Int. J. Production Economics 157 (2014) 170–176

Contents lists available at ScienceDirect

Int. J. Production Economics

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

Simulation optimization for the stochastic economic lot scheduling problem with sequence-dependent setup times

Nils Löhndorf [a,*], Manuel Riel, Stefan Minner [b]

[a] Vienna University of Economics and Business, Vienna, Austria
[b] Technische Universität München, Munich, Germany

ARTICLE INFO

Article history:
Received 29 September 2012
Accepted 8 May 2014
Available online 28 May 2014

Keywords:
Inventory
Multi-product
Lot-sizing and scheduling
Stochastic demand
Sequence-dependent setups
Simulation optimization

ABSTRACT

We consider the stochastic economic lot scheduling problem (SELSP) with lost sales and random demand where switching between products is subject to sequence-dependent setup times. We propose a solution based on simulation optimization using an iterative two-step procedure which combines global policy search with local search heuristics for the traveling salesman sequencing subproblem. To optimize the production cycle, we compare two criteria: minimizing total setup times and evenly distributing setups to obtain a more regular production cycle. Based on a numerical study, we find that a policy with a balanced production cycle leads to lower cost than other policies with unbalanced cycles.

1. Introduction

The integration of lot-sizing, scheduling and safety stock planning is a major challenge in multi-product inventory optimization under stochastic demands. Customized manufacturing and the use of the same manufacturing facilities by multiple products require investments into flexible resources which need to be managed effectively with respect to utilization and inventories. Increasing product variety poses an additional challenge. Especially in the process industry, lot-sizes and production sequences play an important role where setup times considerably decrease facility utilization (see, e.g., Kallrath, 2002). The mainstream planning approach for these problems is to decompose the integrated into several sequential planning problems, typically a deterministic lot-sizing and scheduling problem modeled as a mixed-integer program and a safety stock planning problem using stochastic models. There exist only few approaches that follow an integrated approach, e.g., the (dynamic) capacitated lot-sizing problem under random demands (Helber et al., 2013) or the stochastic economic lot-scheduling problem (SELSP) we consider in this paper. Vaughan (2007) compares different scheduling and sequencing options for the stochastic economic lot scheduling problem. The literature on the SELSP is reviewed in Sox et al. (1999) and Winands et al. (2011).

One dominant assumption in many existing approaches is that setup costs and especially setup times are independent of the chosen production sequence. However, there are several practical problems where sequence-dependency plays a major role. First, there can be significant setup times between different product families (major setups), but small or negligible setup times between individual products belonging to the same family (see e.g. Karalli and Flowers, 2006). Second, setup times between products may be asymmetric, which includes the special case of one-way setup times, for instance due to different cleaning requirements or added machinery tools. Furthermore, the amount of required setup time between items of a certain family might be identical, or linearly (progressively, degressively) increasing. Overviews on available approaches for lot-sizing and sequencing problems with sequence-dependent setup costs and times are available in Allahverdi et al. (1999, 2008). Dobson (1992) analyzes the deterministic ELSP with sequence dependent setup times using the time-varying lot-size approach which was refined recently in Shirodkara et al. (2011). Liberopoulos et al. (2013) analyze the case of a stochastic economic lot-scheduling problem with restrictions in the production sequence.

Our model and the analysis are based on Löhndorf and Minner (2013), who analyze a similar problem under the assumption that setup times are independent of the sequence. They model the problem as a Semi-Markov Decision Problem and compare Approximate Dynamic Programming with direct policy search for fixed-cycle and base-stock policies using simulation optimization. Due to the complexity of the problem, several (meta-)heuristics combined with simulation have been proposed. Wagner and Smits (2004) suggest a local search approach. PaterninaArboleda and Das (2005) develop a multi-agent reinforcement learning approach. Kämpf and Köchel (2006) use a genetic algorithm-based simulation

* Corresponding author.
E-mail addresses: nils.loehndorf@wu.ac.at (N. Löhndorf), m@nuelriel.com (M. Riel), stefan.minner@tum.de (S. Minner).

http://dx.doi.org/10.1016/j.ijpe.2014.05.006
0925-5273/

optimization approach to find the parameters of structured policies. For a more detailed literature review, we refer to Löhndorf and Minner (2013). The sequence dependence of setup times adds another dimension of complexity to the problem. Even the deterministic version of the problem, the sequence dependent economic lot-scheduling problem (SD-ELSP), requires solution of a traveling salesman problem as a subproblem. Due to the overall problem complexity, we resort to a black box approach based on simulation optimization for simultaneously finding production cycles, base-stock-levels, and production frequencies. The use of simulation increases the scope of our method, as it allows practitioners to describe a production process using a simulation model, which offers a detailed but yet user-friendly representation of the actual production process. Based on this approach, we demonstrate that straightforward solutions to the SD-ELSP or SELSP are insufficient by comparing different production policies in a numerical study. Furthermore, we investigate the influence of different setup time characteristics on the performance of the policies. We consider three different policies which all have in common that the global policy search sets the base stock levels and that production follows a fixed pre-defined production cycle.

The paper is organized as follows: in Section 2 we introduce the problem and in Section 3 we first sketch the solution methodology and then the proposed manufacturing policies to be parameterized by simulation optimization. Section 4 reports the results of a numerical study and Section 5 summarizes the main findings and future research opportunities.

