

Department of Industrial Engineering and Management Sciences

Northwestern University, Evanston, Illinois 60208-3119, U.S.A.

Working Paper No. 03-010

Facility Location in Supply Chain Design

Mark S. Daskin

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

Lawrence V. Snyder

Rosemary T. Berger

Department of Industrial and Systems Engineering
Lehigh University
200 West Packer Avenue
Bethlehem, PA 18015-1582

December 2003

This research has been supported in part
by the National Science Foundation under grant DMI-0223277.

ABSTRACT

In this chapter we outline the importance of facility location decisions in supply chain design. We begin with a review of classical models including the traditional fixed charge facility location problem. We then summarize more recent research aimed at expanding the context of facility location decisions to incorporate additional features of a supply chain including LTL vehicle routing, inventory management, robustness, and reliability.

Introduction

The efficient and effective movement of goods from raw material sites to processing facilities, component fabrication plants, finished goods assembly plants, distribution centers, retailers and customers is critical in today's competitive environment. Approximately 10% of the gross domestic product is devoted to supply-related activities (Simchi-Levi, Kaminsky, and Simchi-Levi, 2003, p. 5). Within individual industries, the percentage of the cost of a finished delivered item to the final consumer can easily exceed this value. Supply chain management entails not only the movement of goods but also decisions about (1) where to produce, what to produce, and how much to produce at each site, (2) what quantity of goods to hold in inventory at each stage of the process, (3) how to share information among parties in the process and finally, (4) where to locate plants and distribution centers.

Location decisions may be the most critical and most difficult of the decisions needed to realize an efficient supply chain. Transportation and inventory decisions can often be changed on relatively short notice in response to changes in the availability of raw materials, labor costs, component prices, transportation costs, inventory holding costs, exchange rates and tax codes. Information sharing decisions are also relatively flexible and can be altered in response to changes in corporate strategies and alliances. Thus, transportation, inventory, and information sharing decisions can be readily re-optimized in response to changes in the underlying conditions of the supply chain. Decisions about production quantities and locations are, perhaps, less flexible, as many of the costs of production may be fixed in the short term. Labor costs, for example, are often dictated by relatively long-term contracts. Also, plant capacities must often be taken as fixed in the short-term. Nevertheless, production quantities can often be altered in the intermediate term in response to changes in material costs and market demands.

Facility location decisions, on the other hand, are often fixed and difficult to change even in the intermediate term. The location of a multibillion-dollar automobile assembly plant cannot be changed as a result of changes in customer demands, transportation costs, or component prices. Modern distribution centers with millions of dollars of material handling equipment are also difficult, if not impossible, to relocate except in the long term. Inefficient locations for production and assembly plants as well as distribution centers will result in excess costs being incurred throughout the lifetime of the facilities, no matter how well the production plans, transportation options, inventory management, and information sharing decisions are optimized in response to changing conditions.

However, the long-term conditions under which production plants and distribution centers will operate is subject to considerable uncertainty at the time these decisions must be made. Transportation costs, inventory carrying costs (which are affected by interest rates and insurance costs), and production costs, for example, are all difficult to predict. Thus, it is critical that planners recognize the inherent uncertainty associated with future conditions when making facility location decisions.

Vehicle routing and inventory decisions are generally secondary to facility location in the sense that facilities are expensive to construct and difficult to modify, while routing and inventory decisions can be modified periodically without difficulty. Nevertheless, it has been shown empirically for both location/routing and location/inventory problems that the facility location decisions that would be made in isolation are different from those that would be made taking into account routing or inventory. Similarly, planners are often reluctant to consider robustness and reliability at design time since disruptions may be only occasional; however, large improvements in reliability and robustness can often be attained with only small increases in the cost of the supply chain network.

In this chapter we review several traditional facility location models, beginning with the classical fixed charge location model. We then show how the model can be extended to incorporate additional facets of the supply chain design problem, including more accurate representations of the delivery process, inventory management decisions, and robustness and reliability considerations.

The Fixed Charge Facility Location Problem

The fixed charge facility location problem is a classical location problem and forms the basis of many of the location models that have been used in supply chain design. The problem can be stated simply as follows. We are given a set of customer locations with known demands and a set of candidate facility locations. If we elect to locate a facility at a candidate site, we incur a known fixed location cost. There is a known unit shipment cost between each candidate site and each customer location. The problem is to find the locations of the facilities and the shipment pattern between the facilities and the customers to minimize the combined facility location and shipment costs subject to a requirement that all customer demands be met.

Specifically, we introduce the following notation:

Inputs and Sets:

I	Set of customer locations, indexed by i
J	Set of candidate facility locations, indexed by j
h_i	demand at customer location $i \in I$
f_j	fixed cost of locating a facility at candidate site $j \in J$
c_{ij}	unit cost of shipping between candidate facility site $j \in J$ and customer location $i \in I$

Decision Variables:

$$X_j = \begin{cases} 1 & \text{if we locate at candidate site } j \in J \\ 0 & \text{if not} \end{cases}$$

$$Y_{ij} = \text{fraction of the demand at customer location } i \in I \text{ that is served by a facility at site } j \in J$$

With this notation, the fixed charge facility location problem can be formulated as follows (Balinski, 1965):

$$\text{Minimize } \sum_{j \in J} f_j X_j + \sum_{j \in J} \sum_{i \in I} h_i c_{ij} Y_{ij} \quad (1)$$

$$\text{Subject to } \sum_{j \in J} Y_{ij} = 1 \quad \forall i \in I \quad (2)$$

$$Y_{ij} - X_j \leq 0 \quad \forall i \in I; \forall j \in J \quad (3)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (4)$$

$$Y_{ij} \geq 0 \quad \forall i \in I; \forall j \in J \quad (5)$$

The objective function (1) minimizes the sum of the fixed facility location costs and the transportation or shipment costs. Constraint (2) stipulates that each demand node is fully assigned. Constraint (3) states that a demand node cannot be assigned to a facility unless we open that facility. Constraint (4) is a standard integrality constraint and constraint (5) is a simple non-negativity constraint.

The formulation given above assumes that facilities have unlimited capacity; the problem is sometimes referred to as the uncapacitated fixed charge location problem. It is well known that at least one optimal solution to this problem involves assigning all of the demand at each customer location $i \in I$ fully to the nearest open facility site $j \in J$. In other words, the assignment variables, Y_{ij} , will naturally take on integer values in the solution to this problem. Many firms insist on or strongly prefer such *single sourcing* solutions as they make the management of the supply chain considerably simpler. Capacitated versions of the fixed charge location problem do not exhibit this property; enforcing single sourcing is significantly more difficult in this case (as discussed below).

A number of solution approaches have been proposed for the uncapacitated fixed charge location problem. Simple heuristics typically begin by constructing a feasible solution by greedily adding or dropping facilities from the solution until no further

improvements can be obtained. Maranzana (1964) proposed a neighborhood search improvement algorithm for the closely related P -median problem (Hakimi, 1964, 1965) that exploits the ease in finding optimal solutions to 1-median problem: it partitions the customers by facility and then finds the optimal location within each partition. If any facility changes, the algorithm repartitions the customers and continues until no improvement in the solution can be found. Teitz and Bart (1968) proposed an exchange or “swap” algorithm for the P -median problem that can also be extended to the fixed charge facility location problem. Hansen and Mladenović (1997) proposed a variable neighborhood search algorithm for the P -median problem that can also be used for the fixed charge location problem. Clearly, improvement heuristics designed for the P -median problem will not perform well for the fixed charge location problem if the starting number of facilities is sub-optimal. One way of resolving this limitation is to apply more sophisticated heuristics to the problem. Al-Sultan and Al-Fawzan (1999) applied tabu search (Glover, 1989, 1990; Glover and Laguna, 1997) to the uncapacitated fixed charge location problem. The algorithm was tested successfully on small- to moderate-sized problems.

