Chapter 10

ROUTING PROPANE DELIVERIES

Moshe Dror

Abstract This chapter is about solving the problem of propane deliveries. It is commonly viewed as a representative problem of a much larger family of hard problems of considerable practical significance. This problem has been on the "front burner" of the logistics academic and practitioners community for over twenty years. In this chapter I attempt to describe the practices of a propane distribution company and to summarize the literature on the more general topic of inventory routing. It is one person's point of view and I apologize *ex ante* for my unavoidable biases.

1. Introduction

The outline of this chapter is as follows:

- Personal experience, the basic problem and its variants,
- Real-life examples, starting with Bell et al. (1983),
- The initial analysis of Federgruen and Zipkin (1984),
- The propane distribution as in Dror (1983), and extensions,
- The new wave: Kleywegt et al. (2002, 2004) and Adelman (2003a,b, 2004),
- Summary

1.1 Personal experience

One day in the Fall of 1995, I had the opportunity to spend a day in the "passenger seat" of a propane delivery truck observing a propane delivery operation in a rural area of upstate New York 'first-hand.' More than 10 years before that, in the Spring of 1983, I completed a Ph.D. thesis on this very topic. However, I was only given the chance to observe a "real-life" delivery operation for one day in 1995. While writing my thesis, I did visit the headquarters of a large propane distribution company 3 – 4 times; I talked to it's managers, collected data from the operational offices of one district, but did not drive a delivery truck until Fall of 1995. Talking to the driver of the delivery truck for the whole day, visiting the different customers, and observing the operations first hand was a new experience. On that day we visited customers who had almost empty tanks and customers whose tanks were full, and a number of customers in between the two extremes. After leaving the depot, before starting on their routes, the propane delivery drivers gather (unofficially) at a local diner for a morning coffee and a short chat. They discuss their respective routes and the customers on their routes. Past experiences, road conditions, and advice are freely exchanged. Only after that early gathering do the drivers go ahead and deliver propane to their customers. I believe that this kind of first-hand viewing of an operation is very valuable for understanding and subsequent analysis of logistics operations. I learned a lot that day.

1.2 The basic problem of propane distribution

The basic propane truck delivery operations are usually conducted in rural areas. Propane (and similar liquid or gas products) is used to heat individual houses and other facilities, which are not connected by a delivery (rigid-pipe) network. In densely populated cities economies of scale dictate a different mode of operation via a connected delivery (pipeline) network which provides the commodity (propane) on demand to its customers just as electricity and water is delivered. Rural propane customers are dispersed in a certain geographical area serviceable from a central facility (a depot) located in their area and are serviced from this single depot by delivery trucks (special tank-trucks filled with propane). The delivery trucks are usually of a few (2 – 3) fixed capacity types. Each customer has a propane tank located on his property (next to his house). The tank usually belongs to the propane company servicing the customer on a long-term basis. The customer tanks come in a number (5 to 8) of different sizes. A propane service contract requires the company to maintain a sufficient level of propane in the customer's tank at any time. Once the propane is delivered to a customer, a payment invoice is issued requiring the customer to pay within a week or so. In essence, once the propane is delivered it belongs to the customer. In this setting the subsequent inventory holding costs for the propane are incurred by the customer. It is quite common to hear customers complain about the company filling their tanks to capacity just before summer — a period of very low propane consumption. Even in sparse rural areas there might be pockets of a few locally concentrated users connected by a rigid pipe network to a single propane storage facility. In this case, each customer is equipped with a metering device and billed periodically for his/her consumption. Thus, the propane in the tank incurs a holding cost absorbed by the company. From an academic perspective one might want to view propane distribution as what is called in business practice as *vendor managed inventory* replenishment (VMI). In this chapter we do not try to link propane distribution, modelling, and the literature, with the more general area of VMI.

In a medium size propane service district (2,000 – 5,000 customers, though 10,000 customer districts can also be found) a company uses (owns) 3 to 7 delivery trucks. Each morning, the truck drivers are given a list of customers (their location, tank sizes, names, if necessary special individual characteristics, etc.) and simple dispatching instructions and off they go on their delivery routes. At times, they are unknowingly assigned a customer whose tank is still full and subsequently a service visit is wasted. Presently, without the more prevalent information technology, the exact demand (tank emptiness) becomes known only when the driver checks the tank's propane level on his arrival. Truck dispatching is based on partial information, experience, and a demand forecast which are all used to generate daily dispatching lists. Each day (5 days in every regular week) a preselected subset of customers is replenished. The basic refill policy is to fill-up customers' tanks to their capacity on each replenishment (service) visit. Since the individual customers' consumption rates are only estimated (they are viewed as random variables), this gives rise to stock-out events which necessitate "emergency" deliveries, which represent propane deliveries in response to customer's request because of an empty tank. These emergency deliveries have to be performed 7 days a week, weekends and holidays included, and are quite costly. In one Pennsylvania district of slightly over 2,000 customers, from which initial data was collected for the 1983 study (Dror, 1983), there were about 100 stock-out related deliveries in a period of three months. From the propane company point of view, it is pertinent to operate efficiently with a long time view perspective. That is, the company would like to design the logistic operations which minimize its long run cost of delivering propane to a given set of customers (a district). It is quite common for such US companies to own (operate) 300 – 500 districts. In some larger districts, propane companies may own separate propane storage facilities called satellite depots (see Bard et al., 1998). These satellite depot facilities serve as intermittent refill points for the delivery trucks enabling the trucks to extend their delivery routes without returning to the home depot for refill. It is a quite complex routing setting with many parameters which are only partially known at the beginning of the work day and vary over time. Recently, electronic automatic measuring and reporting devices have been introduced into the propane tanks. These devices relay tank propane level information to the district office. Thus, wasteful visits to customers which do not require certain minimal volume propane delivery could be avoided. I do not know how common this automated reporting technology is. Even if this technology is presently quite common, it only helps in deciding who not to service on a given day but does not eliminate many of the difficulties inherent in efficient propane distribution planning. Key question: What does the district know about the status and evolution of a system (the customers, the delivery trucks, and the depots), and when do they know it?

2. The industrial gases inventory case

A motivating example from Adelman (2003), attributed to Bell et al. (1983), is very useful for introducing some basic difficulties regarding the inventory routing decisions. This example with deterministic daily demands is restated in Figure 10.1 (where the internode distances are in miles) and in the corresponding table below.

Customer	Tank Capacity	Daily Demand
A	5000 gallons	1000 gallons per day
B	3000	3000
C	2000	2000
D	4000	1500

Assuming vehicles of 5000 gallon capacity, a simple inventory routing solution is to replenish customers A and B together every day, and cus-

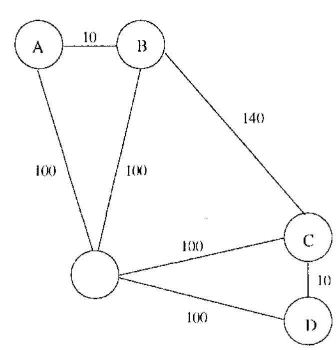

Figure 10.1. A simple example with 4 customers

tomers C and D together every day. The daily "cost" of this solution is 420 miles and uses two vehicles.

