Operations Research Letters 43 (2015) 359–364

Contents lists available at ScienceDirect

Operations Research Letters

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

A polynomial time algorithm for convex cost lot-sizing problems

Z. Melis Teksan, Joseph Geunes *

Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL, 32611, USA

ARTICLE INFO

Article history:
Received 4 December 2014
Received in revised form
25 March 2015
Accepted 28 March 2015
Available online 17 April 2015

Keywords:
Convex cost lot sizing
Polynomial time algorithms
Production and inventory systems
Supply chain management

ABSTRACT

This paper provides a polynomial-time algorithm for economic lot-sizing problems with convex costs in the production and inventory quantities. The resulting algorithm is based on a primal–dual approach that takes advantage of the problem's special structure. This approach improves upon existing results in the literature, which are either pseudo-polynomial or focus on special cases. We apply the approach to a production planning problem with price-dependent supply, leading to an improved bound on the algorithm's running time for a special case.

1. Introduction

The model. We consider the classic discrete-time, finite-horizon economic lot-sizing problem with nondecreasing and convex costs in the production quantities and inventory levels in each period. This problem considers a set of consecutive demand periods and seeks to meet deterministic demand for a product in each period $t = 1, 2, \dots, T$ without lost sales or backlogging, at a minimum total cost over the planning horizon of length T. The total cost incurred over the horizon consists of those costs associated with producing the product in any period as well as the cost of holding inventory of the product between periods. At the beginning of each period, the production quantity in the period is added to remaining inventory from the prior period with zero lead time, and the sum of these two quantities must be at least as great as the demand in the period (with the difference comprising the amount of inventory that will remain at the end of the period). Because all demands and costs are assumed deterministic, all production decisions may be made in advance of the first planning period, and the model can in principle accommodate any finite, deterministic production planning lead time. We assume throughout that no capacity limit exists on the production quantity or the inventory held in any period.

Of particular importance in approaching the solution of problems in this class are the nature and structure of the production and inventory holding costs. Because a wide range of practical settings exist with economies of scale in production, the cost of production in a period is often modeled as a nondecreasing and concave function of the production quantity, while conventional approaches typically treat the cost of holding inventory as a linear function of the inventory amount at the end of the period. This leads to a total cost function that may be expressed as the sum of concave functions of the production quantities, implying a concave total cost function. Conversely, production processes and inventory costs with diseconomies of scale may often be modeled using convex functions of the production quantities and inventory levels. This paper considers this latter class of problems with nondecreasing convex cost functions.

New results. Our main contribution lies in providing a polynomial-time algorithm for the lot-sizing problem we have described with general nondecreasing convex production and holding cost functions. We discuss two ways to implement our approach. The first is an iterative numerical approach that runs in $\mathcal{O}(T^2 \max\{\log T, \log \mathcal{S} \log \mathcal{M}/\epsilon\})$ time. Here T denotes the number of time periods and $\mathcal{S}$ provides an upper bound on the number of non-differentiable points of the cost to supply demand in period s using production in period t for any (t, s) pair with $1 \le t \le s \le T$. The value of $\mathcal{M}$ is an upper bound on the marginal production cost in any period, and ϵ denotes a stopping criterion for a bisection search routine in the proposed algorithm. As we will see in Section 2, for the case in which all of the cost functions are piecewise linear and convex, this iterative approach requires $\mathcal{O}(T^2 \max\{\log T, (\log \mathcal{S})^2\})$ time. The second implementation we provide requires repeated solution of a system of equations with at most $T+1$ equations and $T+1$ variables. Assuming this system of equations can be solved in $\mathcal{O}(\phi(T))$ time, where $\phi(T)$ is a polynomial function of T, this solution approach requires $\mathcal{O}(T \max\{T \log T, \phi(T)\})$ time.

* Corresponding author.
E-mail addresses: zmteksan@ufl.edu (Z.M. Teksan), geunes@ise.ufl.edu (J. Geunes).

http://dx.doi.org/10.1016/j.orl.2015.03.009
0167-6377/

When all costs are differentiable and quadratic, then $\mathcal{O}(\phi(T)) = \mathcal{O}(T)$, and this expression becomes $\mathcal{O}(T^2 \log T)$. To the best of our knowledge, the literature does not contain a special-purpose algorithm for this general problem class that runs in polynomial time in the worst case. Although the resulting problem is a convex optimization problem, the application of a general purpose nonlinear programming solver to problems with special structure may result in unnecessarily long solution times, and the associated running time would be weakly polynomial, regardless of the structure of the cost functions. The algorithm we provide is based on a primal–dual solution approach derived from analyzing the special structure of the problem's generalized Karush Kuhn–Tucker (KKT) conditions, which are necessary and sufficient for optimality. We specialize the resulting algorithm for application to a production planning context involving price-dependent supply components, where component supply is linearly increasing in price. The resulting problem is a special case of the general lot-sizing problem with nondecreasing and convex production costs, where the production costs take a quadratic form and the associated complexity is $\mathcal{O}(T^2 \log T)$.

Directly-related results. Veinott [11] considered a single-stage, dynamic lot-sizing problem in which production and inventory costs are piecewise-linear and convex. He designed a parametric-programming-based procedure in which the solution for a problem with a fixed parameter set is built upon the solution of another problem with a similar parameter set. Veinott [11] assumed all parameters were integer valued, and his procedure resulted in a pseudo-polynomial solution algorithm. The time complexity the algorithm is $\mathcal{O}(TD)$, where $D = \sum_{t=1}^{T} d_t$ denotes the sum of all demands over the planning horizon. Florian et al. [5] referred to Veinott's procedure as the most attractive approach to solve lot sizing problems with convex production and inventory costs and without fixed setup costs, even though the problem is demonstrated to be no harder than linear programming, which is polynomially solvable.

