize (19) over the location variables, X_j , and the assignment variables, Y_{ij} . In the absence of the non-linear term in the assignment variables, solving the problem is simply a matter of computing, $V_j = f_j + \sum_{i \in I} \min(0, \hat{d}_{ij} - \lambda_i)$ for each candidate site $j \in J$, and setting $X_j = 1$ for those candidate sites for which $V_j \leq 0$. (If all V_j values are positive, we identify the smallest positive V_j and set the corresponding $X_j = 1$.) The assignment variables are then easy to determine. One simply sets $Y_{ij} = 1$ if $\hat{d}_{ij} - \lambda_i \leq 0$ and $X_j = 1$, and $Y_{ij} = 0$ otherwise. However, the presence of the non-linear term makes finding an appropriate value of V_j , the value of including candidate site j in the solution, more difficult. To do so, we need to be able to solve a subproblem of the following form for each candidate DC:

$$\text{SP}(j) \quad \tilde{V}_j = \min \sum_{i \in I} b_i Z_i + \sqrt{\sum_{i \in I} c_i Z_i} \quad (20)$$

$$\text{subject to } Z_i \in \{0, 1\}, \quad \forall i \in I, \quad (21)$$

where $b_i = \hat{d}_{ij} - \lambda_i$ and $c_i = \hat{K}_j^2 \mu_i \geq 0$. In (20)–(21) we have replaced the assignment variables Y_{ij} by Z_i to simplify the notation since $\text{SP}(j)$ is specific to DC j . Note that the fixed facility cost is not included in \tilde{V}_j and will be added later.

In [32], we show that this subproblem can be solved by the following procedure.

Step 1. Partition the set I into three sets: $I^+ = \{i: b_i \geq 0\}$, $I^0 = \{i: b_i < 0 \text{ and } c_i = 0\}$ and $I^- = \{i: b_i < 0 \text{ and } c_i > 0\}$. (We note that under our assumption of

the variance being proportional to the mean, the set I^0 will generally be empty. We include it below for completeness.)

Step 2. Sort the elements of I^- so that $b_1/c_1 \leq b_2/c_2 \leq \dots \leq b_n/c_n$, where $n = |I^-|$.

Step 3. Compute the partial sums

$$S_m = \sum_{i \in I^0} b_i Z_i + \sqrt{\sum_{i \in I^0} c_i Z_i} + \sum_{i=1, i \in I^-}^m b_i Z_i + \sqrt{\sum_{i=1, i \in I^-}^m c_i Z_i}.$$

Step 4. Select the value of m that results in the minimum value of S_m .

Since the major step in this algorithm is the sorting of the elements of I^- in step 2, the algorithm has complexity $O(|I| \log |I|)$. This problem must be solved for each $j \in J$, so solving (19) for given values of the Lagrange multipliers has complexity $O(|J||I| \log |I|)$. The proof of the optimality of this procedure relies on the concavity of the square root function. The interested reader is referred to [32] or [31].

Solving subproblem $SP(j)$ for each j enables us to solve the Lagrangian problem (19) and (8)–(10) for fixed values of the Lagrange multipliers in a manner identical to that used for the uncapacitated fixed charge problem. We add f_j to the optimal objective function value of $SP(j)$ to obtain the value of using a DC at candidate site j in the Lagrangian solution. We then select all DCs with negative values. In the case in which all DCs have non-negative values, we select the DC with the smallest non-negative (generally positive) value. This is equivalent to adding the redundant constraint (to the original problem) that $\sum_{j \in J} X_j \geq 1$ meaning that we have to select at least one DC to enable flow from the plant(s) to the retailers. For each opened DC, the corresponding values of the assignment variables are identical to the optimal Z_i values in subproblem $SP(j)$; for unopened DCs (those for which $X_j = 0$ in the Lagrangian solution), $Y_{ij} = 0, \forall i \in I$.

Having solved the Lagrangian problem, we need to find the optimal Lagrange multipliers. We do so using a standard subgradient optimization procedure [15,16]. The optimal value of (19) is a lower bound on the objective function (12).

