THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

A mixed integer linear programming heuristics for computing nonstationary (s,S) policy parameters

Citation for published version:
Rossi, R, Kilic, O & Tarim, SA 2014, 'A mixed integer linear programming heuristics for computing nonstationary (s,S) policy parameters', Paper presented at MAGO 2014 - XII Global Optimisation Workshop (GOW), Malaga, Spain, 1/09/14 - 4/09/14.

Link:
Link to publication record in Edinburgh Research Explorer

Document Version:
Early version, also known as pre-print

Publisher Rights Statement:
© Rossi, R., Kilic, O., & Tarim, S. A. (2014). A mixed integer linear programming heuristics for computing nonstationary (s,S) policy parameters. Paper presented at MAGO 2014 - XII Global Optimisation Workshop (GOW), Malaga, Spain.

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.

OPEN ACCESS

Download date: 25. Apr. 2024

Proceedings of MAGO 2014, pp. 1 – 4.

A mixed integer linear programming heuristics for computing nonstationary (s,S) policy parameters

Roberto Rossi,[1] Onur A. Kilic,[2] and S. Armagan Tarim[2]

[1] *Business School, University of Edinburgh, Edinburgh, United Kingdom,* roberto.rossi@ed.ac.uk

[2] *Institute of Population Studies, Hacettepe University, Ankara, Turkey* {onuralp,armagan.tarim}@hacettepe.edu.tr

Abstract In this work we present a novel MILP based heuristics for computing nonstationary (s,S) policy parameters. This approach presents advantages with respect to other existing methods, since it is easy to implement and features narrower optimality gaps.

Keywords: stochastic lot sizing, (s,S) policy, nonstationary demand, mixed integer linear programming, heuristics

1. Introduction

The stochastic lot sizing problem consists in controlling an inventory system facing random demand over a given planning horizon. The decision maker faces inventory holding costs, if she orders too much; and backorder penalty costs, if she orders too little and demand fulfilment is delayed until the next replenishment arrives. Each time production runs there are fixed and variable production/ordering costs that must be accounted for while controlling the system. The structure of the optimal control policy to this problem has been characterised — under very mild assumptions — over fifty years ago [8]. This control policy, named (s,S), is surprisingly simple; this policy monitors the inventory position, i.e. on hand stock minus backorders plus incoming orders, and issues an order to bring the inventory position up to S whenever the inventory position falls below s.

As pointed out by [4] incorporating more realistic assumptions about product demand constitutes an important research direction in inventory theory. The ability to model and control a nonstationary demand process is essential in practical settings, since only very few businesses actually face stationary demand, while most products are subject to demand processes that evolve over time with frequent changes in their directions and rates of growth or decline.

When demand is nonstationary an (s,S) policy is still cost optimal. However, computing optimal control parameters for this policy constitutes a hard combinatorial task. Standard pseudo-polynomial dynamic programming (DP) algorithms can only tackle small instances. This motivates the investigation of effective heuristics. To date, there are only two established heuristics for computing optimal control policy parameters under nonstationary demand [1, 2]. Unfortunately, these heuristics present a number of drawbacks. They are not easy to implement, since they require dedicated code. Furthermore, in a recent study [3], their respective optimality gap on a large test bed has been found to average 4% to 5%. The same work also demonstrated that approaches such as [9, 6], despite implementing heuristics for control policies that are theoretically inferior to a nonstationary (s,S) policy, feature much lower optimality gaps, i.e. around 1.5%, on the same test bed. This demonstrates that further research is needed to develop more effective heuristics for computing nonstationary (s,S) control policy parameters.

In this work we develop an MILP based heuristics for computing nonstationary (s,S) policy parameters. The key insight upon which sour approach is based comes from the study in [3], which showed that a nonstationary (R,S) policy often performs very close to optimal. The idea is then to use an existing MILP model for computing nonstationary (R,S) policy parameters [5] as a proxy to determine near optimal (s,S) policy parameters. Our heuristics is easy to implement, since it is based solely on a standard MILP model and on a simple binary search procedure. It performs better than other existing approaches, featuring an average optimality gap of 0.2% on our preliminary tests.

2. The stochastic lot sizing problem

The finite-horizon single-item single-stocking location nonstationary stochastic lot sizing problem as introduced in [8] can be formalised as follows. We consider a finite planning horizon of n periods. Customer demand d_t in each period $t = 1, \ldots, n$ is a random variable with known probability distribution. There are three types of costs: a nonlinear purchasing or ordering cost c(z), where z is the amount purchased, which takes the general form

$$c(z) = \begin{cases} K + vz & \text{if } \ z > 0 \\ 0 & \text{otherwise} \end{cases}$$

where K and v denote the fixed and variable purchasing/ordering cost components, respectively; a holding cost of h is paid of each unit of inventory carried from one period to the next; and a shortage cost p which is paid for each unit of demand backordered at the end of a period. Holding and shortage costs are charged at the end of a period. Ordering costs are charged when a purchase is made. Without loss of generality, see [8], delivery of an order is immediate.

