Designing the Distribution Network for an Integrated Supply Chain

Jia Shu * and Jie Sun †

Abstract

We consider an integrated distribution network design problem in which all the retailers face uncertain demand. The risk-pooling benefit is achieved by allowing some of the retailers to operate as distribution centers (DCs) with commitment in service level. The target is to minimize the expected total cost resulted from the DC location, transportation, and inventory. We formulate it as a two-stage nonlinear discrete stochastic optimization problem. The first stage decides which retailers to be selected as DCs and the second stage deals with the costs of DC-retailer assignment, transportation, and inventory. In the literature, the similar models require the demands of all retailers in each scenario to have their variances identically proportional to their means. In this paper, we remove this restriction. We formulate the problem by using a set-covering model, and solve the problem by a column generation approach. With a variable fixing technique, we are able to efficiently solve problems of moderate-size (up to one hundred retailers and nine scenarios). The solution technique exploits only the concavity of the risk-pooling cost structure and can therefore be used in solving more general problems.

1 Introduction

In a highly competitive environment, today's companies must consider their decisions in inventory, transportation, and facility location in an integrated manner. In this paper, we study a mathematical model for integrated distribution network design, which minimizes the expected total cost, including the facility set-up cost, the inventory cost, and the transportation cost over all considered scenarios. This target naturally leads to a two-stage stochastic optimization model. In the first stage, we decide which retailers should be selected as the distribution centers (DCs). In the second stage, after the random parameters are realized, we decide

*Industrial and Manufacturing Engineering Department, Wayne State University, USA

†Department of Decision Sciences and Singapore-MIT Alliance, National University of Singapore. This author's research is partially supported by NUS grants RP314000-042/057-112.

- which retailers to be assigned to a given DC in each scenario;
- how much safety stock to maintain at a given DC so as to protect against the uncertainty in the demands in each scenario;
- how much working inventory to keep at a given DC so as to satisfy the demands of the served retailers in each scenario.

These requirements are often encountered in practice, e.g., in the ALKO Inc. case from a widely used textbook on supply chain management by Chopra and Meindl (2001). Daskin et al. (2002) also mentioned the possible application of this type of models in e-commerce.

Shen (2000), Shen et al. (2003) and Daskin et al. (2002) discuss a single-echelon supply chain network design problem. They use Lagrangian-relaxation and column generation algorithms to handle a nonlinear integer optimization problem. They show that their model can be solved efficiently when the demand at each DC is Poisson or deterministic. Shu, Teo and Shen (2005) and Shu (2003) extended the model to arbitrary demand. The papers of Teo and Shu (2004) and Romeijn, Shu and Teo (2005) considered infinite horizon multi-echelon network design problems in both deterministic and stochastic demand settings. However, to the best of our knowledge, very few research works in this area jointly consider location, inventory, and transportation costs.

Another important feature of our model is to properly incorporate the uncertainties into the choice of the DCs. In the literature of facility location, it is a common practice to consider the possible change in the parameters by a scenario based approach. The pioneering work includes Sheppard (1974), Berman and Krass (2001), and Owen and Daskin (1998), which developed into the study of the stochastic facility location problem. The main concern of those works are the location cost and the distribution cost. As a result, the inventory related cost are often ignored or simplified.

Snyder et al. (2005) provided a stochastic version for the location model with risk pooling by Shen (2000), Shen et al. (2003) and Daskin et al. (2002). A Lagrangian-relaxation based algorithm is adopted and a variable fixing approach is used to reduce the problem size. Similar to the deterministic version, however, there is a rather restrictive assumption on the demand pattern, i.e., the demand should either be deterministic or have a variance identically proportional to its mean for all the retailers.

This paper removes the restrictive demand pattern requirement of Snyder et al. (2005). We reformulate the problem as a set-covering model, solve the pricing problem in $O(n^2 \log n)$ time, and speed up the column generation process using the variable fixing routines.

2 Model Formulation

We propose a two-stage model with discrete scenarios to hedge against the uncertainties. This model is general and flexible enough to be used in other problem settings. For example, it can be interpreted as a multi-period or a multi-commodity version of the model studied in Shu, Teo and Shen (2005).

Stage 1

For a given set I of retailers, we need to determine which retailers to be operated as DCs, which naturally introduces the first stage decision variables. For any retailer $j \in I$, we have

- $X_j = \begin{cases} 1, & \text{if the retailer at location } j \text{ is set up as a DC,} \\ 0, & \text{otherwise.} \end{cases}$

Clearly, the cost associated with this decision is the location cost, i.e., if a DC is set up at location j, the incurred cost is

- f_j: fixed (annual) cost of locating a DC at location j, $\forall j \in I$.

Stage 2

Let P be the set of scenarios to be considered in the second stage. For each scenario $p \in P$, we want to decide which DC to serve each retailer $i, i \in I$. In our model, each retailer can only be served by one DC (single sourcing) for the ease of implementation, i.e., we assume the demand for all the retailers cannot be split in any scenario. Therefore, for each $i \in I$, $j \in I$ and $p \in P$, we have a binary variable

- $Y_{ijp} = \begin{cases} 1, & \text{if DC } j \text{ is used to serve retailer } i \text{ in scenario } p; \\ 0, & \text{otherwise.} \end{cases}$

In different scenarios, the demand patterns faced by each retailer are different. Let us define that each retailer i has a demand pattern with

- μ_{ip}: mean (daily) demand at retailer i in scenario p, $\forall i \in I, \forall p \in P$;
- σ_{ip}^2: variance of (daily) demand at retailer i in scenario p, $\forall i \in I, \forall p \in P$;

in scenario p.

