Performance Measurement for Inventory Routing

Martin Savelsbergh* Jin-Hwa Song*

*School of Industrial & Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0203
{mwps,jhsong}@isye.gatech.edu

1 Introduction

Vendor managed inventory resupply (VMI) has become a popular strategy to reduce inventory holding and/or distribution costs. In environments where VMI partnerships are in effect, the vendor is allowed to choose the timing and size of deliveries. In exchange for this freedom, the vendor agrees to ensure that its customers do not run out of product. Realizing the cost savings opportunities of VMI partnerships, however, is not an easy task, particularly with a large number and variety of customers. The inventory routing problem (IRP) seeks to do exactly that: determining a distribution strategy that minimizes *long term* distribution costs. A large body of literature on the IRP exists.

We do not focus on developing distribution strategies, but instead on measuring the effectiveness of distribution strategies. A popular performance measure used in practice to evaluate distribution strategies in an environment where VMI partnerships are in effect is the volume delivered per mile. As the volume that needs to be delivered by the vendor over a given period of time is determined by the total usage of its customers, and not under the control of the vendor, the vendor strives to minimize the total mileage required to deliver product. However, volume per mile by itself is not a meaningful number, because it is impacted by many factors, such as the geography of customer locations and customer usage patterns, but it is valuable for comparing performance in consecutive periods of time. If a company has a stable customer set and customer usage patterns do not fluctuate much, then an increase (decrease) in volume per mile indicates that distribution planning is improving (worsening).

The above discussion shows that volume per mile is a useful measure for monitoring relative distribution strategy performance. However, volume per mile cannot be used to determine, in an absolute sense, the quality of a distribution strategy. We develop a methodology that allows the computation of tight lower bounds on the total mileage required to deliver product over a period of time (and thus upper bounds on volume per mile). As a result, companies will be able to gain insight into the effectiveness of their distribution strategy.

2 A Simple Bound

Consider the following variant of the inventory routing problem. A single product has to be distributed from a single facility to a set I of n customers over a period of time of length T. Each customer $i \in I$ has the capability to maintain a local inventory of product up to a maximum of C_i. In the period of interest customer i consumes an amount u_i of product. A fleet of homogeneous vehicles, with capacity Q, is available for the distribution of the product. We assume an unlimited supply of product and an unlimited number of vehicles in the fleet. We denote the travel distance between two locations i and j by t_{ij}. The objective is to determine the minimum total distance required to satisfy all demand. Observe that when $C_i \geq Q \ \forall i \in I$, then the optimal distribution strategy is to always deliver a full truck load to a customer right when the customer's storage tank becomes empty. The resulting total distance is $\sum_{i \in I} \frac{u_i}{Q} 2t_{0i}$, where 0 denotes the plant. Therefore, a simple lower bound on the minimum total distance required to satisfy all demand is obtained by assuming that all customers' storage capacities are greater than the truck capacity, i.e.,

$$LB_1\colon \sum_{i \in I} \frac{u_i}{Q} 2t_{0i}$$

3 An Improved Bound

In practice, deliveries to customers with storage capacity less than the truck's capacity, i.e., $C_i < Q$, are usually combined with other deliveries to ensure a high utilization of the truck's capacity.

Define a feasible delivery pattern $P_j = (d_{j1}, d_{j2}, ..., d_{jn})$ to be a delivery pattern that satisfies $\sum_{i \in I} d_{ji} \leq Q$ and $0 \leq d_{ji} \leq C_i \ \forall i \in I$. Let $\delta(P_j) = \{i \in I : d_{ji} > 0\}$ denote the set of customers visited in delivery pattern P_j. The cost of delivery pattern P_j, denoted as $c(P_j)$, is the value of an optimal solution to the traveling salesman problem involving the plant and the customers in $\delta(P_j)$. Let $\mathcal{P}$ be the set of all feasible delivery patterns and let x_j be a decision variable indicating how many times delivery pattern P_j is used. Then the optimal objective function value of the following linear program, called the pattern selection LP, provides a lower bound on the total distance required to satisfy the demand

$$\begin{aligned} D^* = \quad & min && \sum_{P_j \in \mathcal{P}} c(P_j) x_j \\ & s.t. && \sum_{P_j \in \mathcal{P}} d_{ji} x_j \geq u_i, \quad \forall i \in I \\ & && x_j \geq 0 \end{aligned}$$

There are two major obstacles to using this linear program:

- The number of feasible delivery patterns is prohibitively large.

