

An Inventory Routing Problem

A. Campbell
L. Clarke
M. Savelsbergh

April 28, 1999

1 Introduction

PRAXAIR is a large industrial gases company with about 60 production facilities and over 10,000 customers across North America. PRAXAIR recently negotiated a policy with its customers in which PRAXAIR is in charge of managing its customers' inventories. Customers will no longer be calling PRAXAIR to request a delivery. Instead, PRAXAIR will determine who receives a delivery each day and what the size of that delivery will be. PRAXAIR will use gauge readings received from remote telemetry units as well as regular customer phone calls to monitor and forecast product inventories. The distribution planning problems associated with such vendor managed resupply policies are known as inventory routing problems.

Inventory routing problems are very different from vehicle routing problems. Vehicle routing problems occur when customers place orders and the delivery company, on any given day, assigns the orders for that day to routes for trucks. In inventory routing problems, the delivery company, not the customer, decides how much to deliver to which customers each day. There are no customer orders. Instead, the delivery company operates under the restriction that its customers are not allowed to run out of product. Another difference is the planning horizon. Vehicle routing problems typically deal with a single day, with the only requirement being that all orders have to be delivered by the end of the day. Inventory routing problems deal with a longer horizon. Each day the delivery company makes decisions about which customers to visit and how much to deliver to each of them, while keeping in mind that decisions made today impact what has to be done in the future. The objective is to minimize the total cost over the planning horizon while making sure no customers run out of product. The flexibility to decide when customers receive a delivery and how large these deliveries will be may significantly reduce distribution costs. However, this flexibility also makes it very difficult to determine a good, much less an optimal, cost effective distribution plan. When the choice becomes which of the customers to serve each day (PRAXAIR has over 10,000 customers) and how much to deliver to them, the choices become virtually endless.

Vendor managed resupply policies can be used in many situations. In some instances, the use of such a policy is natural, such as when the “customers” are really part of the same company. In others, the use of a vendor managed resupply policy is often the result of lengthy negotiations with customers who have for years followed a policy in which they call in their orders. Examples of industries where vendor managed resupply policies are being used or considered include, the petrochemical industry (gas stations), the grocery industry (supermarkets), the soft drink industry (vending machines), and the automotive industry (parts distribution). The number of industries using vendor managed resupply policies is increasing rapidly. An important reason for this is technology. For a variety of industries/products, the monitoring technology that existed several years ago was not sophisticated enough to make a vendor managed resupply system possible. The only way to check a customer’s inventory for many types of products has been for the vendor to call the customer and for the customer to go look at the meter on the tank, to count the number of items in the vending machine, etc. Now the use of remote telemetry units, scanners, computers and modems allows the monitoring of inventory levels directly by the vendor, opening up new opportunities for vendor managed resupply policies.

The remainder of the chapter is organized as follows. In Section 2, we formally introduce the inventory routing problem. In Section 3, we discuss the two-phase approach we have chosen to solve instances of the IRP. In Section 6, we present the results of some computational experiments.

2 Problem Definition

The inventory routing problem (IRP) is concerned with the repeated distribution of a single product from a single facility, to a set N of customers over a planning horizon of length T , possibly infinity. Customer i consumes the product at a rate u_i (volume per day) and has the capability to maintain a local inventory of the product up to a maximum of C_i . The inventory at customer i is I_i^0 at time 0. A fleet M of homogeneous vehicles, with capacity Q , is available for the distribution of the product. The objective is to minimize the average daily distribution cost during the planning period without causing stockouts at any of the customers. Vehicles are allowed to make multiple trips per day. Three decisions have to be made:

1. When to serve a customer?
2. How much to deliver to a customer when served?
3. Which delivery routes to use?

Real-life inventory routing problems are obviously stochastic. No customer will use product the same way every single day. In many situations, however, usage is relatively predictable and customers generally use about the same amount each day if we look at their total usage for several days in a row. Therefore,

solution approaches developed for the IRP as defined above will still provide useful planning tools.

3 Solution Approach

Even though the IRP is a long-term problem, almost all proposed solution approaches solve only a short-term version of the problem to make it easier. In early work, short-term was often just a single day, but in later work this was expanded to several days. Besides the number of days modeled, key features that distinguish different solution approaches include how the long-term effect of short-term decisions is modeled, and how it is determined which customers are included in the short-term problem. For a review of some of these approaches, we refer to Campbell et al. [2].