The work by Kian et al. [7] is also closely related to ours, as they analyzed a single-stage, uncapacitated economic lot-sizing problem with fixed setup costs and variable costs in each period that are convex in the production quantity (taking the form of a polynomial function of the production quantity). They derived several key optimality conditions for this problem class, as well as a dynamic programming solution algorithm that is exponential in the length of the time horizon. They presented an exact solution algorithm for the lot-sizing problem with zero setup costs and production costs taking the form of polynomial convex functions, and stated that the worst-case time complexity of such an algorithm would be $\mathcal{O}(T^2)$. In Section 3, we will explain why this bound applies to the number of subplans that must be considered, and not to the problem's overall worst-case complexity.

Related work. The problem we consider falls into the class of single-stage, dynamic lot-sizing problems, which has been studied by many researchers, starting with the seminal work of Wagner and Whitin [12], who first considered the problem with concave production costs. Brahimi et al. [1] provided an extensive review of uncapacitated and capacitated versions of the single-stage, dynamic lot-sizing problem. The uncapacitated version of the classical single-stage, dynamic lot-sizing problem with concave costs is polynomially solvable, whereas the general capacitated version of the problem with concave costs is $\mathcal{NP}$-Hard, although the uniform-capacity version can be solved in polynomial time via dynamic programming (see Florian and Klein [4]). Many studies have considered variations of the single-stage, dynamic lot-sizing problem (see e.g. Van Hoesel and Wagelmans [10], and Van den Heuvel and Wagelmans [9]). In these studies, all costs are assumed to be concave in the production and inventory levels.

The literature on lot-sizing problems containing convex production costs is reasonably sparse, with a few notable exceptions. Erenguc and Aksoy [2] considered a single-item, capacitated dynamic lot-sizing problem with fixed production setup costs and linear inventory costs, while variable production costs were piecewise-linear and convex in the production quantity in a period. They used a branch-and-bound algorithm for this problem, which contains neither a convex nor concave objective function. Shaw and Wagelmans [8] developed a pseudo-polynomial dynamic program to solve a capacitated single-item lot-sizing problem with piecewise-linear production costs. Their algorithm can be utilized to solve problems with piecewise-linear and convex production costs, although it does not require any special structure for the piecewise-linear cost function. Feng et al. [3] developed an $\mathcal{O}(T \log T)$ algorithm for the single-item lot-sizing problem with constant capacity, convex inventory costs, and non-increasing fixed order costs.

Paper organization. The rest of this paper is organized as follows. Section 2 provides a general model formulation, description of the Karush Kuhn–Tucker (KKT) conditions for the problem, and development of a polynomial-time algorithm, along with a comparison of this approach to the one provided by Veinott in [11]. Section 3 discusses the production planning problem with price-dependent supply, and illustrates the application of the model to a practical special case that falls within the general problem class we consider. Section 4 contains brief concluding remarks.

2. Problem formulation and solution method

The lot-sizing problem requires meeting a set of demands d_t, for $t = 1, \ldots, T$ without shortages at a minimum total production and inventory holding cost over the horizon of length T. Letting x_t denote the production quantity in period t, it will be convenient and useful to keep track of the amount produced in period t in order to meet demand in period τ, $x_{t\tau}$, where $x_t = \sum_{\tau=t}^{T} x_{t\tau}$ for $t = 1, \ldots, T$. We assume that the cost to produce x_t units in period t is a nondecreasing convex function $f_t(x_t)$. In addition, $h_t(i_t)$ denotes a nondecreasing and convex inventory holding cost function, which depends on the inventory at the end of period t, denoted by i_t. The inventory remaining at the end of period t can be equivalently written as $i_t = \sum_{\tau=1}^{t}\sum_{i=\tau}^{T} x_{\tau i} - \sum_{\tau=1}^{t} d_\tau$; thus, we can alternatively write h_t as a function of the production variables using the expression $h_t\left(\sum_{\tau=1}^{t}\sum_{i=\tau}^{T} x_{\tau i}\right)$, where we have suppressed the dependence of h_t on cumulative demand up to period t for notational convenience. We assume that each of the functions f_t and h_t is everywhere locally Lipschitz continuous. The convex cost lot-sizing problem may then be formulated as follows.

$$\text{P}: \text{ Minimize } \sum_{t=1}^{T}\left\{f_t\left(\sum_{\tau=t}^{T} x_{t\tau}\right) + h_t\left(\sum_{\tau=1}^{t}\sum_{i=\tau}^{T} x_{\tau i}\right)\right\} \tag{1}$$

$$\text{Subject to}: \sum_{t=1}^{\tau} x_{t\tau} = d_\tau, \quad \tau = 1, \ldots, T, \tag{2}$$

$$x_{t\tau} \geq 0, \quad \tau = 1, \ldots, T, \text{ and } t \leq \tau. \tag{3}$$

As the sum of convex functions, the objective function of problem P is convex; this combined with the linear constraint set implies that the generalized KKT conditions are necessary and sufficient for optimality (see Hiriart-Urruty [6]).

Generalized KKT conditions. To characterize the generalized KKT conditions, let $\partial f_t(X^t)$ denote the generalized gradient of f_t at X^t, where $X^t = (x_{tt}, x_{tt+1}, \ldots, x_{tT})$. If f_t is differentiable at X^t, then $\partial f_t(X^t)$ consists of a singleton equal to the partial derivative with respect to any element of X^t at X^t; otherwise, $\partial f_t(X^t)$ corresponds to the set of subgradients at X^t. Similarly, let $\bar{X}^t = (X^1, X^2, \ldots, X^t)$, and let $\partial h_t(\bar{X}^t)$ denote the generalized gradient of h_t at $\bar{X}^t$.