- The calculation of the cost of each delivery pattern involves the solution of a traveling salesman problem.

Below, we discuss how these obstacles can be handled in practice. We start by showing that a much smaller set of feasible delivery patterns can be considered when solving the linear program.

Definition 1. *(Base Pattern) A feasible delivery pattern P is a base pattern if at most one customer, say k, in $\delta(P)$ receives a delivery quantity less than $min(C_k, Q)$, and, in that case, the delivery quantity is $Q - \sum_{i \in \delta(P) \setminus \{k\}} C_i$.*

Theorem 1. *The base patterns are sufficient to find an optimal solution to the Pattern Selection LP.*

Now that we have significantly reduced the number of delivery patterns, we turn our attention to the number of customers visited in a delivery pattern as that impacts the effort required to compute the cost of a delivery pattern.

For any natural number k, let $C'_i = \frac{Q}{k}$ if $C_i < \frac{Q}{k}$, $C'_i = C_i$ if $\frac{Q}{k} \leq C_i \leq Q$, and $C'_i = Q$ if $C_i > Q$. Observe that with these modified storage capacities a base pattern contains at most k customers. Let LB_k denote the optimal value of the pattern selection LP with base patterns based on the modified storage capacities. It is easy to see that LB_k provides a lower bound on D^* for every k and that $LB_1 \leq LB_2 \leq LB_3 \leq ....$ Finally, when $\frac{Q}{k} \leq min\{C_1, C_2, ..., C_n\}$, then $LB_k = D^*$.

For any natural number k, we can also compute an upper bound UB_k on D^*, as follows. We let UB_k be the optimal objective function value of the pattern selection LP in which we only consider base patterns with at most k customers. It is easy to see that $UB_1 \geq UB_2 \geq UB_3 \geq ...$ and that when $k \geq \left\lceil \frac{Q}{min\{C_1, C_2, ..., C_n\}} \right\rceil$, then $UB_k = D^*$.

Our computational experiments have shown that tight bounds are obtained for values $k = 3$ and $k = 4$, in the sense that the gap between LB_k and UB_k is very small. Furthermore, for values $k = 3$ and $k = 4$, the traveling salesman problems that have to be solved involve at most 4 and 5 cities, respectively, and thus can be solved relatively easily by enumeration. Our computational experiments have also shown that even though we have significantly reduced the number of delivery patterns in the pattern selection LP by restricting ourselves to base patterns, the number of base patterns can still be huge (22,575,528 base patterns were generated to compute UB_4 for one of our larger instances). To be able to handle such large instances effectively, we have developed two additional techniques.

So far, we have only exploited feasibility considerations to reduce the set of delivery patterns that need to be considered. Next, we will show how optimality considerations can be exploited effectively to reduce the set of delivery patterns that need to be considered. Consider a base pattern $P = \{d_1, d_2, \ldots, d_n\}$ and the following linear program, called the dominance LP,

$$
\begin{aligned}
z = \quad min \quad & \sum_{\{j:\delta(P_j) \subsetneq \delta(P)\}} c(P_j)\lambda_j \\
s.t. \quad & \sum_{\{j:\delta(P_j) \subsetneq \delta(P)\}} d_{ji}\lambda_j \geq d_i, \quad \forall i \in \delta(P)
\end{aligned}
$$

$$\lambda_j \geq 0$$

If $z \leq c(P)$, then the base patterns with $\lambda_j > 0$ collectively dominate the base pattern P and base pattern P can be eliminated from the pattern selection LP. However, even though the size of a dominance LP is small, setting up and solving it for every base pattern to determine if the base pattern is dominated is computationally prohibitive. Therefore, we rely on easily computable upper bounds on the optimal value of a dominance LP for dominance testing; if $z \leq z_{UB} \leq c(P)$, where z_{UB} denotes an upper bound on z, then the base pattern is dominated and can be eliminated. We compute upper bound z_{UB} by restricting our attention to carefully selected subsets of patterns.

