THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Replenishment planning for stochastic inventory systems with shortage cost

Citation for published version:

Rossi, R, Prestwich, S, Armagan Tarim, S & Hnich, B 2007, Replenishment planning for stochastic inventory systems with shortage cost. in P Van Hentenryck & L Wolsey (eds), *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems: 4th International Conference, CPAIOR 2007, Brussels, Belgium, May 23-26, 2007. Proceedings.* vol. 4510 LNCS, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), Springer-Verlag GmbH, pp. 229-243. https://doi.org/10.1007/978-3-540-72397-4_17

Digital Object Identifier (DOI):

[10.1007/978-3-540-72397-4_17](https://doi.org/10.1007/978-3-540-72397-4_17)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Replenishment Planning for Stochastic Inventory Systems with Shortage Cost

Roberto Rossi,¹ S. Armagan Tarim,² Brahim Hnich,³ and Steven Prestwich¹

Cork Constraint Computation Centre, University College, Cork, Ireland¹
{r.rossi,s.prestwich}@4c.ucc.ie

Department of Management, Hacettepe University, Turkey²
armagan.tarim@hacettepe.edu.tr

Faculty of Computer Science, Izmir University of Economics, Turkey³
brahim.hnich@ieu.edu.tr

Abstract. One of the most important policies adopted in inventory control is the (R,S) policy (also known as the “replenishment cycle” policy). Under the non-stationary demand assumption the (R,S) policy takes the form (R_n, S_n) where R_n denotes the length of the n^{th} replenishment cycle, and S_n the corresponding order-up-to-level. Such a policy provides an effective means of damping planning instability and coping with demand uncertainty. In this paper we develop a CP approach able to compute optimal (R_n, S_n) policy parameters under stochastic demand, ordering, holding and shortage costs. The convexity of the cost-function is exploited during the search to compute bounds. We use the optimal solutions to analyze the quality of the solutions provided by an approximate MIP approach that exploits a piecewise linear approximation for the cost function.

1 Introduction

Much of the inventory control literature concerns the computation of optimal replenishment policies under demand uncertainty. One of the most important policies adopted is the (R,S) policy (also known as the *replenishment cycle* policy). In this policy a replenishment is placed every R periods to raise the inventory position to the order-up-to-level S . This provides an effective means of damping planning instability (deviations in planned orders, also known as *nervousness*) and coping with demand uncertainty. As pointed out by Silver et al. ([8], pp. 236–237), (R,S) is particularly appealing when items are ordered from the same supplier or require resource sharing. In these cases all items in a coordinated group can be given the same replenishment period. Periodic review also allows a reasonable prediction of the level of the workload on the staff involved, and is particularly suitable for advanced planning environments. For these reasons (R,S) is a popular inventory policy.

An important class of stochastic production/inventory control problems assumes a non-stationary demand process. Under this assumption the (R,S) policy takes the non-stationary form (R_n, S_n) where R_n denotes the length of the n^{th}

replenishment cycle and S_n the corresponding order-up-to-level. To compute the *near* optimal policy parameters for (R_n, S_n) , Tarim and Kingsman [4] propose a mixed integer programming (MIP) formulation using a piecewise linear approximation to a complex cost function.

This paper focuses on the work of Tarim and Kingsman, in which a finite-horizon, single-installation, single-item (R_n, S_n) policy is addressed. They assume a fixed procurement cost each time a replenishment order is placed, whatever the size of the order, and a linear holding cost on any unit carried over in inventory from one period to the next. Instead of employing a service level constraint — the probability that at the end of every time period the net inventory will not be negative is at least a certain value (see Tarim and Kingsman [3] for (R_n, S_n) under a service level constraint) — their model employs a penalty cost scheme. They propose a certainty-equivalent formulation of the above problem in the form of a MIP model. So far no CP approach has been proposed for (R_n, S_n) under a penalty cost. In fact, as shown in [4], the cost structure is complex in this case and it differs significantly from the one under a service level constraint. In [2] the authors proposed a CP model under a service level constraint. In this paper it was shown that not only CP is able to provide a more compact formulation than the MIP one, but that it is also able to perform faster and to take advantage of dedicated pre-processing techniques that reduce the size of decision variable domains. Moreover dedicated cost-based filtering techniques were proposed in [1] for the same model, these techniques are able to improve performances of several orders of magnitude.

