Cost-based Filtering for Stochastic Inventory Systems with Shortage Cost

Student: Roberto Rossi[1]
Supervisors: S. Armagan Tarim[2], Brahim Hnich[3] and Steven Prestwich[1]

Cork Constraint Computation Centre, University College, Cork, Ireland[1]
{rrossi,s.prestwich}@4c.ucc.ie
Department of Management, Hacettepe University, Turkey[2]
armagan.tarim@hacettepe.edu.tr
Faculty of Computer Science, Izmir University of Economics, Turkey[3]
brahim.hnich@ieu.edu.tr

Abstract. Cost-based filtering is a promising technique able to improve search performance in combinatorial optimization problems. Such a technique has already been successfully applied to stochastic inventory control problems. We focus on the class of production/inventory control problems that considers a single product and a single stocking location, given a stochastic demand with a known non-stationary probability distribution. An exact CP approach has been recently introduced to find optimal policy parameters for this problem under ordering, holding and shortage cost. We extend such a CP approach using cost-based filtering. Our algorithm can efficiently solve to optimality instances of realistic size, often with no search effort at all. An extended version of this work is presented in the application track of the technical programme.

1 Problem definition and (R_n, S_n) policy

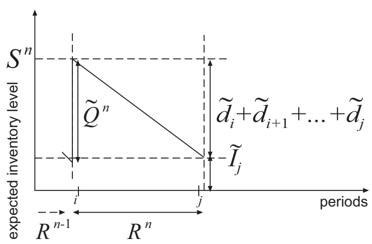

Fig. 1. (R_n,S_n) policy. $\tilde{d}_i + \tilde{d}_{i+1} + \ldots + \tilde{d}_j$ is the expected demand over R_n; $\tilde{I}_j = S_n - \tilde{d}_i + \tilde{d}_{i+1} + \ldots + \tilde{d}_j$ is the expected closing inventory level for R_n.

Inventory control is a topic that has been widely studied in the operational research community. A complete survey over the complexity of deterministic inventory control problems is given in [5]. An important class of production/inventory control problems assumes a non-stationary stochastic demand process. A general *stochastic programming* [1] model for this problem is formulated in [2]. In order to provide a solution for the stochastic production/inventory control a *policy of response* has to be chosen. For a detailed discussion on inventory control policies refer to [12]. A widely studied policy is the (R,S) policy [3]. In this policy a replenishment is placed every R periods to raise the inventory to the order-up-to-level S. This provides an effective means of damping planning instability – deviations in planned orders, also known as *nervousness* (de

Kok and Inderfurth [4], Heisig [8]) – and coping with demand uncertainty. As pointed out by (Silver et al. [12], 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. In (Janssen and de Kok [9]) a two-supplier periodic model is discussed where one supplier delivers a fixed quantity while the amount delivered by the other is governed by an (R,S) policy. Under non-stationary demand assumption this 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 (Fig. 1). We consider a demand d_t in each period $t \in \{1, \ldots, N\}$, which is a normally distributed random variable with known probability density function $g_t(d_t)$. This 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.

2 CP implementation

The *deterministic equivalent* (*stochastic programming* and deterministic equivalent modeling are discussed in [1]) CP formulation proposed in [11] is

$$\min \; E\{TC\} = C \tag{1}$$

subject to

$$\text{objConstraint}\left(C, \tilde{I}_1, \ldots, \tilde{I}_N, \delta_1, \ldots, \delta_N, d_1, \ldots, d_N, a, h, s, v\right) \tag{2}$$

and, for $t = 1 \ldots N$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} \geq 0 \tag{3}$$

$$\tilde{I}_t + \tilde{d}_t - \tilde{I}_{t-1} > 0 \Rightarrow \delta_t = 1 \tag{4}$$

$$\tilde{I}_t \in \mathbb{Z}, \quad \delta_t \in \{0, 1\} \tag{5}$$

Each decision variable $\tilde{I}_t$ represents the expected closing inventory level at the end of period t; it is possible to compute bounds for the domains of these variables. Each $\tilde{d}_t$ represents the expected value of the demand in a given period t according to its probability density function $g_t(d_t)$. The binary decision variables δ_t state whether a replenishment is fixed for period t ($\delta_t = 1$) or not ($\delta_t = 0$). The objective function (1) minimizes the expected total cost over the planning horizon. objConstraint$(C, \tilde{I}_1, \ldots, \tilde{I}_N, \delta_1, \ldots, \delta_N, d_1, \ldots, d_N, a, h, s, v)$, whose propagation logic is described in [11], dynamically computes expected closing-inventory-levels and it assigns to C the expected total cost related to an assignment for replenishment decisions, depending on the demand distribution in each period and on the given combination for problem parameters a, h, s and v.

