Decision support for vendor managed inventory replenishment

Martin Savelsbergh, Georgia Institute of Technology, Atlanta, USA

The business practice called *vendor managed inventory replenishment* (VMI) involves a central decision maker controlling the replenishment of inventory at a number of locations, using a fleet of vehicles that may belong to the vendor or that may be rented. The central decision maker can be the supplier, and the inventory can be kept at independent customers, or the central decision maker can be a manager responsible for inventory replenishment at a number of warehouses or retail outlets belonging to the same company.

VMI differs from conventional inventory management in the following way. With conventional inventory management, the customers monitor their own inventory levels, and when a customer thinks that it is time to reorder, an order for a quantity of the product is placed at the supplier. The supplier receives these orders from the customers, prepares the product for delivery, and makes deliveries using the fleet of vehicles. Conventional inventory management has several disadvantages. It is typical for orders not to arrive uniformly over time, and as a result the supplier's resources, such as the production and storage facilities, as well as transportation resources, cannot be utilized well over time. Also, during times of high demand, the supplier's resources may not be sufficient to satisfy all the customers' demands, with the result that some customers may experience shortages.

With VMI, the supplier monitors the inventory at the customers. This is made possible with modern measurement and communication equipment. The supplier is responsible for maintaining a sufficient inventory level at each customer, and decides which customers should be replenished at which times, and with how much product. With VMI, the supplier has the benefit of more information than with conventional inventory management, such as the current (and possibly past) inventory levels at all the customers, and the customers' usage behavior. It is thus not surprising that VMI has several advantages for the supplier over conventional inventory management. First, VMI may lead to reduced production and inventory costs through a more uniform utilization of resources. Second, VMI may reduce transportation costs, since by proactive planning based on the additional available information instead of reactive response to customers' orders as they arrive, it may be possible to increase the frequency of low-cost full truckload shipments and decrease the frequency of high-cost less-than-truckload shipments. Furthermore, it may be possible to use more efficient routes by coordination of the replenishment at different customers close to each other. Advantages of VMI for customers are more reliable product availability, and the fact that customers have to devote fewer resources to monitoring their inventory levels and placing orders than under conventional inventory management.

The inventory routing problem (IRP) addresses the coordination of inventory replenishment and transportation. Specifically, we study the problem of determining strategies for the distribution of a single product from a single supplier to multiple customers. For this purpose, the supplier operates a fleet of vehicles. It is assumed that product usage at the customers is known (or that the probability distributions of customer usage are known). The supplier makes decisions regarding which customers' inventories to replenish, how much to deliver at each customer, how to combine customers into vehicle routes, and which vehicle routes to assign to each of the vehicles. Such a decision is called an itinerary. The set of feasible itineraries is determined by constraints on the travel times and work hours of vehicles and drivers, delivery time windows at the customers, the storage capacities and current inventory levels of customers, and other constraints dictated by the application. It may be feasible for a vehicle to perform more than one route per day. The objective is to maximize the profit over a given planning horizon (or the expected discounted profit over an infinite horizon), incorporating sales revenues, production costs, transportation costs, inventory holding costs, and shortage penalties.

Our work on this problem is motivated by our collaboration with a producer and distributor of air products. The company operates several plants and produces a variety of products, such as liquid nitrogen and oxygen. The company's bulk customers have their own storage tanks at their sites, which are replenished by tanker trucks under the company's control. Most of the bulk customers participate in the company's VMI program. The inventory levels at the bulk customers are measured by remote telemetry units. Such a device measures the quantity of the product in the storage tank, and is connected through a modem and the telephone network to the company's computer. A telemetry unit can be set to periodically measure the inventory level and send the information to the company's computer, and the computer can also query the telemetry unit at any time, so that the decision maker can obtain inventory information whenever needed. For the most part each customer and each vehicle is allocated to a specific plant, so that the overall problem decomposes according to individual plants. Also, to improve safety and reduce contamination, each vehicle and each storage tank at a customer is dedicated to a particular type of product. Hence the problem also decomposes according to type of product. It seems that the most questionable assumptions are that vehicles and drivers are available at the beginning of each day, mostly because of the notorious unpredictability of driver availability, and that the probability distributions of the customers' demands are known to the supplier and do not change over time. In practice, these probability distributions have to be estimated from data, and of course the probability distributions change over time. Fortunately, in this particular case, a large amount of data is available, and the demand characteristics of consumers do not seem to change rapidly over time. (However, there are significant differences between demand on weekdays and weekends.)

A two-phase approach has been developed for em the deterministic variant of the IRP. Phase I considers a coarse approximation of the problem, with daily decisions, over a k-day planning horizon. Phase II considers a model with decision accuracy in terms of minutes rather than days for the first j days.