Erlenkotter (1978) proposed the well-known DUALOC procedure to find optimal solutions to the problem. Galvão (1993) and Daskin (1995) review the use of Lagrangian relaxation algorithms in solving the uncapacitated fixed charge location problem. When embedded in branch and bound, Lagrangian relaxation can be used to solve the fixed charge location problem optimally (Geoffrion, 1974). The reader interested in a more comprehensive review of the uncapacitated fixed charge location problem is referred to either Krarup and Pruzan (1983) or Cornuéjols, Nemhauser and Wolsey (1990).

One natural extension of the problem is to consider capacitated facilities. If we let b_j be the maximum demand that can be assigned to a facility at candidate site $j \in J$, formulation (1)-(5) can be extended to incorporate facility capacities by including the following additional constraint:

$$\sum_{i \in I} h_i Y_{ij} - b_j X_j \leq 0 \quad \forall j \in J \quad (6)$$

Constraint (6) limits the total assigned demand at facility $j \in J$ to a maximum of b_j . From the perspective of the integer programming problem, this constraint obviates the need for constraint (3) since any solution that satisfies (5) and (6) will also satisfy (3). However, the linear programming relaxation of (1)-(6) is often tighter if constraint (3) is included in the problem.

For fixed values of the facility location variables, X_j , the optimal values of the assignment variables can be found by solving a traditional transportation problem. The embedded transportation problem is most easily recognized if we replace $h_i Y_{ij}$ by Z_{ij} , the quantity shipped from distribution center j to customer i . The transportation problem for fixed facility locations is then

$$\text{Minimize} \quad \sum_{j \in J} \sum_{i \in I} c_{ij} Z_{ij} \quad (7)$$

$$\text{Subject to} \quad \sum_{j \in J} Z_{ij} = h_i \quad \forall i \in I \quad (8)$$

$$\sum_{i \in I} Z_{ij} \leq b_j \hat{X}_j \quad \forall j \in J \quad (9)$$

$$Z_{ij} \geq 0 \quad \forall i \in I; \forall j \in J \quad (10)$$

where we denote the fixed (known) values of the location variables by \hat{X}_j .

The solution to the transportation problem (7)-(10) may involve fractional assignments of customers to facilities. This means that the solution to the problem with

the addition of constraint (6) will not automatically satisfy the single sourcing condition, as does the solution to the uncapacitated fixed charge location problem in the absence of this constraint. To restore the single sourcing condition, we can replace the fractional definition of the assignment variables by a binary one:

$$Y_{ij} \begin{cases} 1 & \text{if demands at customer site } i \in I \text{ are served by a facility at candidate site } j \in J \\ 0 & \text{if not} \end{cases}$$

The problem becomes considerably more difficult to solve since there are now far more integer variables. For given facility sites, even if we ignore the requirement that each demand node is served exactly once, the resulting problems become knapsack problems, which can only be solved optimally in pseudo-polynomial time (as opposed to the transportation problem, which can be solved in polynomial time).

Daskin and Jones (1993) observed that, in many practical contexts, the number of customers is significantly greater than the number of distribution centers that will be sited. As such, each customer represents a small fraction of the total capacity of the distribution center to which it is assigned. Also, if the single sourcing requirement is relaxed, the number of multiply sourced customers is less than or equal to the number of distribution centers minus one. Thus, relatively few customers will be multiply-sourced in most contexts. They further noted that warehouse capacities, when measured in terms of annual throughput as is commonly done, are rarely known with great precision, as they depend on many factors, including the number of inventory turns at the warehouse. (We return to the issue of inventory turns below when we outline an integrated location/inventory model.) They therefore proposed a procedure for addressing the single

sourcing problem that involves (1) ignoring the single sourcing constraint and solving the transportation problem, (2) using duality to find alternate optima to the transportation problem that require fewer customers to be multiply sourced, and (3) allowing small violations of the capacity constraints to identify solutions that satisfy the single sourcing requirement. In a practical context involving a large national retailer with over 300 stores and about a dozen distribution centers, they found that this approach was perfectly satisfactory from a managerial perspective.

In a classic paper, Geoffrion and Graves (1974) extend the traditional fixed charge facility location problem to include shipments from plants to distribution centers and multiple commodities. They introduce the following additional notation:

Inputs and Sets:

- K Set of plant locations, indexed by k
- L Set of commodities, indexed by l
- D_{li} demand for commodity $l \in L$ at customer $i \in I$
- S_{lk} supply of commodity $l \in L$ at plant $k \in K$
- $\underline{V}_j, \bar{V}_j$ minimum and maximum annual throughput allowed at distribution center $j \in J$
- v_j variable unit cost of throughput at candidate site $j \in J$
- c_{lkji} unit cost of producing and shipping commodity $l \in L$ between plant $k \in K$, candidate facility site $j \in J$ and customer location $i \in I$

Decision Variables:

- $Y_{ij} \begin{cases} 1 & \text{if demands at customer site } i \in I \text{ are served by a facility at candidate site } j \in J \\ 0 & \text{if not} \end{cases}$
- Z_{lkji} quantity of commodity $l \in L$ shipped between plant $k \in K$, candidate facility site $j \in J$ and customer location $i \in I$

With this notation, Geoffrion and Graves formulate the following extension of the fixed charge facility location problem.

$$\text{Minimize} \quad \sum_{j \in J} f_j X_j + \sum_{j \in J} v_j \left(\sum_{l \in L} \sum_{i \in I} D_{li} Y_{ij} \right) + \sum_{l \in L} \sum_{k \in K} \sum_{j \in J} \sum_{i \in I} c_{lkji} Z_{lkji} \quad (11)$$

$$\text{Subject to} \quad \sum_{i \in I} \sum_{j \in J} Z_{lkji} \leq S_{lk} \quad \forall k \in K; \forall l \in L \quad (12)$$

$$\sum_{k \in K} Z_{lkji} = D_{li} Y_{ij} \quad \forall l \in L; \forall j \in J; \forall i \in I \quad (13)$$

$$\sum_{j \in J} Y_{ij} = 1 \quad \forall i \in I \quad (14)$$

$$\underline{V}_j X_j \leq \sum_{i \in I} \sum_{l \in L} D_{li} Y_{ij} \leq \bar{V}_j X_j \quad \forall j \in J \quad (15)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (16)$$

$$Y_{ij} \in \{0,1\} \quad \forall i \in I; \forall j \in J \quad (17)$$

$$Z_{lkji} \geq 0 \quad \forall i \in I; \forall j \in J; \forall k \in K; \forall l \in L \quad (18)$$

The objective function (11) minimizes the sum of the fixed distribution center (DC) location costs, the variable DC costs and the transportation costs from the plants through the DCs to the customers. Constraint (12) states that the total amount of commodity $l \in L$ shipped from plant $k \in K$ cannot exceed the capacity of the plant to produce that commodity. Constraint (13) says that the amount of commodity $l \in L$ shipped to customer $i \in I$ via DC $j \in J$ must equal the amount of that commodity produced at all plants that is destined for that customer and shipped via that DC. This constraint stipulates that demand must be satisfied at each customer node for each commodity and also serves as a linking constraint between the flow variables (Z_{lkji}) and the assignment variables (Y_{ij}). Constraint (14) is the now-familiar single-sourcing constraint. Constraint (15) imposes lower and upper bounds on the throughput processed at each distribution center that is used. This also serves as a linking constraint (e.g., it replaces constraint (3)) between the location variables (X_j) and the customer assignment variables (Y_{ij}). Alternatively, it can be thought of as an extension of the capacity constraint (6) above.

