Delivery Cost Approximations for Inventory Routing Problems in a Rolling Horizon Framework [*]

Patrick Jaillet [†]

Department of Management Science and Information Systems
The University of Texas at Austin, Austin, TX 78712

Jonathan F. Bard [‡]

Graduate Program in Operations Research
Mechanical Engineering Department
The University of Texas at Austin, Austin, TX 78712

Liu Huang

DSC Communications, Austin, Texas 78758-4017

Moshe Dror

College of Business and Public Administration
University of Arizona, Tucson, AZ 85721

Abstract

The inventory routing problem considered in this paper is concerned with the repeated distribution of a commodity, such as heating oil, over a long period of time to a large number of customers. The problem involves a central depot as well as various satellite facilities which the drivers can visit during their shift to refill their vehicles. The customers maintain a local inventory of the commodity. Their consumption varies daily and cannot be predicted deterministically. In case of a stockout, a direct delivery is made and a penalty cost is incurred. In this paper we present incremental cost approximations to be used in a rolling horizon framework for the problem of minimizing the total expected annual delivery costs.

Keywords: delivery cost approximations, inventory routing problem, stochastic demand, rolling horizon

Introduction

Inventory routing problems (IRPs) usually arise when consumers rely on a central supplier to provide a given commodity on a repeated basis. For the central supplier, these problems usually involve the daily specification of a sequence of locations that must be visited by a fleet of delivery vehicles for restocking purpose. Whenever a customer unexpectedly stocks out, a special delivery must be made immediately, with a high incurred cost.

Although IRPs have appeared in the literature as early as in the 1970s (e.g., Beltrami and Bodin, 1974 and Russell and Igo, 1979), much of the research is more recent (e.g., Anily and Federgruen, 1990 and 1993, Bell et al., 1983, Bard et al., 1998, Berman and Larson, 1999, Chan, Federgruen, and Simchi-Levi, 1998, Chan, Muriel, and Simchi-Levi, 1998, Chien, Balakrishnan, and Wong, 1989, Christofides and Beasley, 1984, Dror and Ball, 1987, Dror, Ball, and Golden, 1985, Federgruen and Simchi-Levi, 1995, Federgruen and Zipkin, 1984, Fisher et al., 1983, Gaudioso and Paletta, 1992, Golden, Assad, and Dahl, 1984, Kleywegt, Nori, and Savelsbergh, 1998 and 2000, Kumar, Schwarz, and Ward, 1995, Reiman, Rubio, and Wein, 1999, Rubio, 1995, Trudeau and Dror, 1992, and Webb and Larson, 1995).

For typical IRPs, a customer's consumption rate is difficult to predict with certainty and can only be represented at best by a random variable with known probability distribution. If one is willing to assume such probability distributions, a natural objective is to minimize the total expected annual delivery cost including the delivery cost of stockouts. Planning the entire annual distribution scheme in advance would, however, be unreliable and prone to many needed adjustments. Moreover, for computational feasibility, actual routing plans can only be computed for a much shorter time period than an annual horizon.

In Bard et al., 1998 we proposed to deal with these issues by introducing a comprehensive decomposition scheme (based on a rolling horizon framework), in which vehicle routing problems are repeatedly solved over a two-week moving period. More specifically, for a given two-week planning horizon, the method first identifies customers that are to be visited during this time. After an adjustment is made to balance daily demand and accommodate weekend days, customers scheduled for the

first week only are routed. The two-week planning horizon is then shifted by a week and the process is repeated. The overall scheme was tested on data sets generated from field experience with a national liquid propane distributor. One of the unique aspects of the week-long routing subproblems we considered is the presence of satellite facilities where vehicles can be reloaded and customer deliveries continued until the closing time is reached. In Bard et al., 1998, we proposed and tested three heuristics for solving these problems. In another paper Bard et al., 1998b, we have proposed and tested a branch and cut algorithm for solving similar routing problems, called VRP with satellite facilities.

This paper is a close companion to Bard et al., 1998 and contains the main justifications behind the identification processes (selection of the customers to be routed on a specific day of a given week) as described above. Interested readers should refer to Bard et al., 1998 for any details and clarifications on the overall approach and on the use of the material developed in this paper. The next section contains mathematical derivations and approximations. Section 2 is concerned with some numerical testing based on Monte Carlo simulation.