In this paper, we give an *exact* formulation of the (R_n, S_n) inventory control problem via constraint programming, instead of employing a piecewise linear approximation to the total expected cost function. This exact CP formulation provides an optimal solution to (R, S) policy. Our contribution is two-fold: we can now obtain provably optimal solutions, and we can gauge the accuracy of the piecewise linear approximation proposed by Tarim and Kingsman.

2 Problem definition and (R_n, S_n) policy

The demand d_t in period t is considered to be a normally distributed random variable with known probability density function (PDF) $g_t(d_t)$, and is assumed to occur instantaneously at the beginning of each period. The mean rate of demand may vary from period to period. Demands in different time periods are assumed to be independent. A fixed holding cost h is incurred on any unit carried over in inventory from one period to the next. Demands occurring when the system is out of stock are assumed to be back-ordered and satisfied as soon as the next replenishment order arrives. A fixed shortage cost s is incurred for each unit of demand that is back-ordered. A fixed procurement (ordering or set-up) cost a is incurred each time a replenishment order is placed, whatever the size of the order. In addition to the fixed ordering cost, a proportional direct item cost v is incurred. For convenience, and without loss of generality, the initial inventory level is set to zero and the delivery lead-time is not incorporated. It is assumed

that negative orders are not allowed, so that if the actual stock exceeds the order-up-to-level for that review, this excess stock is carried forward and does not return to the supply source. However, such occurrences are regarded as rare events and accordingly the cost of carrying the excess stock is ignored. The above assumptions hold for the rest of this paper.

The general multi-period production/inventory problem with stochastic demands can be formulated as finding the timing of the stock reviews and the size of non-negative replenishment orders, X_t in period t , minimizing the expected total cost over a finite planning horizon of N periods:

$$\min E\{TC\} = \int_{d_1} \int_{d_2} \dots \int_{d_N} \sum_{t=1}^N (a\delta_t + vX_t + hI_t^+ + sI_t^-) g_1(d_1) \dots g_N(d_N) d(d_1) \dots d(d_N) \quad (1)$$

subject to

$$X_t > 0 \Rightarrow \delta_t = 1 \quad (2)$$

$$I_t = \sum_{i=1}^t (X_i - d_i) \quad (3)$$

$$I_t^+ = \max(0, I_t) \quad (4)$$

$$I_t^- = -\min(0, I_t) \quad (5)$$

$$X_t, I_t^+, I_t^- \in \mathbb{Z}^+ \cup \{0\}, \quad I_t \in \mathbb{Z}, \quad \delta_t \in \{0, 1\} \quad (6)$$

for $t = 1 \dots N$, where

d_t : the demand in period t , a normal random variable with PDF $g_t(d_t)$,

a : the fixed ordering cost,

v : the proportional direct item cost,

h : the proportional stock holding cost,

s : the proportional shortage cost,

δ_t : a $\{0,1\}$ variable that takes the value of 1 if a replenishment occurs in period t and 0 otherwise,

I_t : the inventory level at the end of period t , $-\infty < I_t < +\infty$, $I_0 = 0$

I_t^+ : the excess inventory at the end of period t carried over to the next period,
 $0 \leq I_t^+$,

I_t^- : the shortages at the end of period t , or magnitude of negative inventory
 $0 \leq I_t^-$,

X_t : the replenishment order placed and received in period t , $X_t \geq 0$.

The proposed non-stationary (R,S) policy consists of a series of review times and associated order-up-to-levels. Consider a review schedule which has m reviews over the N period planning horizon with orders arriving at $\{T_1, T_2, \dots, T_m\}$, $T_j > T_{j-1}$. For convenience $T_1 = 1$ is defined as the start of the planning horizon and $T_{m+1} = N + 1$ the period immediately after the end of the horizon.

In [3], the decision variable X_{T_i} is expressed in terms of a new variable $S_t \in \mathbb{Z}$, where S_t may be interpreted as the opening stock level for period t , if there is no replenishment in this period (i.e. $t \neq T_i$ and $X_t = 0$) and the order-up-to-level for the i -th review period T_i if there is a replenishment (i.e. $t = T_i$ and $X_t > 0$). According to this transformation the expected cost function, Eq. (1), is written as the summation of m intervals, T_i to T_{i+1} for $i = 1, \dots, m$, defining $D_{t_1, t_2} = \sum_{j=t_1}^{t_2} d_j$:

$$\min E\{TC\} = \sum_{i=1}^m \left(a\delta_{T_i} + \sum_{t=T_i}^{T_{i+1}-1} E\{C_{T_i, t}\} \right) + vI_N + v \int_{D_{1,N}} D_{1,N} \times g(D_{1,N}) d(D_{1,N}), \quad (7)$$