A short-term approach has the tendency to defer as many deliveries as possible to the next planning period, which may lead to an undesirable situation in the next planning period. Therefore, the proper projection of a long-term objective into a short-term planning problem is essential. It needs to capture the costs and benefits of delivering to a customer earlier than necessary. Our focus has been on developing a flexible system capable of handling large instances that properly balances short-term and long-term goals and that considers all the key factors, i.e., geography, inventory, capacity, and usage rate. We wanted also to create a system that would consider routing customers together on a day where none of them are at the point of run out, but where they combine to make a good full truckload delivery. We found that most systems reduce the problem by starting with only the “emergency” customers, never putting together certain combinations that make sense with regard to location and delivery size. The basis for our system is a two-phase solution approach. In the first phase, we determine which customers receive a delivery on each day of the planning period and decide on the size of the deliveries. In the second phase, we determine the actual delivery routes and schedules for each of the days.

As mentioned earlier, real-life inventory routing problems are stochastic. Therefore, any distribution plan covering more than a couple of days will never be executed completely as planned. Actual volumes delivered differ from planned volumes because usage rates deviate from their forecasts, planned driving time is off due to traffic congestion, etc.. Therefore, any planning system needs to be flexible. It needs to take advantage of the latest changes in the data. Given this, our approach is to embed our two-phase solution approach in a rolling horizon framework. We always construct a distribution plan for a month to reflect the long-term nature of the planning problem, but we expect to implement only the first few days. We repeat this as often as necessary using the latest information available.

3.1 Phase I: An Integer Programming Model

At the heart of the first phase is an integer program. Central to the model are two quantities: $L_i^t = \max(0, tu_i - I_i^0)$, a lower bound on the total volume that has to be delivered to customer i by the end of day t , and $U_i^t = tu_i + C_i - I_i^0$, 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 , then 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 have that

$$L_i^t \leq \sum_{1 \leq s \leq t} d_i^s \leq U_i^t \quad \forall i \in N, t = 1, \dots, T.$$

To model the resource constraints with some degree of accuracy and to have a meaningful objective function, we found it necessary to explicitly use delivery routes. However, when we refer to a “route”, we are really referring to a set of customers without enforcing a specific ordering among the customers in the set. We estimate the distance required to visit the customers in the set by the length of the optimal traveling salesman tour through all 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 . 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 ($x_r^t = 0$). The total volume that can be delivered on a single day is limited by a combination of capacity and time constraints. Since vehicles are allowed to make multiple trips per day, we cannot simply limit the total volume delivered on a given day to be the sum of the vehicle capacities. To be more precise, the resource constraints can be modeled by

$$\sum_{i:i \in r} d_{ir}^t \leq Qx_r^t \quad \forall r \in R, t = 1, \dots, T,$$

and

$$\sum_{r:r \in R} T_r x_r^t \leq |M| \quad t = 1, \dots, T.$$

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 basic Phase I integer programming model is given by

$$\min \sum_t \sum_r c_r x_r^t$$

$$L_i^t \leq \sum_{1 \leq s \leq t} d_i^s \leq U_i^t \quad \forall i \in N, t = 1, \dots, T,$$

$$\sum_{i:i \in r} d_{ir}^t \leq Qx_r^t \quad \forall r \in R, t = 1, \dots, T,$$

$$\sum_{r:r \in R} T_r x_r^t \leq |M| \quad t = 1, \dots, T.$$

The first variation of the basic model handles fixed and variable stop times at the customers as well as a vehicle reloading time at the facility. The duration of a route T_r can be modified to include not only the estimated time to drive the distance between the customers on the route, but also a fixed stop time for each customer and an initial fill time for the vehicle required before the route can start. Dispense time at a customer clearly cannot be included in T_r *a priori* because it depends on the size of the delivery. Therefore, we must alter the resource constraint as follows, where F is the percentage of the day required to dispense each unit of product

$$\sum_{r:r \in R} (T_r x_r^t + \sum_{i:i \in r} F d_{ir}^t) \leq |M| \quad t = 1, \dots, T.$$

The second variation handles operating modes of customers. Operating mode refers to the start and end time of customer usage on each day of the week. Before, we assumed that each customer i uses product 24 hours per day everyday. Operating modes are important. When a customer does not use product on the weekend, for example, it has a big impact on properly timing the deliveries. Operating modes can be handled easily by appropriately modifying the lower and upper bound parameters. The values of the lower and upper bounds on day t now depend on where in the week days 1 through t fall.

The third variant handles time windows at customers. An operating mode restricts when a customer uses product. A time window restricts when a customer can receive a delivery. Time windows may be day dependent as well. To handle time windows, the lower and upper bound parameters need to be modified again, but in a slightly different way. Now the lower bound L_i^t needs to be defined as the total volume that has to be delivered to customer i by the closing of the time window on day t to allow customer i to last until the opening of the time window on day $t + 1$ (or the opening of the time window on the first available day for the next delivery if no deliveries can be made on day $t + 1$). The upper bound U_i^t is now defined as the largest volume that customer i can receive by the close of the delivery window on day t .