An improved routing solution consists of a cycling replenishment pattern which repeats every two days. On the first day use only one vehicle and deliver 3000 gallons to customer B and 2000 gallons to customer C, travelling 340 miles. On the second day two vehicles are used. Vehicle 1 delivers 2000 gallons to customer A and 3000 to customer B. Vehicle 2 delivers 2000 to customer C and 3000 to customer D. Each vehicle travels on that day 210 miles. Thus, the average daily distance travelled over two day period is 380 miles. This is 10% lower than the first solution. It is interesting to note that even though this solution has been known since 1983, Adelman (2003a) is the first to derive it analytically and prove its optimality!

The above example illustrates that finding an optimal replenishment solution even in a simple (deterministic) setting can be quite difficult. However, a number of successful distribution solutions have been developed over the years. One of the first success stories is that of industrial gases distribution systems developed by Marshall L. Fisher (Fisher et al., 1982) and his associates and is described below in some detail.

2.1 Industrial gases delivery system

The inventory management of industrial gases introduces real-life logistics issues very similar to the inventory management for propane deliveries. We repeat the operational description from Bell et al. (1983).

In the industrial gases case the main products are oxygen, nitrogen, hydrogen, argon, and carbon monoxide. Essentially, liquid oxygen and nitrogen are manufactured in highly automated plants. The plants serve as supply depots where liquified gases are stored at a temperature less than -320°F. The liquified gases are distributed in cryogenic bulk tankers to industrial users and hospitals. Storage tanks at customer sites are monitored by the supplier under long-term contracts. Similarly to propane, the supplier of liquid gases delivers the product at his discretion with the guarantee of continuous availability. In 1982, one company — the Air Products corporation, employed 340 trucks which travelled over 22 million miles a year. Distribution efficiency is the main competitive tool differentiating among the producers since the manufacturing costs among different companies are about the same and lower distribution costs allow lower pricing or higher profit margins. The decisions taken in distribution operations set the customers' tanks inventory levels, by determining how much to deliver, how to combine the different loads on a truck and how to route the truck. That is, inventory management at customer locations is integrated with vehicle routing. A single liquid gas plant distribution problem may involve several hundred customers and about 20 trucks. Complicating factors include estimating customer usage rates which vary considerably over time. Inventory must be maintained above a specified safety-stock level and customers are not open for delivery every day of the week or during every hour of the day (this varies across customers). Trucks also differ in their capacity and operating costs which might even change by state boundaries because of different state laws from one state to another.

There are numerous other driving cost related characteristics which need to be accounted for in real-life dispatching. We will not go into this here and the more specific details of the system are described in Bell et al. (1983). As we will see in the case of propane delivery, liquid gas delivery also requires careful forecasting of the rate at which each customer is consuming its product and the calculation of the "best" time to deliver, in terms of cost and delivery feasibility. What is usually known in terms of consumption rates is the inventory levels which are recorded before each delivery. In the case of liquid gases some customers are contacted (telephoned) from time to time to establish exact inventory levels to facilitate forecasting and dispatching. However, when deciding on vehicle routing sequences, it is comforting to note that feasible routes contain between two to four customers only. That is, even when dispatching 10 to 30 trucks daily, efficient routes are not very difficult to construct. In fact, the system described in Bell et al. (1983), is designed to produce a distribution schedule for the next two to five days. To select the delivery routes, first, a set of possible routes is generated with the sequencing order of customer stops. However, the delivery amount is not specified in the route generation stage. Since the number of customers on a route is small, "the number of technically feasible routes is not unreasonably large." A large mixed-integer programming model is solved each time which selects the routes from the set of externally generated potential routes, and determines the vehicle, the time each route starts, and the amount to be delivered to each customer on the route. We do not repeat here the mathematical formulation of the route selection model. However, we note that it incorporates parameters which represent the effect of short term delivery decisions on the events beyond the horizon of the model which is two to five days. Otherwise, a short term solution would "paint" the long term efficiency objective into a bad corner. One main difference between delivery of industrial liquid gases and propane is that in the propane case the policy is to fill the customer's tank to capacity on each service visit. In addition, delivery routes usually have between 4–12 customers making the route construction scheme more difficult.

In the next section we focus on one of the first academic attempts to model inventory distribution.

3. The initial analysis

The first, more "mathematical" analysis of the inventory routing problem is contained in the paper by Federgruen and Zipkin (1984). In that paper Federgruen and Zipkin examine:
"the combined problem of allocating scarce resources available at some central depot among several locations (or "customers"), each experiencing a random demand pattern, while deciding which deliveries are to be made by each set of vehicles and in what order."

It sounds very promising and is viewed as an extension of the standard vehicle routing problem where the *"deliveries serve to replenish the inventories to levels that appropriately balance inventory carrying and shortage costs, but thereby incur transportation costs as well."*

Essentially, the problem is examined from the point of view of inventory management in multiple locations with the added complication of routing — constructing delivery routes for a fleet of capacitated vehicles. The inventory status information for each location is assumed to be available at the beginning of the day and delivery quantities together with routes for each vehicle are then computed. The deliveries are executed and then the actual demand is observed with its resulting subsequent holding and shortage penalties. There is no requirement of visiting all the customers. As the authors state *"ours is the first attempt to integrate the allocation and routing problems into a single model."* The importance of such integration and analysis is very nicely motivated by Herron (1979). Federgruen and Zipkin (1984) present a very direct model which views each customer's inventory from the perspective of the newsvendor problem. That is, there are zero delivery costs to a customer and the deliveries to the customer(s) are driven by the shortage costs. For completeness we restate the mathematical formulation below slightly changing the original notation.

paragraphConstants

NV = number of vehicles
$\hat{n} =$ number of locations, with 0 indicating the depot location
$Q_v =$ capacity of vehicle v
$c_{ij} =$ cost of direct travel from location i to location j
$F_i(\cdot) =$ cumulative distribution function of the one period demand in location i, assumed strictly increasing
$h_i^+ =$ inventory carrying cost per unit in location i
$h_i^- =$ shortage cost per unit in location i

$I_i =$ initial inventory at location i
$A =$ total amount of product available at the central depot.

Variables.

$y_{ik} = 1$, if delivery point i is assigned to route (vehicle) k, and is 0 otherwise.

$x_{ijk} = 1$, if vehicle k travels directly from location i to location j, and is 0 otherwise.

q_i is the amount delivered to location i. Note that in the spirit of classical VRP formulations, at most one vehicle visits any given location.