5.2. Finding an upper bound

At each iteration of the Lagrangian procedure, we find an upper bound as follows. We initially fix the DC locations at those sites for which $X_j = 1$ in the current Lagrangian solution. Then we assign retailers to DCs in a two-phased process. First, for each retailer i for which $\sum_{j \in J} Y_{ij} \geq 1$ (each retailer assigned to at least one open DC in the Lagrangian solution), we assign the retailer i to the DC j for which $Y_{ij} = 1$ and that increases the cost the least based on the assignments made so far. In the test cases reported below, retailers are processed in order of non-increasing mean demand. In the second phase, we process retailers i for which $\sum_{j \in J} Y_{ij} = 0$ (retailers that were not assigned to any open DC in the Lagrangian solution). We assign each such retailer to the open DC which increases the total cost the least based on the assignments made so far. Note that the key difference between phases 1 and 2 is that in phase 1 we limit the possible assignments of a retailer to those to which the retailer was tentatively assigned

in the Lagrangian procedure thereby utilizing information provided by the Lagrangian solution. In phase 2, however, we are dealing with retailers that were unassigned in the Lagrangian solution. Hence, for these retailers, we consider all possible assignments to open DCs.

We note that the Theorem indicates that there is a service region around each DC. In particular, if node c is optimally assigned to DC k then any customer j on the path from k to c is also assigned to k . This property could be exploited to develop an improved heuristic algorithm. However, given the strong computational results outlined below, we have not implemented such an algorithm. Instead, we leave the development and implementation of such an algorithm for future work.

5.3. Retailer reassignments

If the value of the upper bound that we compute using the procedure outlined above is better than the best known upper bound, we try to improve the bound further by considering all possible single retailer moves from the DC to which a retailer is currently assigned to another DC. The *value* of the reassignment is generally equal to the sum of the changes in the second and third terms of (12). However, if reassigning retailer i from a DC at site j to another DC will remove all of the assigned demand at site j , then the value of the move is augmented by the fixed cost, f_j , of that DC (since we can remove the site). If the DC is forced into the solution by an additional constraint imposed by the branch and bound algorithm in which the Lagrangian procedure is embedded, the DC cannot be removed from the solution and we do not add the fixed cost into the value of the proposed retailer reassignment. We do not, however, actually remove an open DC from consideration until no improving reassignments can be found. Thus, the value of reassigning retailer i to an open DC with no currently assigned demand equals the sum of the changes in the second and third terms of (12) *minus* the fixed cost for the DC since we would need to re-open the site.

For each retailer, the best possible reassignment is found and performed before finding reassignments for the next retailer. We continue looping through all retailers until we complete one entire loop without finding an improving reassignment. At that point, an open DC with no assigned demand is removed and its fixed cost is actually subtracted from the upper bound on the solution.

5.4. DC exchange algorithm improvements

After the termination of the Lagrangian procedure, we apply a variant of the exchange algorithm proposed by Teitz and Bart [34] for the P -median problem. For each DC in the current solution, we find the best substitute DC that is not in the current solution. For each such potential exchange, retailers are assigned in a greedy manner to the DC which increases the cost the least based on the assignments made so far. Thus, this process is like the second phase assignment process in obtaining the upper bound at each iteration of the Lagrangian procedure in that a retailer can be assigned to *any* open DC. If a DC exchange is found that improves the solution, we make the exchange; otherwise we

proceed to the next open DC to try to find an improving exchange involving that DC. If any improving exchanges are found, we try single retailer reassignments (as discussed in section 5.3) to the best DC configuration we have found and then restart the search for improving exchanges. If a pass through all possible exchanges is made without finding an improving exchange, the exchange algorithm terminates.

5.5. Variable fixing

In addition to the DC exchange heuristic outlined above, at the end of the Lagrangian procedure at the root node, we employ a variable fixing technique. In essence, this can be thought of as performing branch and bound on all of the DC locations, without revising any of the Lagrange multipliers. It uses the following rules:

Node exclusion rule. If candidate DC j is not currently part of the solution to the Lagrangian problem (i.e., $X_j = 0$ in the optimal solution to (19) at some iteration) and if $LB + f_j + \tilde{V}_j > UB$, then candidate DC j cannot be part of the optimal solution, where LB and UB are the current lower and upper bounds on the solution, respectively.