3.2 Phase I: Solving the Integer Programming Model

The integer programming model presented above is not very practical for two reasons: the huge number of possible delivery routes and, although to a lesser extent, the length of the planning horizon. To make the integer program computationally tractable we consider a small (but good) set of delivery routes and aggregate periods toward the end of the planning horizon.

3.2.1 Clusters

Our approach to reduce the number of routes is based on allowing customers to be on a route together only if they are in the same *cluster*. A cluster is a group of customers that can be served cost effectively by a single vehicle for a long period of time. The cost of a cluster is an approximation of the distribution cost for serving the customers in the cluster for a month. The cost of serving a cluster does not only depend on the geographic locations of the customers in the cluster, but also on whether the customers in the cluster have compatible inventory capacities and usage rates. Therefore, to evaluate the cost of a cluster, we need a model that considers all of these factors.

The following approach is used to identify a good set of disjoint clusters covering all customers:

1. Generate a large set of possible clusters.
2. Estimate the cost of serving each cluster.
3. Solve a set partitioning problem to select clusters.

Observe that the selection of clusters only has to be done once as a preprocessing step before the actual planning starts. It does not have to be re-run before every execution of the Phase I integer program. In practice it makes sense to re-cluster when new customers have been added or there have been significant changes to the data.

Since we generate a large number of clusters to choose from, we need a costing procedure that is fast, but able to provide an accurate estimate of the cost of serving the cluster. We decided to use a simple integer program with key features represented.

3.2.2 Aggregation/Relaxation

Given that our two-phase solution approach will be embedded in a rolling horizon framework, the emphasis should be on the quality and detail of the decisions concerning the first few days of the plan. This provides us with an excellent opportunity to reduce the size of the integer program by aggregating days towards the end of the planning period.

For the first k days, we will still have route selection variables for each day, but for the days after that, we will have route selection variables covering periods of several days. Instead of making a decision on whether to execute each route on days 8 to 14 individually, for example, we now decide how many times each of the routes will be executed during the whole week instead. Several aggregation schemes were tested. We found that considering weeks rather than days towards the end of the planning horizon still does a good job of preserving the costs associated with the effect of short-term decisions on the future and yields a significant reduction in CPU time. Therefore, the daily variables associated with these later days are replaced by weekly variables. Upper and lower bounds are altered accordingly as well.

A further simplification is obtained by relaxing the integrality restrictions on the variables representing the weekly decisions. Therefore, the only binary variables appearing in the integer program will be those representing route selections for the first k days.

3.3 Phase II: Scheduling

A solution to the integer program of Phase I specifies the volumes to deliver to each customer for the next k days. It does not specify departure times and customer sequences for the different vehicles. Therefore, we still need to construct vehicle routes and schedules.

Since the delivery volumes specified by the solution to the integer program may not fit before a specific time of the day and may need to be received before a certain later time to prevent run out, these deliveries have self imposed time windows. Therefore, to convert the information provided by the solution to the integer program to daily vehicle routes and schedules, we can solve a sequence of vehicle routing problems with time windows.

However, such an approach does not capitalize on the flexibility inherent in the inventory routing problem. The delivery volumes specified by the solution to the integer program are good from a long-term perspective; they may not be good from a short-term perspective. Therefore, we treat the delivery volumes and timing specified by the solution to the integer programs as suggestions. We try to follow these suggestions as closely as possible, since this helps to achieve our long-term goals, but we allow small deviations when it helps to construct better short-term plans. To be more precise, we construct vehicle routes and schedules for two consecutive days, where we force the total volume delivered to a customer over the two days to be greater than or equal to the total delivery volume specified by the solution to the integer program for these two days, but we do not enforce specific delivery volumes on individual days. In this way, we stay close to the delivery volumes suggested by the integer program, which is 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. Deliveries can be split into smaller pieces, delivering one part on the first day and the second part on the second day if this works out to be better, for example when resources are very tight on one of the days. This flexibility is even more important when we consider the fact that in practice a few customers may not follow a vendor managed resupply policy and may call in orders that need to be added to the daily routing and scheduling problem. With new orders and new accurate up-to-date information on customer inventory levels, it may make sense to shift around some of the deliveries over the next couple of days.

Because of customer usage and customer inventory capacities, there may be customers that require a delivery on both days or even multiple times a day. Consequently, in our two day routing and scheduling problem, we can distinguish two types of customers: customers that require multiple deliveries over the two days and customers that require only one.

We have developed and implemented an insertion heuristic for this two day

routing and scheduling problem. The heuristic is a logical progression of commonly used techniques in insertion heuristics for the vehicle routing problem with time windows, see for example Solomon [5] and Kindervater and Savelsbergh [4].

In the description of the heuristic, we assume, for ease of presentation, that there are no operating modes and no time windows restricting when deliveries can take place. Both complications can easily be handled. We also do not discuss explicitly the use of fixed stop times and unloading times, though both can be included in the travel time value used here.