The term $v \int_{D_{1,N}} D_{1,N} \times g(D_{1,N}) d(D_{1,N})$ is constant and can therefore be ignored in the optimization model. $E\{C_{T_i, t}\}$ of Eq. (7) is defined as:

$$\int_{-\infty}^{S_{T_i}} h(S_{T_i} - D_{T_i, t}) g(D_{T_i, t}) d(D_{T_i, t}) - \int_{S_{T_i}}^{\infty} s(S_{T_i} - D_{T_i, t}) g(D_{T_i, t}) d(D_{T_i, t}). \quad (8)$$

As stated in [4], $E\{C_{T_i, t}\}$ is the expected cost function of a single-period inventory problem where the single-period demand is $D_{T_i, t}$. Since S_{T_i} may be interpreted as the order-up-to-level for the i -th review period T_i and $S_{T_i} - D_{T_i, t}$ is the end of period inventory for the “single-period” with demand $D_{T_i, t}$, the expected total subcosts $E\{C_{T_i, t}\}$ are the sums of single-period inventory costs where the demands are the cumulative demands over increasing periods. By dropping the T_i and t subscripts in Eq. (8) we obtain the following well-known expression for the expected total cost of a single-period newsvendor problem:

$$E\{TC\} = h \int_{-\infty}^S (S - D) g(D) d(D) - s \int_S^{\infty} (S - D) g(D) d(D) \quad (9)$$

where we consider two cost components: holding cost on the positive end of period inventory and shortage cost for any back-ordered demand. Let $G(\cdot)$ be the cumulative distribution function of the demand in our single-period newsvendor problem. A known result in inventory theory (see [17]) is convexity of Eq. (9). The so-called *Critical Ratio*, $\frac{s}{s+h}$, can be seen as the service level β (i.e. probability that at the end of the period the inventory level is non-negative) provided when we fix the order-up-to-level S to the optimal value S^* that minimizes expected holding and shortage costs (Eq. (9)). By assuming $G(\cdot)$ to be strictly increasing, we can compute the optimal order-up-to-level as $S^* = G^{-1}\left(\frac{s}{s+h}\right)$.

2.1 Stochastic cost component in single-period newsvendor

We now aim to characterize the cost of the policy that orders S^* units to meet the demand in our single-period newsvendor problem. Since the demand

D is assumed to be normal with mean μ and standard deviation σ , then we can write $D = \mu + \sigma Z$, where Z is a standard normal random variable. Let $\Phi(z) = \Pr(Z \leq z)$ be the cumulative distribution function of the standard normal random variable. Since $\Phi(\cdot)$ is strictly increasing, $\Phi^{-1}(\cdot)$ is uniquely defined. Let $z_\beta = \Phi^{-1}(\beta)$, since $\Pr(D \leq \mu + z_\beta \sigma) = \Phi(z_\beta) = \beta$, it follows that $S^* = \mu + z_\beta \sigma$. The quantity z_β is known as the safety factor and $S^* - \mu = z_\beta \sigma$ is known as the safety stock. It can be shown [17] that

$$\int_{S^*}^{\infty} (S^* - D)g(D)d(D) = E\{D - S^*\}^+ = \sigma E\{Z - z_\beta\}^+ = \sigma[\phi(z_\beta) - (1 - \beta)z_\beta] \quad (10)$$

where $\phi(\cdot)$ is the PDF of the standard normal random variable. Let $E\{S^* - D\}^+ = \int_{-\infty}^{S^*} (S - D)g(D)d(D)$, it follows

$$\begin{aligned} E\{TC(S^*)\} &= h \cdot E\{S^* - D\}^+ + s \cdot E\{D - S^*\}^+ = \\ &= h \cdot (S^* - \mu) + (h + s)E\{D - S^*\}^+ = \\ &= h z_\beta \sigma + (h + s)\sigma E\{Z - z_\beta\}^+ = \\ &= h z_\beta \sigma + (h + s)\sigma[\phi(z_\beta) - (1 - \beta)z_\beta] = \\ &= (h + s)\sigma\phi(z_\beta) \end{aligned} \quad (11)$$

The last expression $(h + s)\sigma\phi(z_\beta)$ holds only for the optimal order-up-to-level S^* that provides the service level $\beta = \left(\frac{s}{s+h}\right)$ computed from the *critical ratio* (CR). Instead, expression

$$h z_\alpha \sigma + (h + s)\sigma[\phi(z_\alpha) - (1 - \alpha)z_\alpha] \quad (12)$$

can be used to compute the expected total cost for any given level S such that $\alpha = \Phi\left(\frac{S - \mu}{\sigma}\right)$. In Fig. 1 we plot this cost for a particular instance as a function of the opening inventory level S .