In addition to the constraints above, Geoffrion and Graves allow for linear constraints on the location and assignment variables. These can include constraints on the minimum and maximum number of distribution centers to be opened, relationships between the feasible open DCs, more detailed capacity constraints if different commodities use different amounts of a DC's resources, and certain customer service constraints. The authors apply Benders decomposition (Benders, 1962) to the problem after noting that, if the location and assignment variables are fixed, the remaining problem breaks down into $|L|$ transportation problems, one for each commodity.

Geoffrion and Graves highlight eight different forms of analysis that were performed for a large food company using the model arguing, as do Geoffrion and Powers (1980), that the value of a model such as (11)-(18) extends far beyond the mere solution of a single instance of the problem to include a range of sensitivity and what-if analyses.

Integrated Location/Routing Models

An important limitation of the fixed charge location model, and even the multi-echelon, multi-commodity extension of Geoffrion and Graves, is the assumption that full truckload quantities are shipped from a distribution center to a customer. In many contexts, shipments are made in less-than-truckload (LTL) quantities from a facility to customers along a multiple-stop route. In the case of full truckload quantities, the cost of delivery is independent of the other deliveries made, whereas in the case of LTL quantities, the cost of delivery depends on the other customers on the route and the sequence in which customers are visited. Eilon, Watson-Gandy and Christofides (1971) were among the

first to highlight the error introduced by approximating LTL shipments by full truckloads. During the past three decades, a sizeable body of literature has developed on integrated location/routing models.

Integrated location/routing problems combine three components of supply chain design: facility location, customer allocation to facilities and vehicle routing. Many different location/routing problems have been described in the literature, and they tend to be very difficult to solve since they merge two NP-hard problems: facility location and vehicle routing. Laporte (1988) reviews early work on location routing problems; he summarizes the different types of formulations, solution algorithms and computational results of work published prior to 1988. More recently, Min, Jayaraman, and Srivastava (1998) develop a hierarchical taxonomy and classification scheme that they use to review the existing location routing literature. They categorize papers in terms of problem characteristics and solution methodology. One means of classification is the number of layers of facilities. Typically, three-layer problems include flows from plants to distribution centers to customers, while two-layer problems focus on flows from distribution centers to customers.

An example of a three-layer location routing problem is the formulation of Perl (1983) and Perl and Daskin (1985); their model extends the model of Geoffrion and Graves to include multiple stop tours serving the customer nodes but it is limited to a single commodity. Perl defines the following additional notation:

Inputs and Sets:

- P set of points = $I \cup J$
- d_{ij} distance between node $i \in P$ and node $j \in P$
- v_j variable cost per unit processed by a facility at candidate facility site $j \in J$

t_j	maximum throughput for a facility at candidate facility site $j \in J$
S	set of supply points (analogous to plants in the Geoffrion and Graves model), indexed by s
c_{sj}	unit cost of shipping from supply point $s \in S$ to candidate facility site $j \in J$
K	set of candidate vehicles, indexed by k
σ_k	capacity of vehicle $k \in K$
τ_k	maximum allowable length of a route served by vehicle $k \in K$
α_k	cost per unit distance for delivery on route $k \in K$

Decision Variables:

$$Z_{ijk} \begin{cases} 1 & \text{if vehicle } k \in K \text{ goes directly from point } i \in P \text{ to point } j \in P \\ 0 & \text{if not} \end{cases}$$

$$W_{sj} \quad \text{quantity shipped from supply source } s \in S \text{ to facility site } j \in J$$

With this notation (and the notation defined previously), Perl (1983) formulates the following integrated location routing problem:

$$\text{Minimize} \quad \sum_{j \in J} f_j X_j + \sum_{s \in S} \left\{ \sum_{j \in J} c_{sj} W_{sj} \right\} + \sum_{j \in J} v_j \left\{ \sum_{i \in I} h_i Y_{ij} \right\} + \sum_{k \in K} \alpha_k \left\{ \sum_{j \in P} \sum_{i \in P} d_{ij} Z_{ijk} \right\} \quad (19)$$

$$\text{Subject to:} \quad \sum_{k \in K} \sum_{j \in P} Z_{ijk} = 1 \quad \forall i \in I \quad (20)$$

$$\sum_{i \in I} h_i \left\{ \sum_{j \in P} Z_{ijk} \right\} \leq \sigma_k \quad \forall k \in K \quad (21)$$

$$\sum_{j \in P} \sum_{i \in P} d_{ij} Z_{ijk} \leq \tau_k \quad \forall k \in K \quad (22)$$

$$\sum_{i \in V} \sum_{j \in \bar{V}} \sum_{k \in K} Z_{ijk} \geq 1 \quad \forall \text{ subsets } V \subset P \text{ such that } J \subset V \quad (23)$$

$$\sum_{j \in P} Z_{ijk} - \sum_{j \in P} Z_{jik} = 0 \quad \forall i \in P; \forall k \in K \quad (24)$$

$$\sum_{j \in J} \sum_{i \in I} Z_{ijk} \leq 1 \quad \forall k \in K \quad (25)$$

$$\sum_{s \in S} W_{sj} - \sum_{i \in I} h_i Y_{ij} = 0 \quad \forall j \in J \quad (26)$$

$$\sum_{s \in S} W_{sj} - t_j X_j \leq 0 \quad \forall j \in J \quad (27)$$

$$\sum_{m \in P} Z_{imk} + \sum_{h \in P} Z_{jhk} - Y_{ij} \leq 1 \quad \forall j \in J; \forall i \in I; \forall k \in K \quad (28)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (29)$$

$$Y_{ij} \in \{0,1\} \quad \forall i \in I; \forall j \in J \quad (30)$$

$$Z_{ijk} \in \{0,1\} \quad \forall i \in P; \forall j \in P; \forall k \in K \quad (31)$$

$$W_{sj} \geq 0 \quad \forall s \in S ; \forall j \in J \quad (32)$$

The objective function (19) minimizes the sum of the fixed facility location costs, the shipment costs from the supply points (plants) to the facilities, the variable facility throughput costs and the routing costs to the customers. Constraint (20) requires each customer to be on exactly one route. Constraint (21) imposes a capacity restriction for each vehicle, while constraint (22) limits the length of each route. Constraint (23) requires each route to be connected to a facility. The constraint requires that there be at least one route that goes from any set V (a proper subset of the points P that contains the set of candidate facility sites) to its complement, thereby precluding routes that only visit customer nodes. Constraint (24) states that any route entering node $i \in P$ also must exit that same node. Constraint (25) states that a route can operate out of only one facility. Constraint (26) defines the flow into a facility from the supply points in terms of the total demand that is served by the facility. Constraint (27) restricts the throughput at each facility to the maximum allowed at that site and links the flow variables and the facility location variables. Thus, if a facility is not opened, there can be no flow through the facility, which in turn (by constraint (26)) precludes any customers from being assigned to the facility. Constraint (28) states that if route $k \in K$ leaves customer node $i \in I$ and also leaves facility $j \in J$, then customer $i \in I$ must be assigned to facility $j \in J$. This constraint links the vehicle routing variables (Z_{ijk}) and the assignment variables (Y_{ij}). Constraints (29)-(32) are standard integrality and non-negativity constraints.

Even for small problem instances, the formulation above is a difficult mixed integer linear programming problem. Perl solves the problem using a three-phased heuristic. The first phase finds minimum cost routes. The second phase determines

which facilities to open and how to allocate the routes from phase one to the selected facilities. The third phase attempts to improve the solution by moving customers between facilities and re-solving the routing problem with the set of open facilities fixed. The algorithm iterates between the second and third phases until the improvement at any iteration is less than some specified value. Wu, Low, and Bai (2002) propose a similar two-phase heuristic for the problem and test it on problems with up to 150 nodes.