The flexibility to change delivery volumes makes checking the feasibility of insertions much more complex than in the VRP. For example, the insertion of a customer on a route can affect the delivery volume of another customer on an earlier or later route for the same vehicle, which can affect the size and timing of other deliveries for the customers on that route, etc.

To be able to evaluate the feasibility of an insertion, we maintain several quantities related to deliveries to customers already scheduled. Consider a delivery to customer i on route r . The predecessor on the route is denoted by $p(i)$ and the successor on the route is denoted by $s(i)$. The total volume to be delivered to customer i over the two days prescribed by the solution to the Phase I integer program is d_i . We consider a day as ranging from time 0 to 1 for convenience. There is a slight difference for customers that need multiple deliveries over the two days, but the basic quantities we maintain are the following:

1. The minimum delivery volume q_{ri}^{min} .

$$q_{ri}^{min} = d_i.$$

2. The earliest time a delivery can be made t_{ri}^{early}

$$t_{ri}^{early} = \max(t_{rp(i)}^{early} + tt_{p(i),i}, (q_{ri}^{min} - C_i + I_i)/u_i)$$

where $t_{r0}^{early} = t_r^{earlystart}$, the earliest time the route can start, and $tt_{j,k}$ is the travel time from customer j to k . The first term of the maximum represents the time to get to customer i from $p(i)$ and the second term represents the time that the minimum delivery volume can fit at customer i .

3. The latest time a delivery can be made t_{ri}^{late}

$$t_{ri}^{late} = \min(t_{rs(i)}^{late} - tt_{i,s(i)}, I_i/u_i)$$

where $t_{r(n+1)}^{late} = t_r^{lateend}$, the latest time route r can end. The first term of the minimum represents the latest departure time from i to be able to reach $s(i)$ by the latest time for its delivery. The second term represents the time when customer i runs out of product.

4. The maximum delivery volume q_{ri}^{\max}

$$q_{ri}^{\max} = \min(Q - \sum_{j \neq i \in r} q_{rj}^{\min}, C_i, d_i, C_i - I_i + u_i t_{ri}^{\text{late}})$$

The first term of the minimum is the capacity remaining in the vehicle if we assume all other customers on the route will receive their minimum delivery volumes, the second and third terms are obvious, and the fourth term represents the volume that will fit at the latest time a delivery can be made.

Because vehicles can drive multiple routes per day we also maintain several quantities for each route.

1. The earliest time the route can start $t_r^{\text{earlystart}}$
2. The latest time the route can start $t_r^{\text{latestart}}$
3. The earliest time the route can end t_r^{earlyend}
4. The latest time a route can end t_r^{lateend}

Given these quantities the feasibility of an insertion is checked as follows. First, we check whether the minimum delivery volume fits in the vehicle given the other planned deliveries. Next, we compute the earliest time and the latest time a delivery can take place. If the earliest delivery time is greater than the latest delivery time, the insertion is infeasible. Using the latest delivery time, we compute the maximum delivery size. If it is smaller than the minimum delivery size, the insertion is infeasible. If the insertion passes both of these tests, it is feasible.

If an insertion is feasible, the cost of the insertion is evaluated. The cost of an insertion is a weighted sum of several components. The first component is the increase in distance and the second component is an approximation of the minimum increase in waiting time if the insertion is carried out. The third component is a charge for making routes “inflexible”. In the final two day plan, we like to have near capacity routes. Therefore, we want to discourage the construction of routes with a small difference $t_r^{\text{latestart}} - t_r^{\text{earlystart}}$ and a large difference $Q - \sum_{j \in r} q_{rj}^{\max}$, since it is unlikely that such routes can be extended to near capacity routes. A charge is incurred if the insertion forces a route to have a gap between earliest and latest starting time that is less than x minutes and a total maximum delivery volume that is less than $y\%$ of capacity. The charge is inversely related to the size of the gap.

For each delivery to a customer, we maintain the cheapest feasible insertion and the second cheapest feasible insertion, if it exists. Since we can always construct a feasible route with just a delivery to a single customer, there exists at least one feasible insertion.

All that remains to complete the description of the insertion heuristic is to specify how we select the deliveries to be inserted in each iteration. Note that we select deliveries rather than customers, because customers may require multiple deliveries over the two days. We use the following selection rule:

1. If there are deliveries that cannot be inserted into any existing route, then among those deliveries select the one with the most expensive route for itself.
2. If all deliveries can be inserted into at least one existing route, then select the one with the largest difference between the cost of its cheapest and second cheapest insertion.

The first part of the rule captures the idea that if there are deliveries that cannot be inserted in the current set of routes, we know that we have to create at least one more route, so we might as well do it now. The second part of the rule captures the idea of trying to insert a delivery well and before all of its good potential insertion points become infeasible.