1 Expected Total Delivery Cost and Optimal Deliveries for a Customer Over a Given Time Period

1.1 Problem Statement

Our approach attempts to incorporate long-term delivery costs into shorter planning horizons. In doing so, we develop "optimal" *a priori* strategies based on expected consumption of customers, as opposed to reactive strategies based on accurate current knowledge of all local inventories. We are assuming that the central supplier has no reliable monitoring of local inventories. Strategies, based on accurate real-time monitoring of local inventories, are feasible with the adoption of advanced technologies; however, the quantification of the (long-term) benefits associated with the use of such real-time information is far from obvious and an important research topic in

its own, but beyond the scope of the current paper.

We are interested in evaluating the long-term strategies of repeatedly serving a given customer. To do this we assume that the cost for each scheduled delivery at a given customer is a fixed (but customer dependent) given constant b. Note that this is an approximation since the cost due to the service of a customer is in fact dependent on the set of other customers to be served that day (and thus can vary depending on who else is served). In Bard et al., 1998 we provide examples on how to estimate b in practice for a given customer. We assume that the customer's tank capacity is T. We also assume that the daily demand of a customer never exceeds T (i.e., a customer never requires more than one delivery per day).

We define S as the cost of stockout, where $S > b$. Note that we assume S to be fixed (not dependent on the volume of stockout). This is acceptable in case the demand of the product is best approximated by a continuous rate of consumption (as opposed to a discrete lumpy demand); also this assumes that as soon as the tank becomes empty an immediate, and thus costly, special delivery is made. We also assume that if a delivery is scheduled on a specific day and no stockout happened before that day, then no stockout penalty can occur on the delivery day (in other words, deliveries are always made "first thing in the morning").

1.2   Preliminaries

Let us first consider the simple situation of a customer whose daily demand over a given time period is constant and known, say μ ($\mu \leq T$). Assume also that the tank is initially full. Based on the assumption of a constant demand it is clear that the optimal policy (minimizing the total delivery cost) is to fill up the tank every $\lfloor T/\mu \rfloor$ days (i.e., just before the customer runs out of stock). For an n-day period the total delivery cost (not including the initial delivery cost associated with the initial full tank) would then be, for $n \geq 1$:

$$TC_n = \left\lfloor \frac{n-1}{\lfloor T/\mu \rfloor} \right\rfloor b.$$

If over the given time period, the daily demand is not constant, but takes different rates for different subperiods (e.g., due to seasonality), one can divide the overall

period into shorter periods for which constant consumption rate is valid and then patch all solutions together. In the general case for which the daily demand differs for each day of the n-day period in a known fashion, say $\mu_1, \mu_2, \ldots, \mu_n$, an optimal policy is still to serve the customer just before he/she runs out of stock. For the n-day period, let $1 < m_1 < m_2, \cdots < m_{k(n)} \le n$ be the optimal days for a visit. These are defined as follows for $1 \le j \le k(n)$:

$$\sum_{i=1}^{m_j-1} \mu_i \le jT \text{ and } \sum_{i=1}^{m_j} \mu_i > jT.$$

The overall total delivery cost is then

$$TC_n = k(n)b.$$

1.3 Random demand and infinite horizon

Let us now concentrate on the effect of true random fluctuations. Assume that the daily demand corresponds to a stochastic process $(U_i)_{1 \le i \le n}$, where U_i is the demand on day i. Assuming that seasonality and known variations have first been "removed" (see above), it is then natural to assume that the U_i's are i.i.d. random variables with a finite expected value μ. Again, we assume that $Prob(U_i > T) = 0$ and that the tank is initially full. As mentioned in the introduction, we assume that the central supplier cannot monitor what is left in the tank at any time, except at the time of a delivery. This precludes any strategies who would schedule deliveries based on current remaining level of inventories. Under these conditions, we have the following proposition.

Proposition 1. *In case of an infinite horizon, there exists an optimal policy which (i) fills up the tank during each delivery and (ii) following any prior (scheduled or stockout) delivery, plans the next delivery d^* days after. The optimal replenishment interval d^* is a constant chosen so as to minimize the expected daily cost.*

Proof. Since the cost of serving the customer is b for each scheduled visit and S for each stockout visit, irrespective of the quantity served, serving the maximum amount during every delivery cannot be sub-optimal.