Like the three-layer formulation of Perl, two-layer location routing formulations (e.g., Laporte, Nobert and Pelletier (1983), Laporte, Nobert and Arpin (1986) and Laporte, Nobert and Taillefer (1988)) usually are based on integer linear programming formulations for the vehicle routing problem (VRP). Flow formulations of the VRP often are classified according to the number of indices of the flow variable: $X_{ij} = 1$ if a vehicle uses arc (i,j) or $X_{ijk} = 1$ if vehicle k uses arc (i,j). The size and structure of these formulations make them difficult to solve using standard integer programming or network optimization techniques. Motivated by the successful implementation of exact algorithms for set-partitioning-based routing models, Berger (1997) formulates a two-layer location/routing problem that closely resembles the classical fixed charge facility location problem. Unlike other location/routing problems, she formulates the routes in terms of paths, where a delivery vehicle may not be required to return to the distribution center after the final delivery is made. The model is appropriate in situations where the deliveries are made by a contract carrier or where the commodities to be delivered are perishable. In the latter case, the time to return from the last customer to the distribution center is much less important than the time from the facility to the last customer. Berger defines the following notation:

Inputs and Sets:

P_j set of feasible paths from candidate distribution center $j \in J$

c_{jk} cost of serving the path $k \in P_j$

a_{ik}^j 1 if delivery path $k \in P_j$ visits customer $i \in I$; 0 if not

Decision Variables:

$V_{jk} \begin{cases} 1 & \text{if path } k \in P_j \text{ is operated out of distribution center } j \in J \\ 0 & \text{if not} \end{cases}$

Note that there can be any number of restrictions on the feasible paths in set P_j ; in fact, the more restrictive the conditions imposed on P_j are, the smaller the cardinality of P_j is. Restricting the total length of the paths, Berger formulates the following integrated location/routing model:

$$\text{Minimize } \sum_{j \in J} f_j X_j + \sum_{j \in J} \left\{ \sum_{k \in P_j} c_{jk} V_{jk} \right\} \quad (33)$$

$$\text{Subject to: } \sum_{j \in J} \left\{ \sum_{k \in P_j} a_{ik}^j V_{jk} \right\} = 1 \quad \forall i \in I \quad (34)$$

$$V_{jk} - X_j \leq 0 \quad \forall j \in J; \forall k \in P_j \quad (35)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (36)$$

$$V_{jk} \in \{0,1\} \quad \forall j \in J; \forall k \in P_j \quad (37)$$

The objective function (33) minimizes the sum of the facility location costs and the vehicle routing costs. Constraint (34) requires each demand node to be on one route. Constraint (35) states that a route can be assigned only to an open facility. Constraints (36) and (37) are standard integrality constraints.

Although the similarity between this location routing model and the classical fixed charge location model (1)-(5) is striking, this model is much more difficult to solve for two reasons. First, the linear programming relaxation provides a weak lower bound.

The linear programming relaxation typically has solutions in which the path variables are assigned very small fractional values and the location variables are assigned fractional variables large enough only to satisfy constraints (35). To strengthen significantly the linear programming relaxation, constraints (35) can be replaced by the following constraints:

$$\sum_{k \in P_j} a_{ik}^j V_{jk} - X_j \leq 0 \quad \forall i \in I; \forall j \in J \quad (38)$$

Consider a customer node $i \in I$ that is served (in part) using routes that emanate from facility $j \in J$. The first term of (38) is the sum of all route assignment variables that serve that customer and that are assigned to that facility. (In the linear programming relaxation, these assignment variables may be fractional). Thus, this sum can be thought of as the fraction of demand node $i \in I$ that is served out of facility $j \in J$. The constraint requires the location variable to be no smaller than the largest of these sums for customers assigned (in part) to routes emanating from the facility.

Second, there is an exponential number of feasible paths associated with any candidate facility, so complete enumeration of all possible columns of the problem is prohibitive. Instead, Berger develops a branch-and-price algorithm, which uses column generation to solve the linear programs at each node of the branch-and-bound tree. The pricing problem for the model decomposes into a set of independent resource-constrained shortest path problems.

The development and the use of location routing models have been more limited than both facility location and vehicle routing models. In our view, the reason is that it is difficult to combine, in a meaningful way, facility location decisions, which typically are strategic and long-term, and vehicle routing decisions, which typically are tactical and

short-term. The literature includes several papers that attempt to accommodate the fact that the set of customers to be served on a route may change daily, while the location of a distribution center may remain fixed for years. One approach is to define a large number of customers and to introduce a probability that each customer will require service on any day. Jaillet (1985, 1988) introduces this concept in the context of the probabilistic traveling salesman problem. Jaillet and Odoni (1988) provide an overview of this work and related probabilistic vehicle routing problems. The idea is extended to location routing problems in Berman, Jaillet and Simchi-Levi (1995). Including different customer scenarios, however, increases the difficulty of the problem, so this literature tends to locate a single distribution center. In our view, the problem of approximating LTL vehicle tours in facility location problems without incurring the cost of solving an embedded vehicle routing or traveling salesman problem remains an open challenge worthy of additional research.

Integrated Location/Inventory Models

The fixed charge location problem ignores the inventory impacts of facility location decisions; it deals only with the tradeoff between facility costs, which increase with the number of facilities, and the average travel cost, which decreases approximately as the square root of the number of facilities located (call it N). Inventory costs increase approximately as the square root of N . As such, they introduce another force that tends to drive down the optimal number of facilities to locate. Baumol and Wolfe (1958) recognized the contribution of inventory to distribution costs over forty years ago when they stated, “standard inventory analysis suggests that, optimally, important inventory

components will vary approximately as the square root of the number of shipments going through the warehouse.” (p. 255) If the total number of shipments is fixed, the number through any warehouse is approximately equal to the total divided by N . According to Baumol and Wolfe, the cost at each warehouse is then proportional to the square root of this quantity. When the cost per warehouse is multiplied by N , we see that the total distribution cost varies approximately with the square root of N . This argument treats the cost of holding working or cycle stock; Eppen (1979) argued that safety stock costs also increase as the square root of N (assuming equal variance of demand at each customer and independence of customer demands).

While the contribution of inventory to distribution costs has been recognized for many years, only recently have we been able to solve the resulting non-linear models that result from incorporating inventory decisions in facility location models. Shen (2000) and Shen, Coullard, and Daskin (2003) introduced a location model with risk pooling (LMRP). The model minimizes the sum of fixed facility location costs, direct transportation costs to the customers (which are assumed to be linear in the quantity shipped), working and safety stock inventory costs at the distribution centers and shipment costs from a plant to the distribution center (which may include a fixed cost per shipment). The last two quantities—the inventory costs at the distribution centers and the shipment costs of goods to the distribution centers—depend on the allocation of customers to the distribution centers. Shen introduces the following additional notation:

Inputs and Sets:

- μ_i, σ_i^2 mean and variance of the demand per unit time at customer $i \in I$
- c_{ij} a term that captures the annualized unit cost of supplying customer $i \in I$ from facility $j \in J$ as well as the variable shipping cost from the supplier to facility $j \in J$

- ρ_j a term that captures the fixed order costs at facility $j \in J$ as well as the fixed transport costs per shipment from the supplier to facility $j \in J$ and the working inventory carrying cost at facility $j \in J$
- ω_j a term that captures the lead time of shipments from the supplier to facility $j \in J$ as well as the safety stock holding cost