Just like in the newsvendor inventory model, the inventory cost function $C_i(\cdot)$ and its derivative $C_i'(\cdot)$, for $i = 1, \ldots, \hat{n}$, are given by

$$C_i(q_i) = \int_{I_i+q_i}^{\infty} h_i^-(\xi - I_i - q_i)dF_i(\xi) + \int_0^{I_i+q_i} h_i^+(I_i + q_i - \xi)dF_i(\xi)$$
$$C_i'(q_i) = (h_i^+ + h_i^-)F_i(I_i + q_i) - h_i^-$$

Now the mathematical formulation expressing a single period cost minimization is stated as follows:

$$\min \sum_{i,j,k} c_{ij}x_{ijk} + \sum_i C_i(q_i) \tag{10.1}$$

subject to the following constraints

$$\sum_i q_i y_{ik} \leq Q_k, \qquad k = 1, \ldots, \text{NV}; \tag{10.2}$$

$$\sum_i q_i \leq A; \tag{10.3}$$

$$\sum_{k=1}^{\text{NV}} y_{0k} = \text{NV}; \tag{10.4}$$

$$\sum_{k=0}^{\text{NV}} y_{ik} = 1, \qquad i = 1, \ldots, \hat{n}; \tag{10.5}$$

$$\sum_i x_{ijk} = y_{jk}, \qquad j = 0, \ldots, \hat{n}; k = 1, \ldots, \text{NV}; \tag{10.6}$$

$$\sum_j x_{ijk} = y_{ik}, \qquad i = 0, \ldots, \hat{n}; k = 1, \ldots, \text{NV}; \tag{10.7}$$

$$\sum_{(i,j)\in S\times S} x_{ijk} \leq |S| - 1, \quad S \subseteq \{1, \ldots, \hat{n}\}, 2 \leq |S| \leq \hat{n} - 1, \quad k = 1, \ldots, \text{NV}; \tag{10.8}$$

$$x_{ijk} \in \{0.1\}, \qquad i, j = 0, \ldots, \hat{n}; k = 1, \ldots, \text{NV}; \tag{10.9}$$

$$y_{ik} \in \{0,1\}, \qquad i = 0, \ldots, \hat{n}; k = 1, \ldots, \text{NV}; \tag{10.10}$$

$$q_i \geq 0, \qquad i = 1, \ldots, \hat{n}. \tag{10.11}$$

This formulation has a mixture of 0-1 integer linear VRP type constraints and nonlinear constraints (10.2). It is a single period formulation which generates deliveries driven essentially by the expected shortage costs. For the propane delivery setting the formulation unnecessarily assumes limited supply at the depot (constraint (10.3)) but does not contain the tank capacity constraints at the customer locations. It also incorporates an inventory holding cost per unit per unit time — which is not directly applicable to the propane case. More importantly, it charges inventory shortage costs per unit per unit time. In the propane case the shortage costs might be best represented by some step function representing customer specific cost of emergency delivery. However this is a very nice model and for a fixed vector y it decomposes into simpler problems. On the one hand we get the inventory allocation problem, and on the other NV-TSPs, one for each route-vehicle. It is an attractive approach but unfortunately not appropriate for the propane delivery long term optimization problem. The primary reason is that this model is a single period optimization which does not project short-term decisions into long term cost implications. Thus, it might attempt to myopically "paint" a sequence of one period solutions into a long-term bad solution. We will return to this point later when we compare this model with the later model taken from Dror and Ball (1987).

4. The initial propane delivery model (Dror and Ball, 1987)

We first describe a number of very simple principles guiding the efficiency of propane deliveries.

- Visit a customer as infrequently as possible. This translates simply to delivering as much as possible to a given customer on each visit. In other words, if it is feasible to deliver as much as the customer's tank capacity on each visit, then do so.
- If it does not cost extra to visit a customer then replenish him/her. That is, if you can save a future service visit which has a positive cost by delivering early at no (or small) cost, then go ahead and replenish.
- Replenishing earlier reduces the risk of stock-out and increases the present value of cash-flow (see Dror and Trudeau, 1996).

To develop the above principles more formally we describe a basic analysis of a single customer with a fixed sized tank (size T), a cost of refill b (b_i for customer i), daily consumption rate μ (deterministic for now) and initial inventory I_0. Assume the analysis for an n-day period (for "large" n). Note that the tank is refilled on each visit and the customer is serviced when the tank inventory reaches zero.

In this case, the (optimal) cost of service visits is:

$$\frac{b(n\mu - I_0)}{T}.$$

If we plan deliveries for the next m ($\ll n$) days, and select this planning horizon m small enough such that no customer will need more than one replenishment during the next m days, then for each customer we will have to examine two possible cases:

(1) If $I_0 - m\mu < 0$, this customer must be replenished during the next m days, otherwise a stockout occurs.

(2) If $I_0 - m\mu \geq 0$, the customer need not be replenished during the next m days.

If case 1 occurs, then the optimal policy would dictate that the replenishment take place on day $t^* = I_0/\mu$, allowing for non-integer t^* values. Clearly, if case 2 occurs it is best not to replenish the customer in this current m day period. This single customer analysis is very basic and does not communicate any problem dynamics. What if the capacity of the system is insufficient to replenish all the customers whose t^* day falls on a specific day during the current m-day period, but is sufficient to replenish all the customers which have their best delivery day t^* fall during some day of the current m-day period? Some of these customers will have to be replenished on a day different than their corresponding t^*. Thus, just for that reason we have to calculate for each customer the marginal cost over n days, denoted $c(t)$, of replenishment on day t deviating from day t^*. That is, $c(t^*) \equiv 0$, and $c(t) > 0$, $t \neq t^*$. There are other important reasons for evaluating this marginal cost $c(t)$, for instance balancing the work load over time. Another quantity is calculated for the customers who do not need to be replenished on day t during the current m-day period. This quantity is denoted by $g(t)$ and represents the decrease in future costs (over an n day horizon) if the customer is replenished during the current m-day period at no cost instead of being replenished at his "best" day in a future m-day period at cost b. Below we repeat the calculation from Dror and Ball (1987).

If replenishment is executed on day t, then the closing inventory I_c is defined by

$$I_c = T - (m - t)\mu.$$

Now a simple difference calculation for $c(t)$ and $g(t)$ is as follows:

$$\begin{aligned} b\frac{(n\mu-(T-(m-t)\mu))}{T} - b\frac{(n\mu-(T-(m-t^{*})\mu))}{T} &= \frac{b}{T}(I_c-\mu t) \\ &= c(t) \\ b\frac{(n\mu-(I_c-m\mu))}{T} - b\frac{(n\mu-(T-(m-t)\mu))}{T} &= b\frac{(T-I_c+t\mu)}{T} \\ &= b-c(t)=g(t) \end{aligned}$$

The above simple deterministic single customer analysis is extended to a more realistic stochastic model and later $c(t)$ and $g(t)$ are used as a cost coefficient in a multi-customer setting. The major weakness of the above analysis lies in its assumption that we know b—the cost of visiting a customer, and that this value, even though customer specific, remains constant for all the replenishments in the n-period. Clearly, this is not entirely true in real-life propane distribution. We will return to this important point later on.