Fig. 1. Single-period holding and shortage cost as a function of the opening inventory level S . The demand is normally distributed with mean 200 and standard deviation 20. Holding cost is 1, shortage cost is 10.

2.2 Stochastic cost component in multiple-period newsvendor

The considerations in the former sections refer to a single-period problem, but they can be easily extended to a replenishment cycle $R(i, j)$ that covers the period span $i, \dots, j$. The demand in each period is normally distributed with PDF $g_i(d_i), \dots, g_j(d_j)$. The cost for the multiple periods' replenishment cycle, when ordering costs are neglected, can be expressed as

$$E\{TC\} = \sum_{k=i}^j \left(h \int_{-\infty}^S (S - d_{i,k}) g_{i,k}(d_{i,k}) d(d_{i,k}) - s \int_S^{\infty} (S - d_{i,k}) g_{i,k}(d_{i,k}) d(d_{i,k}) \right) \quad (13)$$

Since demands are independent and normally distributed in each period, the term $g_{i,j}(d_{i,j})$ (that is the p.d.f. for the overall demand over the period span $\{i, \dots, j\}$) can be easily computed (see [12]) once the demand in each period $d_i, \dots, d_j$ are known. It is easy to apply the same rule as before and compute the second derivative of this expression:

$$\frac{d^2}{dS^2} E\{TC\} = \sum_{k=i}^j (h \cdot g_{i,k}(S) + s \cdot g_{i,k}(S)) \quad (14)$$

which is again a positive function of S , since $g_{i,k}(S)$ are PDFs and both holding and shortage cost are assumed to be positive. The expected cost of a single replenishment cycle therefore remains convex in S regardless of the periods covered. Unfortunately it is not possible to compute the CR as before, using a simple algebraic expression to obtain the optimal S^* which minimizes the expected cost. But since the cost function is convex, it is still possible to compute S^* efficiently. Eq. (12) can be extended in the following way to compute the cost for the replenishment cycle $R(i, j)$ as a function of the opening inventory level S :

$$\sum_{k=i}^j (h z_{\alpha(i,k)} \sigma_{i,k} + (h + s) \sigma_{i,k} [\phi(z_{\alpha(i,k)}) - (1 - \alpha(i,k)) z_{\alpha(i,k)}]) \quad (15)$$

where $G_{i,k}(S) = \alpha(i, k)$ and $z_{\alpha(i,k)} = \Phi^{-1}(\alpha(i, k))$. Therefore we have $j - i + 1$ cost components: the holding and shortage cost at the end of period $i, i + 1, \dots, j$. In Fig. 2 we plot this cost for a particular instance as a function of the opening inventory level S . For each possible replenishment cycle we can efficiently compute the optimal S^* that minimizes such a cost function, using gradient based methods for convex optimization such as Newton's method. Notice that the complete expression for the cost of replenishment cycles that start in period $i \in \{1, \dots, N\}$ and end in period N is

$$\sum_{k=i}^N (h z_{\alpha(i,k)} \sigma_{i,k} + (h + s) \sigma_{i,k} [\phi(z_{\alpha(i,k)}) - (1 - \alpha(i,k)) z_{\alpha(i,k)}]) + v \left(S - \sum_{k=i}^N d_k \right) \quad (16)$$

Fig. 2. Three periods holding and shortage cost as a function of the opening inventory level S . The demand is normally distributed in each period with mean respectively 150, 100, 200, the coefficient of variation is 0.1. Holding cost is 1, shortage cost is 10.

In fact for this set of replenishment cycles we must also consider the unit cost component. Once S^* is known, by subtracting the expected demand over the replenishment cycle we obtain the optimal expected buffer stock level $b(i, j)$ required for such a replenishment cycle in order to minimize holding and shortage cost. Notice that every other choice for buffer stock level will produce a higher expected total cost for $R(i, j)$.

