

An $O(T^3)$ algorithm for the economic lot-sizing problem with constant capacities

Citation for published version (APA):

van Hoesel, C. P. M., & Wagelmans, A. P. M. (1993). An $O(T^3)$ algorithm for the economic lot-sizing problem with constant capacities. (Memorandum COSOR; Vol. 9337). Technische Universiteit Eindhoven.

Document status and date:

Published: 01/01/1993

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

EINDHOVEN UNIVERSITY OF TECHNOLOGY
Department of Mathematics and Computing Science

Memorandum COSOR 93-37

**An $O(T^3)$ algorithm for the economic
lot-sizing problem with constant capacities**

C.P.M. van Hoesel

A.P.M. Wagelmans

Eindhoven, November 1993
The Netherlands

Eindhoven University of Technology
Department of Mathematics and Computing Science
Probability theory, statistics, operations research and systems theory
P.O. Box 513
5600 MB Eindhoven - The Netherlands

Secretariat: Dommel building 0.03
Telephone: 040-47 3130

ISSN 0926 4493

An $O(T^3)$ algorithm for
the economic lot-sizing problem
with constant capacities

C.P.M. van Hoesel * A.P.M. Wagelmans †

June 1993

Abstract

We develop an algorithm that solves the constant capacities economic lot-sizing problem with concave production costs and linear holding costs in $O(T^3)$ time. The algorithm is based on the standard dynamic programming approach which requires the computation of the minimal costs for all possible subplans of the production plan. Instead of computing these costs in a straightforward manner, we use structural properties of optimal subplans to arrive at a more efficient implementation. Our algorithm improves upon the $O(T^4)$ running time of an earlier algorithm.

Key words: Economic lot-sizing, complexity.

AMS Subject classification: 90B.

*Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands. E-mail: HOESEL@BS.WIN.TUE.NL

†Erasmus University Rotterdam, P.O. Box 1738, 3000 DR Rotterdam, The Netherlands. E-mail: WAGELMAN@CVX.EUR.NL

1 Introduction

In the single-item capacitated economic lot-sizing problem there is demand for a single item in T consecutive periods. The demand in a certain period may be satisfied by production in that period or from inventory that has been produced in earlier periods. It is assumed that there is no inventory at the beginning of period 1 and that no inventory should be left at the end of period T . Furthermore, capacity constraints on the production levels have to be taken into account. The total costs associated with a production plan depend on the production and inventory levels. A fixed set-up cost is incurred when production takes place in a certain period. In addition there are production costs which are a function of the production level. Finally, there are holding costs, which are a function of the inventory level at the end of the period. The objective is to find a feasible production plan that minimizes total costs.

In most models that have been studied in the literature, the cost functions are assumed to be concave or linear. Under these assumptions, many uncapacitated models are polynomially solvable. For instance, if all cost functions are linear, then the uncapacitated version of the above problem is solvable in $O(T \log T)$ time (cf. Aggarwal and Park [1], Federgruen and Tzur [5], Wagelmans et al. [10]). Polynomial algorithms also exist for many other uncapacitated lot-sizing problems with linear costs (cf. Van Hoesel et al. [8]). The uncapacitated problem with concave production and holding costs is solvable in $O(T^2)$ time (cf. Veinott [9]).

For capacitated problems the situation is quite different. Florian et al. [7] and Bitran and Yanasse [2] have shown that the single item capacitated economic lot-sizing problem is NP-hard, even in many special cases. Bitran and Yanasse also designed a classification scheme for capacitated lot-sizing problems with linear production and holding costs. They use the four field notation $\alpha/\beta/\gamma/\delta$, where α , β , γ and δ represent the set-up cost, unit holding cost, unit production cost and capacity type, respectively. Each of the parameters α , β and γ can take on one of the values G , C , ND , NI or Z . G means that the parameter follows an arbitrary pattern over time, whereas C , ND , NI and Z indicate constant, non-decreasing, non-increasing respectively zero parameter values. δ can take on the values G , C , ND or NI ; in case there are no capacity restrictions, the fourth field is omitted.

A very successful DP approach to solve the most general problem, $G/G/G/G$, has recently been proposed by Chen et al. [3]. We also refer to that paper for a discussion of other work that has been done on NP-hard versions of the capacitated economic lot-sizing problem.