Suppose that DC j serves all the retailers in S, $S \subseteq I$ in scenario p. The associated cost includes the following components:

- $\sum_{i \in S} \alpha_{ijp}\mu_{ip}$: a term that is linear in μ_{ip}. It can be interpreted as the transportation cost from DC j to retailer i, for all $i \in S$. In this setting, α_{ijp} corresponds to the unit transportation cost from DC j to retailer i in scenario p.

- $G_{jp}(\sum_{i \in S} \mu_{ip})$: a term that is concave and nondecreasing in the expected throughput assigned. This concave term captures the economies of scale, e.g., it can be regarded as the total inventory and ordering cost for the DC j when using the EOQ model or in the case of ALKO.

- $H_{jp}(\sum_{i \in S} \sigma_{ip}^2)$: a term that is concave and nondecreasing in the total demand variance experienced by the DC j in scenario p. This term reflects the risk-pooling effects by consolidating demand at a centralized location. The safety stock cost at the DC j is usually approximated by a concave function of the total demand variance for all the retailers it serves.

As a common approach in the scenario-based model, each scenario $p \in P$ is assigned a probability P_p, and the objective is to minimize the first stage cost combined with the expectation of the second stage cost. Then, our integrated distribution network design model can be formulated as:

$$\min \quad \sum_{j \in I} f_j X_j + \sum_{p \in P} \sum_{j \in I} P_p \Bigg\{ \Bigg(\sum_{i \in I} (\alpha_{ijp}\mu_{ip}) Y_{ijp} \Bigg) + G_{jp} \Bigg(\sum_{i \in I} \mu_{ip} Y_{ijp} \Bigg) + H_{jp} \Bigg(\sum_{i \in I} \sigma_{ip}^2 Y_{ijp} \Bigg) \Bigg\} \tag{1}$$

$$\text{subject to} \qquad \sum_{j \in I} Y_{ijp} = 1, \qquad \forall i \in I, \quad \forall p \in P \tag{2}$$

$$Y_{ijp} - X_j \leq 0, \qquad \forall i, j \in I, \quad \forall p \in P \tag{3}$$

$$Y_{ijp} \in \{0,1\}, \qquad \forall i, j \in I, \quad \forall p \in P \tag{4}$$

$$X_j \in \{0,1\}, \qquad \forall j \in I \tag{5}$$

When there is only one scenario, the first two terms in (1) and the constraints (2)-(5) consist of the formulation of the uncapacitated facility location model, which is already NP-hard. Our model considers multiple scenarios and includes two nonlinear terms with binary variables and therefore tends to be more difficult to solve.

An Example: The Risk-Pooling Model

Here we use the model described in Snyder et al. (2005), which is a special case of our model, to illustrate the generic model (1)-(5). The reader might refer to Snyder et al. (2005) for a more detailed derivation. The decisions to be made are exactly the same as that in the generic model (1)–(5). Besides the demand parameters, the probabilities of the scenarios, and the DC location costs, the other inputs in their model are summarized in the following:

- β: weight factor associated with the shipment cost
- θ: weight factor associated with the inventory cost
- d_{ijp}: per unit shipment cost from DC-based retailer j to retailer i, $\forall i, j \in I, \forall p \in P$
- a_{jp}: per unit shipment cost from an external supplier to DC j, $\forall p \in P$
- h_p: inventory holding cost per unit of product per year $\forall p \in P$
- α: desired percentage of retailers orders satisfied (fill rate)
- z_α: standard normal deviate such that $P(z \leq z_\alpha) = \alpha$
- L_{jp}: lead time in days from the supplier to DC j, $\forall p \in P$
- g_{jp}: fixed shipment cost from an external supplier to DC j, $\forall p \in P$
- F_{jp}: fixed cost of placing an order at DC j, $\forall j \in I, \forall p \in P$

Note that our solution approach allows α, β, and θ to be scenario-dependent. For ease of exposition, we might assume that they are the same for all scenarios in our analysis.

The model Snyder et al. (2005) considers two kinds of inventory at the DCs: the working inventory to satisfy the demand of retailers and the safety stock to protect against the uncertainties in demand faced by the retailers. It is assumed that all the retailers keep minimum inventory and thus their inventory costs are ignored. However, the location cost of the DCs, the transportation cost from the supplier to the DCs and from the DCs to the retailers, and the ordering cost for the DCs are all taken into account. As a result, the cost components include:

From Stage 1

- DC location cost: $\sum_j f_j X_j$, which is the first term in (1).