6

Since the underlying stochastic process is stationary, and since we have an infinite horizon, delivery days correspond to renewal events associated to a simple stationary renewal process, and an optimal a priori planning strategy (i.e., minimizing the expected daily delivery cost) will always schedule the next delivery a constant number of days after any given delivery. □

1.4 Optimal policies over a finite horizon

1.4.1 Constant replenishment interval strategies

Because (regular or stockout) delivery costs are fixed, irrespective of the amount delivered, we can safely assume that, even for a finite horizon, any valid strategies will fill up the tank at each delivery.

For an n-day period, an intelligent strategy should however adapt itself and choose the next scheduled visit based not only on the tank size and probability distribution of the daily demand of the customer, but also on the number of days remaining until the end of the period. In other words an optimal general policy may lead to non-identical replenishment intervals over a given n-day period.

However, because of the conditions of our problem (e.g. each refill costs the same, irrespective of the amount delivered), we believe that such an adaptive strategy would not bring a significant difference (in the total expected delivery cost) over a simpler (and analytically tractable) non-adaptive (i.e. with constant replenishment intervals) strategy. This is why we will restrict ourselves to the following set of acceptable policies:

Policy Restriction. For a given n-day period, feasible policies will be chosen among the restricted set of constant replenishment interval strategies, i.e., strategies whose scheduled replenishment intervals are constant over the time period.

A consequence of this restriction is that any optimal such policies over a n-day period should lead to a constant replenishment interval $d^*(n)$ which should in fact be independent of n, when n is large enough. We will test this approximation in Section 2.

1.4.2 Optimal constant replenishment interval strategies

Define $ETC_n(d)$ to be the expected total cost over an n-day period of a constant replenishement interval strategy of d days.

As before, we assume an initial full tank. Let us derive a detailed analytical formula for $ETC_n(d)$. Since we assume that any scheduled deliveries are always done "first thing in the morning" and that the daily demand can never exceed the tank size T, there will never be any stockout if $d = 1$.

Assume now that $d > 1$ and that we are at the beginning of day 1 with a full tank (the next scheduled visit is then on day $d + 1$). If on one day i during the next $d - 1$ days there is a stockout then one makes a delivery on that day and the next scheduled visit is updated to $d + i$ and so on. We then have a *renewal* stochastic process where the renewal points are the times of refill (regular or due to stockout).

Let p_j be the probability that a stockout first appears on the j^{th} day following day 1, $1 \leq j \leq d - 1$. This probability is given as

$$
\begin{aligned}
p_j &= Prob(D_j \leq T \text{ and } D_{j+1} > T) \\
&= \int_0^T P(U_{j+1} > T - x) f_{D_j}(x) dx,
\end{aligned}
$$

where $D_k = \sum_{i=1}^k U_i$ is the total consumption from day 1 to k, and $f_{D_k}(x)$ is the density function of the distribution of D_k.

Let $p = p_1 + p_2 + \cdots + p_{d-1}$ be the probability that there is (at least) a stockout in the next $d - 1$ days, and $p_d \triangleq 1 - p$ be the probability of no stockout during that period. We then have the following relationships:

For $d \geq n \geq 2$

$$
ETC_n(d) = \sum_{j=1}^{n-1} p_j (ETC_{n-j}(d) + S),
$$

where $ETC_1(d) = 0$.

For $n \geq d + 1$

$$
ETC_n(d) = \sum_{j=1}^{d-1} p_j (ETC_{n-j}(d) + S) + (1 - p)(ETC_{n-d}(d) + b).
$$

Define $y_n = ETC_n(d)$ (y_n is of course dependent on d, but for notational simplicity

we will not make this dependence explicit). One can rewrite the previous relationships as follows:

$$
y_n = \begin{cases} \sum_{j=1}^{n-1} p_j y_{n-j} + S \sum_{j=1}^{n-1} p_j & \text{if } 2 \leq n \leq d \\[2ex] \sum_{j=1}^{d} p_j y_{n-j} + pS + (1-p)b & \text{if } n \geq d+1 \end{cases} \tag{1}
$$

with the initial condition $y_1 = 0$. From this initial condition and recursions (1), one can numerically compute y_j for all j. However, we would like to obtain an analytical expression for y_n, especially when n is large compared to d. To do so, let us consider the recursions when $n \geq d+1$, i.e.,

$$
y_n - \sum_{j=1}^{d} p_j y_{n-j} = pS + (1-p)b, \tag{2}
$$

and the associated homogeneous equation

$$
y_n - \sum_{j=1}^{d} p_j y_{n-j} = 0. \tag{3}
$$

The corresponding characteristic function (see e.g., Kelley and Peterson, 1991) for the definition of a characteristic function and its use for solving homogeneous linear recursions) is given by

$$
x^d - \sum_{j=1}^{d} p_j x^{d-j} = 0. \tag{4}
$$

Let us characterize the roots of equation (4).