Node inclusion rule. If candidate DC j is part of the solution to the Lagrangian problem (i.e., $X_j = 1$ in the optimal solution to (19) at some iteration) and if $LB - (f_j + \tilde{V}_j) > UB$, then candidate DC j must be part of the optimal solution. Note that when node j is part of the best-known solution, $f_j + \tilde{V}_j$ will generally be negative so that the left hand side of the inequality above will generally be greater than LB .²

5.6. Branch and bound

It may be that after the exchange algorithm and the variable forcing routine, either the lower bound equals the upper bound or there are *no* unforced DC location variables. In that case, the solution corresponding to the upper bound is optimal, so the lower bound is set to the upper bound and the algorithm terminates. In our experience this often occurs as indicated below. If, however, the lower bound is less than the upper bound and some candidate DC locations are not forced in or out of the solution, we employ branch and bound. Note that we have implemented branch and bound only on the location variables and not on the assignment variables. In all our computational experience with the model we have never found a case in which there was a need to branch on the assignment variables.

At each node of the branch and bound tree, we first try to branch on a location variable corresponding to a facility that is in the best-known feasible solution (i.e., the solution corresponding to the best-known upper bound). From among these locations, we select the DC (free) location with the largest assigned demand. If all DC locations in the solution corresponding to the best upper bound are forced into the solution, then we

² The authors gratefully acknowledge Larry Snyder for pointing out a minor error in the statement of the node exclusion and inclusion rules in an earlier version of the paper.

branch on the first un-forced candidate location in the list of candidate locations. Once a location has been identified for branching, we first force the site out of the solution and then force the site into the solution. Node selection is done in a depth-first manner.

6. Computational results

In this section, we outline computational results from two experiments. The first was designed to test the algorithm's computational capabilities and to compare the algorithm with a set partitioning approach to the same problem proposed by Shen [31] and Shen, Coullard and Daskin [32]. For that experiment we employed two datasets: one had 88 retail locations and the other had 150 retail locations. The second experiment was based on a modification of the 150-node data set and was designed to assess the sensitivity of the results with respect to changes in the relative importance of the transport and inventory terms. We also used this dataset to assess the impacts of significant reductions in the fixed order costs as might result from the use of e-commerce technologies. In all cases, each retail location was also a candidate DC location. Also, in all cases, we set d_{ij} , the unit cost of shipping from candidate DC j to retailer i , to the great circle distance between these locations.

For all cases, the parameters for the Lagrangian procedure are shown in table 3.

Two datasets were used for the initial set of tests. They are minor modifications of the 88 and 150-node datasets given in [7]. For the 88-node dataset – representing the 50 largest cities in the 1990 U.S. census along with the 48 capitals of the continental U.S. – the mean demand was obtained by dividing the population data by 1000 and rounding the result to the nearest integer. Fixed facility location costs were obtained by dividing the facility location costs in [7] by 100. For the 150-node dataset – representing the 150 largest cities in the continental U.S. for the 1990 census – the mean demand was obtained in the same manner. The fixed facility costs were all set to 100, one-one-thousandth of the value in the dataset given by Daskin [7]. These changes were made to allow us to deal with smaller numbers.

Table 4 presents the coefficients used in all the runs for both datasets.

Table 5 presents the results for the two datasets. As the transport costs increase (as β goes up), the number of DCs goes up. Conversely, as inventory costs increase (as θ goes up), the number of DCs goes down. Also, when inventory considerations dominate (θ is very large relative to β), it is sometimes optimal to assign one or more

Table 3
Parameters for Lagrangian relaxation runs.

Parameter	Value
Maximum number of iterations	400
Minimum alpha multiplier	0.00000001
Maximum number of iterations before halving alpha	12
Crowder's damping factor	0.3
Initial Lagrange multiplier value	$10\bar{\mu} + 10f_j$

Table 4
Model coefficients for 88 and 150-node test problems.