From each scenario of Stage 2

- Average transportation cost from DCs to retailers: $\beta \sum_{j\in I} \sum_{i\in I} \mu_{ip}(d_{ijp} + a_{jp})Y_{ijp}$. It is easy to see the term $\beta(d_{ijp} + a_{jp})$ corresponds to the factor α_{ijp} in (1).
- Average working inventory cost and ordering cost: $\sum_{j\in I} \sqrt{2\theta h_p(F_{jp} + \beta g_{jp}) \sum_{i\in I} \mu_{ip} Y_{ijp}}$. For any DC j, this term is clearly approximated by the solution to the EOQ equation with ordering cost $F_{jp} + \beta g_{jp}$, holding cost θh_p and demand $\sum_{i\in I} \mu_{ip} Y_{ijp}$. It represents the term $G_{jp}(\cdot)$ in (1).

- Average safety stock cost: $\theta h_p z_\alpha \sqrt{L_{jp}} \sum_{j\in I} \sqrt{\sum_{i\in I} \sigma_{ip}^2 Y_{ijp}}$. This component captures the risk-pooling effects, and obviously is a special case for the term $H_{jp}(\cdot)$ in (1).

To sum up, the objective function is

$$\sum_{j\in I} f_j X_j + \sum_{p\in P} P_p \left(\beta \sum_{j\in I} \sum_{i\in I} \mu_{ip}(d_{ijp}+a_{jp}) Y_{ijp} + \sum_{j\in I} \sqrt{2\theta h_p (F_{jp} + \beta g_{jp}) \sum_{i\in I} \mu_{ip} Y_{ijp}} + \theta h_p z_\alpha \sum_{j\in I} \sqrt{\sum_{i\in I} L_{jp} \sigma_{ip}^2 Y_{ijp}} \right).$$

Since this problem shares the same constraints of the generic model, the risk-pooling model can be formulated as

$$\min \quad \sum_{j\in I} f_j X_j + \sum_{p\in P} \sum_{j\in I} P_p \left(\sum_{i\in I} \hat{d}_{ijp} Y_{ijp} + K_{jp} \sqrt{\sum_{i\in I} \mu_{ip} Y_{ijp}} + q_{jp} \sqrt{\sum_{i\in I} \sigma_{ip}^2 Y_{ijp}} \right) \tag{6}$$

$$\text{subject to} \quad \sum_{j\in I} Y_{ijp} = 1, \qquad \forall i \in I, \quad \forall p \in P \tag{7}$$

$$Y_{ijp} - X_j \le 0, \qquad \forall i, j \in I, \quad \forall p \in P \tag{8}$$

$$Y_{ijp} \in \{0, 1\}, \qquad \forall i, j \in I, \quad \forall p \in P \tag{9}$$

$$X_j \in \{0, 1\}, \qquad \forall j \in I, \tag{10}$$

where

$$\begin{aligned} \hat{d}_{ijp} &= \beta \mu_{ip}(d_{ijp} + a_{jp}), \\ K_{jp} &= \sqrt{2\theta h_p (F_{jp} + \beta g_{jp})}, \\ q_{jp} &= \theta h_p z_\alpha \sqrt{L_{jp}}. \end{aligned}$$

3 Set-Covering Formulation

In the literature, it is a common practice to approximate the nonlinear integer programming models by linearizing the objective function and to solve the problem by a linear integer programming approach (cf. Erlebacher and Meller (2000)). Sometimes the nonlinear integer programming problem can also be solved directly using the Lagrangian relaxation approach (cf. Snyder et al. (2005)) or heuristic approaches (cf. Sun and Gu (2002)). In this section, we reformulate our problem as a set-covering model, which replaces the nonlinear objective function with a linear objective function of an exponential number of variables, and naturally fits into the framework of the column generation procedure. The set-covering model is known to be equivalent ("dual") to the Lagrangian relaxation approach (cf. Bertsimas and Tsitsiklis (1997)) in theory but it comes with significant computational advantages in our case.

For any scenario, a feasible solution to our decision problem consists of a partition of the set I of retailers into subsets, $R_1, R_2, \ldots, R_n$, together with one designated DC for each of these n sets. Let $\mathcal{R}$ be the collection of all subsets of the set I, i.e., $\mathcal{R} = \{S : S \subseteq I\}$. Let $C_{j,R_1,R_2,\ldots,R_{|p|}}$ be the expected total cost associated with DC j if it serves all retailers in R_p in scenario p for any scenario $p \in P$, i.e.,

$$C_{j,R_1,R_2,\ldots,R_{|P|}} = f_j + \sum_{p\in P} P_p \Big\{ \sum_{i\in R_p} \alpha_{ijp}\mu_{ip} + G_{jp}\Big(\sum_{i\in R_p} \mu_{ip}\Big) + H_{jp}\Big(\sum_{i\in R_p} \sigma_{ip}^2\Big)\Big\}.$$