Lemma 1. *1 is a root of (4) with multiplicity one. Any other root r is such that $|r| \leq 1$, where $|r|$ is the absolute value of r if r is real, and the modulus of r if r is complex.*

Proof. Since $\sum_{i=1}^{d} p_i = 1$, 1 is root of (4). One can thus rewrite (4) as

$$
(x-1)[x^{d-1} + (1-p_1)x^{d-2} + \cdots + (1-p_1-p_2-\cdots-p_{d-1})] = 0,
$$

which shows that 1 is a root with multiplicity one only. (The second factor on the left hand side of the equation above is a polynomial of degree $d-1$ of which 1 is not a root. This can be seen by simple substitution.)

9

Any other root r will have $|r| \leq 1$. Indeed if $|r| > 1$, then $|r^d| > |r^i|$ for all $1 \leq i \leq d$. But this is clearly impossible since (4) implies that $|r^d| \leq \sum_{i=1}^{d} p_i |r^{d-i}| \leq \max_{1 \leq i \leq d} |r^{d-i}|$. $\qquad\square$

Since the general solutions y_n of (3) are given by weighted sums of the roots of the characteristic equations raised to the n^{th} power, the implication of Lemma 1 is that $y_n = f(n, d) + \alpha(d)$ where $f(n, d)$ goes to zero exponentially fast as n goes to infinity ($f(n, d)$ is the part corresponding to the weighted sum of roots r such that $r < 1$, each raised to the n^{th} power).

Let us now look at a particular solution to (2). Because 1 is a root of order 1 in the characteristic function, a particular solution will be of the form $y_n = \beta(d)n$, where $\beta(d)$ is a constant. Replacing y_n by $\beta(d)n$ in (2), we get $\beta(d) = (pS + (1 - p)b) / \sum_{j=1}^{d} jp_j$.

We just proved the following result:

Theorem 1. *The expected total cost of filling up a customer's tank using a constant replenishment period strategy of d days over an n-day period ($n \geq d + 1$) is given by*

$$ETC_n(d) = f(n, d) + \alpha(d) + \beta(d)n, \qquad (5)$$

where

$$\beta(d) = \frac{pS + (1 - p)b}{\sum_{j=1}^{d} jp_j}, \qquad (6)$$

and where $f(n, d)$ goes to zero exponentially fast as n goes to infinity.

This leads to our final approximation result.

Corollary 1. *An optimal constant replenishment period strategy of serving a customer over a large n-day period will correspond to choosing d^* so as to minimize $\beta(d)$.*

Proof. From Theorem 1, an optimal policy will correspond to finding d so as to minimize $ETC_n(d) = f(n, d) + \alpha(d) + \beta(d)n$. When n is large enough, $\beta(d)n$ is the main component of the total cost. $\qquad\square$

Notes:

1. From (5) we have $\lim_{n\to\infty} ETC_n(d)/n = \beta(d)$. From an application of a discrete version of the renewal reward theorem (see for example Ross, 1997 for an elementary treatment of renewal reward theory), we know that the expression for $\beta(d)$ should have a simple interpretation. Indeed the numerator of $\beta(d)$ is the expected cost of a delivery (i.e. the expected "cost of a cycle") and the denominator is the expected number of days between deliveries (i.e. the expected "duration of a cycle").

2. The expression given in (6) could naturally be obtained directly by the renewal reward theorem for the infinite horizon case. Our derivation can precisely measure the approximation made from a finite to an infinite horizon period.

3. As an illustration, when $d = 2$, equation (4) has two roots, 1 and $-(1-p)$, and the exact expression for y_n, $n \geq 3$, is given by

$$y_n = \frac{pS - b}{(2-p)^2}\left[-(1-p)\right]^n - \frac{pS + b(1-p)(3-p)}{(2-p)^2} + \frac{pS + (1-p)b}{2-p}n.$$

When n is large, approximating this expression by the linear term $\frac{pS+(1-p)b}{2-p}n$ overestimates the true total cost by a constant $\frac{pS+b(1-p)(3-p)}{(2-p)^2}$.