These rules are first applied to the deliveries to customers that require multiple deliveries over the two days. The idea is that these deliveries will be the most difficult to schedule feasibly, so we need to handle these first. When all of these are scheduled, these same insertion rules are then applied to the remaining deliveries.

After a feasible schedule is created, we run one more heuristic, a delivery amount optimization routine, which finalizes the schedule. It reviews the current schedule, decides which of the customers will have their delivery amounts increased and by how much, and determines at what time within the final feasible time ranges the deliveries will take place.

The insertion heuristic described above is embedded into a greedy randomized adaptive search procedure (GRASP) [3]. A GRASP combines a greedy heuristic with randomization. Whenever the heuristic selects the next delivery to be inserted, it will pick randomly from the q best choices, where q is pre-specified. This allows the algorithm to make choices that do not seem to be the best at the time, but may provide better opportunities later. In a GRASP framework, the heuristic is executed many times and the best plan obtained is picked.

4 Computational Experience

In this section, we present the results of various computational experiments that demonstrate the viability and value of the presented approach and illustrate many of the complexities of inventory routing problems.

4.1 Instances

For our computational experiments we have used real-life data from two of PRAXAIR's production facilities. We have chosen these two production facil-




Figure 1: Facility A and its customers




Figure 2: Delivery frequencies of customers of facility A

ities because the characteristics of the set of customers they serve are quite different in terms of geography, tank capacities, and usage rates.

Production facility A serves 50 customers that are fairly spread out, covering a mostly rural area with some small clusters of customers near cities. The facility is located in the northwestern corner of the state, not in the center. (See Figure 1.) Customers are between 4 minutes to 4 and a half hours driving time from the facility, with an average of 3 hours. The average driving time between two customers is 2 hours and 10 minutes. Of the 50 customers, 72 % require less than 1 delivery per week, 16% require between 1 and 2, 8% require between 2 and 3, and 4% require between 3 and 6 deliveries per week. (See Figure 2.) With respect to tank capacities, 22% of the customers can receive a delivery of more than a truckload, but 58% cannot receive even half of a truckload. (See Figure 3; the heavy line indicates vehicle capacity.)

Production facility B serves 87 customers spread over a large geographic area in the northern United States. The customers are concentrated heavily in the




Figure 3: Tank capacities of customers of facility A




Figure 4: Facility B and its customers

middle of the area, where the facility is also located, and become less concentrated as the distance from the center increases. (See Figure 4.) Customers are between 6 minutes and 10 hours driving time from the facility, with an average of 2 hours and 20 minutes. The average driving time between two points is 3 hours and 40 minutes. In terms of usage, 90% need less than 1 delivery per week, 8% need between 1 and 2, and only 2% require more than 1 delivery per week. (See Figure 5.) Furthermore, 21% can receive a delivery of a truckload and 41 % cannot receive half of a truckload. (See Figure 6.)

Roughly 75% of the customers at both plants use product 24 hours a day, 7 days a week. Of the customers that are not constant users, many change how they use product depending on the day of the week. Most use product roughly the same way Monday through Friday, but often only 8 to 10 hours per day. The usage pattern usually changes on the weekend, with many of these customers not using product at all on Sundays and less than half of a weekday amount on Saturdays.

Other relevant information used in our computational experiments is that the time of a delivery is calculated as $.5 + (\text{vehicle pump rate})(\text{quantity delivered})$,




Figure 5: Delivery frequencies for customers of facility B




Figure 6: Tank capacities of customers of facility B

that it takes 1 hour to reload a vehicle at the facility before it can depart again on another route, that all vehicles drive at a speed of at 45 mph, and that deliveries can be made 24 hours a day.

Finally, the initial inventory for all customers was chosen randomly, with the restriction that the inventory level should be sufficient to last the customer until the first time a vehicle would be able to arrive at the customer to refill its tank.

4.2 Solution quality

A solution to the inventory routing problem for a given planning period specifies which vehicles are visiting which customers on each day of the planning period, in what order the deliveries are being made, and how much is delivered to each customer. However, even with all this information it is still nontrivial to evaluate the quality of the solution. Since the IRP is really an infinite horizon problem, we have only specified the first part of a solution. For example, if we consider a planning period of two weeks, as we will do in our computational experiments, it is not obvious how to compare two solutions and claim that one is better than the other. If the total distance traveled in one solution is less than in the other solution, this represents a smaller driving costs. However, if, in the solution

with a higher total distance traveled, only full truckload deliveries are made, how can we say this solution is worse? It utilizes the trucks extremely well and may end in a state that is a much better starting point for the deliveries that have to be made in the following weeks.

Therefore, in addition to looking at the obvious statistics, such as the number of trips, the number of stops, the total volume delivered, and the total distance traveled, we also look at several other statistics to evaluate the quality of a solution for a two week planning period. Some of these statistics are used by PRAXAIR to evaluate their own performance, others are proposed in the literature, and some we just found to be interesting.