With this notation, Shen formulates the LMRP as follows:

$$\text{Minimize } \sum_{j \in J} f_j X_j + \sum_{j \in J} \sum_{i \in I} c_{ij} \mu_i Y_{ij} + \sum_{j \in J} \rho_j \sqrt{\sum_{i \in I} \mu_i Y_{ij}} + \sum_{j \in J} \omega_j \sqrt{\sum_{i \in I} \sigma_i^2 Y_{ij}} \quad (39)$$

$$\text{Subject to } \sum_{j \in J} Y_{ij} = 1 \quad \forall i \in I \quad (2)$$

$$Y_{ij} - X_j \leq 0 \quad \forall i \in I; \forall j \in J \quad (3)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (4)$$

$$Y_{ij} \geq 0 \quad \forall i \in I; \forall j \in J \quad (5)$$

The first term of the objective function (39) represents the fixed facility location costs. The second term captures the cost of shipping from the facilities to the customers as well as the variable shipment costs from the supplier to the facilities. The third term represents the working inventory carrying costs which include any fixed (per shipment) costs of shipping from the supplier to the facilities. The final term represents the safety stock costs at the facilities. Note that the objective function is identical to that of the fixed charge location problem (1) with the addition of two non-linear terms, the first of which captures economies of scale regarding fixed ordering and shipping costs and the second of which captures the risk pooling associated with safety stocks. Also note that the constraints of the LMRP are identical to those of the fixed charge location problem.

Shen (2000) and Shen, Coullard, and Daskin (2003) recast this model as a set covering problem where the sets contain customers to be served by facility $j \in J$. As in

Berger's location/routing model, the number of possible sets is exponentially large. Thus, they propose solving the problem using column generation. The pricing problems are non-linear integer programs, but their structure allows for a low-order polynomial solution algorithm. Shen assumes that the variance of demand is proportional to the mean. If demands are Poisson, this assumption is exact and not an approximation. With this assumption, he is able to collapse the final two terms in the objective function into one term. The resulting pricing problems can then be solved in $O(|I|\log|I|)$ time for each candidate facility and in $O(|J||I|\log|I|)$ time for all candidate facilities at each iteration of the column generation algorithm. Shu, Teo, and Shen (2001) showed that the pricing problem with two square root terms (i.e., without assuming that the variance-to-mean ratio is constant for all customers) can be solved in $O(|I|^2 \log|I|)$ time. Daskin, Coullard, and Shen (2003) developed a Lagrangian relaxation algorithm for this model and found it to be slightly faster than the column generation method.

One of the important qualitative findings from Shen's model is that, as inventory costs increase as a percentage of the total cost, the number of facilities located by the LMRP is significantly smaller than the number that would have been sited by the uncapacitated fixed charge location model, which ignores the risk pooling effects of inventory management. Shen and Daskin (2003) have extended the model above to account for customer service considerations. As customer service increases in importance, the number of facilities used in the optimal solution grows, eventually approaching and even exceeding the number used in the uncapacitated fixed charge model.

Several joint location/inventory models appeared in the literature prior to Shen's work. Barahona and Jensen (1998) solve a location problem with a fixed cost for stocking a given product at a DC. Erlebacher and Meller (2000) use various heuristic techniques to solve a joint location-inventory problem with a highly non-linear objective function. Teo, Ou, and Goh (2001) present a $\sqrt{2}$ -approximation algorithm for the problem of choosing DCs to minimize location and inventory costs, ignoring transportation costs. Nozick and Turnquist (2001a, 2001b) present models that, like Shen's model, incorporate inventory considerations into the fixed charge location problem; however, they assume that inventory costs are linear, rather than concave, and DC-customer allocations are made based only on distance, not inventory.

Ozsen, Daskin, and Coullard (2003) have extended the LMRP to incorporate capacities at the facilities. Capacities are modeled in terms of the maximum (plausible) inventory accumulation during a cycle between order receipts. This model is considerably harder to solve than is its uncapacitated cousin. However, it highlights an important new dimension in supply chain operations that is not captured by the traditional capacitated fixed charge location model. In the traditional model, capacity is typically measured in terms of throughput per unit time. However, this value can change as the number of inventory turns per unit time changes. Thus, the measure of capacity in the traditional model is often suspect. Also, using the traditional model, there are only two ways to deal with capacity constraints as demand increases: build more facilities or reallocate customers to more remote facilities that have excess capacity. In the capacitated version of the LMRP, a third option is available, namely ordering more frequently in smaller quantities. By incorporating this extra dimension of choice, the

capacitated LMRP is more likely to reflect actual managerial options than is the traditional fixed charge location model.

To some extent, merging inventory management with facility location decisions suffers from the same conceptual problems as merging vehicle routing with location. Inventory decisions, as argued above, can be revised much more frequently than can facility location decisions. Nevertheless, there are three important reasons for research to continue in the area of integrated inventory-location modeling. First, early results suggest that the location decisions that are made when inventory is considered can be radically different from those that would be made by a procedure that fails to account for inventory. Second, as indicated above, the capacitated LMRP better models actual facility capacities than does the traditional fixed charge location model, as it introduces the option of ordering more often to accommodate increases in demand. Third, we can solve fairly large instances of the integrated location-inventory model outlined above. In particular, the Lagrangian approach can often solve problems with 600 customers and 600 candidate facility sites in a matter of minutes on today's desktop computers.

Planning Under Uncertainty

Long-term strategic decisions like those involving facility locations are always made in an uncertain environment. During the time when design decisions are in effect, costs and demands may change drastically. However, classical facility location models like the fixed charge location problem treat data as though they were known and deterministic, even though ignoring data uncertainty can result in highly suboptimal

solutions. In this section, we discuss approaches to facility location under uncertainty that have appeared in the literature.

Most approaches to decision making under uncertainty fall into one of two categories: stochastic programming or robust optimization. In stochastic programming, the uncertain parameters are described by discrete scenarios, each with a given probability of occurrence; the objective is to minimize the expected cost. In robust optimization, parameters may be described either by discrete scenarios or by continuous ranges; no probability information is known, however, and the objective is typically to minimize the worst-case cost or regret. (The regret of a solution under a given scenario is the difference between the objective function value of the solution under the scenario and the optimal objective function value for that scenario.) Both approaches seek solutions that perform well, though not necessarily optimally, under any realization of the data. We provide a brief overview of the literature on facility location under uncertainty here. For a more comprehensive review, the reader is referred to Owen and Daskin (1998) or Berman and Krass (2001).

Sheppard (1974) was one of the first authors to propose a stochastic approach to facility location. He suggests selecting facility locations to minimize the expected cost, though he does not discuss the issue at length. Weaver and Church (1983) and Mirchandani, Oudjit, and Wong (1985) present a multi-scenario version of the P -median problem. Their model can be translated into the context of the fixed-charge location problem as follows. Let S be a set of scenarios. Each scenario $s \in S$ has a probability q_s of occurring and specifies a realization of random demands (h_{is}) and travel costs (c_{ijs}). Location decisions must be made now, before it is known which scenario will occur.