1.5 Incremental Costs

In the previous section we have shown that the optimal (constant replenishment interval) strategy to serve a given customer corresponds to planning on replenishing the tank every $d^*(n)$ days, where $d^*(n)$ is the value of d that minimizes $ETC_n(d) = f(n,d) + \alpha(d) + \beta(d)n$.

We would like now to estimate the impact on the overall total cost (over an n-day period) of switching a given scheduled delivery of a customer to another day. This incremental cost is an essential input in Ball et al., 1998 used in the second phase of the customer selection (balancing the pre-selected customers among the days of the planning current week).

As before, we assume that we are at the beginning of day 1 with a full tank. Suppose now that on the first visit the optimal a priori strategy is not followed and the customer is scheduled to be replenished on day $j + 1$, $j \neq d^*$. Suppose also

11

that after the first visit, the optimal replenishment period of d^* is resumed. The difference between the expected total cost incurred by following this strategy and the optimal expected total cost will hereafter be called the incremental cost and denoted by c_j. By definition $c_{d^*} = 0$. Let us derive the general incremental cost c_j. Let $X(j)$ be the random variable representing the number of days until the next actual delivery assuming a scheduled replenishment interval of j days. (In case of a stockout $X(j) < j$, otherwise $X(j) = j$.) We then have

Theorem 2. *The incremental cost c_j of serving a customer after j days instead of d^* days is given by*

$$c_j = g(n, j, d^*) + E[X(j)](\beta(j) - \beta(d^*)), \tag{7}$$

where, for any $k \geq 2$,

$$E[X(k)] = k(1 - \sum_{i=1}^{k-1} p_i) + \sum_{i=1}^{k-1} ip_i, \tag{8}$$

and where $g(n, d^*, d^*) = 0$ and $g(n, j, d^*)$ goes to zero exponentially fast as n goes to infinity.

Proof. Let $SUB_n(j, d^*)$ be the expected total cost corresponding to the suboptimal strategy. We have

$$
\begin{aligned}
SUB_n(j, d^*) &= \sum_{i=1}^{j-1} p_i(ETC_{n-i}(d^*) + S) + (1 - \sum_{i=1}^{j-1} p_i)(ETC_{n-j}(d^*) + b) \\
&= \alpha(d^*) + \beta(d^*)n + E[X(j)](\beta(j) - \beta(d^*)) + h(n, j, d^*),
\end{aligned}
$$

where $h(n, j, d^*)$ (sum of the terms $f(n - i, d^*)$ as defined in (5)) goes to zero exponentially fast as n goes to infinity. By writing $c_j = SUB_n(j, d^*) - ETC_n(d^*)$ and defining $g(n, j, d^*) = h(n, j, d^*) - f(n, d^*)$, we get the result. □

Notes:

1. When n is large enough, the incremental cost given in (7) can be approximated as follows:

$$c_j \approx E[X(j)](\beta(j) - \beta(d^*)). \tag{9}$$

12

2. Note that if the expected total cost $ETC_n(d)$ is approximated by its linear component $\beta(d)n$ to start with, then $\beta(d)$ can be interpreted as the daily cost rate associated with the strategy of filling the customer every d days. Under this interpretation one can then approximate $ETC_n(j, d^*)$ as follows:

$$SUB_n(j, d^*) \approx \beta(j)E[X(j)] + \beta(d^*)(n - E[X(j)]).$$

The first component is the expected number of days with a suboptimal cost of $\beta(j)$ per day, and the second term corresponds to the remaining total expected cost under the optimal policy. This leads to the same approximation of c_j as given in (9).

2 Monte-Carlo Simulation and Testing

The purpose of this section is to use Monte-Carlo simulation to (i) validate one consequence of Policy Restriction of Section 1.4.1, and (ii) validate the limiting approximation formula (9) for the computation of the incremental costs.