F	Fixed order cost	10
$g_j, \forall j \in \mathbf{J}$	Fixed shipping cost from supplier to DC	10
$a_j, \forall j \in \mathbf{J}$	Unit shipping cost from supplier to DC	5
γ	Variance to mean ratio	1.00
z_α	Service level parameter	1.96
L	Lead time	1
h	Unit holding cost	1
χ	Days per year	1

retailers to a DC that is other than the least cost DC in terms of local delivery cost (d_{ij}). For example, for the final case (150 nodes, $\beta = 0.001$ and $\theta = 1$), three retailers were assigned to DCs other than the one with the smallest local delivery cost. Reassigning each to the least local delivery cost DC would increase the cost by \$3.37 or 0.027%. The local delivery costs decrease by almost \$11 out of roughly \$4,870 (0.22%), while the safety stock, fixed order costs, and working inventory costs each increase by about \$4.7 or 0.245%. Thus, while it is optimal to assign these three retailers to DCs other than the ones that could provide the least cost local delivery service, reassigning them to the least-cost DCs would not increase the total cost very much.

Table 5 also compares the computation times for the algorithm presented above with those obtained in [32] using a set partitioning approach. Times obtained for our model are on a Dell Latitude CPx computer running at 650 MHz using Windows 98. The program was written in Delphi 5. The Shen, Coullard, and Daskin times are on a Sun Spark Station running the SunOS 4.1.3u5 operating system. Our computation times are consistently lower than those obtained using the column generation approach. However, our times tend to grow with β while the column generation times decrease with β . Thus, for very large values of β relative to θ , it might be better to use the column generation approach. In many cases, we did not need to use branch and bound at all; when we did very few nodes needed to be evaluated. Finally, the variable forcing rules were exceptionally effective, forcing almost all nodes in or out of the solution at the root node. This is in part due to the very tight bounds obtained at the root node.

Finally, we consider a different variant of the 150-node dataset given in [7]. In this case, we divided the demands by 250 and truncated the result to obtain the mean daily demand. Table 6 gives the model coefficients for this set of runs, which were designed to better replicate real-world conditions. Two supplier locations were considered – Chicago and Phoenix – and the fixed shipping cost from the supplier to a candidate DC was taken to be function of the distance from the closer supplier to the candidate DC site as shown in table 6.

Figure 3 shows the optimal solution for a base case of $\beta = 0.00025$ and $\theta = 0.1$. The supplier in Chicago serves 6 of the 10 DCs and about 73% of the total demand. Figure 4 breaks down the total cost of approximately \$4.49 million into its constituent parts. Note that in this case, the transport costs from the suppliers to the DCs as well as the safety stock costs – all buried in the “other” category – are negligible. Figures 5

Table 5
Results for 88 and 150 node datasets.

No. nodes	Beta	Theta	Objective function	# of DCs	# Retailers assigned to non-closest DC	Iter.	B&B nodes	Time (sec.)	# DCs forced IN at root node	# DCs forced OUT at root node	Shen et al. time (sec.)	Shen et al. columns
88	0.001	0.1	13,229.55	9	0	192	1	2.08	9	79	626	12,246
88	0.002	0.1	19,975.37	11	0	308	1	3.40	10	77	65	3,417
88	0.003	0.1	25,306.68	15	0	205	1	2.69	15	73	22	2,352
88	0.004	0.1	28,752.64	21	0	501	3	5.87	13	62	15	1,409
88	0.005	0.1	31,390.69	23	0	263	1	3.57	23	65	8	846
88	0.001	0.1	13,229.55	9	0	192	1	2.09	9	79	626	12,246
88	0.002	0.2	20,491.17	10	0	250	1	2.69	10	77	89	4,185
88	0.005	0.5	33,794.94	22	0	193	1	2.74	22	65	7	906
88	0.005	0.1	31,390.69	23	0	263	1	3.62	23	65	8	846
88	0.005	0.5	33,794.94	22	0	193	1	2.75	22	65	7	906
88	0.005	1	35,876.10	21	0	120	1	1.27	21	67	12	1,222
88	0.005	5	47,348.38	17	1	121	1	1.93	17	71	60	4,579
88	0.005	10	57,959.54	12	2	161	1	1.98	12	76	145	9,158
88	0.005	20	74,760.97	9	2	379	1	3.96	9	79	869	20,114
150	0.0004	0.01	3,977.27	15	0	448	3	13.62	11	134	252	9,624
150	0.0006	0.01	4,867.76	21	0	403	3	13.90	21	128	180	8,075
150	0.0008	0.01	5,580.71	26	0	403	3	15.32	26	123	55	3,351
150	0.001	0.01	6,162.99	28	0	405	3	15.87	28	121	29	2,318
150	0.0005	0.01	4,459.32	18	0	423	3	13.68	15	130	277	10,907
150	0.001	0.02	6,410.09	28	1	400	1	15.55	28	122	45	2,730
150	0.002	0.04	8,988.62	41	0	171	1	4.78	41	109	22	1,286
150	0.001	0.01	6,162.99	28	0	405	3	15.81	28	121	29	2,318
150	0.001	0.1	7,508.49	26	1	241	1	10.94	26	124	102	4,233
150	0.001	0.5	10,175.71	21	2	400	1	13.74	21	129	261	11,795
150	0.001	1	12,380.63	15	3	444	3	13.46	14	133	730	21,944