4.1 The stochastic single customer

Let r_t denote the amount of propane consumed by customer on day t. Normally, we do not know the value of r_t. We do not know its exact value for past days, which is less important, and, even more so, we do not know its value for future days. That is, r_t, $t=1,2,\ldots$ are random variables. For simplicity we assume that r_ts are independent identically distributed random variables for each t (consider that the seasonality effects are removed) with mean μ, variance σ^2, and cumulative distribution function $F(\cdot)$. The randomness (and variability) of consumption makes the replenishment scheduling a risky proposition. Guessing that there is enough propane in a customer's tank when there is not usually results in a high cost emergency replenishment. Guessing that there is little left in the tank when there is a lot left results in an almost equally costly visit. Thus, it is of value to calculate the replenishment day which balances the risk of the two cost penalties and at the same time accounts for future implications of an expected delivery volume that is less than the tank capacity. We describe below this calculation assuming that the customer's tank is full on day 1.

Let $R_t=\sum_{i=1}^{t} r_i$ denote the cumulative consumption over a t day period. Let $P_S(t)$ denote the probability that a stockout occurs on day t given that the tank has not been refilled prior to day t. Thus, assuming that $\mu<T$,

$$P_S(t)=\text{Prob}\{R_{t-1}\le T<R_t\}=\text{Prob}\{R_{t-1}\le T\}-\text{Prob}\{R_t\le T\}$$

$$= F^{(t-1)}(T) - F^{(t)}(T)$$

where $F^{(k)}(T)$ is the k-fold convolution of F (set $F^{(0)} \equiv 0$).

Now we are in a position to write down the expression for the expected cost during the next $n + m$ days associated with a delivery on day t denoted by $E(t)$. Denote by $k^* = \max\{k : \text{such that } k\mu \leq T\}$,

$$\begin{aligned} E(t) &= \sum_{i=1}^{t-1} \big(S + c(i)\big) P_S(i) + \left(1 - \sum_{i=1}^{t-1} P_S(i)\right) c(i) \\ &= \sum_{i=1}^{t-1} \big(S + c(t) - c(i)\big) P_S(i) - c(t) \end{aligned}$$

for $1 \leq t \leq k^*$ (or $k^* + 1$).

What is particularly interesting is that in Dror and Ball (1987), it has been proven that $E(t)$, $1 \leq t \leq k^*$ is a strictly convex function by proving that $P_S(t) > P_S(t-1)$, $2 \leq t \leq k^*$, for r_ts normally distributed with coefficient of variation ≤ 1. Moreover, in Kreimer and Dror (1990), this result was strengthened by proving that the relation $P_S(t) > P_S(t-1)$, $2 \leq t \leq k^*$ holds for a number of other interesting distributions. In Dror (2002), the relation $P_S(t) > P_S(t-1)$, $2 \leq t \leq k^*$ (monotonicity) was stated formally as a more general conjecture.

In summary, the result is that $E(t)$, when viewed as a continuous function of t, is convex; thus it achieves its minimum at a single point (or at most 2 points as a discrete function). That is, let $E(t^*) = \min\{E(t) : 1 \leq t \leq k^*\}$ determines the "best" (minimal expected cost) day for replenishment — t^*. It is appropriate to note that a similar analysis (with similar results) has been conducted by Jaillet et al. (2002).

4.2 The propane routing model

The notation is similar but not identical to the presentation of the model by Federgruen and Zipkin (1984).

Constants.

NV = number of vehicles
$M =$ the set of customers, with 0 indicating the depot location
$Q =$ capacity of a vehicle (homogenous vehicles)
$c_{ij} =$ cost of servicing customer i then travelling from i to j
$m =$ number of days in planning period.

The quantities defined in the previous subsection now become customer specific, so that we have T_i, b_i, μ_i, S_i, $c_i(t)$, $g_i(t)$, I_{0i}, and t_i^*, defined for all $i \in M$. Customer i's expected demand on day t is denoted by $q_i(t)$ and equal to $T_i - I_{0i} - \mu_i t$. Since not all customers in M need to be replenished during the current planning period, we partition the customers into two subsets. Let $\widehat{M} = \{i \in M$ be such that $t_i^* \leq m\}$ as the customers who must be replenished during the current planning period, and $M^c = M \setminus \widehat{M}$ the rest of customers. In addition, to simplify the formulation we denote by TSP(N) a travelling salesman problem solution for customers in $N \subset M$.

Variables.

$y_{iwt} =$ 1, if customer i is assigned to route (vehicle) w on day t, and is 0 otherwise.

Now the mathematical formulation expressing a single period cost minimization is stated as follows:

$$\min \sum_{w=1}^{\mathrm{NV}} \sum_{t=1}^{m} \Big(\mathrm{TSP}(N_{wt}) + \sum_{i \in \widehat{M}} c_i(t) y_{iwt} - \sum_{i \in M^c} g_i(t) y_{iwt} \Big) \tag{10.12}$$

subject to the following constraints

$$\sum_{w=1}^{\mathrm{NV}} \sum_{t=1}^{t_i^*} y_{iwt} = 1, \qquad \forall\, i \in \widehat{M}, \tag{10.13}$$

$$\sum_{w=1}^{\mathrm{NV}} \sum_{t=1}^{t_i^*} y_{iwt} \leq 1, \qquad \forall\, i \in M^c, \tag{10.14}$$

$$\sum_{i \in M} q_i(t) y_{iwt} \leq Q, \qquad w = 1, \ldots, \mathrm{NV}; t = 1, \ldots, m \tag{10.15}$$

$$N_{wt} = \{i : y_{iwt} = 1\}, \qquad w = 1, \ldots, \mathrm{NV}; t = 1, \ldots, m \tag{10.16}$$

$$y_{iwt} \in \{0, 1\}, \qquad \forall\, i, w, t \tag{10.17}$$

The y_{iwt} variables indicate for customer i the replenishment day and the replenishment vehicle. We artificially require that customers be replenished before or at their best day t^*. Customers who do not have their best day fall in the current m-period, do not have to be replenished (10.14). Other than the term for TSP(N) in the objective function followed by the appropriate set partition in (10.16), the formulation resembles that of the generalized assignment problem. The stochasticity is captured by the t^*'s and the dynamics (long-term implications) by the $c_i(t)$'s and $g_i(t)$'s. The "big" problem is that of calculating the individual b_i values required for calculation of $c_i(t)$'s ($g_i(t)$'s). Dror and Ball (1987) offered only an approximation of unproven quality, but their computational tests compared very favorably with the real-life results (Trudeau and Dror (1992)). For practitioners, a solution system based on this approach is best described in Dror and Trudeau (1988). This (10.12)–(10.17) mathematical formulation is similar in spirit to the formulation from Bell et al. (1983). There are however a number of differences. It is not a set-covering approach. That is, it is not a scheme to cover a given set of customers by selecting routes, each containing a subset of customers, from a large family of externally generated routes. It is a customer selection approach which selects subsets of customers together with the days in which to replenish these customer subsets (see also Dror et al., 1985, 1986). In addition, the amounts delivered are determined by the replenishment day since the policy is always to fill-up the tank, and the delivery implications are explicitly projected forward. In Bell et al. (1983) the future implications of a present delivery are not clearly spelled out.