2. Model

We consider the continuous time stochastic economic lot scheduling problem with $n \in \{1, 2, \dots, N\}$ products. There is a single machine that can only manufacture one product at a time. If the machine state changes from one product to another, it has to be set up. This requires a deterministic, sequence-dependent setup time s_{nm} to change over from product n to product m. We further assume that the setup status is preserved over an idle period. The production of one unit of product n requires a deterministic production time p_n. During a setup or the production of a single item, interruption is not permitted. Inventories are subject to holding cost h_n per item and unit of time and cannot exceed a maximum inventory level $\overline{y}_n$. Demand for each product n follows a compound renewal process with inter-arrival distribution F_n^A and demand size distribution F_n^D. Inter-arrival times and demand size are independent for each product and across products. As in Altiok and Shiue (1995) and Krieg and Kuhn (2002), we assume that unsatisfied customer demand is lost at cost v_n per item, but we allow for partial fulfillment of an order.

The problem can be modeled and for very small instances with few products be solved using stochastic dynamic programming, see e.g. Graves (1980) or Löhndorf and Minner (2013).

While the model can be easily extended to handle setup cost in addition to setup times, in line with Krieg and Kuhn (2002), we assume that setup cost is negligible, since "often no explicit setup cost is incurred, and the latter is used merely to represent opportunity cost of setup times" (Federgruen and Katalan, 1996).

3. Solution method

3.1. Global policy search

As in Löhndorf and Minner (2013), we propose to directly search for optimal parameters of simple policies for production control. To guide the search for the optimal parameter vector, we use the CMA-ES algorithm (Hansen and Ostermeier, 2001). CMA-ES generates new candidate vectors from a multivariate normal distribution, i.e., $\mathcal{N}(\mu^x, \text{diag}(\sigma^x))$, which serves as an internal model of promising search steps. Throughout the search process, the algorithm updates the distribution's means and covariances to increase the likelihood of previously successful search steps.

(1) Input arguments: initial guess μ^x, trust region σ^x
(2) Do for $i = 1, 2, \dots, I$
  (2.1) Get $(x^1, ..., x^K) \leftarrow G^M(K)$ from internal model
  (2.2) Do for $k = 1, 2, \dots, K$
    (2.2.1) Do for $t = 1, 2, \dots, T$
      (2.2.1.1) Compute $(c_t, \tau_t, S) \leftarrow S^M(S, \pi(S; x^k))$
    (2.2.2) Compute $r^k \leftarrow \sum_{t=e+1}^{T} c_t \left(\sum_{t=1}^{T} \tau_t\right)^{-1}$
  (2.3) Update internal model $U^M((x^1, ..., x^K), (r^1, ..., r^K))$
(3) Return best solution x^*

Fig. 1. Generic policy search for production control.

Fig. 1 outlines a generic formulation of the CMA-ES algorithm. Denote $\pi(\cdot; x)$ as a control policy which is characterized by a (continuous) parameter vector x, and denote S^M as a sample from the state transition function of the SELSP which, for a given state S and decision $\pi(S; x)$, returns a realization of the immediate cost c, the sojourn time τ and the successor state S'. The objective of the algorithm is to search for an x that minimizes the expected average cost. The algorithm is initialized with a *guess* of the best solution, μ^x, as well as a *trust* region, σ^x, in which the solution is likely to be found. The main loop consists of three steps: (2.1) generation of a set of K candidate solutions $(x^1, \dots, x^K)$ from the internal model G^M which controls the search process; (2.2) simulating the transition process for T periods and recording the average cost for each candidate policy; (2.3) updating the internal model using the sampled information.

3.2. Production policies

Denote $Y = \{Y_1, \dots, Y_N\}$ as the set of base-stock levels and $Q = \{Q_1, \dots, Q_J\} \in \mathbb{Q}$ as the production sequence of length J, where $\mathbb{Q}$ is defined as the power set of Q. For a given position j in sequence Q and order-up-to levels Y, the production policy is given by

$$\pi(S; x) = \begin{cases} 0 & \text{if } y_n = Y_n \quad \forall\ n, \\ Q_{z(j)} & \text{otherwise,} \end{cases} \tag{1}$$

where the recursive function z is defined as

$$z(j) = \begin{cases} j & \text{if } y_k < Y_k : k = Q_j, \\ z(j \bmod J + 1) & \text{otherwise.} \end{cases} \tag{2}$$

For a given position j, the function returns the next position in the sequence for which $y_n < Y_n$, where the modulus ensures that the production cycle is repeated as soon as $j = J$.

In contrast to Löhndorf and Minner (2013), sequence-dependent setup times have to be taken into account when constructing a production sequence from policy parameters. We therefore apply an iterative two-step procedure which combines the heuristic (local) search with the global policy search to jointly optimize base-stock levels, as well as the production sequence.

3.2.1. Common cycle policy (CCP)

The simplest production policy with a fixed production sequence is the *common cycle policy*, where each product is produced exactly once during a cycle. The optimal production sequence is set in advance by finding the sequence with the minimum total setup time. Since the production sequence remains constant, the global policy search merely has to set the base-stock levels. To find a production sequence which minimizes the total