Table 6
Model coefficients for final set of test problems.

F	Fixed order cost	1000
$g_j, \forall j \in J$	Fixed shipping cost from supplier to DC	$40 + 0.1(\text{distance})$
$a_j, \forall j \in J$	Unit shipping cost from supplier to DC	0.5
γ	Variance to mean ratio	1.00
z_α	Service level parameter	1.96
L	Lead time (days)	7
h	Unit holding cost	25
χ	Days per year	250

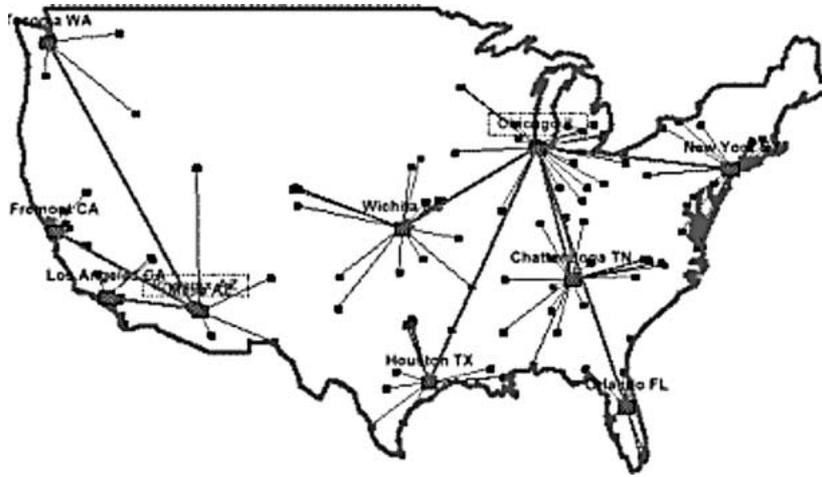


Figure 3. Solution for base case.

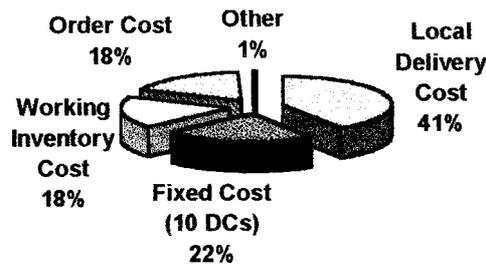


Figure 4. Distribution of costs for base case.

and 6 show the sensitivity of the cost and number of facilities to changes in the transport and inventory cost weights. As the relative importance of inventory costs goes up, the number of DCs located goes down and as the importance of transport costs goes up, the number of DCs increases. The figures also compare this model with the traditional uncapacitated fixed charge (UFC) location model. Not surprisingly, the total cost for the UFC model is always below that of the model presented above which includes additional

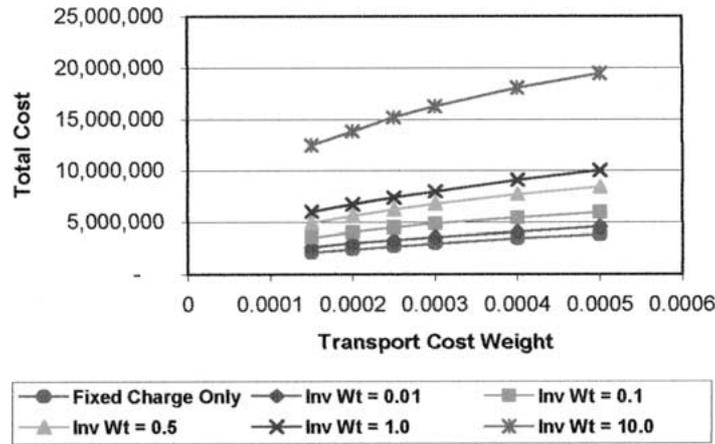


Figure 5. Sensitivity of cost to changes in transport and inventory weights.