However, customers may be assigned to facilities after the scenario is known, so the Y variables are now indexed by a third subscript, s . The objective is to minimize the total expected cost. The stochastic fixed charge location problem is formulated as follows:

$$\text{Minimize} \quad \sum_{j \in J} f_j X_j + \sum_{s \in S} \sum_{j \in J} \sum_{i \in I} q_s h_{is} c_{ijs} Y_{ijs} \quad (40)$$

$$\text{Subject to} \quad \sum_{j \in J} Y_{ijs} = 1 \quad \forall i \in I; \forall s \in S \quad (41)$$

$$Y_{ijs} - X_j \leq 0 \quad \forall i \in I; \forall j \in J; \forall s \in S \quad (42)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (43)$$

$$Y_{ijs} \geq 0 \quad \forall i \in I; \forall j \in J; \forall s \in S \quad (44)$$

The objective function (40) computes the total fixed cost plus the expected transportation cost. Constraint (41) requires each customer to be assigned to a facility in each scenario. Constraint (42) requires that facility to be open. Constraints (43) and (44) are integrality and non-negativity constraints. The key to solving this model and the P -median-based models formulated by Weaver and Church (1983) and Mirchandani, Oudjit, and Wong (1985) is recognizing that the problem can be treated as a deterministic problem with $|I||S|$ customers instead of $|I|$.

Snyder, Daskin, and Teo (2003) consider a stochastic version of the LMRP. Other stochastic facility location models include those of Louveaux (1986), França and Luna (1982), Berman and LeBlanc (1984), Carson and Batta (1990), and Jornsten and Bjorndal (1994).

Robust facility location problems tend to be more difficult computationally than stochastic problems because of their minimax structure. As a result, the literature on

robust facility location generally falls into one of two categories: analytical results and polynomial-time algorithms for restricted problems like 1-median problems or P -medians on tree networks (see Chen and Lin (1998), Burkhard and Dollani (2001), Vairaktarakis and Kouvelis (1999), and Averbakh and Berman (2000)) and heuristics for more general problems (Serra, Ratick, and ReVelle (1996), Serra and Marianov (1998), and Current, Ratick, and ReVelle (1997)).

Solutions to the stochastic fixed charge problem formulated above may perform well in the long run but poorly in certain scenarios. To address this problem, Snyder and Daskin (2003b) combine the stochastic and robust approaches by finding the minimum-expected-cost solution to facility location problems subject to an additional constraint that the relative regret in each scenario is no more than a specified limit. They show empirically that by reducing this limit, one obtains solutions with substantially reduced maximum regret without large increases in expected cost. In other words, there are a number of near-optimal solutions to the fixed charge problem, many of which are much more robust than the true optimal solution.

Location Models with Facility Failures

Once a set of facilities has been built, one or more of them may from time to time become unavailable – for example, due to inclement weather, labor actions, natural disasters, or changes in ownership. These facility “failures” may result in excessive transportation costs as customers previously served by these facilities must now be served by more distant ones. In this section, we discuss models for choosing facility locations to minimize fixed and transportation costs while also hedging against failures within the system. We call the ability of a system to perform well even when parts of the system have failed the “reliability” of the system. The goal, then, is to choose facility locations that are both inexpensive and reliable.

The robust facility location models discussed in the previous section hedge against uncertainty in the problem data. By contrast, reliability models hedge against uncertainty in the solution itself. Another way to view the distinction in the context of supply chain design is that robustness is concerned with “demand-side” uncertainty (uncertainty in demands, costs, or other parameters), while reliability is concerned with “supply-side” uncertainty (uncertainty in the availability of plants or distribution centers).

The models discussed in this section are based on the fixed charge location problem; they address the tradeoff between *operating cost* (fixed location costs and day-to-day transportation cost—the classical fixed charge problem objective) and *failure cost* (the transportation cost that results after a facility has failed). The first model considers the *maximum* failure cost that can occur when a single facility fails, while the second model considers the *expected* failure cost given a fixed probability of failure. The strategy behind both formulations is to assign each customer to a *primary* facility (which

serves it under normal conditions) and one or more *backup* facilities (which serve it when the primary facility has failed). Note that although we refer to primary and backup *facilities*, “primariness” is a characteristic of assignments, not facilities; that is, a given facility may be a primary facility for one customer and a backup facility for another.

In addition to the notation defined earlier, let

$$Y_{ijk} = \begin{cases} 1 & \text{if facility } j \in J \text{ serves as the primary facility and facility } k \in J \text{ serves as the secondary} \\ & \text{facility for customer } i \in I \\ 0 & \text{if not} \end{cases}$$

and let V be a desired upper bound on the failure cost that may result if a facility fails.

Snyder (2003) formulates the maximum-failure-cost reliability problem as follows:

$$\text{Minimize} \quad \sum_{j \in J} f_j X_j + \sum_{i \in I} \sum_{j \in J} \sum_{k \in J} h_i c_{ij} Y_{ijk} \quad (45)$$

$$\text{Subject to} \quad \sum_{j \in J} \sum_{k \in J} Y_{ijk} = 1 \quad \forall i \in I \quad (46)$$

$$\sum_{k \in J} Y_{ijk} \leq X_j \quad \forall i \in I; \forall j \in J \quad (47)$$

$$Y_{ijk} \leq X_k \quad \forall i \in I; \forall j \in J; \forall k \in J \quad (48)$$

$$\sum_{i \in I} \sum_{k \in J} \sum_{l \in J} h_i c_{ik} Y_{ikl} + \sum_{i \in I} \sum_{k \in J} h_i c_{ik} Y_{ijk} \leq V \quad \forall j \in J \quad (49)$$

$$Y_{ijj} = 0 \quad \forall i \in I; \forall j \in J \quad (50)$$

$$X_j \in \{0,1\} \quad \forall j \in J \quad (51)$$

$$Y_{ijk} \geq \{0,1\} \quad \forall i \in I; \forall j \in J; \forall k \in J \quad (52)$$

The objective function (45) sums the fixed cost and transportation cost to customers from their primary facilities. (The summation over k is necessary to determine

the assignments, but the objective function does not depend on the backup assignments.) Constraint (46) requires each customer to be assigned to one primary and one backup facility. Constraints (47) and (48) prevent a customer from being assigned to a primary or a backup facility, respectively, that has not been opened. (The summation on the left-hand side of (47) can be replaced by Y_{ijk} without affecting the IP solution, but doing so considerably weakens the LP bound.) Constraint (49) is the reliability constraint and requires the failure cost for facility j to be no greater than V . The first summation computes the cost of serving each customer from its primary facility if its primary facility is not j , while the second summation computes the cost of serving customers assigned to j as their primary facility from their backup facilities. Constraint (50) requires a customer's primary facility to be different from its backup facility, and constraints (51) and (52) are standard integrality and non-negativity constraints. This model can be solved for small instances using an off-the-shelf IP solver, but larger instances must be solved heuristically.

The expected-failure-cost reliability model (Snyder and Daskin, 2003a) assumes that multiple facilities may fail simultaneously, each with a given probability q of failing. In this case, a single backup facility is insufficient, since a customer's primary and backup facilities may both fail. Therefore, we define

$$Y_{ijr} = \begin{cases} 1 & \text{if facility } j \in J \text{ serves as the level-} r \text{ facility for customer } i \in I \\ 0 & \text{if not} \end{cases}$$

A "level- r " assignment is one for which there are r closer facilities that are open. If $r = 0$, this is a primary assignment; otherwise it is a backup assignment. The objective is to minimize a weighted sum of the operating cost (the fixed charge location problem objective) and the expected failure cost, given by

$$\sum_{i \in I} \sum_{j \in J} \sum_{r=0}^{|J|-1} h_i c_{ij} q^r (1-q) Y_{ijr} .$$

Each customer i is served by its level- r facility (call it j) if the r closer facilities have failed (this occurs with probability q^r) and if j itself has not failed (this occurs with probability $1 - q$). The full model is omitted here. This problem can be solved efficiently using Lagrangian relaxation.

Few firms would be willing to choose a facility location solution that is, say, twice as expensive as the optimal solution to the fixed charge problem just to hedge against occasional disruptions to the supply chain. However, Snyder and Daskin (2003a) show empirically that it often costs very little to “buy” reliability: like robustness, reliability can be improved substantially with only small increases in cost.