Central to the Phase I integer program are two parameters: $LL_i^t = \max(0, tU_i - I_i^0)$, a lower bound on the total volume that has to be delivered to customer i by the end of day t, and $UL_i^t = tU_i + C_i - I_i^0$, an upper bound on the total volume that can be delivered to customer i by the end of day t. Let d_i^t represent the delivery volume to customer i on day t. To ensure that no stockout occurs at customer i and to ensure that we do not exceed the inventory capacity at customer i, we need to require

$$LL_i^t \leq \sum_{1 \leq s \leq t} d_i^s \leq UL_i^t \quad \forall i \quad \forall t. \tag{1}$$

To model resource constraints with some degree of accuracy and to have a meaningful objective function, we found it necessary to introduce the notion of "delivery routes" even in Phase I. However, it should be noted that when we refer to a "route" in the integer program, we will be referring only to a set of customers without enforcing a specific ordering among them. We estimate the distance required to visit the customers in the set to be the length of the optimal traveling salesman tour through the customers. Now, let R be the set of delivery routes, let T_r denote the duration of route r (as a fraction of a day), and let c_r be the cost of executing route r (which we assume depends on the duration of the route r). Furthermore, let x_r^t be a 0-1 variable indicating whether route r is used on day t ($x_r^t = 1$) or not ($x_r^t = 0$) and d_{ir}^t be a continuous variable representing the amount of product delivered to customer i on route r on day t (replacing the use of the variable d_i^t). The total volume that can be delivered on a single day is limited by a combination of capacity and time constraints. The resource constraints can be modeled by (recalling that our fleet has M vehicles of capacity Q)

$$\sum_{i: i \in r} d_{ir}^t \leq Q x_r^t \quad \forall r \in R \quad \forall t \tag{2}$$

and

$$\sum_{r: r \in R} T_r x_r^t \leq M \quad \forall t. \tag{3}$$

These constraints ensure that we do not exceed the vehicle capacity on any of the selected routes and that the time required to execute the selected routes does not exceed the time available.

The high-level base plan produced in Phase I does not specify departure times and delivery sequences for the different vehicles. We still need to construct vehicle routes and schedules. The delivery quantities and times specified by the solution to the Phase I integer program are good from a long-term perspective; they may not be as good from a short-term perspective. Therefore, we treat the delivery sizes and times specified by the solution to the integer programs as suggestions. We follow these suggestions as closely as possible, since this helps to achieve our long-term goals, but some changes are often needed. Since some constraints common in practice cannot be enforced at the daily level in the Phase I integer program, the routes chosen may either be infeasible when creating the detailed schedule for the jth day or not as efficient as other potential schedules. Also in constructing the delivery schedule, we may be able to deliver more to the customers than the quantity specified by the Phase I integer program. Since this creates fuller truckload routes (and thus more revenue to the vendor) and does not lead to increased long term costs, we want to take advantage of these opportunities where possible. To be more precise, we construct vehicle routes and schedules for j consecutive days, where we force the total volume delivered to a customer over the j days to be greater than or equal to the total delivery volume specified by the solution to the integer program for these days, but we do not enforce specific delivery quantities to occur on specific days. In this way, we should stay close to the delivery quantities suggested by the Phase I integer program, which are good from a long-term perspective, but we introduce some flexibility in the daily routing and scheduling, which is good from a short-term perspective. We have developed and implemented an insertion heuristic for the j-day routing and scheduling problem that takes advantage of the inherent flexibility of the problem.

We formulate *the stochastic variant of the IRP* as a discrete time Markov decision process (MDP) with the following components:

1. The state x is the current inventory at each customer, and X_t denotes the state at time t.

2. The action space $\mathcal{A}(x)$ for each state x is the set of all itineraries that satisfy the work load constraints, such that the vehicles' capacities are not exceeded, and the customers' storage capacities are not exceeded after deliveries. Let $A_t \in \mathcal{A}(X_t)$ denote the itinerary chosen at time t.

3. The known joint probability distribution F of customer demands gives a known Markov transition function Q, according to which transitions occur.

4. A reward of $r_n(d)$ is earned for delivery of quantity d at customer n. A transportation cost of c_{ij} is incurred each time arc (i, j) is traversed by a vehicle. Inventory holding costs $h_n(x + d)$ is incurred if the inventory at customer n at the beginning of the time period is x and quantity d is delivered at customer n. A penalty of $p_n(s)$ is incurred when customer n has a shortage of s during the time period. Let $g(x, a)$ denote the expected single stage net reward if the process is in state x at time t, and itinerary $a \in \mathcal{A}(x)$ is implemented.