We redefine the decision variables $Z_{j,R_1,R_2,\ldots,R_{|P|}}$, for all $j \in I$, $R_1, R_2, \ldots, R_{|P|} \in \mathcal{R}$. $Z_{j,R_1,R_2,\ldots,R_{|P|}}$ is equal to 1 if DC j serves all retailers in R_p in scenario p for any scenario $p \in P$; and 0 otherwise. With this notation, the set-covering reformulation of problem (1)–(5) becomes:

$$\begin{aligned}
(\mathcal{M}_{\mathcal{R}}) \qquad \min \quad & \sum_{j\in I}\sum_{R_1\in\mathcal{R}}\sum_{R_2\in\mathcal{R}}\cdots\sum_{R_{|P|}\in\mathcal{R}} C_{j,R_1,R_2,\ldots,R_{|P|}} Z_{j,R_1,R_2,\ldots,R_{|P|}} \\
\text{subject to} \quad & \sum_{j\in I}\sum_{R_p\subseteq\mathcal{R}: i\in R_p} Z_{j,R_1,R_2,\ldots,R_{|P|}} \geq 1, \qquad \forall i \in I, p \in P, \\
& Z_{j,R_1,R_2,\ldots,R_{|P|}} \in \{0,1\}, \qquad \forall j \in I, R_1, R_2, \ldots, R_{|P|} \in \mathcal{R}.
\end{aligned}$$

We can solve the linear relaxation of $(\mathcal{M}_{\mathcal{R}})$ by the column generation approach. At each iteration, we solve the LP with a partial set of columns (the so-called master LP problem), which gives us an optimal solution $\bar{Z}_{j,R_1,R_2,\ldots,R_{|P|}}$, $j \in I$, $R_1, R_2, \ldots, R_{|P|} \in \mathcal{R}$, and the corresponding optimal LP dual solution $\bar{\pi}_{ip}$, $i \in I, p \in P$. To check whether we obtain the optimal solution to the linear relaxation of $(\mathcal{M}_{\mathcal{R}})$, we need to know whether the reduced cost

$$C_{j,R_1,R_2,\ldots,R_{|P|}} - \sum_{p\in P}\sum_{i\in R_p} \bar{\pi}_{ip} \geq 0$$

is nonnegative for each $j \in I$ and $R_1, R_2, \ldots, R_{|P|} \in \mathcal{R}$. If the answer is yes, then $\bar{Z}$ is an optimal solution to the linear relaxation of $(\mathcal{M}_{\mathcal{R}})$. Otherwise, some column $(j, R_1, R_2, \ldots, R_{|P|})$ must have a negative reduced cost, so we add that column $(j, R_1, R_2, \ldots, R_{|P|})$ to the master LP, and begin the next iteration.

An efficient way to check the reduced cost for each column is to solve the following integer programming problem for each fixed j:

$$\begin{aligned}
(\mathcal{P}_j) \qquad \min \quad & f_j + \sum_{p\in P} P_p \Big\{ \sum_{i\in I} \frac{(P_p\alpha_{ijp}\mu_{ip} - \bar{\pi}_{ip})}{P_p} Y_{ijp} + G_{jp}\Big(\sum_{i\in I}\mu_{ip}Y_{ijp}\Big) + H_{jp}\Big(\sum_{i\in I}\sigma_{ip}^2 Y_{ijp}\Big)\Big\} \\
\text{subject to} \quad & Y_{ijp} \in \{0,1\}, \quad \forall i \in I,
\end{aligned}$$

which we call the *pricing problem*. Obviously, if the optimal value of $\mathcal{P}_j$ is nonnegative for all $j \in I$, then all reduced costs must be nonnegative. When there exists some j with the optimal value of $\mathcal{P}_j$ to be negative, we find the column to be added in the master problem.

4 Solution Procedure

In this section we first show how to solve the pricing problem efficiently. We then show how to speed up the column generation procedure using a variable fixing technique.

4.1 The Pricing Problem

In the pricing problem $(\mathcal{P}_j)$, we have a constant f_j and a summation of some separable functions. Therefore, we can remove f_j from the objective and decompose $(\mathcal{P}_j)$ into $|P|$ problems of $(\mathcal{P}'_{jp})$ which corresponds to scenario $p \in P$:

$$
\begin{array}{lll}
(\mathcal{P}'_{jp}) \qquad & \omega^*_{jp} = \min & \displaystyle\sum_{i \in I} a_i z_i + G_{jp}\left(\sum_{i \in I} b_i z_i\right) + H_{jp}\left(\sum_{i \in I} c_i z_i\right) \\
& \text{subject to} & z_i \in \{0,1\}, \quad \forall i \in I,
\end{array}
$$

where

$$
\begin{aligned}
a_i &= (P_p \alpha_{ijp} \mu_{ip} - \bar{\pi}_{ip})/P_p, \\
b_i &= \mu_{ip}, \\
c_i &= \sigma^2_{ip}, \\
z_i &= Y_{ijp}.
\end{aligned}
$$

Let z^* and ω^*_{jp} be the optimal solution and the optimal objective value to $(\mathcal{P}'_{jp})$ respectively. Clearly, $\sum_{p \in P} P_p \omega^*_{jp} + f_j$ is the optimal value of $(\mathcal{P}_j)$, which corresponds to the reduced cost for all the columns with designated DC j. Therefore, to efficiently solve $(\mathcal{P}'_{jp})$ is essential for the column generation procedure. Note that the pricing problem introduced in Snyder et al. (2005) is a special case of $(\mathcal{P}'_{jp})$.