An *upper bound* for the value of the opening inventory level in each period $t \in \{1, \dots, N\}$ can be computed by considering the buffer stock $b(1, N)$ required to optimize the convex cost of a single replenishment cycle $R(1, N)$ that covers the whole planning horizon. Then for each period $t \in \{1, \dots, N\}$, $\max(S_t) = \sum_t^N \tilde{d}_t + b(1, N)$. A *lower bound* for the value of the expected closing inventory level in each period $t \in \{1, \dots, N\}$, i.e. opening inventory level minus expected demand, can be computed by considering every possible buffer stock $b(i, j)$ required to optimize the convex cost of a single replenishment cycle $R(i, j)$, independently of the other cycles that are planned. The lower bound will be the minimum value among all these possible buffer values for $j \in \{1, \dots, N\}$ and $i \in \{1, \dots, j\}$.

3 Deterministic equivalent CP formulation

Building on the considerations above it is easy to construct a *deterministic equivalent* CP formulation for the non-stationary (R_n, S_n) policy under stochastic demand, ordering cost, holding and shortage cost. (For a detailed discussion on deterministic equivalent modeling in stochastic programming see [14]).

In order to correctly compute the expected total cost for a replenishment cycle $R(i, j)$ with opening inventory level S_i , we must build a special-purpose constraint *objConstraint*($\cdot$) that dynamically computes such a cost by means of

an extended version of Eq. (15)

$$C(S_i, i, j) = a + \sum_{k=i}^j (hz_{\alpha(i,k)}\sigma_{i,k} + (h+s)\sigma_{i,k}[\phi(z_{\alpha(i,k)}) - (1-\alpha(i,k))z_{\alpha(i,k)}]) \quad (17)$$

that considers the ordering cost. Then the expected total cost for a certain replenishment plan will be computed as the sum of all the expected total costs for replenishment cycles in the solution, plus the respective ordering costs. $objConstraint(\cdot)$ also computes the optimal expected buffer stock level $b(i, j)$ for every replenishment cycle $R(i, j)$ identified by a partial assignment for $\delta_{k \in \{1, \dots, N\}}$ variables. A *deterministic equivalent* CP formulation is

$$\min E\{TC\} = C \quad (18)$$

subject to

$$objConstraint(C, \tilde{I}_1, \dots, \tilde{I}_N, \delta_1, \dots, \delta_N, d_1, \dots, d_N, a, h, s) \quad (19)$$

and for $t = 1 \dots N$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \quad (20)$$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} > 0 \Rightarrow \delta_t = 1 \quad (21)$$

$$\tilde{I}_t \in \mathbf{Z}, \quad \delta_t \in \{0, 1\} \quad (22)$$

Each decision variable $\tilde{I}_t$ represents the expected closing inventory level at the end of period t ; bounds for the domains of these variables can be computed as explained above. Each $\tilde{d}_t$ represents the expected value of the demand in a given period t according to its PDF $g_t(d_t)$. The binary decision variables δ_t state whether a replenishment is fixed for period t ($\delta_t = 1$) or not ($\delta_t = 0$).

Eq. (20) enforces a no-buy-back condition, which means that received goods cannot be returned to the supplier. As a consequence of this the expected inventory level at the end of period t must be no less than the expected inventory level at the end of period $t-1$ minus the expected demand in period t . Eq. (21) expresses the replenishment condition. We have a replenishment if the expected inventory level at the end of period t is greater than the expected inventory level at the end of period $t-1$ minus the expected demand in period t . This means that we received some extra goods as a consequence of an order.

The objective function (18) minimizes the expected total cost over the given planning horizon. $objConstraint(\cdot)$ dynamically computes buffer stocks and it assigns to C the expected total cost related to a given assignment for replenishment decisions, depending on the demand distribution in each period and on the given combination for problem parameters a, h, s . In order to propagate this constraint we wait for a partial assignment involving $\delta_t, t = 1, \dots, N$ variables. In particular we look for an assignment where there exists some i s.t. $\delta_i = 1$, some $j > i$ s.t. $\delta_{j+1} = 1$ and for every $k, i < k \leq j, \delta_k = 0$. This will

Fig. 3. A replenishment cycle $R(i, j)$ is identified by the current partial assignment for δ_i variables.

uniquely identify a replenishment cycle $R(i, j)$ (Fig. 3). There may be more replenishment cycles associated to a partial assignment. If we consider each $R(i, j)$ identified by the current assignment, it is easy to minimize the convex cost function already discussed, and to find the optimal expected buffer stock $b(i, j)$ for this particular replenishment cycle independently on the others. By doing this for every replenishment cycle identified, two possible situations may arise: the buffer stock configuration obtained satisfies every inventory conservation constraint (Eq. (20)), or for some couple of subsequent replenishment cycles this constraint is violated (Fig. 4). Therefore we observe an expected negative order quantity. If the latter situation arises we can adopt a fast convex optimization