5. The objective is to maximize the expected total discounted value over an infinite horizon. Let $\alpha \in [0, 1)$ denote the discount factor. Let $V^*(x)$ denote the optimal expected value given that the initial state is x, i.e.,

$$V^*(x) \equiv \sup_{\{A_t\}_{t=0}^{\infty}} E\left[\sum_{t=0}^{\infty} \alpha^t g(X_t, A_t) \middle| X_0 = x\right] \tag{4}$$

Solving the MDP exactly involves solving the following optimality equation, for each state $x \in \mathcal{X}$:

$$V^*(x) = \sup_{a \in \mathcal{A}(x)} \left\{ g(x, a) + \alpha \int_{\mathcal{X}} V^*(y)\, Q[dy \mid x, a] \right\} \tag{5}$$

This requires the following major computational tasks to be performed.

1. Estimation of the optimal value function V^*. Because V^* appears in the left hand side and right hand side of (5), most algorithms for computing V^* involves the computation of successive approximations to $V^*(x)$ for every $x \in \mathcal{X}$. Clearly, this is practical only if the state space $\mathcal{X}$ is small. In the case of the IRP, the size of the state space is exponential in the number of customers. Thus, even if the inventory levels are discretized, the state space $\mathcal{X}$ is far too large to compute $V^*(x)$ for every $x \in \mathcal{X}$ if there are more than about four customers. Our approach is to approximate $V^*(x)$ by a function $\hat{V}(x)$ that is much easier to compute. The approximation method is briefly discussed below.

2. Estimation of the expected value (integral) in (5). For many applications, this is a high dimensional integral, which requires a lot of computational effort to compute accurately. In the case of the IRP, the number of dimensions is equal to the number of customers, which can be as much as several hundred. Conventional numerical integration methods are not practical for the computation of such high dimensional integrals. We estimate the expected value with a sample average, and use common random numbers and orthogonal arrays as variance reduction techniques to improve the accuracy of the estimates.

3. The maximization problem on the right hand side of (5) has to be solved to determine the optimal action for each state. This optimization problem is very hard, because the NP-hard vehicle routing problem is a special case. Besides choosing which customers to visit and the vehicle routes, maximization problem (5) also has to choose a delivery quantity at each visited customer, taking into account not only travel costs, but also expected shortage costs, holding costs, and the values of future states. We propose a construction heuristic for solving this problem, and present computational results.

The approximating function $\hat{V}$ of the optimal value function V^* is constructed as follows. A number of subsets of customers are chosen. These subsets cover the set of all customers, and may overlap. All singletons are included in the chosen subsets. One guideline for choosing the subsets is that a single vehicle should be able to visit all customers in the subset in one day. An MDP is defined for each subset, that mimics the overall process as experienced by the subset of customers. These MDPs are small enough to be solved exactly using standard techniques. For any given state $x \in \mathcal{X}$ of the overall process, the approximating value $\hat{V}(x)$ is the optimal value of an optimization problem, that selects a partition among the subsets of customers to combine the optimal values of the subset MDPs. In general, this optimization problem can be formulated as an integer program. With direct deliveries, the problem reduces to a nonlinear knapsack problem, which is easy to solve with dynamic programming. With at most two customers per subset, the problem is approximated with a matching problem.

The approximation method has been implemented, and initial results are promising.

References

[1] A. Kleywegt, V. Nori, and M.W.P. Savelsbergh (2002). Dynamic Programming Approximations for a Stochastic Inventory Routing Problem. Transportation Science, to appear.

[2] A. Kleywegt, V. Nori and M.W.P. Savelsbergh. The Stochastic Inventory Routing Problem with Direct Deliveries. Transportation Science 36 (2002) 94-118.

[3] A. Campbell and M.W.P. Savelsbergh (2002). A Decomposition Approach for the Inventory Routing Problem. Transportation Science, to appear.

[4] A. Campbell and M.W.P. Savelsbergh (2002). Efficiently Handling Practical Complexities in Insertion Heuristics. Transportation Science, to appear.

[5] A. Campbell and M.W.P. Savelsbergh (2001). Delivery Volume Optimization. Transportation Science, to appear.

[6] A. Campbell, L. Clarke and M.W.P. Savelsbergh. Inventory Routing in Practice. The Vehicle Routing Problem (P. Toth, D. Viego, Eds.), SIAM monographs on discrete mathematics and applications (2002) 309-330.

[7] A. Campbell, L. Clarke, A. Kleywegt and M.W.P. Savelsbergh. The Inventory Routing Problem. T. Crainic and G. Laporte (eds.). Fleet Management and Logistics. Kluwer Academic Publishers (1997) 95-113.