Proposition 1 *For any $(\mathcal{P}'_{jp})$, $j \in I$ and $p \in P$, there exists an optimal solution z^* such that $a_i < 0$ for any $z^*_i = 1$, $i \in I$.*

Proof: Let z^* be an optimal solution to $(\mathcal{P}'_{jp})$. Suppose that there exists an k, $k \in I$ with $z^*_k = 1$ and $a_k \geq 0$. Since $b_i, c_i \geq 0$, G_{jp} and H_{jp} are concave and nondecreasing, we have

$$
\sum_{i \in I} a_i z^*_i + G_{jp}\left(\sum_{i \in I} b_i z^*_i\right) + H_{jp}\left(\sum_{i \in I} c_i z^*_i\right) \geq \sum_{i \in I, i \neq k} a_i z^*_i + G_{jp}\left(\sum_{i \in I, i \neq k} b_i z^*_i\right) + H_{jp}\left(\sum_{i \in I, i \neq k} c_i z^*_i\right).
$$

Therefore, z' such that $z'_i = z^*_i$ for any $i \in I \setminus \{k\}$ and $z'_k = 0$, must be an optimal solution. By repeating this process, we can obtain an optimal solution z'' with $a_i < 0$ for any $z''_i = 1$. $i \in I$.

□

As a result of Proposition 1, we may restrict our search for the optimal solution in the set I^-, which is defined by $I^- := \{i | a_i < 0, i \in I\}$. Note that there are $2^{|I^-|}$ choices of z. To find an efficient algorithm, we identify a nice structure of the solution z^* (cf. Chakravarty et al. (1985) and Shu, Teo and Shen (2005)).

Shu, Teo and Shen (2005) showed that the above problem can be solved efficiently by separating the points $(\frac{b_i}{a_i}, \frac{c_i}{a_i})$ with a line in a 2-dimensional plane, which is the *primal approach*. Each separation corresponds to an extreme point of the associated zonotope and the total number of different partitions are small. By using the point-line duality, Shu (2003) gave the *dual approach* for this problem. By further exploiting the special structure, Shu, Teo and Shen (2005) proposed a more *straightforward algorithm* for the above problem, which is the third approach. All three algorithms implicitly utilize the fact that the set of all lines in 2-dimesional plane has low VC-dimension and the computational complexity is $O(n^2 \log n)$.

4.2 Variable Fixing

We should note that the problem $(\mathcal{P}'_{jp})$ needs to be solved for each retailer j, $j \in I$ and each scenario p, $p \in P$. Although the solution algorithm for $(\mathcal{P}'_{jp})$ is $O(n^2 \log n)$, this approach still significantly slows down the column generation routine (cf. Shu, Teo and Shen (2005)). A possible method to speed up this process is to decide whether a retailer can be a DC in the optimal solution early in the column generation routine (variable fixing). In this section, we show how the primal and dual information obtained in the column generation procedure can be used to achieve effective variable fixing for our problem.

Recall that the set-covering model we are trying to solve is in the form:

$$\begin{aligned}
\min \quad & \sum_{j\in I} \sum_{R_1 \in \mathcal{R}} \sum_{R_2 \in \mathcal{R}} \cdots \sum_{R_{|P|} \in \mathcal{R}} C_{j,R_1,R_2,\ldots,R_{|P|}} Z_{j,R_1,R_2,\ldots,R_{|P|}} \\
\text{subject to} \quad & \sum_{j\in I} \sum_{R_p \subseteq \mathcal{R}: i \in R_p} Z_{j,R_1,R_2,\ldots,R_{|P|}} \geq 1, \qquad \forall i \in I, p \in P, \\
& Z_{j,R_1,R_2,\ldots,R_{|P|}} \in \{0,1\}, \qquad \forall j \in I, R_1, R_2, \ldots, R_{|P|} \in \mathcal{R}.
\end{aligned}$$

After each iteration of the column generation routine, we have:

- A set of dual prices $\bar{\pi}_{ip}, i \in I, p \in P$.
- A set of primal feasible (fractional) solution $Z_{j,R_1,R_2,\ldots,R_{|P|}}, j \in I, R_1, R_2, \ldots, R_{|P|} \in \mathcal{R}$.
- After solving the pricing problem (one for each potential DC location), we obtain the reduced cost $r_j \equiv \min_{R_1,R_2,\ldots,R_{|P|}} \left(C_{j,R_1,R_2,\ldots,R_{|P|}} - \sum_{p\in P} \sum_{i \in R_p} \bar{\pi}_{ip} \right), j \in I$. Note that some of the r_j's may be non-negative.

Let Z_{IP} and Z_{LP} denote the optimal integral and fractional solution to the set-covering problem.

Proposition 2 $\sum_{j:r_j\leq 0} r_j + \sum_p \sum_j \bar{\pi}_{jp}$ *is a lower bound to* Z_{IP}.