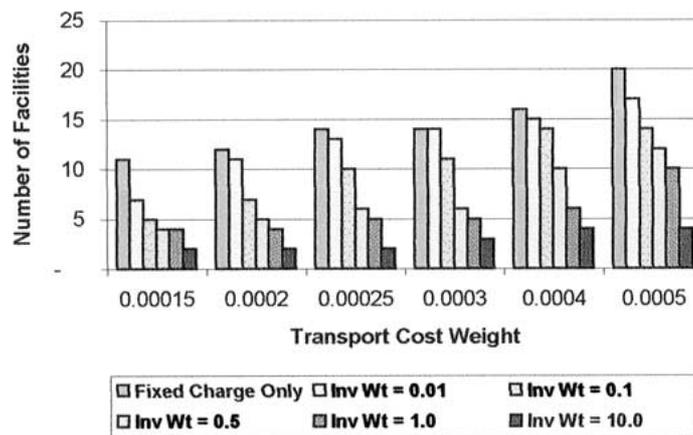


Figure 6. Sensitivity of number of DCs to changes in transport and inventory weights.

cost components. Also, due to the concavity of the additional cost terms, the UFC consistently locates more DCs than does the model above.

Finally, we consider how the solution might change if the fixed cost of placing an order decreased significantly from \$1000 per order to \$10 per order as might be the case with e-commerce technologies. The total cost decreases from \$4.49 million to \$2.9 million. Interestingly, it is now optimal to open 14 DCs instead of the 10 found in the base case. This is reassuring since the demand-weighted average local delivery distance to a retailer from a DC will decrease from 127.1 to 88.8 miles. It is reassuring that the average distance decreases in this case, since e-commerce is often associated with expectations of improved customer service. Figure 7 shows the new solution. The supplier in Chicago serves 9 of the DCs and about 71% of the demand. Compared to the cost breakdown in the base case (figure 4), in this case, 48% of the total cost is in facility

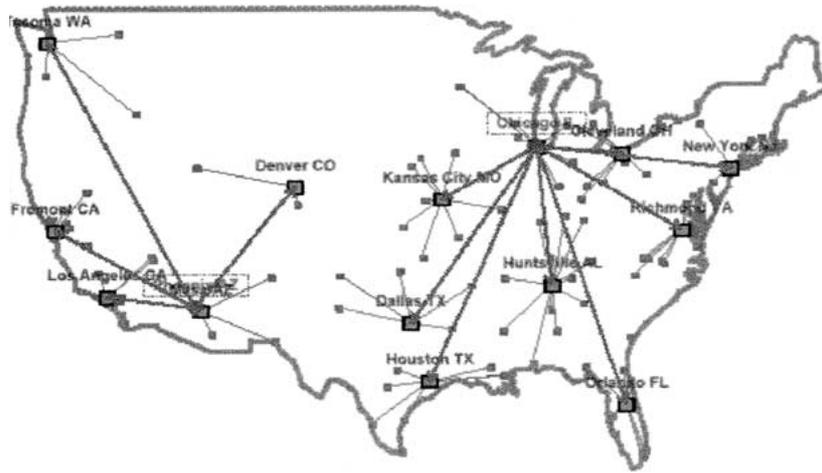


Figure 7. Solution with reduced fixed ordering costs.

costs and another 44.3% is in local delivery costs. Order costs and working inventory costs each constitute 3.3% of the total costs while safety stock costs and shipment costs from the supplier to the DCs are again negligible.