5. The Markov decision process approach for inventory routing

Clearly propane delivery routing is merely one representative of a large class of practical significant problems. Yet due to the inherent combinatorial and stochastic nature of this class, it remains notoriously intractable. Formulating the control problem as a Markov decision process represents an attractive modelling approach which captures most of the system dynamics intrinsic to propane delivery. Following Minkoff (1993), there have been a number of attempts to do just that. Markov decision process modelling of inventory routing has taken-off in the work of Kelywegt, et al. (2002, 2003). However, the concomitant contribution by Adelman (2003a,b, 2004) is the most promising solution approach yet. We attempt below to provide a brief summary of the main ideas in these works.

5.1 The Markov decision process model (MDP)

The Markov decision process model is stated as follows (We modify the models of Minkoff, 1993; Kleywegt et al., 2002, slightly to unify notation and assumptions.):

(1) The state variable $I = (I_1, \ldots, I_{\bar{m}})$, where M is the customer set and $\bar{m} = |M|$, represents the current amount of inventory at each customer. The constant vector $T = (T_1, \ldots, T_{\bar{m}})$ represents the customers' tank capacities. Thus, the inventory can vary (continuously or discretely) in the product state space $\mathcal{I}$ bounded below by the zero vector and above by the vector of tank capacities. Let $I_t = (I_{1t}, \ldots, I_{\bar{m}t}) \in \mathcal{I}$ denote the inventory state at time t.

(2) Given a state vector $I \in \mathcal{I}$, denote by $A(I)$ the set of all feasible decisions. A decision $a \in A(I_t)$ in time t selects (i) the subset of customers for replenishment, and (ii) the vehicles' replenishment routes. Note that the amount to be replenished can be either a part of the decision, or, like in a partially observable MDPs, the outcome of customer inventory level observed on delivery if we always refill customer's tank. In the second case, the decision a will have to contain an estimate of what should be the replenishment volume. However, the actual delivery value might be quite different. Let $a_t \in A(I_t)$ be the decision chosen at time t. In our propane Markov decision model we assume that the exact demand is revealed only when the vehicle arrives at customer location and the policy is to fill-up the tank.

(3) The system's randomness is expressed in terms of the daily consumption rate $r = (r_1, \ldots, r_{\bar{m}})$. That is, the amount that can be delivered to customer i at time t (the demand q_{it}), equals $T_i - I_{it}$, which is a random variable dependent on r_i's since the last replenishment. The amount delivered to customer i (denoted by $d_{it}(a)$ in deference to q_{it}) by executing the policy a on day t can be either zero, a predetermined quantity d_{it}, or $T_i - I_{it}$ (if a replenishment always fills-up the tank). Let $U = \{I_{t+1} \in R_+^{\bar{m}} : \big((I_{1t} - r_{1t} + d_{1t}(a)), \ldots, (I_{\bar{m}t} - r_{\bar{m}t} + d_{\bar{m}t}(a))\big)\}$. The known joint probability distribution F of customers demands $q_t = (q_{1t}, \ldots, q_{\bar{m}t})$ gives us a known transition function in the form of a conditional probability distribution. That is, for any state $I \in \mathcal{I}$, and decision $a \in A(I)$, we have

$$\text{Prob}\{U \mid I_t, a] = F(U \mid I_t, a).$$

(4) Let $\zeta(I, a)$ denote the expected single stage net reward (cost) if the process is in state I at time t, and decision $a \in A(I)$ is implemented. Note that not only the exact customer demands are random variables

but also the costs of the corresponding routing solution is a random variable since we do not know this cost until we execute the route and incur the additional recourse routing costs in response to route failures (see Trudeau and Dror, 1992).

(5) The objective is to maximize the expected total discounted value (or the present value of the cash flow), over an infinite (or finite "long" n-day) horizon. The decisions in time t, $a(t)$, are restricted to the feasible sets $A(I_t)$ for each t and depend only on the history $(I_0, a_0, \ldots, I_{t-1}, a_{t-1}, I_t)$ of the process up to time t. Let Π be the set of policies which depend on the history up to time t. Let $\alpha \in [0, 1)$ denote the discount factor. Let $\nu^*(I)$ denote the optimal expected value given the initial state is I, then

$$\nu^*(I) \equiv \sup_{\pi \in \Pi} E^{\pi} \left[\sum_{t=1}^{\infty} \alpha^t \zeta(I_t, a_t) | I_0 = I \right]$$

Following standard text book analysis (see Bertsekas and Shreve, 1978), a stationary deterministic policy π selects a decision $\pi(I) \in A(I)$ based only on the current state I. In principle, under some conditions, one can solve the above system by dynamic programming, computing the optimal value function ν^* and an optimal policy π^*. However, for the problem described here as the propane inventory problem, this is clearly impractical. The state space $\mathcal{I}$ is much too big (uncountable). The dimensionality is too high. The subproblems which need to be solved are NP-hard, etc. See the detailed arguments in Kleywegt et al. (2002, 2004). On the surface, this modelling approach seems to lead nowhere. Still, as a mathematical model it has the ability to represent the intrinsic problem details in a clear manner. Minkoff (1993) and Kleywegt et al. (2002, 2004) both attempted the ambitious undertaking of "salvaging" this Markov decision process approach to obtain reasonable solutions for inventory routing. (See also Berman and Larson (2001), for a different modelling approach.) In essence, Minkoff (1993) and Kleywegt et al. (2002, 2004), solution approach partitions the set of customers and estimates parameters for each subset by simulation. The optimal value function ν^* is approximated by $\hat{\nu}$ by choosing a collection of subsets (of size 1 or 2) of customers that partition the customer set. The approximate function $\hat{\nu}$ is computed for each subset and the sum over the subsets constitutes the approximate value. To simplify matters, Kleywegt et al. (2002, 2004) discretize their inventory demand state space. In all fairness, in Kleywegt et al. the focus is on designing vehicle routes which are limited to one or at most two customers, and the customers stockouts are due to lack of available vehicles. Since in our experience with propane delivery, vehicle availability was never the reason for stockouts, we do not consider these limitations here. There are however a number of questions regarding the modelling and solution methodology of Kleywegt et al. (2002, 2004). In addition, it is not clear how Kleywegt et al. solutions compare with real-life inventory routing since they did not conduct computational study which compares their results with real world data. However, many of the questions/reservations regarding their model are subsequently addressed in the work of Adelman (2003a, 2004) which we describe next.