Proof: In the optimal IP solution, by concavity of the objective function, it is easy to check that $C_{j,R_1,R_2,\ldots,R_{|P|}} + C_{j,R'_1,R'_2,\ldots,R'_{|P|}} \geq C_{j,R_1\cup R'_1,R_2\cup R'_2,\ldots,R_{|P|}\cup R'_{|P|}}$. Hence, for every j, $j \in I$, the inequality

$$\sum_{R_1}\sum_{R_2}\cdots\sum_{R_{|P|}} Z_{j,R_1,R_2,\ldots,R_{|P|}} \leq 1$$

is valid for the problem. Therefore, we can add the above inequality, one for each j, to the set-covering model, to obtain a stronger LP relaxation. The Lagrangian dual of the new LP relaxation is thus equivalent to

$$L(\lambda) = \sum_p\sum_j \lambda_{jp} + \min\Big\{ \Big(\sum_j\sum_{R_1}\cdots\sum_{R_{|P|}}(C_{j,R_1,\ldots,R_{|P|}} - \sum_{p\in P}\sum_{i\in R_p}\lambda_{ip})Z_{j,R_1,\ldots,R_{|P|}}\Big) : \\ 1 \geq Z_{j,R_1,\ldots,R_{|P|}} \geq 0, \sum_{R_1}\cdots\sum_{R_{|P|}} Z_{j,R_1,\ldots,R_{|P|}} \leq 1,\ \forall\ j\Big\}.$$

The problem decomposes for each retailer j, and hence $Z_{IP} \geq \max_\lambda L(\lambda) \geq L(\bar{\pi}) = \sum_p\sum_j \bar{\pi}_{jp} + \sum_{j:r_j\leq 0} r_j$. □

We can now find a sufficient condition such that a retailer j^* with $r_{j^*} > 0$ can never be included in the optimal solution. Let UB be an upper bound for Z_{IP}.

Proposition 3 *If* $\sum_{j:r_j\leq 0} r_j + \sum_p\sum_j \bar{\pi}_j + r_{j^*} > UB$, *then retailer* j^* *will never be used as a DC in the optimal solution to the (integral) set-covering problem.*

Proof: Suppose that DC j^* is opened in an optimal solution. Then Z_{IP} remains unchanged if we impose the additional condition: $\sum_{R_1}\cdots\sum_{R_{|P|}} Z_{j^*,R_1,\ldots,R_{|P|}} = 1$ to the existing set of constraints. The Lagrangian dual, in this case, reduces to

$$L'(\lambda) = \sum_p\sum_j \lambda_{jp} + \min\Big\{ \Big(\sum_j\sum_{R_1}\cdots\sum_{R_{|P|}}(C_{j,R_1,\ldots,R_{|P|}} - \sum_{p\in P}\sum_{i\in R_p}\lambda_{ip})Z_{j,R_1,\ldots,R_{|P|}}\Big) : \\ 1 \geq Z_{j,R_1,\ldots,R_{|P|}} \geq 0, \sum_{R_1}\cdots\sum_{R_{|P|}} Z_{j,R_1,\ldots,R_{|P|}} \leq 1,\ \forall\ j, \sum_{R_1}\cdots\sum_{R_{|P|}} Z_{j^*,R_1,\ldots,R_{|P|}} = 1\Big\}.$$

Hence $Z_{IP} \geq \max_\lambda L'(\lambda) \geq L'(\bar{\pi}) = \sum_p\sum_j \bar{\pi}_{jp} + \sum_{j:r_j\leq 0} r_j + r_{j^*}$. On the other hand, we have $Z_{IP} \leq UB$. This gives rise to a contradiction. □

Once we determine that the retailer j^* will never be used as a DC in the optimal solution, then we do not need to solve the pricing problem corresponding to j^* any more in the rest of the column generation procedure. Moreover, all columns that has been generated using j^* as DC can also be deleted from the set-covering model, which further improves the effectiveness of the variable fixing routine.

The variable fixing technique depends largely on the quality of the upper bound UB. If in a certain stage of the column generation process, all $Z_{j,R_1,\ldots,R_{|P|}}$ in the optimal solution are integral, then the optimal value of $\sum_{j,R_1,\ldots,R_{|P|}} C_{j,R_1,\ldots,R_{|P|}} Z_{j,R_1,\ldots,R_{|P|}}$ obtained in that iteration is an upper bound to Z_{IP}. Unfortunately this is not true for all instances. As in Shu et al. (2005), we generate an upper bound for the IP by generating a feasible solution in the following way:

- Let Z^* be the optimal LP solution obtained by solving the problem using a partial set of columns.

- Order the retailer-scenario combinations according to non-decreasing value of demand.

- For a given retailer-scenario combination (say, for retailer i and scenario p) on the list, if there exists $Z^*_{j,R_1,\ldots,R_{|P|}} = 1$ and $i \in R_p$ for some DC j, then retailer i is served by DC j in scenario p. Otherwise, there exists a set of DCs S such that there exists $Z^*_{j,R_1,\ldots,R_{|P|}} > 0$ and $i \in R_p$ for any $j \in S$. We serve retailer i in scenario p using the DC that will lead to the least total cost, and remove the retailer-scenario combination for retailer i and scenario p from the list.

- Repeat the previous step until the list is empty.

In this way, we can generate a feasible solution to the distribution network design problem, which can be used as an upper bound to perform variable fixing in the column generation routine.