2.1 Cost-based filtering by relaxation

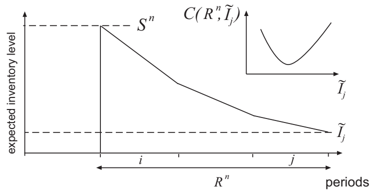

Fig. 2. Convexity of the expected total cost associated to a given replenishment cycle R^n covering periods $\{i, \ldots, j\}$. The expected total cost is a function of $\tilde{I}_j$ (expected closing-inventory-level).

Cost-based filtering is an elegant way of combining techniques from CP and Operations Research (OR) [6]. OR-based optimization techniques are used to remove values from variable domains that cannot lead to better solutions. This type of domain filtering can be combined with the usual CP-based filtering methods and branching heuristics, yielding powerful hybrid search algorithms. In [14] the authors adopt a relaxation proposed by Tarim in [13] for the CP model that computes (R^n,S^n) policy parameters under service level constraints. When the relaxed model is solved it provides good bounds for the original problem. Furthermore the relaxed problem is a Shortest Path Problem that can be solved in polynomial time. Therefore it is easy to obtain good bounds at each node of the search tree. In the same work the authors also explain how it is possible to take into account a partial assignment for replenishment decisions $\delta_1, \ldots, \delta_N$ and for expected closing-inventory-levels $\tilde{I}_1, \ldots, \tilde{I}_N$ when the relaxed problem is constructed, so that the effect of these assignments is reflected on the bound that is obtained by solving the relaxed problem. As shown in [14], the CP model proposed for computing (R^n,S^n) policy parameters under service level constraints can be reduced to a Shortest Path Problem if the inventory conservation constraint and the replenishment condition constraint, that is constraint 3 and 4 in our model under shortage cost scheme, are relaxed for replenishment periods. That is for each possible pair of replenishment cycles $\langle R(i, k-1), R(k, j)\rangle$ where $i, j, k \in \{1, \ldots, N\}$ and $i < k \leq j$, the relationship between the opening inventory level of $R(k, j)$ and the closing inventory level of $R(i, k-1)$ is not considered. The same approach can be translated to the CP model for (R^n,S^n) under shortage cost scheme. In [11] the authors provided a general function $C(i, j, \tilde{I}_j)$ to compute the expected total cost of replenishment cycle $R(i, j)$, when an expected closing-inventory-level $\tilde{I}_j$ is held in period j. Furthermore they proved that this function is convex in $\tilde{I}_j$ (Fig. 2). If we consider each replenishment cycle $R(i, j)$ independently, we can efficiently compute the optimal expected closing-inventory-level that minimizes the expected total cost associated to such a cycle using gradient based methods for convex optimization. In this way we obtain a set $\mathcal{S}$ of $N(N+1)/2$ possible replenishment cycles and respective order-up-to-levels. Our new problem is to find an optimal set $\mathcal{S}^* \subset \mathcal{S}$

of consecutive disjoint replenishment cycles that covers our planning horizon at the minimum cost. In [14] it was shown that the optimal solution to this relaxation is given by the shortest path in a graph from a given initial node to a final node where each arc represents a specific cost. We now adapt their approach to our model that employs a shortage cost scheme. If N is the number of periods in the planning horizon of the original problem, we introduce $N+1$ nodes. Since we assume, without loss of generality, that an order is always placed at period 1, we take node 1, which represents the beginning of the planning horizon, as the initial one. Node $N+1$ represents the end of the planning horizon. For each possible replenishment cycle $R(i, j-1)$ such that $i, j \in \{1, \dots, N+1\}$ and $i < j$, we introduce an arc (i, j) with associated cost $Q(i, j) = C(i, j-1, \tilde{I}^*_{j-1})$, where $\tilde{I}^*_{j-1}$ is the closing-inventory-level that minimizes the convex cost of replenishment cycle $R(i, j-1)$. Since we are dealing with a one-way temporal feasibility problem, when $i \geq j$, we introduce no arc.