5.2 Price-directed Markov models

Adelman (2004) states his decision model clearly: "The dispatcher chooses nonnegative integer-valued replenishment quantities $q =< q_1, q_2, \ldots, q_{\bar{m}} >$ with q_i equal to the quantity replenished at i, $i = 1, \ldots, \bar{m}$." For simplicity, we can adopt this notion of deciding the propane replenishment quantities a priori regardless of the amount I_{it} realized at time t in location i and the actual demand at i (remember that we do not know the exact inventory levels before service and therefore do not know how much is needed to fill the tank at i). The state space is as before the product space of estimated (and known) inventory levels $\mathcal{I}$. After estimating an inventory state $I \in \mathcal{I}$ the dispatcher selects the subset of customers who will be replenished in the current period. As before (for instance, Dror and Ball, 1987), it is assumed that no customer will be replenished more than once in a period. The customers are partitioned into non-empty (disjoint) subsets $M = \{M_1 \cup \cdots \cup M_K\}$, where K is the number of subsets ($K \leq \bar{m}$) including the subset of customers who are not to be replenished in the current period, say M_K. Note that K and the particular partition are part of the action $a \in A(I)$. The idea is that the customers in each subset M_j, $j = 1, \ldots, K-1$ are replenished together (the same vehicle trip) in the current period (in Adelman, 2004, a period is a day). Based on the present state I, the corresponding action space $A(I)$ consists of determining the partition number K, the partition $M = \{M_1 \cup \cdots \cup M_K\}$, and the vector q. The vehicle capacity constraints specify that $\sum_{j \in M_i} q_j \leq Q$, $i = 1, \ldots, K-1$. In addition, the components of the replenishment vector q as a function of the state I have to confirm to the customer tank constraints. That is, $q_i(I) \leq \max\{0, T_i - I_i\}$, $i = 1, \ldots, \bar{m}$. In fact, one can replace the actual tank capacities T_i with artificial tank capacities T_i', or vehicle capacity Q with artificial vehicle capacity $Q' < Q$, as in a chance constrained models, to control the route failure probability for each subset of customers (see Trudeau and Dror, 1992).

After executing action $a \in A(I_t)$, the system observes (on delivery) a partial realization of demand. That is, the system observes only the demand quantities of the customers who were replenished at period t (day t). However, Adelman (2004) like Kleywegt et al. (2002, 2004) models the MDP as if the entire vector of demands $d = \langle d_1, \ldots, d_{\bar{m}} \rangle$ is observed after the decision a is taken. Here, we follow their modelling approach with respect to the probability distribution of the demand vector. That is, let $\eta(d)$ denote the probability that the demand equal d where $d_i \in D_i$, and D_i is a finite set of nonnegative integers.

Once the demand is realized, costs are computed. That is, given an action $a \in A(I_t)$, we obtain a partition of M as $M = M_1(a) \cup \cdots \cup M_{K(a)}(a)$ and the corresponding cost equal to $\sum_{i=1}^{K(a)-1} C_i\big(M_i(a)\big)$, where the cost for replenishing a given subset $M_i(a)$ is $C_i\big(M_i(a)\big)$ — the cost of the replenishment route (a TSP route) through M_i. Clearly, if our convention is that the subset $M_{K(a)}(a)$ does not get replenished in the current period, then the cost $C_{K(a)}\big(M_{K(a)}(a)\big) \equiv 0$. In addition to the delivery (routing) cost, Adelman (2004) also uses a traditional linear form to account for inventory holding and shortage costs in each location in the form of $g_i(I_i, q_i, d_i) = h_i(I_i + q_i - d_i)^+ + b_i\big(d_i - (I_i + q_i)\big)^+$.

Adelman (2004) derives an infinite horizon, expected discounted cost MDP which requires the dispatcher to find an optimal expected cost minimizing policy. After deriving the optimality equations (following Puterman, 1994) for finding an optimal policy that is Markovian, stationary, and deterministic, a linear program is proposed to solve the problem. Again, because of the huge size of the subsequent model, approximation solution schemes must be proposed. That is, the optimality equations are:

$$\nu^*(I) = \min_{a \in A(I)} \left\{ \sum_{i \in M} g_i\big(I_i, q_i(a)\big) + \sum_{j=1}^{K(a)-1} C_j\big(M_j(a)\big) + \alpha \sum_{I' \in \mathcal{I}} p(I'|I, a)\nu^*(I') \right\}, \quad \forall\, I \in \mathcal{I}$$

where $g_i\big(I_i, q_i(a)\big)$ is the expected holding and stockout cost for item i given the current state I_i and $q_i(a)$ is replenished. The linear program is:

$$\mathrm{LP}_0 = \max_{\nu} \sum_{I \in \mathcal{I}} s(I)\nu(I)$$

$$\nu(I) \leq \sum_{i \in M} g_i\big(I_i, q_i(a)\big) + \sum_{j=1}^{K(a)} C_j(M_j(a)) + \alpha \sum_{I' \in \mathcal{I}} p(I' \mid I, a)\nu^*(I'), \quad \forall\, a \in A(I); I \in \mathcal{I}$$

where $s(I) > 0$ can be arbitrary positive constants for all $I \in \mathcal{I}$.

Substituting $\nu(I)$ by the sum of customer dependent value functions $V_i(I)$, that is, $\nu(I) = \sum_{i \in M} V_i(I), \forall\, I \in \mathcal{I}$ we can rewrite the above linear program.

An important modelling novelty introduced in Adelman's (2004) math programming based solution scheme, is his approximation of $C_i(M_i)$ from below with $\sum_{j \in M_i} W_j(q_j)$, where $W_j(q_j)$ represents the allocated cost of replenishing customer j with quantity q_j. Without recasting the full analysis of Adelman (2004), we note that the inventory replenishment solution uses the $W_j(q_j)$ in a similar role to the customer specific b_j value in Dror and Ball (1987). The optimal $W_j^*(q_j)$ values have to satisfy cost allocation efficiency conditions. That is, $\sum_{j \in M_i} W_j^*(q_j) = C(M_i)$ = the cost of the TSP tour through the subset of customers M_i including the depot. Thus, one approximating model proposed by Adelman (2004) is

$$\mathrm{LP}_{\mathrm{app}} = \max_{V,W} \sum_{i \in M} \sum_{I_i \in \mathcal{I}} s(I_i) V_i(I_i)$$

$$V_i(I_i) \leq g_i(I_i, q_i) + W_i(q_i) + \alpha \sum_{I_i' \in \mathcal{I}} p_i(I_i' \mid I_i, q_i) V_i(I_i'), \quad \forall\, q_i \in \Upsilon(I_i), I \in \mathcal{I}$$

$$\sum_{i \in M'} W_i(q_i) \leq C(M'), \quad \forall M' \in M, q \in \Upsilon(M')$$

where $\Upsilon(M') = \{q_i, i \in M' : \sum_{i \in M'} q_i \leq Q \cdot \mathrm{NV}\}, \forall M' \subset M$, and $C(M')$ is the cost on an optimal VRP solution.