If we could not relocate or add to the set of open DCs used in the base case, but could simply reoptimize the retailer assignments and inventory management policies, the costs would be 4.1% higher than those for the solution shown in figure 7. If we were constrained to using the 10 DCs in the base case, but could add DCs, it would be optimal to use 3 additional DCs located in Cleveland, OH, Denver, CO, and Richmond, VA. The total cost in this case is only 1.1% greater than that of the optimal solution in figure 7.

7. Conclusions and directions for future work

We have presented a facility location model that incorporates working and safety stock inventory costs at the distribution centers as well as the economies of scale that exist in the transport costs from suppliers to DCs. The model is identical to that proposed by Shen [31] and Shen, Coullard and Daskin [32]. Both references solve the problem by first recasting it as a set partitioning problem and then solving the resulting model using column generation. In this paper, we presented a Lagrangian solution algorithm for the case in which the ratio of the variance of demand at the retailers to the mean demand is the same for all retailers. A number of improvement heuristics were outlined for the problem. Also, two variable forcing rules that are applied at the root node of the branch and bound tree were discussed.

The algorithm was tested on two datasets consisting of 88 and 150 nodes, respectively. The computational results compare favorably with those of the set partitioning approach proposed by Shen [31] and Shen, Coullard and Daskin [32]. In many cases, branch and bound was not needed because the bounds at the root node proved the optimality of the solution and/or because we could force all of the candidate DCs into or

out of the solution. In a final set of tests, we included explicit supplier locations and modified the inputs to better reflect actual conditions. When fixed order costs are significantly reduced, the number of facilities located increases. Despite our best efforts to reflect realistic conditions, it is important to test the data with actual corporate data to confirm that the model is accurately capturing the relevant costs.

A number of extensions should be considered. First, we need to develop ways of solving the problem when the ratio of the variance to the mean is not identical for all retailers. Second, we hope to consider the cases with multiple items as well as a constraint on the maximum allowable inventory at a DC and a constraint on the maximum demand that can be served by a supplier. Third, it might be possible to incorporate local delivery cost estimates that better reflect less-than-truckload routing since such approximations often involve the square root of demand [6]. Finally, we would like to extend the model to account for different future demand or travel cost scenarios. We are working in all of these areas.

Acknowledgments

This work was supported by NSF grants DMI-96-34750 and DMI-98-12915. This support is gratefully acknowledged. The authors would also like to acknowledge the contributions of the anonymous referees whose comments improved the paper significantly.

References

- [1] S. Axsater, Using the deterministic EOQ formula in stochastic inventory control, *Management Science* 42 (1996) 830–834.
- [2] M. Balinski, Integer programming: Methods, uses, computation, *Management Science* 13 (1965) 253–313.
- [3] R.T. Berger, C.R. Coullard and M.S. Daskin, Modeling and solving location-routing problems with route-length constraints, Working paper, Northwestern University (1998).
- [4] O. Berman, P. Jaillet and D. Simchi-Levi, Location-routing problems with uncertainty, in: *Facilities Location*, ed. Z. Drezner (Springer, 1995) pp. 427–452.
- [5] L.M.A. Chan, A. Federguen and D. Simchi-Levi, Probabilistic analysis and practical algorithms for inventory routing models, *Operations Research* 46 (1998) 96–106.
- [6] C. Daganzo, *Logistics Systems Analysis*, Lecture Notes in Economics and Mathematical Systems, eds. M. Beckmann and W. Krelle (Springer, Berlin, 1991).
- [7] M.S. Daskin, *Network and Discrete Location: Models, Algorithms, and Applications* (Wiley, New York, 1995).
- [8] M.S. Daskin and S.H. Owen, Location models in transportation, in: *Handbook of Transportation Science*, ed. R. Hall (Kluwer Academic, Norwell, MA, 1999) pp. 311–360.
- [9] Z. Drezner, ed., *Facility Location: A Survey of Applications and Methods* (Springer, New York, 1995).
- [10] G. Eppen, Effects of centralization on expected costs in a multi-location newsboy problem, *Management Science* 25(5) (1979) 498–501.
- [11] S.J. Erlebacher and R.D. Meller, The interaction of location and inventory in designing distribution systems, *IIE Transactions* 32 (2000) 155–166.
- [12] D. Erlenkotter, A dual-based procedure for uncapacitated facility location, *Operations Research* 14 (1978) 361–368.