A popular statistic used in industry is *average volume per mile*. This statistic averages the volume per mile of all the trips, where the volume per mile of a trip is what we expect it to be, namely the total volume delivered on a trip divided by the total distance traveled on the trip. It is easy to see that this number is very sensitive to the distance of customers to the facility and therefore does not seem to provide reliable information in an averaged form. For example, if we consider a trip to a customer 4 miles away from the facility where a full truckload is delivered, the volume per mile is equal to truckload divided by 8. If we consider another trip to a customer 40 miles away from the facility where we also deliver a full truckload, the volume per mile is equal to truckload divided by 80. The average of these two volumes per mile ($\frac{11}{160}$ truckload) does not provide much information.

A more sensible statistic, especially over a period of several days, is *total volume per mile*, defined as the total volume delivered to all customers over the time period considered divided by the total distance traveled over this time period. Since we are looking only at the first piece of a long-term problem, it makes sense, in this first piece, to deliver more product than required to ensure that customers will not run out, if it can be done at a relative small cost, i.e., a small increase in distance traveled. A large value of total volume per mile indicates that you are successful at doing so.

A third statistic, inspired by the “weighted delivery radius” discussed in [1], is *weighted volume per mile*. The weighted volume per mile for a trip with n customers is computed as:

$$\frac{d_1 \cdot tt_{0,1} + d_2 \cdot tt_{0,2} + \dots + d_n \cdot tt_{0,n}}{\text{total round trip distance}}$$

The intuition behind this statistic is revealed when we look at the values it gives for the example given above. The value it gives for both trips is 0.5 truckload (which is the largest value possible). It says that both trips are equally good, in fact as good as possible, since the best you can do when serving a customer for a long period of time is to deliver full truckloads. When a trip contains several stops to deliver a full truckload or when a trip does not deliver a full truckload, the value of this statistic will go down. The other benefit of this statistic is that it still provides relevant information when it is averaged over a number of trips (assuming all vehicle capacities are equal).

Other statistics that are also important to consider include the *average inventory level before delivery* and *average inventory level after delivery*. Obviously, higher values are preferred, especially with respect to the average inventory level before delivery, because a high value indicates that you are less likely to experience stockouts due to fluctuations in usage rates. Furthermore, the average vehicle utilization is an interesting statistic. You like this value to be high, but not at all costs. You do not want to drive many extra miles just to ensure a high vehicle utilization. (In practice, there is a strong belief that every vehicle should leave the facility fully loaded and return empty. In part though, this is motivated by the inherent stochasticity which sometimes allows for larger than expected deliveries.)

Finally, we may also want to look at the number of vehicles used. We do not want to put too much importance on this statistic in our experiments though. In the long-term, eliminating a vehicle represents significant savings for a company, but in the short-term, we cannot really argue that one solution is really better than another just because it uses one less vehicle.

In our tests, we used the number of vehicles used in practice as the maximum number of vehicles available. We operated under the assumption that the number used in practice was necessary (long-term) and that therefore minimizing this number (short-term) does not make sense. If everything else is equal, though, this may be used as criterion for preferring one solution over another.

4.3 Alternate heuristic

To be able to compare the quality of the solutions produced by our proposed approach to current industrial practices, we implemented a solution approach based on the rules-of-thumb and ideas most commonly used in practice. After many discussion with the planners at PRAXAIR, we identified the following rules:

1. Create trips around customers that must receive a delivery on the day being considered. Fill up that customer to capacity and then add nearby customers to the trip if there is remaining inventory in the vehicle.
2. All customers on a trip are filled up to capacity except for the last one.
3. Discourage a vehicle from returning to the facility without delivering its remaining capacity to some customer.
4. Do not create trips involving only customers that do not require a delivery on the day being considered unless there are excess resources that day, and it appears that there will be insufficient resources when the first of these customers requires delivery.

4.4 Computational experiments

The first experiment compares the solutions obtained by our proposed optimization based approach (O) to the solutions obtained by the industry approach (I).

The results can be found in Table 1. In this table, and in all other tables, the top part contains the results for facility A, and the bottom part contains the results for facility B. We present the number of vehicles used (V), the number of trips made (T), the number of stops made (S), the vehicle utilization (U), the total volume delivered (Vol), the total distance traveled (Mile), the volume per mile (V/M), the average volume per mile (aV/M), the average weighted volume per mile (wV/M), the average inventory before delivery (Bef), and the average inventory after delivery.