Fig. 4. The expected total cost of both replenishment cycles is minimized, but the inventory conservation constraint is violated between $R(i, k)$ and $R(k+1, j)$

procedure to compute a feasible buffer stock configuration with minimum cost. The key idea is to identify two possible limit situations: we increase the opening inventory level of the second cycle, thus incurring a higher overall cost for it, to preserve optimality of the first cycle (Fig. 5 - a). Or we decrease the buffer stock of the first replenishment cycle, thus incurring a higher overall cost for it, to preserve optimality of the second cycle cost (Fig. 5 - b). A key observation is that, when negative order quantity scenarios arise, at optimality the closing inventory levels of the first and the second cycle lie in the interval delimited by the two situations described. This directly follows from the convexity of both the cost functions. Moreover the closing inventory level of the first cycle must be equal to the opening inventory level of the second cycle. In fact, if this does not hold, then either the first cycle has a closing inventory level higher than the opening inventory level of the second cycle and the solution is not feasible (Fig.

Fig. 5. Feasible limit situations when negative order quantity scenarios arise

6 - a), or the first cycle has a closing inventory level smaller than the opening inventory level of the second cycle. In the latter case we can obviously decrease the overall cost by choosing a smaller opening inventory level for the second cycle (Fig. 6 - b). The algorithm for computing optimal buffer stock configura-

Fig. 6. Infeasible (a) and suboptimal (b) plans realized when the opening inventory level of the second cycle doesn't equate the closing inventory level of the first cycle

tions in presence of negative order quantity scenarios simply exploits the linear dependency between opening inventory level of the second cycle and closing inventory level of the first cycle. Due to this dependency the overall cost is still convex in $b(i, k)$ (or equivalently in $b(k+1, j)$, since they are linearly dependent) and we can apply any convex optimization technique to find the optimal buffer stock configuration. Notice that this reasoning still holds in a recursive process. Therefore we can optimize buffer stock for two subsequent replenishment cycles, then we can treat these as a new single replenishment cycle, since their buffer stocks are linearly dependent, and repeat the process in order to consider the next replenishment cycle if a negative order quantity scenario arises.

Once buffer stocks are known we can apply Eq. (17) to the opening inventory level $S_i = \tilde{d}_i + \dots + \tilde{d}_j + b(i, j)$ and compute the cost $C(S_i, i, j)$ associated to a given replenishment cycle. Since the cost function in Eq. (17) is convex and we handle negative order quantity scenarios, a lower bound for the expected total cost associated to the current partial assignment for δ_i , $t = 1, \dots, N$ variables is now given by the sum of all the cost components $C(S_i, i, j)$, for each replenishment cycle $R(i, j)$ identified by the assignment. Furthermore this bound is tight

Period	1	2	3	4	5	6	7	8
$\bar{d}_t$	200	100	70	200	300	120	50	100

Table 1. Expected demand values

if all the δ_t variables have been assigned. *objConstraint*($\cdot$) exploits this property in order to incrementally compute a lower bound for the cost of the current partial assignment for δ_t variables. When every δ_t variable is ground, since such a lower bound becomes tight, buffer stocks computed for each replenishment cycle identified can be assigned to the respective I_t variables. Finally, in order to consider the unit variable cost v we must add the term $v \cdot I_N$ to the cycle cost $C(S_i, i, N)$ for $i \in \{1, \dots, N\}$. Therefore the complete expression for the cost of replenishment cycles that start in period $i \in \{1, \dots, N\}$ and end in period N is:

$$C(S_i, i, N) = a + \sum_{k=i}^N (hz_{\alpha(i,k)}\sigma_{i,k} + (h+s)\sigma_{i,k}[\phi(z_{\alpha(i,k)}) - (1 - \alpha(i,k))z_{\alpha(i,k)}]) + v \left(S_i - \sum_{k=i}^N d_k \right) \quad (23)$$