Let y denote the stock level immediately after purchases are delivered, the expected holding and shortage cost for a generic period are given by

$$L(y) = \begin{cases} \displaystyle\int_0^y h(y-\omega)g(\omega)\mathrm{d}\omega + \int_y^\infty p(\omega - y)g(\omega)\mathrm{d}\omega & y \geq 0 \\ \displaystyle\int_0^\infty p(\omega - y)g(\omega)\mathrm{d}\omega & y < 0 \end{cases}$$

where $g_t(\cdot)$ denotes the probability density function of the demand in period t. If the initial inventory at the beginning of the planning horizon is x and $C_n(x)$ represents the expected total cost over the n-periods planning horizon if provisioning is done optimally then $C_n(x)$ satisfies

$$C_n(x) = \min_{y \geq x} \left\{ c(y-x) + L_n(y) + \int_0^\infty C_{n-1}(y-\omega)g_n(\omega)\mathrm{d}\omega \right\}$$

If $y_n(x)$ is the argument minimising the above functional equation, then $y_n(x) - x$ denotes the optimal initial purchase.

3. (s,S) policy

As shown in [8], the optimal control policy for the problem introduced in Section 2 takes a surprisingly simple form. The result stems from the study of the following function

$$G_n(y) = cy + L_n(y) + \int_0^\infty C_{n-1}(y-\omega)g_n(\omega)\mathrm{d}\omega$$

More specifically, Scarf proved that $G_n(y)$ is K-convex.

Definition 1. Let $K \geq 0$, and let $f(x)$ be a differentiable function, $f(x)$ is K-convex if

$$K + f(a + x) - f(x) - af'(x) \geq 0$$

for all positive a and all x. This definition can be extended to a non differentiable function.

It follows that, under general nonstationary settings, the optimal policy can be described via n pairs (s_i,S_i), where s_i denotes the reorder point and S_i the order-up-to-level for period i. In practice, S_n denotes the absolute minimum of $G_n(y)$ and $s_n < S_n$ is the unique value such that $K + G_n(S_n) = G_n(s_n)$. The fact that $G_n(y)$ is K-convex ensures that ripples in the above nonlinear function do not affect the existence of a unique reorder point $s_n \leq S_n$, since their height will never exceed K.

We shall now illustrate graphically the notion of K-convexity on a simple numerical example. Consider a planning horizon of $n = 4$ period and a demand d_t normally distributed in each period t with mean $\mu_t \in \{20, 40, 60, 40\}$, for period $t = 1, \ldots, n$ respectively. The standard deviation σ_t of the demand in period t is equal to $0.25\mu_t$. Other problem parameters are $K = 100$, $h = 1$ and $p = 10$; to better conceptualise the example we let $v = 0$. In Fig. 1 we plot $G_n(y)$ for an initial inventory $y \in (0, 200)$. In period one, when the opening inventory level y falls below

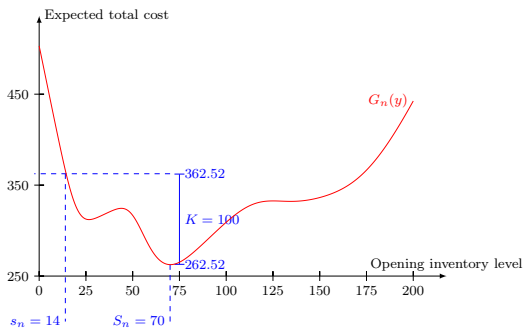

Figure 1. Plot of $G_n(y)$

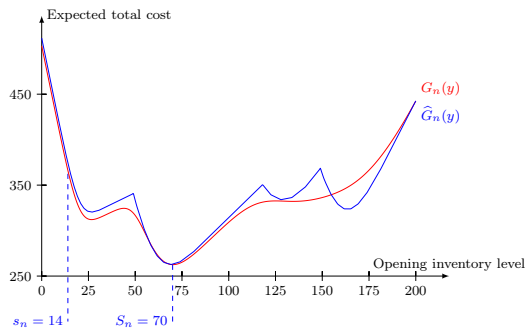

Figure 2. $G_n(y)$ vs $\widehat{G}_n(y)$

14 it is convenient to pay the fixed ordering cost K to increase available inventory to S_n, i.e. $K + G_n(S_n) = G_n(s_n) \leq G_n(y)$. Comparable graphs can be produced for all other periods.