Table 1: Comparison of optimization approach and industry approach

Approach	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
O	3	65	118	95.72	2613027	18841	138.69	410	18856	24.54	81.64
I	4	67	90	89.58	2519989	18988	132.71	309	18357	11.41	81.27
O	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49
I	3	61	93	85.41	28656473	15042	1905.10	7016	206276	9.83	88.97

Our approach clearly outperforms the industry approach. It does better for both facilities on all the important statistics. The difference in the underlying ideas of the two approaches is most clearly observed in the "average inventory before delivery" column. The industry approximation approach is driven by customers that are getting near to run out and that have to be visited, which results in a low "average inventory before delivery", whereas our approach looks further ahead and attempts to identify good opportunities to visit customers before they are getting near to run out.

As we indicated above, we believe that the strength of our approach is that it considers "enough" of the future to make the right decisions. In the next two experiments, we investigate the impact of varying the amount of future considered. In our chosen approach, we consider 5 days in full detail plus 4 weeks in aggregated form beyond this. We note that considering 5 days in full detail is already more than many of the solution approaches proposed in the literature. In Table 2, we show the results when we vary the amount of future considered in aggregated form.

Table 2: Value of the length of the planning horizon

Setting	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
5 days, 0 wk	4	69	119	92.43	2678599	20284	132.05	408	18114	29.87	87.76
5 days, 1 wk	4	67	117	95.04	2674527	19640	136.18	411	18687	25.26	82.18
5 days, 4 wk	3	65	118	95.72	2613027	18841	138.69	410	18856	24.54	81.64
5 days, 6 wk	4	74	125	88.10	2738189	20836	131.42	392	17514	24.26	80.74
5 days, 0 wk	3	64	105	84.58	29773123	14512	2051.62	7087	204340	20.30	92.33
5 days, 1 wk	3	60	102	89.02	29375272	14050	2090.77	7855	218711	18.03	92.16
5 days, 4 wk	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49
5 days, 6 wk	3	57	111	89.40	28026955	13789	2032.56	7496	214607	19.28	85.22

It is interesting to observe the increase in the number of deliveries when 6 weeks are considered. When 6 weeks are considered, a larger portion of the objective function value is representing future costs, and optimizing with this objective apparently allows us to make some unwise and expensive decisions in

the part of the planning period that really counts, i.e., the first five days. If we do not consider any of the future beyond the 5 days, we appear to be missing some beneficial opportunities. There is not much difference, though, between considering 1 week or 4 weeks beyond the 5 days.

Next, we decided to investigate the effect of considering fewer days in full detail. By reducing this number from 5, we make the IPs smaller and therefore easier to solve, but it is not clear what the impact will be on the solutions. In Table 3, we show the results when we vary the number of days considered in full detail.

Table 3: Value of the number of days of considered in full detail

Setting	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
2 days, 4 wk	*										
3 days, 4 wk	4	69	103	90.32	2617434	19856	131.82	283	18188	27.99	91.30
5 days, 4 wk	3	65	118	95.72	2613027	18841	138.69	410	18856	24.54	81.64
2 days, 4 wk	3	66	102	86.19	31285911	14027	2230.41	7298	214827	18.40	94.75
3 days, 4 wk	3	65	103	84.53	30220444	13172	2294.29	8357	219112	17.83	92.99
5 days, 4 wk	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49

As expected, for facility A the quality of the solutions decreases when we consider fewer days in full detail. In fact, when we just consider two days, we are unable to construct a solution in which none of the customers runs out of product during the planning period. In this case, the IP selects delivery amounts for customers that turn out to be impossible to schedule with the routing heuristic, because too many deliveries must occur on a specific day and roughly at the same time. On the other hand, for facility B the best solution is found when we consider only 3 days with full detail. It is interesting to observe that when we consider fewer days, the number of stops decreases. Apparently, when we consider more days in full detail, the IP starts looking for opportunities to make deliveries to customers that only require a delivery a few days out, whereas the IP is unable to do that when fewer days are considered in full detail.

Besides the amount of future considered, the quality of the solution is also impacted by the parameter settings used in the routing and scheduling heuristic and whether or not delivery amount optimization is active. When delivery amount optimization is not active, a delivery amount cannot be set above the amount specified by the IP. In Table 4 we present the results of our approach with (Y) and without (N) delivery optimization. Without delivery optimization, we expect the average vehicle utilization and the total volume to be less.

Table 4: Value of delivery amount optimization

Setting	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
N	3	69	126	87.56	2537554	19929	127.33	330	17077	24.90	78.00
Y	3	65	118	95.72	2613027	18441	138.69	410	18856	24.54	81.64
N	3	61	111	85.81	28789623	15274	1884.88	6612	202495	19.12	85.10
Y	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49

Looking at the summary statistics, the delivery amount optimization clearly

improves truck utilization and also leads to a significantly better total volume per mile and weighted volume per mile. The increase in total volume delivered was slightly less than 3% for facility A and slightly more than 5% for facility B.