Adelman (2004) has shown that $\mathrm{LP}_{\mathrm{app}}$ gives the same results as forcing separable V in LP_0, but LP_{app} is much easier to solve. The optimal vector W^* of $W_i^*(q_i)$'s is coupled with the optimal vector V^*. When the optimal prices $V_i^*(I_i)$ are used to obtain the control solution then Adelman calls it a *price-directed* control policy. We note that our definition of $\Upsilon(M')$ is different than that in Adelman (2004) since the cost function $C(\cdot)$ must also depend on the full vector q, because it is now the solution to a VRP instead of a TSP. However, we believe that the math goes through in this case if we ignore the travel time. There are a number of key technical details in Adelman (2004) which we omit here for the sake of space. Incidently, in Adelman (2003b), Proposition 2 shows that when $C(M')$ is the cost of the optimal VRP solution, it can be decomposed into individual TSP solutions for the purpose of solving the relaxed LP. Based on the computational results, this (Adelman, 2004) solution methodology is proven superior to that of Minkoff (1993) and Kleywegt et al. (2002).

5.3 Cost allocation for subsets and inventory

In Adelman (2003a) the "price-directed" solution methodology for inventory routing receives an additional boost in terms of clarification of ideas, solution philosophy, and results. However, we should note that this paper looks for optimal policies in a deterministic setting like the one in the example described in Figure 1. The key concept in this paper is that of incremental cost when considering, in current time, the replenishment for customer i. That is, the key value which "real-world" dispatcher ought to examine is $C(M_j \cup i) - C(M_j), M_j \subset M, i \notin M_j$, together with the future cost implication of delivering quantity $d_i > 0$ to i.

Since all the costs have to be absorbed by the customers, Adelman's (2003a) analysis requires a cost allocation process which is applied simultaneously to routing and inventory replenishment decisions. (For cost allocation in vehicle routing see Gothe-Lundgren et al., 1996.) The propane delivery problem is formulated as that of minimizing long-run time average replenishment costs. This objective corresponds nicely to the objective of maximizing the long-run average number of units (gallons) delivered per hour of delivery operation which is used in real-life propane distribution. Adelman (2003a) formulates the problem as a control problem using dynamic system equations. Without restating the evolution of the problem modelling and the technical details involved, we note that the main thrust is to reformulate the deterministic control problem as a nonlinear program in which "in the long-run averages, replenishment must equal consumption." Solving the nonlinear program leads to the development of what is called the *price-directed* operating policy which maximizes the net-value of the replenishment. Incidently, Adelman proved that the objective used by Dror and Ball (1987), is also a net-value replenishment maximizing objective justifying its apparent success. Next, we sketch out Adelman's (2003a) modelling approach. We attempt to keep the notational convention of the earlier sections.

Adelman links the initiation of a replenishment action to any subset of customers with an occurrence of one stockout (or more than one, if occurring simultaneously) in the system which triggers a "must" replenishment response. (This is not how propane replenishment systems behave in practice, but it is quite fitting in this setting.) Now the time is measured not in day units, but as the elapsed time between the successive initiations of new replenishment activities. Thus, $\overline{T}_t$ represents the time elapsed between replenishment epochs t and $t+1$, $t = 1, 2, 3, \ldots$, and I_{it} be the inventory level at customer i just before the tth replenishment operations activation. We require that at least one day lapses between consecutive replenishment activities. Given a set M of customers, let $\widehat{M}_{\hat{A}} \subseteq M$ denote a subset of customers, and let the zero-one variable $Z_{\widehat{M}_{\hat{A}},t} = 1$, if customers in $\widehat{M}_{\hat{A}}$ are replenished during epoch t. The corresponding control problem is formulated as follows:

(CONTROL)

$$\inf \lim_{N\to\infty} \sup \frac{\sum_{t=1}^{N} \sum_{\widehat{M}_{\hat{A}} \subseteq M} C(\widehat{M}_{\hat{A}}) Z_{\widehat{M}_{\hat{A}},t}}{\sum_{t=1}^{N} \bar{T}_t} \tag{10.18}$$

$$I_{i,t+1} = I_{i,t} + d_{i,t} - \mu_i \bar{T}_t, \qquad \forall \text{ positive integer } t; \forall i \in M \tag{10.19}$$

$$d_{i,t} \leq (T_i - I_{i,t}) \cdot \sum_{\widehat{M}_{\hat{A}} \subseteq M : i \in \widehat{M}_{\hat{A}}} Z_{\widehat{M}_{\hat{A}},t}, \qquad \forall \text{ positive integer } t; \forall i \in M \tag{10.20}$$

$$\sum_{i \in \widehat{M}_{\hat{A}}} d_{i,t} \leq Q \cdot \text{NV}, \qquad \forall \text{ positive integer } t \tag{10.21}$$

$$\sum_{\widehat{M}_{\hat{A}} \subseteq M} Z_{\widehat{M}_{\hat{A}},t} = 1, \qquad \forall \text{ positive integer } t \tag{10.22}$$

$$Z_{\widehat{M}_{\hat{A}},t} \in \{0,1\}, \qquad \forall \widehat{M}_{\hat{A}} \subseteq M, \forall \text{ positive integer } t \tag{10.23}$$

$$s, I, \bar{T} \geq 0 \tag{10.24}$$

The objective (10.18) minimizes the long-run average replenishment costs. Note that $C(\widehat{M}_{\hat{A}})$ denotes the cost of the corresponding VRP solution through the subset $\widehat{M}_{\hat{A}}$. Constraints (10.19)) state the conservation of inventory for each customer i. Constraints (10.20) insure that the individual tank capacities are respected. Constraints (10.21) make sure that for the replenishments scheduled in an epoch the vehicle fleet capacity is not exceeded. Constraints (10.22) state that exactly one subset is selected for replenishment in each replenishment epoch. The other constraints are just state $0-1$ selection for subsets and nonnegativity of the corresponding vectors. We note that it is straight forward in this formulation to limit the choice of the subsets $\widehat{M}_{\hat{A}}$ which can be considered for replenishment and thus manage the size of the corresponding control problem.

In order to solve the above problem, a nonlinear programming model is proposed which is a relaxation of the original. Denote by $Z_{\widehat{M}_{\hat{A}}}$ a nonnegative decision variable representing the long-run time average rate that the subset $\widehat{M}_{\hat{A}}$ is replenished together. For each such subset $\widehat{M}_{\hat{A}}$ which contains i, let $d_{i,\widehat{M}_{\hat{A}}}$ denote the decision variable representing the average replenishment quantity delivered to i when replenishing the subset $\widehat{M}_{\hat{A}}$. The corresponding program is stated as:

$$\textbf{(NLP)} \qquad \min \sum_{\widehat{M}_{\hat{A}} \subseteq M} C(\widehat{M}_{\hat{A}}) Z_{\widehat{M}_{\hat{A}}} \tag{10.25}$$

$$\sum_{\widehat{M}_{\hat{A}} \subseteq M; i \in \widehat{M}_{\hat{A}}} d_{i,\widehat{M}_{\hat{A}}} Z_{\widehat{M}_{\hat{A}}} = \mu_i, \qquad \forall i \in M \tag{10.26}$$

$$\sum_{i \in \widehat{M}_{\hat{A}}} d_{i,\widehat{M}_{\hat{A}}} \leq Q \cdot \mathrm{NV}, \qquad \forall \widehat{M}_{\hat{A}} \subseteq M \tag{10.27}$$

$$d_{i,\widehat{M}_{\hat{A}}} \leq T_i, \qquad \forall \widehat{M}_{\hat{A}} \subseteq M, i \in \widehat{M}_{\hat{A}} \tag{10.28}$$

$$Z, d \geq 0 \tag{10.29}$$

Adelman (2003a) shows that solutions to (NLP) may not be necessarily implementable because it does not capture all the dynamics in the system.