4. (R,S) policy

A widely adopted control policy, alternative to the (s,S) policy, is the (R,S) policy. In this policy, all replenishment periods must be fixed at the beginning of the planning horizon; however, the decision maker can decide upon the actual order quantity just before issuing a replenishment. Under a nonstationary settings this policy takes the form (δ_i,S_i), where δ_i is a binary variable that is set to 1 if a replenishment is scheduled in period i and S_i denotes the order-up-to-level associated with a replenishment that occurs in period i. [5] developed a mixed integer linear programming model to compute near-optimal (R,S) policy parameters. To model nonlinear expected holding and shortage costs the authors exploit piecewise linear upper and lower bounds of the first order loss function [7]. An interesting feature of the model in [5] is the fact that, despited being explicitly developed for the (R,S) policy, it can nevertheless be used as a "proxy" to the expected total cost of an (s,S) policy, i.e. $G_n(y)$. In fact, we can first observe that the optimal expected total cost and the order-up-to-level for period one returned by the model for an initial stock level of x units are tight approximation to $C_n(x)$ and $y_n(x) = S_n$, respectively. Furthermore, if we set $\delta_1 = 0$ — i.e. we do not schedule any replenishment at the beginning of the planning horizon — since $G_n(y)$ is K-convex, there is a unique reorder point $s_n < S_n$ such that $K + G_n(S_n) = G_n(s_n)$. We can therefore exploit a binary search on $y < S_n$

Table 1. Optimal policy for the numerical example; expected total cost (ETC) estimated at 95% confidence.

	SDP — ETC: (362.2,362.9)		MILP — ETC: (363.0,363.1)	
t	S_t	s_t	S_t	s_t
1	70.0	14.0	70.2	15.0
2	141	29.5	53.9	29.0
3	113	58.0	116	58.1
4	53.5	28.5	53.9	29.0

to find s_n. In the binary search procedure, the cost associated with a given opening inventory level y can be approximated using $\widehat{G}_n(y)$, the solution of the MILP model. We then repeat this procedure to find S_i and s_i for each period $i = 1, \ldots, n$, by analysing $\widehat{G}_n(y), \widehat{G}_{n-1}(y), \ldots$.

In Fig. 2, for the numerical example previously discussed, we plot $\widehat{G}_n(y)$, obtained via the MILP model in [5], when we vary y, i.e. the opening inventory level at the beginning of the planning horizon. We also compare it to the plot of $G_n(y)$ obtained via a standard DP approach. The optimal policy found via DP for the above example is contrasted in Table 1 against the policy obtained via the MILP heuristics.

5. Computational Experience

We conducted a comprehensive numerical analysis on 1152 small instances over an 8-period planning horizon. These instanced have been solved to optimality via DP. Factors that have been varied in our analysis include demand pattern, fixed and proportional ordering cost, holding cost, penalty cost and demand variability. All instances could be solved in about 10 secs by our method, as opposed to about a minute required by the DP code. The average optimality gap observed for our approach is 0.2%. In contrast, [1, 2] feature optimality gaps of 2.09% and 3.52%, respectively; however, these latter heuristics are faster than ours.

References

[1] Ronald G. Askin. A procedure for production lot sizing with probabilistic dynamic demand. AIIE Transactions, 13(2):132–137, June 1981.

[2] Srinivas Bollapragada and Thomas E. Morton. A simple heuristic for computing non-stationary (s, S) policies. Operations Research, 47(4):576–584, 1999.

[3] Gozdem Dural-Selcuk, Onur A. Kilic, S. Armagan Tarim, and Roberto Rossi. A comparison of lot sizing methods for non-stationary stochastic demands, December 2013.

[4] S. C. Graves. A Single-Item Inventory Model for a Non-Stationary Demand Process. Manufacturing & Service Operations Management, 1:50–61, 1999.

[5] Roberto Rossi, Onur A. Kilic, and S. Armagan Tarim. Piecewise linear approximations for the static-dynamic uncertainty strategy in stochastic lot-sizing, December 2013.

[6] Roberto Rossi, S. Armagan Tarim, Brahim Hnich, and Steven Prestwich. Constraint programming for stochastic inventory systems under shortage cost. Annals of Operations Research, 195(1):49–71, July 2012.

[7] Roberto Rossi, S. Armagan Tarim, Steven Prestwich, and Brahim Hnich. Piecewise linear approximations of the standard normal first order loss function. Applied Mathematics and Computation, to appear.

[8] Herbert E. Scarf. Optimality of (s, S) policies in the dynamic inventory problem. In K. J. Arrow, S. Karlin, and P. Suppes, editors, Mathematical Methods in the Social Sciences, pages 196–202. Stanford University Press, Stanford, CA, 1960.

[9] S. Armagan Tarim and Brian G. Kingsman. Modelling and computing (Rn,Sn) policies for inventory systems with non-stationary stochastic demand. European Journal of Operational Research, 174(1):581–599, October 2006.