4 Comparison of the CP and MIP approaches

In [4] Tarim and Kingsman proposed a piecewise linear approximation of the cost function for the single-period newsvendor type model under holding and shortage costs, which we analyzed above. Thus they were able to build a MIP model approximating an optimal solution for the multi-period stochastic lot-sizing under fixed ordering, holding and shortage costs. They gave a few examples to show the effect of higher noise levels (uncertainty in the demand forecasts) on the order schedule. Using the same examples we shall compare the policies obtained using our exact CP approach with their approximation. Depending on the number of segments used in the piecewise approximation, the quality of the solutions obtained can be improved. We shall consider approximations with two and seven segments. The forecast of demand in each period are given in Table 1. We assume that the demand in each period is normally distributed about the forecast value with the same coefficient of variation τ . Thus the standard deviation of demand in period t is $\sigma_t = \tau \cdot \bar{d}_t$. In all cases, initial inventory levels, delivery lead-times and salvage values are set to zero.

In Fig. 7–11 optimal replenishment policies obtained with our CP approach are compared for four different instances, with respect to τ , v , a and s , with the policies provided by the 2-segment (PW-2) and 7-segment (PW-7) approximations. For each instance we compare the expected total cost provided by the exact method with the expected total cost provided by the policies found using approximate MIP models. Since the cost provided by PW-2 and PW-7 is an

approximation, it often differs significantly from the real expected total cost related to policy parameters found by these models. It is therefore not meaningful to compare the cost provided by the MIP model with that of the optimal policy obtained with our CP model. To obtain a meaningful comparison we computed the real expected total cost by applying the exact cost function (Eqs. 17, 23) discussed above to the (R^n, S^n) policy parameters obtained through PW-2 and PW-7. It is then possible to assess the accuracy of approximations in [4]. Fig.

Fig. 7. $h = 1, a = 250, s = 10, v = 0, \tau = 0.0$

7 shows the optimal replenishment policy for the deterministic case ($\tau = 0.0$). The direct item cost (v) is taken as zero. Four replenishment cycles are planned. The (R^n, S^n) policy parameters are $R = [3, 1, 3, 1]$ and $S = [370, 200, 470, 100]$. The total cost for this policy is 1460. Fig. 8 shows an instance where we con-

Fig. 8. $h = 1, a = 250, s = 10, v = 0, \tau = 0.1$

sider low levels of forecast uncertainty ($\tau = 0.1$). In this case both PW-2 and PW-7 perform well compared to our exact CP solutions. Since forecast uncertainty must be considered, all the models introduce buffer stocks. The optimal (R^n, S^n) policy parameters found by our CP approach are $R = [3, 1, 2, 2]$ and $S = [384, 227, 449, 160]$. The PW-2 solution is 1.75% more costly than the exact solution, while the PW-7 solution is slightly more costly than the exact solution.

Period	1	2	3	4	5	6	7	8
d_t	200	100	70	200	300	120	200	300

Table 2. Expected demand values

Fig. 9 shows that as the level of forecast uncertainty increases ($\tau = 0.2$), the

Fig. 9. $h = 1, a = 250, s = 10, v = 0, \tau = 0.2$

quality of the PW-2 solution deteriorates, in fact it is now 3.62% more costly than the exact solution. The optimal (R^n, S^n) policy parameters found by our CP approach are $R = [3, 1, 2, 2]$ and $S = [401, 253, 479, 170]$. In contrast the PW-7 solution is still only slightly more costly than the exact solution. As noted

Fig. 10. $h = 1, a = 350, s = 50, v = 0, \tau = 0.3$

in [4] the quality of the approximation decreases for high ratios s/h . In Fig. 10 we consider $s/h = 50$ and a different demand pattern. The forecast of demand in each period are given in Table 2. Now the PW-2 solution is 6.66% more costly than the exact approach, while the PW-7 solution is 1.03% more costly. The optimal (R^n, S^n) policy parameters found by our CP approach are $R = [3, 1, 2, 1, 1]$ and $S = [483, 324, 592, 324, 486]$. In Fig. 11 we consider the same instance but

a direct item cost is now incurred ($v = 15$). The buffer stock held in the last replenishment cycle is affected by this parameter, and is decreased from 186 to 63. The PW-7 policy is now 0.84% more costly than the exact one. For these