- [13] A. Federgruen and D. Simchi-Levi, Analytical analysis of vehicle routing and inventory routing problems, in: *Network Routing*, eds. M. Ball, T. Magnanti, C. Monma and G. Nemhauser, Handbooks in Operations Research and Management Science, Vol. 8 (North-Holland, Amsterdam, 1995) pp. 297–373.
- [14] A. Federgruen and P. Zipkin, A combined vehicle routing and inventory allocation problem, *Management Science* 32 (1984) 1019–1036.
- [15] M.L. Fisher, The Lagrangian relaxation method for solving integer programming problems, *Management Science* 27 (1981) 1–18.
- [16] M.L. Fisher, An applications oriented guide to Lagrangian relaxation, *Interfaces* 15(2) (1985) 2–21.
- [17] S.C. Graves, A.H.G. Rinnooy Kan and P.H. Zipkin, *Logistics of Production and Inventory* (Elsevier, Amsterdam, 1993).
- [18] S. Hakimi, Optimum location of switching centers and the absolute centers and medians of a graph, *Operations Research* 12 (1964) 450–459.
- [19] S. Hakimi, Optimum location of switching centers in a communications network and some related graph theoretic problems, *Operations Research* 13 (1965) 462–475.
- [20] W. Hopp and M.L. Spearman, *Factory Physics: Foundations of Manufacturing Management* (Irwin, Chicago, 1996).
- [21] A.J. Kleywegt, V.S. Nori and M.L.P. Savelsbergh, The stochastic inventory routing problem, Working paper, Georgia Institute of Technology, 2000.
- [22] M. Korkel, On the exact solution of large-scale simple plant location problems, *European Journal of Operations Research* 39 (1989) 157–173.
- [23] G. Laporte, Location-routing problems, in: *Vehicle Routing: Methods and Studies*, eds. B.L. Golden and A.A. Assad (North-Holland, Amsterdam, 1988) pp. 163–197.
- [24] G. Laporte and P.J. Dejax, Dynamic location–routing problems, *Journal of the Operations Research Society* 40(5) (1989) 471–482.
- [25] H. Lee, Information distortion in a supply chain: The bullwhip effect, *Management Science* 43 (1996) 546–558.
- [26] H. Lee, V. Padmanabhan and S. Whang, The bullwhip effect in supply chains, *Sloan Management Review*, Spring (1997) 93–102.
- [27] H. Min, V. Jayaraman and R. Srivastava, Combined location-routing problems: A synthesis and future research directions, *European Journal of Operational Research* 108 (1998) 1–15.
- [28] D.C. Montgomery, G.C. Runger and N.F. Hubele, *Engineering Statistics* (Wiley, New York, 1998).
- [29] S. Nahmias, *Production and Operations Management*, 3rd edn. (Irwin, Chicago, 1997).
- [30] Y. Sheffi, *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods* (Prentice-Hall, Englewood Cliffs, NJ, 1985).
- [31] Z.J. Shen, Efficient algorithms for various supply chain problems, Ph.D. dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL (2000).
- [32] Z.J. Shen, C.R. Coullard and M.S. Daskin, A joint location-inventory model, to appear in *Transportation Science*.
- [33] D. Simchi-Levi, P. Kaminsky and E. Simchi-Levi, *Designing and Managing the Supply Chain: Concepts, Strategies and Case Studies* (Irwin McGraw Hill, Boston, MA, 2000).
- [34] M.B. Teitz and P. Bart, Heuristic methods for estimating generalized vertex median of a weighted graph, *Operations Research* 16 (1968) 955–961.
- [35] C.P. Teo, J. Ou and M. Goh, Impact on inventory costs with consolidation of distribution centers, *IIE Transactions* 33 (2001) 99–110.
- [36] S. Viswanathan and K. Mathur, Integrating routing and inventory decisions in one-warehouse multi-retailer multiproduct distribution system, *Management Science* 43 (1997) 294–312.
- [37] Y.-S. Zheng, On properties of stochastic inventory systems, *Management Science* 38(1) (1992) 87–103.
- [38] P.H. Zipkin, *Foundations of Inventory Management* (Irwin, Burr Ridge, IL, 1997).