The cost of the shortest path from node 1 to node $N+1$ in the given graph is a valid lower bound for the original problem, as it is a solution of the relaxed problem. In fact, as shown in [11], the expected total cost function for each replenishment cycle is convex in the expected closing-inventory-level held at the end of the cycle. Therefore in order to meet the violated inventory conservation constraints, if any exists, we will incur an overall higher expected total cost for a given group of replenishment cycles (Fig. 3). Furthermore it is easy to map the optimal solution for the relaxed problem, that is the set of arcs participating to the shortest path, to a solution for the original problem by noting that each arc (i, j) represents a replenishment cycle $R(i, j-1)$. The feasibility of such a solution with respect to the original problem can be checked by verifying that it satisfies every relaxed constraint. If no inventory conservation constraint is violated, it is easy to see that the computed cost is optimal for the given replenishment plan. We will now show how to exploit this *lower bound* in an *optimization oriented* global constraint able to dynamically produce good bounds when a partial solution is provided. A detailed discussion on optimization oriented global constraints can be found in [7]. The costs stored in the connection matrix can be adjusted to reflect the current partial assignment for decision variables δ_t and $\tilde{I}_t$ exactly in the way shown for the service level constrained model [14]. If a given $\tilde{I}_t$, $t \in \{i, \dots, j\}$ is assigned to a value, the expected closing-inventory-level ($\tilde{I}_j$) for the replenishment cycle $R(i, j)$, which covers period t, is uniquely determined and therefore the expected total cost for such a replenishment cycle can be directly computed from $C(i, j, \tilde{I}_j)$. If in a given period t, δ_t is assigned to 0, we remove from the graph every inbound arc to node t and every outbound arc from node t. This

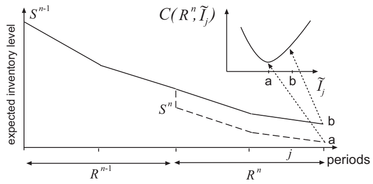

Fig. 3. The optimal expected closing-inventory-level for replenishment cycle R^n considered alone is a, this minimizes the convex cost associated to replenishment cycle R^n. In order to meet the inventory conservation constraint for the stocks carried over from cycle R^{n-1}, the minimum expected closing-inventory-level required is b. Such a value produces a higher expected total cost for R^n.

prevents node t from being part of the shortest path, and hence prevents period t from being a replenishment period. On the other hand, if δ_t is assigned to 1, then we split the planning horizon into two at period t, thus obtaining two new subproblems $\{1,\ldots,t-1\}$ and $\{t,\ldots,N\}$. We can then separately solve these two subproblems by relaxing them and solving the respective Shortest Path Problems. Note that the action of splitting the time span is itself a relaxation; in fact it means overriding constraints (3,4) for period t. It follows that the cost of the overall solution obtained by merging the two subproblem solutions is again a valid lower bound for the original problem. Furthermore it is easy to characterize when such a bound is an exact one. When the solutions of the two subproblems are both feasible with respect to the original model and the condition $\tilde{I}_{t-1} \leq S_t$ is satisfied, the solution obtained merging the two solutions obtained for the independent subproblems is both feasible and optimal for the original problem. We have shown how to act when each of the possible cases, $\delta_i = 1$ and $\delta_i = 0$, is encountered. It is now possible at any point of the search in the decision tree to compute valid lower bounds.

