A state space augmentation algorithm for the replenishment cycle inventory policy

Roberto Rossi [a,*], S. Armagan Tarim [b], Brahim Hnich [c], Steven Prestwich [d]

[a] Logistics, Decision and Information Sciences, Wageningen UR, Hollandseweg 1, 6706 KN Wageningen, The Netherlands
[b] Department of Management, Hacettepe University, Turkey
[c] Faculty of Computer Science, Izmir University of Economics, Izmir, Turkey
[d] Cork Constraint Computation Centre, University College, Cork, Ireland

ARTICLE INFO

ABSTRACT

In this work we propose an efficient dynamic programming approach for computing replenishment cycle policy parameters under non-stationary stochastic demand and service level constraints. The replenishment cycle policy is a popular inventory control policy typically employed for dampening planning instability. The approach proposed in this work achieves a significant computational efficiency and it can solve any relevant size instance in trivial time. Our method exploits the well known concept of state space relaxation. A filtering procedure and an augmenting procedure for the state space graph are proposed. Starting from a relaxed state space graph our method tries to remove provably suboptimal arcs and states (filtering) and then it tries to efficiently build up (augmenting) a reduced state space graph representing the original problem. Our experimental results show that the filtering procedure and the augmenting procedure often generate a small filtered state space graph, which can be easily processed using dynamic programming in order to produce a solution for the original problem.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Inventory theory provides methods for managing and controlling inventories under different constraints and environments. An interesting class of production/inventory control problems is the one that considers the single-location, single-product case under non-stationary stochastic demand and service level constraints. Such a problem has been widely studied because of its key role in practice.

Different inventory control policies can be adopted for the above mentioned problem. For a discussion of inventory control policies, see Silver et al. (1998). One of the possible policies that can be adopted is the replenishment cycle policy, (R,S). A detailed discussion on the characteristics of (R,S) can be found in de Kok (1991). In this policy an order is placed every R periods to raise the inventory level to the order-up-to-level S. This provides an effective means of dampening planning instability (deviations in planned orders, also known as *nervousness* (de Kok and Inderfurth, 1997; Heisig, 2002) and coping with demand uncertainty. As pointed out by Silver et al. (1998, 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 and risk management (Tang, 2006).

Under the non-stationary demand assumption the replenishment cycle policy takes the form (R^n, S^n) where R^n denotes the length of the nth replenishment cycle and S^n the respective order-up-to-level. In this policy, the actual order quantity for replenishment cycle n is determined after the demand in previous periods has been observed. The order quantity is computed as the amount of stock required to raise the closing inventory level of replenishment cycle $n-1$ up to level S^n. In order to provide a solution for our problem under the (R^n, S^n) policy we must populate both the sets $\{R^n | n = 1, \ldots, M\}$ and $\{S^n | n = \{1, \ldots, M\}$, where M denotes the number of replenishment cycles scheduled over a finite planning horizon of N periods.

The problem of populating these sets has been solved to optimality only recently, due to the complexity involved in the modeling of uncertainty and of the policy-of-response. As Silver points out, computing replenishment cycle policy parameters under non-stationary stochastic demand is a computationally hard task (Silver, 1978). Early works in this area adopted heuristic strategies such as those proposed by Silver (1978), Askin (1981), and Bookbinder and Tan (1988). Under some mild assumptions, the first complete solution method for this problem was introduced by Tarim and Kingsman (2004), who proposed a

* Corresponding author. Tel.: +31 317 482321; fax: +31 317 485646.
E-mail addresses: roberto.rossi@wur.nl, robros@gmail.com (R. Rossi),
armtar@yahoo.com (S.A. Tarim), brahim.hnich@ieu.edu.tr (B. Hnich),
s.prestwich@4c.ucc.ie (S. Prestwich).

deterministic equivalent mixed integer programming (MIP) formulation for computing (R^n, S^n) policy parameters. Tempelmeier (2007) extended Tarim and Kingsman's MIP formulation in order to consider different service level measures, such as the "fill rate". Nevertheless, empirical results showed that Tarim and Kingsman's model is unable to solve large instances. Tarim and Smith (2008) therefore introduced a more compact and efficient constraint programming formulation of the same problem that showed a significant computational improvement over the MIP formulation. The constraint programming formulation has been further enhanced by means of dedicated cost-based filtering algorithms developed by Tarim et al. (2009). A stochastic constraint programming (Tarim et al., 2006) approach for computing optimal (R^n, S^n) policy parameters is proposed in Rossi et al. (2008). In this work the authors drop the mild assumptions originally introduced by Tarim and Kingsman and compute optimal (R^n, S^n) policy parameters. Of course, there is a price to pay for dropping Tarim and Kingsman's assumptions, in fact this latter approach is less efficient than the one in Tarim and Smith (2008). Finally, Pujawan and Silver (2008) recently proposed a novel and effective heuristic approach.