Conclusions and Directions for Future Work

Facility locations decisions are critical to the efficient and effective operation of a supply chain. Poorly placed plants and warehouses can result in excessive costs and degraded service no matter how well inventory policies, transportation plans, and information sharing policies are revised, updated, and optimized. At the heart of many supply chain facility location models is the fixed charge location problem. As more facilities are located, the facilities tend to be closer to customers resulting in lower transport costs, but higher facility costs. The fixed charge facility location problem finds the optimal balance between fixed facility costs and transportation costs. Three important extensions of the basic model consider (1) facility capacities and single sourcing requirements, (2) multiple echelons in the supply chain, and (3) multiple products.

The fixed charge location problem, as well as these extensions, assume that shipments from the warehouses or distribution centers to the customers or retailers are made in truckload quantities. In reality, distribution to customers is often performed using less-than-truckload routes that visit multiple customers. This chapter reviewed two different approaches to formulating integrated location/routing models. However, as indicated above, these approaches suffer from the fundamental problem that facility locations are typically determined at a strategic level while vehicle routes are optimized at the operational level. In other words, the set of customers and their demands may change daily resulting in daily route changes, while the facilities are likely to be fixed for years. We believe that additional research is needed to find improved ways of approximating the impact of less-than-truckload deliveries on facility location costs without embedding a vehicle routing problem (designed to serve one realization of customer demands) in the facility location model.

Incorporating inventory decisions in facility location models appears to be critical for supply chain modeling. As early as 1958, researchers recognized that inventory costs would tend to increase with the square root of the number of facilities used. Only recently, however, have non-linear models that approximate this relationship between inventory costs and location decisions been formulated and solved optimally. While we believe that these models represent an important step forward in location modeling for supply chain problems, considerable additional research is needed. In particular, researchers should attempt to incorporate more sophisticated inventory models, including multi-item inventory models and models that account for inventory accumulation at all echelons of the supply chain. Heuristic approaches to the multi-item problem have

recently been proposed by Balcik (2003) and an optimal approach has been suggested by Snyder (2003). The latter model, however, assumes that items are ordered separately, resulting in individual fixed order costs for each commodity purchased.

Finally, since facility location decisions are inherently strategic and long-term in nature, supply chain location models must account for the inherent uncertainty surrounding future conditions. We have reviewed a number of scenario-based location models as well as models that account for unreliability in the facilities themselves. This too is an area worthy of considerable additional research. For example, generating scenarios that capture future uncertainty and the relationships between uncertain parameters is one critical area of research. Reliability-based location models for supply chain management are still in their infancy. In fact, it is not immediately clear how to marry reliability modeling approaches and the integrated location/inventory models we reviewed, since the non-linearities introduced by the inventory terms complicate the computation of failure costs. In this regard, the more general techniques of stochastic programming (Birge and Louveaux, 1997) may ultimately prove fruitful.

References

- Al-Sultan, K. S. and M. A. Al-Fawzan, 1999, "A Tabu Search Approach to the Uncapacitated Facility Location Problem," *Annals of Operations Research*, **86**, pp. 91-103.
- Averbakh, I. and O. Berman, 2000, "Minmax Regret Median Location on a Network Under Uncertainty," *INFORMS Journal on Computing*, **12**:2, pp. 104-110.
- Balcik, B., 2003, "Multi-Item Integrated Location/Inventory Problem," M.S. Thesis, Department of Industrial Engineering, Middle East Technical University.

- Balinski, M. L., 1965, "Integer Programming: Methods, Uses, Computation," *Management Science*, **12**, 253-313.
- Barahona, F. and D. Jensen, 1998, "Plant Location with Minimum Inventory," *Mathematical Programming*, **83**, pp. 101-111.
- Baumol, W. J. and P. Wolfe, 1958, "A Warehouse-Location Problem," *Operations Research*, **6**, pp. 252-263.
- Benders, J. F., 1962. "Partitioning Procedures for Solving Mixed-Variables Programming Problems," *Numerische Mathematik*, **4**, pp. 238-252.
- Berger, R., 1997, "Location-Routing Models for Distribution System Design," Ph.D. Dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.
- Berman, O., and B. LeBlanc, 1984, "Location-Relocation of Mobile Facilities on a Stochastic Network," *Transportation Science*, **18**:4, pp. 315-330.
- Berman, O., P. Jaillet and D. Simchi-Levi, 1995, "Location-Routing Problems with Uncertainty," in *Facility Location: A Survey of Applications and Methods*, Z. Drezner, ed., Springer, New York.
- Berman, O., and D. Krass, 2002, "Facility Location Problems with Stochastic Demands and Congestion," in *Facility Location: Applications and Theory*, Z. Drezner and H. W. Hamacher, eds., Springer, New York, pp. 331-373.
- Birge, J. R. and F. Louveaux, 1997, *Introduction to Stochastic Programming*, Springer, New York.
- Burkhard, R. E. and H. Dollani, 2001, "Robust Location Problems with Pos/Neg Weights on a Tree," *Networks*, **38**:2, pp. 102-113.
- Carson, Y. M. and Y. Batta, 1990, "Locating an Ambulance on the Amherst Campus of the State University of New York at Buffalo," *Interfaces*, **20**:5, pp. 43-49.
- Chen, B. T. and C. S. Lin, 1998, "Minimax-Regret Robust 1-Median Location on a Tree," *Networks*, **31**, 93-103.
- Cornuéjols, G., G. L. Nemhauser and L. A. Wolsey, 1990, "The Uncapacitated Facility Location Problem," in *Discrete Location Theory*, P. B. Mirchandani and R. L. Francis, eds., Wiley, NY, pp. 119-171.
- Current, J., S. Ratick, and C. ReVelle, 1997, "Dynamic Facility Location When the Total Number of Facilities is Uncertain: A Decision Analysis Approach," *European Journal of Operational Research*, **110**:3, pp. 597-609.

- Daskin, M. S., 1995, Network and Discrete Location: Models, Algorithms and Applications, John Wiley and Sons, Inc., New York.
- Daskin, M. S., C. Coullard, and Z.-J. M. Shen, 2002, "An Inventory-Location Model: Formulation, Solution Algorithm and Computational Results," *Annals of Operations Research*, **110**, 83-106.
- Daskin, M.S., P. C. Jones, 1993, "A New Approach To Solving Applied Location/Allocation Problems," *Microcomputers in Civil Engineering*, **8**, pp. 409-421.
- Eilon, S., C. D. T. Watson-Gandy, and N. Christofides, 1971, Distribution Management: Mathematical Modeling and Practical Analysis, Hafner Publishing Co., NY.
- Eppen, G., 1979, "Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem," *Management Science* **25**:5, pp. 498-501.
- Erlebacher, S. J. and R. D. Meller, 2000, "The Interaction of Location and Inventory in Designing Distribution Systems," *IIE Transactions* **32**, pp. 155-166.
- Erlenkotter, D., 1978, "A Dual-Based Procedure for Uncapacitated Facility Location," *Operations Research*, **26**, pp. 992-1009.
- França, P. M. and H. P. L. Luna, 1982, "Solving Stochastic Transportation-Location Problems by Generalized Benders Decomposition," *Transportation Science*, **16**, pp. 113-126.
- Galvão, R. D., 1993, "The Use of Lagrangean Relaxation in the Solution of Uncapacitated Facility Location Problems," *Location Science*, **1**:1, pp. 57-79.
- Geoffrion, A.M., 1974, "Lagrangian Relaxation for Integer Programming," *Mathematical Programming Study*, **2**, pp. 82-114.
- Geoffrion, A. M. and G. W. Graves, 1974, "Multicommodity Distribution System Design by Benders Decomposition," *Management Science*, **20**:5, pp. 822-844.
- Geoffrion, A. M. and R. F. Powers, 1980, "Facility Location Analysis is Just the Beginning (if you do it right)," *Interfaces*, **10**:2, pp. 22-30.
- Glover, F., 1989, "Tabu Search — Part I," *ORSA Journal on Computing*, **1**:3, pp. 190-206.
- Glover, F., 1990, "Tabu Search — Part II," *ORSA Journal on Computing*, **2**:1, pp. 4-32.