3 Experimental Results

In this section we show the effectiveness of our approach. A single problem is considered and the period demands are generated from seasonal data with no trend: $\tilde{d}_t = 50[1 + \sin(\pi t/6)]$. In addition to the "no trend" case (P1) we also consider three others: (P2) positive trend case, $\tilde{d}_t = 50[1 + \sin(\pi t/6)] + t$; (P3) negative trend case, $\tilde{d}_t = 50[1 + \sin(\pi t/6)] + (52 - t)$; (P4) life-cycle trend case, $\tilde{d}_t = 50[1 + \sin(\pi t/6)] + \min(t, 52 - t)$. In each test we assume an initial null inventory level and a normally distributed demand for every period with a coefficient of variation $\sigma_t/\tilde{d}_t$ for each $t \in \{1,\ldots,N\}$, where N is the length of the considered planning horizon. We performed tests using four different ordering cost values $a \in \{50, 100, 150, 200\}$ and two different $\sigma_t/\tilde{d}_t \in \{1/3, 1/6\}$. The planning horizon length takes even values in the range $[20, 38]$. The holding cost used in these tests is $h = 1$ per unit per period. Our tests also consider two different shortage cost values $s = 15$ and $s = 25$. Direct item cost is $v = 2$ per unit produced. All experiments were performed on an Intel(R) Centrino(TM) CPU 1.50GHz with 500Mb RAM. The solver used is Choco [10], an open-source solver developed in Java. The cost-based filtering techniques presented are implemented as dedicated constraints within Choco. The same variable and value selection heuristics used in [14] are employed. The running time limit is set to 5 seconds. Our method was able to solve 315 of the 320 instances considered in the given time limit and almost all the instances were solved in less than a second. When the cost-based filtering method we proposed is not used, the pure CP approach is never able to provide an optimal solution within the given running time limit for every instance. Finally it should be also noted that the worst case running time of our approach over the whole test bed was $6,77$ minutes. Therefore even in the few cases where an optimal solution is not found in a less than 5 seconds, our cost-based filtering techniques provides a reasonable running time.

4 Conclusions

We presented a cost-based filtering technique able to speed up the search for optimal (R^n, S^n) policy parameters under a shortage cost scheme. Our experimental results prove that such a technique brings a significant improvement in the efficiency of the existing CP approach for this problem.

References

1. J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Springer Verlag, New York, 1997.
2. J. H. Bookbinder and J. Y. Tan. Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34:1096–1108, 1988.
3. A. G. de Kok. Basics of inventory management: part 2 The (R,S)-model. Research memorandum, FEW 521, 1991. Department of Economics, Tilburg University, Tilburg, The Netherlands.
4. T. de Kok and K. Inderfurth. Nervousness in inventory management: Comparison of basic control rules. *European Journal of Operational Research*, 103:55–82, 1997.
5. M. Florian, J. K. Lenstra, and A. H. G. Rinooy Kan. Deterministic production planning: Algorithms and complexity. *Management Science*, 26(7):669–679, 1980.
6. F. Focacci, A. Lodi, and M. Milano. Cost-based domain filtering. In *Proceedings of the 5th International Conference on the Principles and Practice of Constraint Programming*, pages 189–203. Springer Verlag, 1999. Lecture Notes in Computer Science No. 1713.
7. F. Focacci and M. Milano. Connections and integrations of dynamic programming and constraint programming. In *Proceedings of the International Workshop on Integration of AI and OR techniques in Constraint Programming for Combinatorial Optimization Problems CP-AI-OR 2001*, 2001.
8. Gerald Heisig. *Planning Stability in Material Requirements Planning Systems.* Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.
9. F. Janssen and T. de Kok. A two-supplier inventory model. *International Journal of Production Economics*, 59:395–403, 1999.
10. F. Laburthe and the OCRE project team. Choco: Implementing a cp kernel. Technical report, Bouygues e-Lab, France, 1994.
11. R. Rossi, S. A. Tarim, B. Hnich, and S. Prestwich. Replenishment planning for stochastic inventory systems with shortage cost. In *Proceedings of the International Conference on Integration of AI and OR techniques in Constraint Programming for Combinatorial Optimization Problems CP-AI-OR 2007*, pages 229–243. Springer Verlag, 2007. Lecture Notes in Computer Science No. 4510.
12. E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling.* John-Wiley and Sons, New York, 1998.
13. S. A. Tarim. *Dynamic Lotsizing Models for Stochastic Demand in Single and Multi-Echelon Inventory Systems.* PhD thesis, Lancaster University, 1996.
14. S. A. Tarim, B. Hnich, R. Rossi, and S. Prestwich. Cost-based filtering for stochastic inventory control. In *Recent Advances in Constraints Joint ERCIM/CoLogNET International Workshop on Constraint Solving and Constraint Logic Programming, CSCLP 2006.* Springer Verlag, 2007. Lecture Notes in Computer Science, to appear.