Our default settings for the GRASP are to run the routing and scheduling heuristic 25 times and to select from among the 3 best choices. To investigate the impact of these settings as well as the importance of randomization, we conducted an experiment in which we executed the heuristic without any randomization (pure greedy) and with different settings for the number of replications. The results are presented in Table 5.

Table 5: Value of randomization

Setting	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
NO	4	71	121	87.82	2618675	21799	120.13	394	17092	23.54	79.87
5	3	67	116	93.70	2636760	19315	136.51	419	18592	24.76	82.77
25	3	65	118	95.72	2613027	18441	138.69	410	18856	24.54	81.64
50	4	65	121	93.22	2544803	19176	132.71	327	17793	23.74	77.21
NO	3	59	99	86.22	27979401	14378	1945.99	6990	200752	15.28	88.34
5	3	62	111	87.99	30004693	15094	1987.86	6773	201000	19.88	89.13
25	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49
50	3	63	111	87.57	30343074	14605	2077.58	7762	211049	19.88	89.61

Without randomization the solution has a noticeably low average vehicle utilization, total volume per mile, and average weighted volume per mile. On the other hand, going to 50 replications does not seem to improve over 25 replications, in fact it does slightly worse. This is possible because the replications are for two days of the schedule at a time. Which schedule is selected affects what deliveries are made, what the customer inventories are at the end of the two days, and thus the input for the next integer program that is solved.

To obtain more insight in the behavior of the GRASP, we kept track of the total distance traveled for all 50 replications at two different points in the two week planning period. The results are plotted in Figure 7.

The criterion used to pick the best solution out of the 25 produced by the GRASP is total travel distance. However, the results may be quite different if we decide to use average weighted volume per mile as the criterion to pick the best solution. Our last computational experiment relating to the GRASP compares the behavior based on different selection criteria. The results are presented in Table 6, where D indicates that schedules are selected based on mileage and WVM indicates that schedules are selected based on average weighted volume per mile.

Table 6: Comparing distance and weighted volume per mile as selection criterion

Setting	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
D	3	65	118	95.72	2613027	18441	138.69	410	18856	24.54	81.64
WVM	4	65	120	96.83	2643371	19341	136.67	423	19018	24.99	82.64
D	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49
WVM	3	58	105	91.34	29138802	14235	2046.98	8291	220091	20.50	91.76

There are various other parameters that can be set in the routing and





Figure 7: Effects of randomization

scheduling heuristic. Some of these help to construct solutions that reflect company policy. For example, in our default approach, we have not penalized waiting time at customers. In practice, though, waiting time is often strongly discouraged or not even allowed. To see the impact of discouraging waiting time on the quality of the solutions, Table 7 presents the solution statistics when we do not penalize waiting time (W) and when we do penalize waiting time (NW).

Table 7: Comparing allowing waiting and penalizing waiting at customers

Setting	V	T	S	U	Vol	Mile	V/M	aV/M	wV/M	Bef	Aft
W	3	65	118	95.72	2613027	18441	138.69	410	18856	24.54	81.64
NW	4	70	119	90.50	2660785	20019	132.91	390	17817	21.26	77.66
W	3	61	106	90.26	30283480	14226	2128.74	8019	215095	19.91	92.49
NW	3	58	103	88.46	28220108	14908	1892.95	6298	197806	18.06	88.17

As we expected, when we allow waiting at customers, we get a higher truck utilization, a better total volume per mile, and average weighted volume per mile. For facility B, the differences are quite significant.

5 Conclusion

We have presented an inventory routing problem and an optimization based approach for its solution. Extensive computational experiments indicate the value and potential of optimization based approaches for this type of complex routing and scheduling problem. The inventory routing problem is of special interest because it integrates two components of supply chain management: inventory control and vehicle routing. This type of integration is essential to improve overall system performance.

References

- [1] W. Bell, L. Dalberto, M. Fisher, A. Greenfield, R. Jaikumar, P. Kedia, R. Mack, and P. Prutzman. Improving the distribution of industrial gases with an on-line computerized routing and scheduling optimizer. *Interfaces*, 13(6):4–23, 1983.
- [2] A. Campbell, L. Clarke, A. Kleywegt, and M. Savelsbergh. Inventory routing. In T. Crainic and G. Laporte, editors, *Fleet Management and Logistics*. Kluwer Academic Publishers, 1998.
- [3] T.A. Feo and M.G.C. Resende. Greedy randomized adaptive search procedures. *Journal of Global Optimization*, 6:109–133, 1995.
- [4] G. A. P. Kindervater and M. W. P. Savelsbergh. Vehicle routing: Handling edge exchanges. In E. H. L. Aarts and J. K. Lenstra, editor, *Local Search in Combinatorial Optimization*, pages 337–360. Wiley, 1997.
- [5] Marius Solomon. Algorithms for the vehicle routing and scheduling problems with time window constraints. *Operations Research*, 35(2):254–265, 1987.