- Glover, F. and M. Laguna, 1997, Tabu Search, Kluwer Academic Publishers, Boston, MA.
- Hansen, P. and N. Mladenović, 1997, "Variable Neighborhood Search for the P -Median," *Location Science*, **5**:4, pp. 207-226.
- Hakimi, S. L., 1964, "Optimum Location of Switching Centers and the Absolute Centers and Medians of a Graph," *Operations Research*, **12**, pp. 450-459.
- Hakimi, S. L., 1965, "Optimum Distribution of Switching Centers in a Communication Network and Some Related Graph Theoretic Problems," *Operations Research*, **13**, pp. 462-475.
- Jaillet, P., 1985, "The Probabilistic Traveling Salesman Problem," Technical Report 185, Operations Research Center, M.I.T., Cambridge, MA.
- Jaillet, P., 1988, "A Priori Solution of a Traveling Salesman Problem in Which a Random Subset of the Customers are Visited," *Operations Research*, **36**, pp. 929-936.
- Jaillet, P. and A. Odoni, 1988, "Probabilistic Vehicle Routing Problems," in Vehicle Routing: Methods and Studies, B.L. Golden and A.A. Assad, eds., North-Holland, Amsterdam, pp. 293-318.
- Jornsten, K. and M. Bjorndal, 1994, "Dynamic Location Under Uncertainty," *Studies in Regional and Urban Planning*, **3**, pp. 163-184.
- Krarup, J. and P. M. Pruzan, 1983, "The Simple Plant Location Problem: Survey and Synthesis," *European Journal of Operational Research*, **12**, 36-81.
- Laporte, G., 1988, "Location Routing Problems," in Vehicle Routing: Methods and Studies, B.L. Golden and A.A. Assad (Editors). North-Holland, Amsterdam, pp. 163-197.
- Laporte, G., Y. Nobert, and D. Arpin, 1986, "An Exact Algorithm for Solving a Capacitated Location-Routing problem," *Annals of Operations Research*, **6**, pp. 293-310.
- Laporte, G., Y. Nobert, and J. Pelletier, 1983, "Hamiltonian Location Problems," *European Journal of Operational Research*, **12**, pp. 82-89.
- Laporte G., Y. Nobert, and S. Taillefer, 1988, "Solving a Family of Multi-Depot Vehicle Routing and Location-Routing problems," *Transportation Science*, **22**, pp. 161-172.
- Louveaux, F. V., 1986, "Discrete Stochastic Location Models," *Annals of Operations Research*, **6**, pp. 23-34.

- Maranzana, F. E., 1964, "On the Location of Supply Points to Minimize Transport Costs," *Operational Research Quarterly*, **15**, pp. 261-270.
- Min, H., V. Jayaraman and R. Srivastava, 1998, "Combined Location-Routing Problems: A Synthesis and Future Research Directions," *European Journal of Operational Research*, **108**, pp. 1-15.
- Mirchandani, P. B., A. Oudjit, and R. T. Wong, 1985, "Multidimensional Extensions and a Nested Dual Approach for the m-Median Problem," *European Journal of Operational Research*, **21**, pp. 121-137.
- Nozick, L. K. and M. A. Turnquist, 2001a, "Inventory, Transportation, Service Quality and the Location of Distribution Centers," *European Journal of Operational Research*, **129**, pp. 362-371.
- Nozick, L. K. and M. A. Turnquist, 2001b, "A Two-Echelon Inventory Allocation and Distribution Center Location Analysis," *Transportation Research Part E*, **37**, pp. 421-441.
- Owen, S. H. and M. S. Daskin, 1998, "Strategic Facility Location: A Review," *European Journal of Operational Research*, **111**, pp. 423-447.
- Ozsen, L., M. S. Daskin and C. R. Coullard, 2003, "Capacitated Facility Location Model with Risk Pooling," submitted for publication.
- Perl, J., 1983, "A Unified Warehouse Location-Routing Analysis," Ph.D. Dissertation, Department of Civil Engineering, Northwestern University, Evanston, IL
- Perl, J. and M. S. Daskin, 1985, "A Warehouse Location-Routing Problem," *Transportation Research*, **19B**:5, pp. 381-396.
- Serra, D. and V. Marianov, 1998, "The p -Median Problem in a Changing Network: The Case of Barcelona," *Location Science*, **6**, pp. 383-394.
- Serra, D., S. Ratick, and C. ReVelle, 1996, "The Maximum Capture Problem with Uncertainty," *Environment and Planning B*, **23**, pp. 49-59.
- Shen, Z. J., 2000, "Efficient Algorithms for Various Supply Chain Problems," Ph.D. Dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208.
- Shen, Z.-J. M., C. Coullard, and M. S. Daskin, 2003, "A joint location-inventory model," *Transportation Science*, **37**(1), pp. 40-55.

- Shen, Z.-J. M. and M. S. Daskin, 2003, "Tradeoffs Between Customer Service and Cost in an Integrated Supply Chain Design Framework," submitted to *Manufacturing and Service Operations Management*.
- Sheppard, E. S., 1974, "A Conceptual Framework for Dynamic Location-Allocation Analysis," *Environment and Planning A*, **6**, pp. 547-564.
- Shu, J., C.-P. Teo, and Z.-J. M. Shen, 2001, "Stochastic Transportation-Inventory Network Design," preprint.
- Simchi-Levi, D., P. Kaminsky, and E. Simchi-Levi, 2003, Designing and Managing the Supply Chain: Concepts Strategies and Case Studies, Second Edition, McGraw-Hill, Irwin, Boston, MA.
- Snyder, L. V., 2003, "Supply Chain Robustness and Reliability: Models and Algorithms," Ph.D. Dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208.
- Snyder, L. V. and M. S. Daskin, 2003a, "Reliability Models for Facility Location: The Expected Failure Cost Case," submitted for publication.
- Snyder, L. V. and M. S. Daskin, 2003b, "Stochastic p -Robust Location Problems," Working Paper.
- Snyder, L. V., M. S. Daskin and C.-P. Teo, 2003, "The Stochastic Location Model with Risk Pooling," submitted for publication.
- Teitz, M. B. and P. Bart, 1968, "Heuristic Methods for Estimating the Generalized Vertex Median of a Weighted Graph," *Operations Research*, **16**, pp. 955-961.
- Teo, C.-P., J. Ou, and M. Goh, 2001, "Impact on Inventory Costs with Consolidation of Distribution Centers," *IIE Transactions*, **33**:2, pp. 99-110.
- Vairaktarakis, G. L. and P. Kouvelis, 1999, "Incorporation Dynamic Aspects and Uncertainty in 1-Median Location Problems," *Naval Research Logistics*, **46**:2, pp. 147-168.
- Weaver, J. R. and R. L. Church, 1983, "Computational Procedures for Location Problems on Stochastic Networks," *Transportation Science*, **17**, pp. 168-180.
- Wu, T.-H., C. Low, and J.-W. Bai, 2002, "Heuristic Solutions to Multi-Depot Location-Routing Problems," *Computers and Operations Research*, **29**, pp. 1393-1415.