Operations Research Letters 42 (2014) 161–165

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

A reformulation for the stochastic lot sizing problem with service-level constraints

Huseyin Tunc[a], Onur A. Kilic[a,*], S. Armagan Tarim[a], Burak Eksioglu[b]

[a] Institute of Population Studies, Hacettepe University, Ankara, Turkey

[b] Department of Industrial and Systems Engineering, Mississippi State University, MS, United States

ARTICLE INFO

Article history:
Received 23 September 2013
Received in revised form
28 January 2014
Accepted 28 January 2014
Available online 7 February 2014

Keywords:
Stochastic lot sizing
Reformulation
Service level

ABSTRACT

We study the stochastic lot-sizing problem with service level constraints and propose an efficient mixed integer reformulation thereof. We use the formulation of the problem present in the literature as a benchmark, and prove that the reformulation has a stronger linear relaxation. Also, we numerically illustrate that it yields a superior computational performance. The results of our numerical study reveals that the reformulation can optimally solve problem instances with planning horizons over 200 periods in less than a minute.

1. Introduction

We address the stochastic lot sizing problem with α service level constraints. A description of the problem is as follows. A retailer faces stochastic demand over a finite planning horizon. Period demands are independent but not necessarily identically distributed. When an order is placed, a fixed ordering cost is incurred. Also, a holding cost is incurred for each unit of inventory carried from one period to the next. Demands realized when the system is out of stock are backordered. The retailer aims to plan her orders so as to minimize the expected total cost over the planning horizon while guaranteeing a pre-defined non-stockout probability in any period. This problem appears in many industrial environments where demand is non-stationary and stochastic (e.g. due to short product life cycles), and supply requires fixed freight fees (e.g. due to distant offshore suppliers). Also, α service level is the very common in environments where stock-outs are rather costly and independent of the magnitude or the duration of the stock-outs (e.g. due to costly stoppages in production) [12].

Bookbinder and Tan [2] proposed a number of strategies for this problem which differ with respect to the time epochs where the decisions on when and how much to order are taken. One of these strategies is the so-called static–dynamic uncertainty strategy where a complete order schedule is fixed at the outset, and order quantities are decided upon following an order-up-to policy at each replenishment epoch after the actual inventory is observed. Because it offers a rigid order schedule, this strategy is particularly appealing in material requirement planning, joint replenishment, and shipment consolidation environments (see [5,12,6,7]). Furthermore, it effectively hedges against demand uncertainty (see [4,17]).

The current paper belongs to the growing literature on the application of static–dynamic uncertainty strategy in stochastic lot sizing. A significant progress has been made in this line of research over the last decades. Silver [10] and Askin [1] studied the problem under a backorder cost scheme. They proposed heuristics building on the well-known Silver and Meal [11] heuristic. Bookbinder and Tan [2] addressed the problem with an α service-level constraint—the very problem considered in the current study. They developed a sequential method where an order schedule is obtained first, and then order-up-to levels are determined. Tarim and Kingsman [14] developed a mixed integer programming (MIP) formulation for the same problem which determines the order schedule and order-up-to levels simultaneously. Tarim and Kingsman [15], Tempelmeier [16], and Rossi et al. [8] adapted this formulation to take into account different service measures. Recently, Rossi et al. [9] and Tarim et al. [13] developed computationally efficient special-purpose algorithms where an optimal solution is obtained by iteratively solving a relaxed version of the problem that is formulated and solved as a shortest-path problem.

The contribution of the current study is to present a reformulation for the static–dynamic uncertainty strategy in the stochastic lot sizing problem with α service level constraints. As opposed to the special-purpose algorithms in the literature (see [9,13]), the reformulation is a deterministic equivalent MIP model. As such, it has

* Corresponding author.
E-mail address: onuralp@hacettepe.edu.tr (O.A. Kilic).

http://dx.doi.org/10.1016/j.orl.2014.01.010
0167-6377/

the advantage of being at users' disposal to be fed into commercial solvers without the need of tailor-made computer programs. We verify the efficiency of the reformulation both analytically and numerically by comparing it against Tarim and Kingsman's [14] formulation—which we use as a reference. We first show that the linear relaxation of the reformulation is stronger. Then, we numerically illustrate that the reformulation is far more time-efficient. The linear relaxation of the reformulation averages an integrality gap of 0.2% for the test instances considered. As a result, the reformulation can solve realistic-sized problem instances to optimality in the order of seconds.

The reminder of the paper is organized as follows. In Section 2, we set the notation and introduce the reference formulation. In Section 3, we develop the reformulation. In Section 4, we provide a comparison of the reference formulation and the reformulation. In Section 5, we report the numerical experiments and obtained results.

2. Notation and Tarim and Kingsman's formulation

We consider an N-period planning horizon. The demands $d_1, d_2, \ldots, d_N$ are independent random variables with known distribution functions. The fixed ordering cost is a per order. The holding cost is h per unit per period. The minimum non-stockout probability that should be attained in any period is α.