5 Computational Results

In this section, we summarize our computational experience with the algorithms outlined in the previous section. All the instances were solved on a P4-1.8G workstation running the Windows XP operating system. We used the risk pooling model as described in Shen et al. (2003) to design the computational experiment for our model. To facilitate proper comparison, we have also imposed the additional assumption, used in Shen et al. (2003) and Snyder et al. (2005), that a DC at retailer j must be used to serve the demand at retailer j. This is not necessarily true in the optimal solution. The solution approach described for the general case can be easily modified to handle this additional assumption.

Stochastic Network Design with Variable Fixing

The algorithm for the network design problem is coded in C++, and the linear programming problem is solved by using CPLEX 7.5 LP Solver.

The means μ_{ip} and variances σ_{ip}^2 of the demands are randomly generated in $[100, 1600]$ for all $i \in I, p \in P$. The holding cost $h = 1$, and the service level $\alpha = 97.5\%$ ($z_\alpha = 1.96$). For all $j \in I, p \in P$, the per unit shipment cost from an external supplier to DC j is $a_{jp} = 5$, the fixed shipment cost from an external supplier to DC j is $g_{jp} = 10$, and the cost of placing an order at DC j is $F_{jp} = 10$. We randomly generate $|P|$ numbers (N_p), and let $P_p = N_p / \sum_{p \in P} N_p$. It is easy to see we can ensure $\sum_{p \in P} P_p = 1$ in this way. Our goal is two fold. First, we want to know the size of problems that can be effectively solved by the algorithm. Second, we test how the solutions depend on the values of β, θ, and the fraction of retailers used as DCs in the solution.

For each of the instances, we first solve the linear programming relaxation of the set-covering model via column generation. The initial set of columns include all singletons. The column labeled "No. of Columns Generated" indicates the total number of columns added during this phase. The column labeled "No. of DCs OUT" indicates the number of retailers ruled out from being possible DCs in the optimal linear programming solution by the variable fixing technique. The resulting final optimal objective value is denoted by Z_{LP}. In most of all instances generated, the corresponding optimal solutions are integral. We denote by Z_H the best upper bound we obtained. In the case where the linear-programming relaxation solution is not integral, Z_H is obtained by applying an integer-programming solver to the final master problem.

To speed up the column generation algorithm, we do not solve the pricing problem at every iteration. Instead, we maintain a column pool and we first look for columns with negative reduced cost in the column pool. If the search is successful, we add these columns to the master problem. We also update the reduced costs of the columns in the column pool and remove those with large positive reduced costs from the pool. We solve the pricing problem only when the column pool is empty. When this happens, the pricing algorithm is run to either prove the optimality of the current solution or to find new columns with negative reduced costs. Only when the column pool is empty shall we do variable fixing procedure again.

Tables 1, 2, and 3 highlight the results of our computational study. For problems with up to 100 retailers and 9 scenarios, we are able to solve them within 15 minutes.

As we can see from these tables, the problem takes longer to solve when β decreases for fixed θ or when θ increases for fixed β. When the transportation costs increase relative to other costs, i.e., β increases, the No. of DCs increases too. While the inventory costs increase relative to other costs, i.e., θ increases, the No. of DCs decreases.

INPUT		OUTPUT					
β	θ	No. of Scenarios	No. of DCs selected	No. of DCs OUT	CPU Time (seconds)	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	3	7	31	21	187	1
0.001	0.1	5	6	20	94	853	1.002
0.001	0.1	9	7	9	137	1,577	1.002
0.005	0.1	3	14	20	3	74	1
0.005	0.1	5	13	14	28	232	1
0.005	0.1	9	14	7	44	471	1
0.005	0.5	3	13	18	10	172	1
0.005	0.5	5	11	11	40	361	1.001
0.005	0.5	9	11	8	61	672	1
0.005	1	3	11	18	4	87	1
0.005	1	5	10	16	32	282	1.002
0.005	1	9	10	11	55	591	1.001
0.005	10	3	6	31	20	191	1
0.005	10	5	7	21	99	973	1
0.005	10	9	7	10	158	1,781	1.001

Table 1: Computational results for the 40 retailers instance.

INPUT		OUTPUT					
β	θ	No. of Scenarios	No. of DCs Selected	No. of DCs OUT	CPU Time (seconds)	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	3	6	64	80	761	1
0.001	0.1	5	8	53	156	1,433	1
0.001	0.1	9	8	33	313	2,714	1.002
0.005	0.1	3	25	36	20	195	1
0.005	0.1	5	24	30	54	466	1
0.005	0.1	9	26	22	131	1,117	1.003
0.005	0.5	3	19	47	33	318	1
0.005	0.5	5	17	40	87	742	1
0.005	0.5	9	18	16	156	1,295	1
0.005	1	3	11	56	62	556	1
0.005	1	5	12	43	112	962	1
0.005	1	9	14	20	218	1,842	1.003
0.005	10	3	7	63	91	858	1.002
0.005	10	5	8	52	171	1,572	1
0.005	10	9	8	35	328	2,981	1.002

Table 2: Computational results for the 80 retailers instance.