In this paper, we build on Tarim and Kingsman's modeling assumptions and we develop a state-of-the-art algorithm for computing optimal (R^n, S^n) policy parameters. Two existing techniques—dynamic programming and state space relaxation—are combined in order to obtain an effective approach for computing (R^n, S^n) policy parameters. Dynamic programming (DP) is an optimization procedure that solves optimization problems by decomposing them into a nested family of subproblems. DP is based on the *principle of optimality* (Bellman, 1957; Dreyfus and Law, 1989) and it has been applied to solve a wide variety of combinatorial optimization problems, as well as optimal control problems. State space relaxation (SSR) considers the DP formulation of a combinatorial optimization problem, and modifies this formulation to obtain a different—and possibly more compact—DP formulation whose optimal solution is a lower bound for the original problem. Christofides et al. (1981) proposed that SSR has been successfully applied to constrained variants of routing problems (see, e.g. Mingozzi et al., 1997; Focacci and Milano, 2001). Roughly speaking, SSR maps the original state space graph to a new state space graph having a smaller number of vertices, and whose shortest path represents a lower bound for the cost of the shortest path in the original state space graph.

In this work, we enhance these known approaches with a novel strategy: we introduce a filtering procedure for the state space graph and an augmenting procedure that is able to build a reduced state space graph for the original problem starting from a filtered state space graph for the relaxed problem. The concept of state space augmentation (Boland et al., 2006) is known in the operations research literature. A dual approach to state space augmentation also exists and is known as decremental SSR (Righini and Salani, 2008). Nevertheless, the idea of filtering a relaxed state space graph is, to the best of our knowledge, a novel contribution. Our experimental results prove the effectiveness of such an approach for computing optimal (R^n, S^n) policy parameters.

The paper is structured as follows. In Section 2 we introduce the problem definition and the modeling assumptions adopted in this work. In Section 3 we describe a DP reformulation for Tarim and Kingsman's model. An SSR for this reformulation is presented in Section 4. A procedure for filtering the relaxed state space graph is presented in Section 5. An augmenting procedure for the relaxed state space graph is described in Section 6. An example that demonstrates the algorithm proposed is given in Section 7. Our computational experience and a comparison with the state-of-the-art approaches for computing replenishment cycle policy parameters are discussed in Section 8. In Section 9 we draw conclusions.

2. Problem definition and modeling assumptions

The single-location, single-product production/inventory control problem under non-stationary stochastic demand and service level constraints is formulated in this paper by using the following inputs and assumptions.

We consider a planning horizon of N periods and a demand d_t for each period $t \in \{1, \ldots, N\}$, which is a non-negative random variable with known probability density function and expected value $\tilde{d}_t$. We assume that the demand occurs instantaneously at the beginning of each time period. The demand is non-stationary, that is it can vary from period to period, demands in different periods are assumed to be independent. 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. The sell-back of excess stock is not allowed, if the actual stock exceeds the order-up-to-level for a given review, this excess stock is carried forward and it is not returned to the supply source. However, as in Bookbinder and Tan (1988), Tarim and Kingsman (2004), Tarim and Smith (2008), and Tempelmeier (2007) such occurrences are regarded as rare events and accordingly the cost of carrying this excess stock and its effect on the service levels of subsequent periods are ignored.

A fixed delivery cost a is incurred for each order. A linear holding cost h is incurred for each unit of product carried in stock from one period to the next. Our aim is to find a replenishment plan that minimizes the expected total cost, which is composed of ordering costs and holding costs, over the N-period planning horizon, satisfying the service level constraints. As a service level constraint we require that, with a probability of at least a given value α, at the end of each period the net inventory will be non-negative. As pointed out in Tempelmeier (2007), since period demands are random, the net inventory may become negative. However, the number of stock-outs is restricted by the service level constraints enforced. While computing holding costs, we will assume, as in Bookbinder and Tan (1988), Tarim and Kingsman (2004), Tarim and Smith (2008), and Tempelmeier (2007), that the service level is set large enough to ensure that the net inventory will be a good approximation of the inventory on hand.

3. A DP formulation for the deterministic equivalent problem

We hereby introduce a deterministic equivalent DP formulation for computing optimal (R^n, S^n) policy parameters.

Definition. A replenishment cycle, $T(i, j)$, is the time span between two consecutive orders/productions occurring in periods i and $j+1, j \geq i$.

Definition. The cycle buffer stock, $b(i, j)$, denotes the minimum expected buffer stock level required to satisfy the required non-stock-out probability during $T(i, j)$.

We define $b(i, j)$, $i = 1, \ldots, N$, $j = i, \ldots, N$, as

$$b(i, j) = G^{-1}_{d_i + d_{i+1} + \cdots + d_j}(\alpha) - \sum_{k=i}^{j} \tilde{d}_k, \tag{1}$$

where $G_{d_i + d_{i+1} + \cdots + d_j}$ is the cumulative probability distribution function of $d_i + d_{i+1} + \cdots + d_j$. It is assumed that G is strictly increasing, hence G^{-1} is uniquely defined. It should be noted that it is possible to consider different service level measures—for instance the "fill rate"—simply by introducing a different definition for the cycle buffer stock (see also Tempelmeier, 2007).

Since N is the number of periods in our planning horizon, this will also be the number of steps in the system. A state s_k at step k represents a possible expected closing-inventory-level, $\tilde{I}_k$, at the end of period k. The decision x_k to be taken at step k is to place an