Next, we observe that a pattern selection LP has a large aspect ratio, i.e., a large ratio of number of columns to number of rows. Linear programs with large aspect ratios occur frequently when set partition or set covering formulations are used to model practical situations, for example in air crew scheduling applications. We have developed a specialized linear programming solver exploiting the fact that most variables will have a zero value in an optimal solution have been developed for such problems. The optimizer solves the LP with only a subset of the full set of variables (assuming a zero solution value for each of the remaining variables). From the solution to this partial LP, the reduced costs of the remaining variables can be computed. Variables with reduced costs less than zero are added to the partial LP and the partial LP is resolved and the process repeats. If no negative reduced cost variables exist, then the current solution is an optimal solution to the full problem.

4 Computational Experiments

We conducted various computational experiments to analyze the effect on the lower bound on the minimum total mileage required to satisfy demand of explicitly taking varying storage capacities into account. The data used in our experiments had usage information for about 2000 customers served from 36 plants (with the smallest plant serving about 10 customers and the largest plant serving about 150 customers). Each customer is supplied from one particular plant. Consequently, we are dealing with independent 36 instances. (The data was provided by Praxair Inc., a producer and distributor of industrial gases and long-time member of the Leaders in Logistics program at Georgia Tech.)

The primary experiment involved computing increasingly tighter lower and upper bounds on D^* the bound on the minimum total mileage required to satisfy demand. The results are displayed in 1.

First, the results shows that limiting ourselves to patterns with at most three or four customers is sufficient to obtain tight bounds on D^*. Second, the results show that allowing more deliveries per trip has a substantial effect on the upper bound, but hardly any effect on the lower bound. The latter result was somewhat counter to our expectations, but has important implications because it suggests that investing in larger storage facilities at customers, which is often discussed as a potential way of reducing distribution costs, may not deliver the desired savings. Finally, by comparing the actual incurred mileage to LB_4 over a period of time,

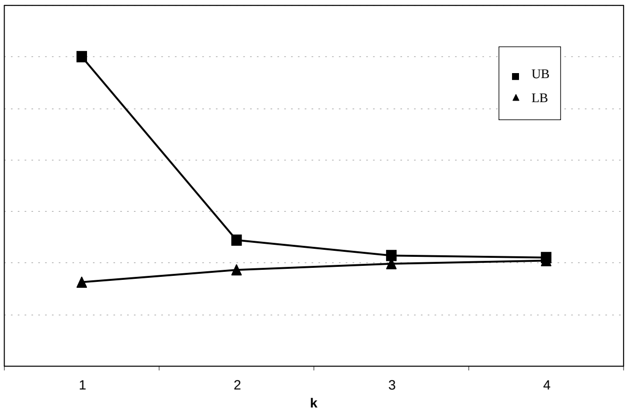

Figure 1: Lower and upper bounds on D^*

Praxair will gain insight in the effectiveness of its distribution strategy and in the potential savings that may result from improvements to its distribution strategy.

Next, we investigated whether the behavior observed for the complete system is also observed at the individual plant level. We found that for bounds LB_4 and UB_4 the largest relative gap is 2.53% for Plant 18 and the smallest relative gap is 0.02% from Plant 1. To understand the cause of the differences, we examined these plants in more detail. Two factors clearly impact the difference between the value of LB_4 and UB_4:

- The number of customers with $C_i < \frac{Q}{4}$
- The number of times we have to make deliveries to customers with $C_i < \frac{Q}{4}$

Note that when all customers served by a plant have $C_i \geq \frac{Q}{4}$, then we have $LB_4 = UB_4$. When we look more closely at Plant 1, we see that when a direct delivery policy would be employed, the number of deliveries is 568.4 (computed as $\sum_i \frac{u_i}{min(C_i,Q)}$). Among these 554.0 correspond to deliveries to customers with $C_i \geq \frac{Q}{4}$, i.e., 97.5% of the total number if deliveries. On the other hand, for Plant 18 the number of deliveries is 163.8 when a direct delivery policy is employed, out of which 100.9 correspond to deliveries to customers with $C_i \geq \frac{Q}{4}$, i.e., only 61.6% of the total number of deliveries.