Next, a dual problem to (NLP) is formulated below with decision variables V_i and data d_i derived from a set $\mathcal{D}_O$.

$$\textbf{(D)} \qquad \max \sum_{i \in M} \mu_i V_i \tag{10.30}$$

$$\sum_{i \in \widehat{M}_{\hat{A}}} d_i V_i \leq C(\widehat{M}_{\hat{A}}), \quad \forall\ (\widehat{M}_{\hat{A}}, d) \in \mathcal{D}_O \tag{10.31}$$

The interpretation of the V_i's is that "at optimality they are the marginal costs, or prices, associated with satisfying constraints (10.26) of (NLP)" and "$\mu_i V_i$ at optimality can be interpreted as the total allocated cost rate for replenishing customer i in the optimal solution to (NLP)."

As far as solution, Adelman (2003a) solves the (NLP) by solving a version of the dual problem (D) by column generation procedure. We do not describe here the technical details and the difficulties involved. With this solution scheme he can prove that the solution for the motivating example (Figure 10.1) is indeed optimal!

The computational study of Adelman's (2003a) price-directed solution methodology demonstrates its viability. It produces results superior to all the previously proposed solution schemes.

6. Summary

This chapter is about solving the problem of propane deliveries. It is commonly viewed as a representative problem of a much larger family of hard problems of considerable practical significance. This problem has been on the "front burner" of the logistics academic and practitioners community for over twenty years. In fact, it was voted as the "most important/interesting" current OR problem in an unofficial gathering of Operations Research professionals which took place in early 1983 at Cornell University. Has it been solved now?

Clearly, Bell et al. (1983) describe a workable solution to the problem. They were able to construct a mathematical optimization module which routinely solved mixed integer programs with 800,000 variables and 200,000 constraints. In 2004, with our present computing power, this model should be able to solve problems 100 times larger. Could they prove optimality of their solution for the example in Figure 10.1? I do not think so. This problem was solved by Adelman (2003a).

In another solution scheme, Dror and Ball (1987) proposed and implemented a solution methodology which routinely solved problems with 5,000 customers. While it was (and may still be) a very promising solution methodology, it did not claim or deliver optimal solutions.

Presently, the work of Adelman (2003a,b, 2004) stands out as the reigning incumbent. Adelman (2003a), describes computational testing of his approach on a number of instances form Praxair, Inc. available online. These computational results are very promising. We would hope that more testing on "real-world" problems and operational implementation would follow.

Final Note. Propane deliveries are made every day in the US, Canada, and Europe (to my knowledge). The real-world operators are making money replenishing customers with propane and are in the market for improved solution methodologies for their operations.

Acknowledgment The author is grateful to Dan Adelman from The University of Chicago, Graduate School of Business, for many comments and constructive critique of this chapter. I claim all the remaining mistakes.

References

Adelman, D. (2003a). Price-directed replenishment of subsets: Methodology and its application to inventory routing. *MSOM*, 5(4):348–371.

Adelman, D. (2003b). Internal Transfer Pricing for a Decentralized Operation with a Shared Supplier. Working paper, October 2003, The University of Chicago, Graduate School of Business.

Adelman, D. (2004). A price-directed approach to stochastic inventory/routing. Forthcoming in *Operations Research.*

Bard, J.F., Huang,L., Jaillet,P., and Dror, M. (1998). A decomposition approach to the inventory routing problem with satellite facilities. *Transportation Science*, 32:189 – 203.

Bell, E.T., Dalberto, L.M., Fisher, M.L., Greenfield, A., Jaikumar, R., Kedia, P., and Prutzman, P. (1983). Improving the distribution of industrial gases with on-line computerized routing and scheduling optimizer. *Interfaces*, 3(6):4 – 23.

Berman, O. and Larson, R.C. (2001). Deliveries in an inventory/routing problem using stochastic dynamic programming. *Transportation Science*, 35:192 – 213.

Bertsekas, D.P. and Shreve, S.E. (1978). *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, New York.

Dror, M. (1983). *The Inventory Routing Problem*. Ph.D. Thesis, University of Maryland at College Park.

Dror, M. (2002). Routing with stochastic demands: A survey. In: M. Dror, P. L'Ecuyer, and F. Szydarovszky (eds.), *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, pages 629 – 653, Kluwer Academic Publishers.

Dror, M. and Ball, M.O. (1987). Inventory/routing: Reduction from annual to a short period problem. *Naval Research Logistics*, 34:891 – 905.

Dror, M., Ball, M.O., and Golden, B. (1985/6). Computational comparison of algorithms for inventory routing. *Annals of Operations Research*, 4:3 – 23.

Dror, M. and Trudeau, P. (1988). Inventory routing: Operational design. *Journal of Business Logistics*, 9:165 – 183.

Dror, M. and Trudeau, P. (1996). Cash flow optimization in delivery scheduling. *European Journal of Operational Research*, 88:504 – 515.

Federgruen, A. and Zipkin, P. (1984). A combined vehicle routing and inventory allocation problem. *Operations Research*, 32:1019 – 1037.

Fisher, M., Greenfield, A., Jaikumar, R., and Kedia, P. (1982). *Real-Time Scheduling of Bulk Delivery Fleet: Practical Application and Lagrangian Relaxation*. Technical report, The Wharton School, University of Pennsylvania, Department of Decision Science.

Gothe-Lundgren, M., Jornsten, K., and Varbrand, P. (1996). On the nucleolus of the basic vehicle routing game. *Mathematical Programming*, 72:83 – 100.

Herron, D. (1979). Managing physical distribution for profit. *Harvard Business Review*, 79:121 – 132.

Jaillet, P., Huang, L., Bard, J.F, and Dror, M. (2002). Delivery cost approximations for inventory routing problems in a rolling horizon framework. *Transportation Science*, 36:292 – 300.

Kleywegt, A.J., Nori, V.S., and Savelsbergh, M.W. (2002). The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36:94 – 115.

Kleywegt, A.J., Nori, V.S., and Savelsbergh, M.W. (2004). Dynamic programming approximations for stochastic inventory routing problem, *Transportation Science*, 38:42 – 70.

Kreimer, J. and Dror, M. (1990). The monotonicity of threshold detection probability in stochastic accumulation processes. *Computers & Operations Research*, 17:63 – 71.

Minkoff, A.S. (1993). A Markov decision model and decomposition heuristics for dynamic vehicle dispatching. *Operations Research*, 41:77 – 90.

Puterman, M.L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley.

Trudeau, P. and Dror, M. (1992). Stochastic inventory routing: Stockout and route failure. *Transportation Science*, 26:172 – 184.