With respect to polynomially solvable special cases of the capacitated economic lot-sizing problem, the following results are known. Bitran and Yanasse showed that $ND/Z/ND/NI$ and $C/Z/C/G$ can be solved in $O(T)$ respectively $O(T \log T)$ time, Chung and Lin [4] gave an $O(T^2)$ algorithm for $NI/G/NI/ND$ and an $O(T^4)$ algorithm for $G/G/G/C$ was presented by Florian and Klein [6]. The latter algorithm also solves the more general constant capacity problem in which the cost functions are concave instead of linear.

In this paper we will show that when the production costs are concave and the holding costs are linear, it is possible to solve the economic lot-sizing problem with constant capacities in $O(T^3)$ time. Hence, for this case we improve upon the Florian–Klein algorithm.

This paper is organized as follows. In Section 2 we introduce notation and make a few preliminary remarks. Section 3 contains a description of the new algorithm and concluding remarks can be found in Section 4.

2 Preliminaries

We will use the following notation:

T : the length of the planning horizon;
 C : the capacity of the production in each period;

Furthermore, for each period $t \in \{1, \dots, T\}$:

d_t : the demand for the item in t ;
 x_t : the production level in t ;
 I_t : the inventory level at the end of t ($I_0 = 0$);
 f_t : the set-up cost in t ;
 h_t : the unit holding cost in t ;
 $p_t(x_t)$: the production costs in period t , a concave function of x_t .

Without loss of generality we may assume:

- (a) For each period t : $d_t \leq C$. If this is not the case, we can move the excess demand in t to the preceding period $t - 1$ without changing the set of feasible solutions.
- (b) The holding costs are all equal to zero. If this is not the case, then an equivalent problem is obtained by omitting the holding costs and redefining the production costs as $\tilde{p}_t(x_t) = p_t(x_t) + \sum_{i=t}^T h_i x_i$ (cf. Wagelmans et al. [10]). Note that we can achieve this only when the original holding costs are linear.

A feasible production plan can be subdivided into several *subplans*. Such a subplan (s, t) ($1 \leq s \leq t \leq T$) consists of a set of consecutive periods, starting with s and ending with t . It is characterized by the fact that there is no incoming inventory and no leaving inventory, i.e., $I_{s-1} = I_t = 0$, and the intermediate inventories are all positive, i.e., $I_s, \dots, I_{t-1} > 0$.

Florian and Klein [6] have shown that there is an optimal schedule in which each subplan contains at most one period in which production is positive and less than capacity, i.e., for such a period τ we have $0 < x_\tau < C$. This property is often referred to as the *fractional production property* and it even holds for general capacities.

The algorithm that will be presented consists of two phases:

Phase 1: Find the minimum cost of the subplan (s, t) , for each pair s, t ($1 \leq s \leq t \leq T$).

Phase 2: Find, in a network with vertices $\{1, \dots, T + 1\}$ and arcs $(s, t + 1)$ ($1 \leq s \leq t \leq T$), the shortest path from vertex 1 to vertex $T + 1$, where the length of arc $(s, t + 1)$ is equal to the minimum cost of subplan (s, t) .

Except for vertex $T + 1$, the vertices on the shortest path found in phase 2 correspond to the first periods of the subplans of an optimal production plan. Given the cost of each subplan, the second phase can be implemented in $O(T^2)$ time, since the network is acyclic and the number of arcs is $O(T^2)$. In the next section we will give an $O(T^3)$ algorithm to solve phase 1.

3 The algorithm

We will show in this section that the optimal cost of each subplan can be calculated in $O(T)$ amortized time. Without loss of generality, we consider subplan $(1, T)$. Since we have a fixed subplan, the production in the “fractional” production period can be determined as $f : 0 \leq f < C$ such that $f + KC = \sum_{t=1}^T d_t$ for some integer K . Hence, K is the number of times we will produce at full capacity in the subplan. For notational convenience, we let $cost(t)$ denote the cost of producing at full capacity in period t , i.e., $cost(t) = f_t + \tilde{p}_t(C)$ ($t \in \{1, \dots, T\}$).

Our algorithm solves the problem iteratively, by calculating for all relevant $t \in \{1, \dots, T\}$ the optimal solution when the fractional period is fixed to t . We will denote the latter problem by $P(t)$. Clearly, we can restrict the fractional production f to periods t with $\sum_{i=t}^T d_i \geq f$, since fractional production in later periods will lead to positive ending inventory in period T , and therefore to infeasible production plans. Therefore, we define l to be the latest period such that $\sum_{i=l}^T d_i \geq f$.