For the first testing, the input of the simulation is the total number of days n for the horizon and the frequency of visit d (in days). The output is the total delivery cost. For a fixed n, we vary d to find the optimal a priori frequency $d^*(n)$. The Policy Restriction implies that $d^*(n)$ should be independent of n, when n is large. The code is a fixed-time-increment simulation in which time is advanced day by day. For each day, one first checks whether there is a scheduled delivery. If yes, the level in the tank is set back to T and the total cost is incremented by b. Then the consumption of the customer is randomly generated from a truncated normal distribution and is deducted from the current level of the tank. If the new tank level is negative then an emergency visit is scheduled, the level of the tank is set back to T and the total cost is incremented by S.

The other testing (incremental cost) is done via a similar simulation, except that the input is n, $d^*(n)$ (as computed above), and j (which represents the initial number of days before the first scheduled visit). In both cases, for a given input, we run 10,000 independent simulations and average the output to obtain an approximation of the expected value under consideration.

We assume that the customer's consumption rate follows a truncated normal distribution $N(\mu, \sigma)$, where negative values of the normal distribution are ignored. We initially assume that $T = 1000$, $\mu = 40$, $\sigma = 10$, $b = 16.27$, and $S = 68$. Note that $T/\mu = 25$ would represent the number of days between each refill, i.e., the optimal a priori frequency if the consumption rate were a constant equal to μ.

2.1   Results for Test 1

By varying the time period n from 0.5 years, 1 year, 2 years, 5 years to 10 years, we found that the optimal a priori policy was always to refill the customer every $d^* = 23$ days. We repeated the experiments with different σ's, from $\sigma = 5$ to $\sigma = 30$, and reached the same conclusion, i.e., that d^* was independent of the time period n.

For $n = 10$ years, Figure 1 presents the simulated expected total cost for d varying from 15 to 30, as well as the calculated expected total cost based on Theorem 1 and approximated by $\beta(d)n$ (i.e., with only the linear part). The fit between the two curves is almost perfect. To apply Theorem 1, the probability p_i's must be evaluated. The easiest way to do this is via Monte-Carlo simulation. To achieve a satisfactory level of precision, about 100,000 independent simulations were conducted for each p_j.

Again we repeated the experiments with $n=10$ and different σ's, ranging from $\sigma = 5$ to $\sigma = 30$. To compare the simulated results with the analytic results, we defined a fitness ratio r as follows. Suppose x_i $(i = 1, \cdots, n)$ is a set of numbers, and z_i $(i = 1, \cdots, n)$ another set, then

$$r = \frac{\sum_{i=1}^{n} |x_i - z_i|}{\sum_{i=1}^{n} x_i/n}$$

is our measure of fitness. Note that the numerator is simply the L_1 measure of discrepancy between the two sets of numbers and the denominator allows for a measure independent of units. The smaller r is the better the fitness between the data sets x and z. In Table 1, we present six cases with σ varying from 5 to 30. We give the minimal expected total cost for each case, using d^* to represent the best frequency and ETC_{min} to indicate the minimal expected total cost. The data show that for all six cases we get the same d^* with both simulation and the analytic method. The fitness ratio r does not exceed 1.07, which indicate a good fit between the two approaches.
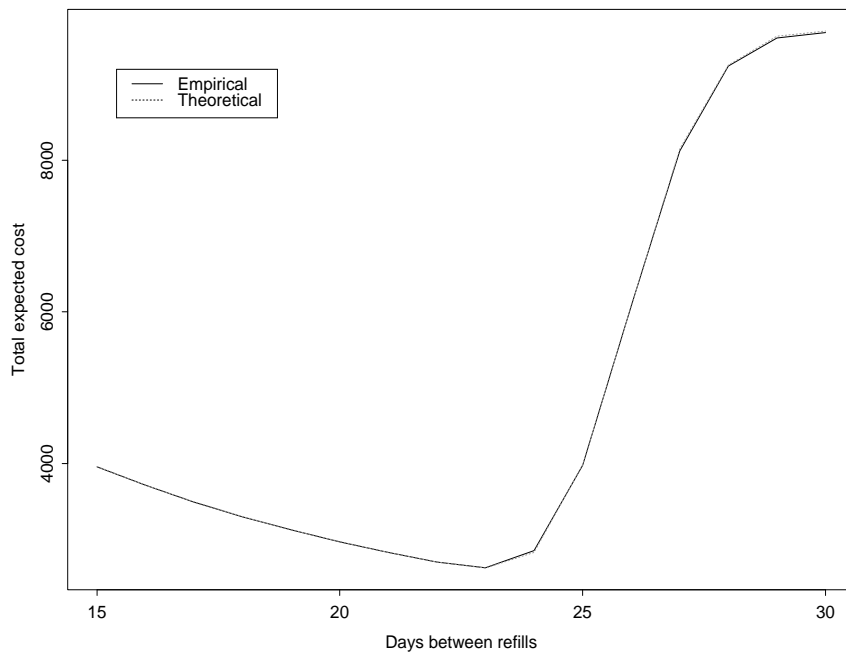
Figure 1: Annual expected total cost