Tarim and Kingsman [14] developed the following MIP formulation for solving the stochastic lot sizing problem under the static–dynamic uncertainty strategy

$$\min \quad \sum_{t=1}^{N}\left(a\delta_t + h\mathsf{E}\{I_t\}\right) \tag{1}$$

$$\text{subject to} \quad \mathsf{E}\{I_t\} = \mathsf{E}\{S_t\} - \mathsf{E}\{d_t\} \quad t \in [1, N] \tag{2}$$

$$\mathsf{E}\{S_t\} \geq \mathsf{E}\{I_{t-1}\} \quad t \in [1, N] \tag{3}$$

$$\mathsf{E}\{S_t\} - \mathsf{E}\{I_{t-1}\} \leq M\delta_t \quad t \in [1, N] \tag{4}$$

$$\mathsf{E}\{I_t\} \geq \sum_{p=1}^{t}\left(G^{-1}_{d_{t-p+1}+\cdots+d_t}(\alpha) - \sum_{k=t-p+1}^{t}\mathsf{E}\{d_k\}\right)P_{tp} \quad t \in [1, N] \tag{5}$$

$$\sum_{p=1}^{t} P_{tp} = 1 \quad t \in [1, N] \tag{6}$$

$$P_{tp} \geq \delta_{t-p+1} - \sum_{k=t-p+2}^{t}\delta_k \quad t \in [1, N],\ p \in [1, t] \tag{7}$$

$$\mathsf{E}\{I_t\}, \mathsf{E}\{S_t\} \geq 0, \qquad \delta_t, P_{tp} \in \{0, 1\} \quad t \in [1, N],\ p \in [1, t] \tag{8}$$

where the decision variables are comprised of

- $\mathsf{E}\{I_t\}$ expected inventory level at the end of period t
- $\mathsf{E}\{S_t\}$ expected post-replenishment inventory level (order-up-to level if an order is scheduled, and expected opening inventory level otherwise) at period t
- δ_t binary variable that takes the value of 1 if an order is scheduled in period t, and 0 otherwise
- P_{tp} binary variable that takes the value of 1 if the most recent order prior to period t was in period $t-p+1$, and 0 otherwise.

Here, M is a sufficiently large positive number. The function $G_{d_{t-p+1}+\cdots+d_t}(\cdot)$ is the cumulative distribution of the random variable $d_{t-p+1}+\cdots+d_t$. It is important to remark that it could be possible to obtain these convolutions in closed form if period demands belong to the same distribution family (e.g. normal and gamma distributions), otherwise it is necessary to employ numerical methods. It is assumed that G is strictly increasing. Thus G^{-1} is uniquely defined. This assumption holds if the probability density function of period demands are strictly positive within their domains — which is the case for most commonly used demand distributions. It is also assumed, without loss of generality, that the initial inventory position I_0 (and hence its expectation $\mathsf{E}\{I_0\}$) is zero, and an order is placed in the first period.

The model can be summarized as follows. The objective function (1) minimizes the expected total inventory ordering and holding costs over the planning horizon. It should be noted that this model is designed for high service level regimes—as is usually the case for α service level models. Hence, when computing the expected holding costs, it is reasonable to approximate the expected inventory on-hand $\mathsf{E}\{\max(0, I_t)\}$, with the expected inventory level $\mathsf{E}\{I_t\}$ (see [2,14]). Constraint (2) coordinates inventory conservation by fixing the expected inventory carried forward to the difference between the post-replenishment inventory level and the expected demand. Constraint (3) rules out expected negative order quantities. Constraint (4) states that the expected post-replenishment inventory level can only exceed the expected opening inventory level if there is a scheduled order. Constraint (5) expresses the service level requirement. Here, the term $G^{-1}_{d-p+1+\cdots+d_t}(\alpha)$ is the minimum post-replenishment inventory level in period $t-p+1$ that is sufficient to provide the desired service level in period t, given the most recent order prior to period t was in period $t-p+1$. Notice that this term can be calculated offline. The constraint, thus, guarantees that the service level is satisfied by bounding the expected inventory level at period t from below, while conditioning on the timing of the most recent order prior to period t. Constraint (6) makes sure that the most recent replenishment period prior to any given period is unique. Constraint (7) states that no order should be placed in between any given period and its designated most recent replenishment period. Finally, (8) sets variable domains.

3. Reformulation

The reformulation builds on the idea of re-defining the problem by means of alternative decision variables which enable us to provide a stronger linear relaxation of the overall problem. In the following, we first introduce the preliminaries regarding the reformulation, and then provide the new MIP model.

Let us consider a case where we are given the timing of two consecutive replenishment periods, say periods i and j ($i < j$), and refer to the time interval $[i, j-1]$ as the replenishment cycle. The expected total cost to be incurred over the replenishment cycle can be written as

$$a + h\sum_{t=i}^{j-1}\left(y - \sum_{k=i}^{t}\mathsf{E}\{d_k\}\right) \tag{9}$$

where y is the post-replenishment inventory level of the ith period. Here, we remark that we approximate the expected inventory on-hand $\mathsf{E}\{\max(0, y - \sum_{k=i}^{t} d_k)\}$, with the expected inventory level $y - \sum_{k=i}^{t}\mathsf{E}\{d_k\}$. This is also the case in Tarim and Kingsman's model. The expression above is increasing on y. Thus, in order to obtain a lower cost, one would pick an order-up-to level as small as possible. However, in order to satisfy the α service level constraint, the order-up-to level should be larger than or equal to $G^{-1}_{d_i+\cdots+d_{j-1}}(\alpha)$. As a result, we can define a lower bound on the expected cycle cost as

$$c_{ij} := a + h\sum_{t=i}^{j-1}\left(G^{-1}_{d_i+\cdots+d_{j-1}}(\alpha) - \sum_{k=i}^{t}\mathsf{E}\{d_k\}\right). \tag{10}$$

Furthermore, building on this lower bound, we can re-write the expected cycle cost given in (9) as

$$c_{ij} + h(j-i)(y - G^{-1}_{d_i+\cdots+d_{j-1}}(\alpha)) \tag{11}$$

for all y such that $y \geq G^{-1}_{d_i+\cdots+d_{j-1}}(\alpha)$. The first part of this expression can be regarded as the static, and the second part as the incremental component of the expected total cycle cost. It is important