In subsection 3.1 we will describe a greedy algorithm that solves $P(t)$ for a fixed value of $t \in \{1, \dots, l\}$ and prove the optimality of this approach. We will also study structural properties of the optimal solution of $P(t)$ and of the optimal solution of a closely related problem. In subsection 3.2 the iterative algorithm will be described. The idea is that the aforementioned structural properties can be exploited such that it is not necessary to perform the greedy algorithm explicitly for every relevant value of t . Finally, in subsection 3.3, some details on data will be added and the claimed running time will be proved.

3.1 Greedy algorithm

Suppose that the fractional period is taken equal to s ($1 \leq s \leq l$), then the corresponding optimal solution can be determined by the following greedy algorithm.

We start with the (infeasible) production plan in which only the fractional production takes place, i.e., the inventory levels are $I_t = -\sum_{i=1}^t d_i$ for $t = 1, \dots, s-1$ and $I_t = f - \sum_{i=1}^t d_i$ for $t = s, \dots, T$. The K periods in which we produce at full capacity are chosen as follows. We consider the periods from 1 to T in increasing order. If $I_t < 0$ when period t is considered, then the cheapest available period in the set $\{1, \dots, t\}$ is chosen as production period (if necessary, we break ties by choosing the earliest cheapest period). Because period t forces us to choose a production period, we call this period a *choice period*. If we decide to produce at capacity in period τ , then the ending inventory of all periods $\{\tau, \dots, T\}$ increases with C . Thus, I_t becomes nonnegative (because $d_t \leq C$ and $I_{t-1} \geq 0$). The greedy aspect is that the decision to produce is always taken as late as possible (thereby maximizing the number of periods available for production) and that, among all periods available for production, the cheapest one is chosen.

The following lemma is easy to prove.

Lemma 1 *Let I_t^0 ($1 \leq t \leq T$) denote the inventory levels at the start of the greedy algorithm after the fractional production has been assigned. Then w_k , the k -th choice period, is the first period $t \in \{1, \dots, T\}$ for which $I_{t-1}^0 \geq -(k-1)C$ and $I_t^0 < -(k-1)C$ ($1 \leq k \leq K$).*

As a corollary we get the following condition.

Feasibility condition

A production schedule is feasible, i.e., $I_t \geq 0$ ($1 \leq t \leq T$), if and only if for every choice period w_k ($1 \leq k \leq K$), there are at least k production periods in $\{1, \dots, w_k\}$.

Every feasible production plan is completely characterized by its full production periods. To compare two feasible production plans S and S' , we can look at the first full production period in which they differ. If that period is earlier in S than it is in S' , then solution S is called *lexicographically earlier* than solution S' .

Clearly, the greedy solution is feasible. The optimality of the greedy approach is stated in the following theorem.

Theorem 2 *The greedy algorithm constructs the lexicographically earliest optimal production plan for $P(t)$, $t \in \{1, \dots, l\}$.*

Proof. Let S be the lexicographically earliest optimal solution. Suppose it is not equal to the solution S_G created by the greedy algorithm.

Let w_1, \dots, w_K be the choice periods for the greedy algorithm, and let the respective full production periods chosen by the greedy algorithm be τ_1, \dots, τ_K . Let k be the minimum index such that τ_k is not in S . There is a period τ' in $\{1, \dots, w_k\}$ that is a production period in S but not in S_G , because otherwise S would have less than k production periods in $\{1, \dots, w_k\}$, violating the feasibility condition.

Consider the following cases.

- (1) If $cost(\tau') < cost(\tau_k)$, then this contradicts the fact that the greedy algorithm chooses the cheapest available production period for w_k .
- (2) If $cost(\tau') = cost(\tau_k)$ and $\tau' < \tau_k$, then this contradicts the fact that the greedy algorithm chooses the earliest period among the cheapest available ones.

The feasibility condition also holds for the solution created from S by replacing τ' by τ_k as a production period. Therefore, we can conclude the following.

- (3) If $cost(\tau') > cost(\tau_k)$, then the solution S can be improved.
- (4) If $cost(\tau') = cost(\tau_k)$ and $\tau' > \tau_k$, then the solution S is not the lexicographically earliest optimal solution.