Cases	Analytical Calculation, x's		Numerical Simulation, z's		Fitness Ratio
σ	d^*	ETC_{min}	d^*	ETC_{min}	r
5	24	2477.43	24	2477.67	0.14
10	23	2618.62	23	2624.04	0.16
15	22	2790.16	22	2791.82	0.23
20	21	2969.13	21	2996.05	0.28
25	19	3210.13	19	3222.46	0.57
30	18	3459.64	18	3493.97	1.07

Table 1: Comparison of numerical results and analytic results

2.2 Results for Test 2

By assuming that $T = 1000$, $\mu = 40$, $\sigma = 10$, $b = 16.27$, and $S = 68$, the optimal a priori frequency is to serve the customer every $d^* = 23$ days (as indicated in Figure 1 above).

In Figure 2, we present, for $n = 10$, the simulated expected incremental costs corresponding to j varying from 15 to 30 (the "empirical" curve). Also we have used (9) to calculate the analytical (theoretical) approximated values for comparison purposes. The fit between simulated and calculated is almost perfect. Additional testings for varying σ yielded similar results.
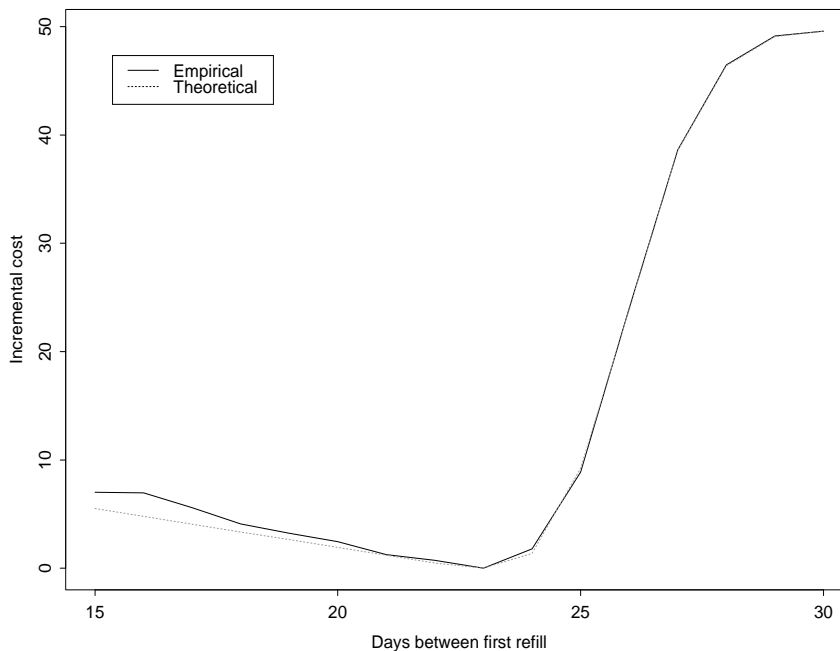

Figure 2: Incremental cost

3   Concluding Remarks

In this paper we have presented the justifications behind the derivations of the incremental costs as used in the rolling horizon framework of Bard et al., 1998. The

essence of our contributions lies in the ability to reduce the problem from an annual time base to a bi-weekly rolling planning period via the approximations of delivery costs.

References

S. Anily, and A. Federgruen, "One Warehouse Multiple Retailer Systems with Vehicle Routing Costs," *Management Science* **36**, 92-114 (1990).

S. Anily, and A. Federgruen, "Two-echelon Distribution Systems with Vehicle Routing Costs and Central Inventories," *Operations Research* **41**, 37-47 (1993).

J. Bard, L. Huang, M. Dror, and P. Jaillet, "A Branch and Cut Algorithm for the VRP with Satellite Facilities," *IIE Transactions* **30**, 831-834 (1998b).

J. Bard, L. Huang, P. Jaillet, and M. Dror, "A Decomposition Approach to the Inventory Routing Problem with Satellite Facilities," *Transportation Science* **32**, 189-203 (1998).