INPUT		OUTPUT					
β	θ	No. of Scenarios	No. of DCs Selected	No. of DCs OUT	CPU Time (seconds)	No. of Columns Generated	Z_H/Z_{LP}
0.001	0.1	3	9	68	126	1,023	1.002
0.001	0.1	5	10	57	289	2,403	1
0.001	0.1	9	8	38	484	4,289	1.001
0.005	0.1	3	35	41	36	304	1
0.005	0.1	5	33	33	99	826	1.002
0.005	0.1	9	36	24	210	1,856	1
0.005	0.5	3	30	49	61	514	1
0.005	0.5	5	29	41	114	974	1
0.005	0.5	9	31	22	252	2,196	1.004
0.005	1	3	16	61	94	797	1
0.005	1	5	17	48	172	1,476	1
0.005	1	9	17	26	354	3,064	1.003
0.005	10	3	9	70	196	1,834	1.002
0.005	10	5	10	53	388	3,472	1
0.005	10	9	10	40	772	6,672	1

Table 3: Computational results for the 100 retailers instance.

6 Conclusions

In this paper we proposed a new model for the integrated distribution network design problem, which optimizes the expected total cost of location, transportation, inventory, and DC-retailer assignment under stochastic and scenario-based parameters.

A prototype of the model was proposed by Snyder et al. (2005), in which a special case of the general model was solved efficiently if in each scenario the demand is deterministic or has a variance identically proportional to its mean for all the retailers. Our model does not requires this condition and is therefore more realistic.

Computational study demonstrated that our approach can efficiently solve moderate-sized network design problem to near optimality. Furthermore, it appears that our solution techniques may be used in more general network design problems since the algorithm only uses the concavity property of the objective function, which suggests that the multi-period or multi-commodity network design problem could also be formulated and solved in this framework.

In our model, a retailer might be assigned to different DCs in different scenarios. In practice, sometimes to switch the DCs serving a retailer incurs high cost. Therefore, it is interesting to study the problem where DC-retailer assignment decision is scenario independent, i.e., DC-retailer assignments are the same across all the scenarios. The pricing problem arises from the column generation algorithm turns out containing multiple square-root components instead of two as in this paper. Designing an efficient algorithm for this problem is both challenging and interesting.

References

[1] Berman O. and D. Krass (2001), "Facility Location Problems with Stochastic Demands and Congestion", *Facility Location: Applications and Theory*, edited by A. Drezner and H.W. Hamacher, Springer-Verlag, New York, pp. 331-373.

[2] Bertsimas, D. and J. N. Tsitsiklis (1997), *Introduction to Linear Optimization*, Athena Scientific, Belmont, Massachusetts.

[3] Chakravarty, A. K., J. B. Orlin and U. G. Rothblum (1985), "Consecutive Optimizers for a Partitioning Problem with Applications to Optimal Inventory Groupings for Joint Replenishment", *Operations Research* **33** 820-834.

[4] Chopra, S. and P. Meindl (2001), *Supply Chain Management: Strategy, Planning and Operation.*

[5] Daskin, M. S., Collette R. Coullard and Z. J. Max Shen (2002), "An Inventory-Location Model: Formulation, Solution Algorithm and Computational Results", *Annals of Operations Research* **110** 83-106.

[6] Erlebacher, S. J. and R. D. Meller (2000), "The Interaction of Location and Inventory in Designing Distribution Systems", *IIE Transactions* **32** 155-166.

[7] Owen, S. H., and M. S. Daskin (1998), "Strategic Facility Location: A Review", *European Journal of Operational Research* **111** 423-447.

[8] Romeijn, H. E., J. Shu and C. P. Teo (2005), "Optimal Design of Two-Echelon Supply Chain Network", working paper.

[9] Shen, Z. J. Max (2000), "Approximation Algorithms for Various Supply Chain Problems", Ph.D. Thesis, Dept. of Industrial Engineering and Management Sciences, Northwestern University.

[10] Shen, Z. J. Max, Collette R. Coullard and Mark S. Daskin (2003), "A Joint Location-Inventory Model", *Transportation Science* **37** 40-55.

[11] Sheppard, E. S. (1974), "A Conceptual Framework for Dynamic Location-Allocation Analysis", *Environment and Planning A* **6** 547-564.

[12] Shu, J. (2003), "Integrated Logistics Network Design Optimization: The Impact of Multi-Echelon Inventory Management", Ph.D. Thesis, High Performance Computation for Engineered Systems, Singapore-MIT Alliance.

[13] Shu, J., C. P. Teo and Z. J. Max Shen (2005), "Stochastic Transportation-Inventory Network Design Problem", *Operations Research* **53** 48-60.

[14] Snyder, L. V., Mark S. Daskin and C. P. Teo (2005), "The Stochastic Location Model with Risk Pooling", submitted for publication.

[15] Sun, J. and Y. Gu (2002), "A Parametric Approach for a Nonlinear Discrete Location Problem", *Journal of Combinatorial Optimization* **6** 119-132.

[16] Teo, C. P., J. H. Ou and K. H. Goh (2001), "Impact on Inventory Costs with Consolidation of Distribution Centers", *IIE Transactions* **33** 99-110.

[17] Teo, C. P. and J. Shu (2004), "Warehouse-Retailer Network Design Problem", *Operations Research* **52** 396-408.

[18] Vapnik, V. N. and A. Ya. Chervonenkis (1971), "On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities", *Theory Probab. Appl.* **16** 264-280.