Fig. 11. $h = 1$, $a = 350$, $s = 50$, $v = 15$, $\tau = 0.3$

instances seven segments usually provides a solution with a cost reasonably close to optimal. In terms of running times, for all these instances both the MIP approximations and the CP model perform very quickly. In our experiments we used ILOG OPL Studio 3.7 to solve the MIP models of [4], and Choco [16] (an open source solver written in Java) to implement our CP model. All experiments were performed on an Intel Centrino 1.5 GHz with 500Mb RAM. Since the planning horizon is short (8 periods), we were able to solve any instance in less than a second. As the planning horizon length increases the pure CP model becomes slower than the MIP one. This is due both to the size of decision variable domains and to the lack of good bounds in the search. We do not discuss efficiency issues in this paper, but we emphasise that a significant reduction in decision variable domain sizes can be achieved in a way similar to the one discussed in [2]. Furthermore it is possible to incorporate in our CP model dedicated cost-based filtering methods [15] based on a *dynamic programming relaxation* [5] that is able to generate good bounds during the search. Such a technique has been already employed under a service level constraint [1] and preliminary results in this direction under a penalty cost suggest that our exact CP model, when enhanced with these dedicated filtering techniques, is able to produce an optimal solution for instances up to 50 periods and more in a few seconds.

5 Conclusions

We presented a CP approach that finds optimal (R_n, S_n) policies under non-stationary demands. Using our approach it is now possible to evaluate the quality of a previously published MIP-based approximation method, which is typically faster than the pure CP approach. Using a set of problem instances we showed that a piecewise approximation with seven segments usually provides good quality solutions, while using only two segments can yield solutions that differ signifi-

cantly from the optimal. In future work we will aim to develop domain reduction techniques and cost-based filtering methods to enhance the performance of our exact CP approach.

Acknowledgements: this work was supported by Science Foundation Ireland under Grant No. 03/CE3/I405 as part of the Centre for Telecommunications Value-Chain-Driven Research (CTVR) and Grant No. 00/PI.1/C075.

References

1. S. A. Tarim, B. Hnich, R. Rossi, S. Prestwich. Cost-Based Filtering for Stochastic Inventory Control. *Lecture Notes in Computer Science*, Springer-Verlag, 2007, to appear.
2. S. A. Tarim, B. Smith. Constraint Programming for Computing Non-Stationary (R,S) Inventory Policies. *European Journal of Operational Research*. to appear.
3. S. A. Tarim, B. G. Kingsman. The Stochastic Dynamic Production/Inventory Lot-Sizing Problem With Service-Level Constraints. *International Journal of Production Economics* 88:105–119, 2004.
4. S. A. Tarim, B. G. Kingsman. Modelling and Computing (R^n, S^n) Policies for Inventory Systems with Non-Stationary Stochastic Demand. *European Journal of Operational Research* 174:581–599, 2006.
5. S. A. Tarim. Dynamic Lotsizing Models for Stochastic Demand in Single and Multi-Echelon Inventory Systems. PhD Thesis, Lancaster University, 1996.
6. J. H. Bookbinder, J. Y. Tan. Strategies for the Probabilistic Lot-Sizing Problem With Service-Level Constraints. *Management Science* 34:1096–1108, 1988.
7. H. M. Wagner, T. M. Whitin. Dynamic Version of the Economic Lot Size Model. *Management Science* 5:89–96, 1958.
8. E. A. Silver, D. F. Pyke, R. Peterson. Inventory Management and Production Planning and Scheduling. John Wiley and Sons, New York, 1998.
9. E. L. Porteus. Foundations of Stochastic Inventory Theory. Stanford University Press, Stanford, CA, 2002.
10. K. Apt. Principles of Constraint Programming. Cambridge University Press, Cambridge, UK, 2003.
11. A. Charnes, W. W. Cooper. Chance-Constrained Programming. *Management Science* 6(1):73–79, 1959.
12. L. Fortuin. Five Popular Probability Density Functions: a Comparison in the Field of Stock-Control Models. *Journal of the Operational Research Society* 31(10):937–942, 1980.
13. I. J. Lustig, J.-F. Puget. Program Does Not Equal Program: Constraint Programming and its Relationship to Mathematical Programming. *Interfaces* 31:29–53, 2001.
14. J. R. Birge, F. Louveaux. Introduction to Stochastic Programming. Springer Verlag, New York, 1997.
15. F. Focacci, A. Lodi, M. Milano. Cost-Based Domain Filtering. Fifth International Conference on the Principles and Practice of Constraint Programming, *Lecture Notes in Computer Science* 1713, Springer Verlag, 1999, pp. 189–203.
16. F. Laburthe and the OCRE project team. Choco: Implementing a CP Kernel. Bouygues e-Lab, France.
17. G. Hadley, T. M. Whitin. Analysis of Inventory Systems. Prentice Hall, 1964.