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**, 4-23 (1983).

E. Beltrami, and L. Bodin, "Networks and Vehicle Routing for Municipal Waste Collection," *Networks* **4**, 65-94 (1974).

O. Berman and R. Larson, "Deliveries in an Inventory/Routing Problem Using Stochastic Dynamic Programming," Working Paper, Massachusetts Institute of Technol-

ogy, Boston, 1999.

A. Chan, A. Federgruen, and D. Simchi-Levi, "Probabilistic Analyses and Practical Algorithms for Inventory-Routing Models," *Operations Research* **46**, 96-107 (1998).

A. Chan, A. Muriel, and D. Simchi-Levi, "Integrating Inventory and Transportation in Supply Chain Management," Working Paper, Department of IE&MS, Northwestern University, Evanston, IL, 1998.

T. Chien, A. Balakrishnan, and R.T. Wong, "An Integrated Inventory Allocation and Vehicle Routing Problem," *Transportation Science* **23**, 67-76 (1989).

N. Christofides, and J. Beasley, "The Period Routing Problem," *Networks* **14**, 237-256 (1984).

M. Dror, and M. Ball, "Inventory/Routing: Reduction from an Annual to a Short-Period Problem," *Naval Research Logistics* **34**, 891-905 (1987).

M. Dror, M. Ball, and B. Golden, "A Computational Comparison of Algorithms for the Inventory Routing Problem," *Annals of Operations Research* **4**, 3-23 (1985).

A. Federgruen, and D. Simchi-Levi, "Analysis of Vehicle Routing and Inventory-Routing Problems," in *Handbooks in Operations Research and Management Science*, M. Ball et al. (eds), 297-371, North Holland, Amsterdam, 1995.

A. Federgruen, and P. Zipkin, "A Combined Vehicle Routing and Inventory Allocation Problem," *Operations Research* **32**, 1019-1037 (1984).

M. Fisher, A. Greenfield, R. Jaikumar, and P. Kedia, "Real Time Scheduling of a Bulk Delivery Fleet: Practical Application of Lagrangian Relaxation," Working Paper 82-10-11, Department of Decision Science, The Wharton School, University of Pennsylvania, Philadelphia, 1983.

M. Gaudioso, and G. Paletta, "A Heuristic for the Periodic Vehicle Routing Problem,"

Transportation Science **26**, 86-92 (1992).

B. Golden, A. Assad, and R. Dahl, "Analysis of a Large Scale Vehicle Routing Problem with an Inventory Component," *Large Scale Systems* **7**, 181-190 (1984).

W. Kelley and A. Peterson, *Difference Equations: An Introduction with Applications*, Academic Press, Boston (1991).

A. Kleywegt, V. Nori, and M. Savelsbergh, "A Computational Approach for the Inventory Routing Problem," Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, 1998.

A. Kleywegt, V. Nori, and M. Savelsbergh, "The Stochastic Inventory Routing Problem with Direct Deliveries," Working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, 2000.

A. Kumar, L. Schwarz, and J. Ward, "Risk-Pooling Along a Fixed Delivery Route Using a Dynamic Inventory-Allocation Policy," *Management Science* **41**, 344-362 (1995).

M. Reiman, R. Rubio, and L. Wein, "Heavy Traffic Analysis of the Dynamic Stochastic Inventory-Routing Problem," *Transportation Science* **33**, 361-381 (1999).

R. Rubio, Dynamic-Stochastic Vehicle Routing and Inventory Problem, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, 1995.

R. Russell, and W. Igo, "An Assignment Routing Problem," *Networks* **9**, 1-17 (1979).

S. Ross, *Introduction to Probability Models*, Academic Press, Inc., Englewood Cliffs, New Jersey (1997).

P. Trudeau, and M. Dror, "Stochastic Inventory Routing: Route Design with Stockouts and Route Failures," *Transportation Science* **26**, 171-184 (1992).

I. Webb, and R. Larson, "Period and Phase of Customer Replenishment: A New Approach to the Strategic Inventory/Routing Problem," *European Journal of Operational Research* **85**, 132-148 (1995).

TABLE AND FIGURE LEGENDS

Table 1: Comparison of numerical results and analytic results

Figure 1: Annual expected total cost

Figure 2: Incremental cost