Hence, the assumption that $\tau' \neq \tau_k$ always leads to a contradiction. We conclude that S_G is equal to S , the lexicographically earliest optimal solution. \square

When referring to *the* optimal solution, we will from now on always mean the lexicographically earliest optimal solution determined by the greedy algorithm.

In the iterative algorithm to be presented in the next subsection, we will not only compute the optimal solutions for the problems $P(t), t \in \{1, \dots, l\}$, but also the optimal solutions of the l problems defined as follows. Let $t \in \{1, \dots, l\}$, then $P(t)'$ is the problem of finding an optimal schedule when f units become available in period t completely for free, i.e., *without costing any money or capacity*. Clearly, a feasible solution for this problem corresponds to a choice of K full production periods. The only difference with the feasible solutions of $P(t)$ is that now also period t is available for full production (at cost $cost(t)$). It follows immediately that the optimal solution can be found by applying the greedy algorithm. Again, when referring to *the* optimal solution of $P(t)'$, we will mean the solution constructed by the greedy algorithm. Furthermore, note that the optimal solution of $P(l)$ is also optimal for $P(l)'$, since production of C units in l would lead to a positive ending inventory in period T .

The following property plays a key role in our iterative algorithm.

Lemma 3 *Let $t \in \{1, \dots, l-1\}$. The optimal solutions of $P(t+1)'$ and $P(t)'$ differ with respect to the full production periods in at most one period. Moreover, if there is a difference, then the optimal solution of $P(t)'$ is obtained from the optimal solution of $P(t+1)'$ by replacing a full production period in $\{1, \dots, t\}$ by a period in $\{t+1, \dots, T\}$.*

Proof. We will prove that the solutions produced by the greedy algorithm in both problems differ in at most one production period, as described in the lemma.

Comparing the initial inventories of $P(t+1)'$ and $P(t)'$ we observe that the inventories are all equal, except that the inventory I_t^0 is f units higher in $P(t)'$ than in $P(t+1)'$. Thus, only for period t it may be possible that the situation with respect to the condition of lemma 1 is different. To be more precise, it is possible that for $P(t+1)'$ the condition is satisfied in t for some k , whereas it is not satisfied in t for $P(t)'$. Thus, if t is a choice period in $P(t+1)'$ it need not be a choice period in $P(t)'$. Instead, there may be a later choice period, say u . All other choice periods are equal for both problems.

Clearly, if the choice periods are equal in both problems, or if the same choices for production in t and u are made by the greedy algorithm, the optimal solutions do not differ. Hence, we only have to examine the case where the choices in t and u differ, say τ' is chosen in t and τ'' is chosen in u . Note that t and u are the first periods where the choices may differ, since the earlier choice periods are equal in both problems.

By definition, τ' is the cheapest available production period in $\{1, \dots, t\}$. Moreover, τ'' is the cheapest available production period in $\{1, \dots, u\}$. It follows that $cost(\tau'') < cost(\tau')$, and $\tau'' > t$.

We will show that in the remainder of the greedy algorithm the number of different production periods for both problems remains at most one, and that the difference is always as specified in the lemma.

As argued before, the choice periods after u are equal for both problems. Let those periods be w_k, \dots, w_K , and consider the production period chosen at w_k .

(a) Suppose that τ' is the period chosen at w_k in problem $P(t)'$.

Because τ' is the cheapest period available for production up to w_k in $P(t)'$, it follows

that τ'' is the cheapest period available up to w_k in $P(t+1)'$. Clearly, from w_k on the partial solutions are equal again.

- (b) Suppose $\tau \neq \tau'$ is the period chosen at w_k in problem $P(t)'$. τ is the cheapest period available for production up to w_k , and it holds that $\tau > t$ (since τ' is the cheapest period available up to t). Moreover, in $P(t+1)'$ it is also the cheapest period for production, unless τ'' is cheaper. However, which of these periods is chosen does not matter. In both cases the difference with the partial solution of $P(t)'$ remains one period, either τ or τ'' , and both are later than t .

If (a) occurs, then it follows immediately that the full production periods of the optimal solutions of $P(t)'$ and $P(t+1)'$ are the same. If (b) occurs, the above argument can be repeated for the later choice periods w_{k+1}, \dots, w_K , and the lemma is proved. \square

The following lemma states a property about the difference between the optimal solution of $P(t)$ and the optimal solution of $P(t)'$. It can be proved using similar arguments as in the proof of Lemma 3.

Lemma 4 *Let $t \in \{1, \dots, l\}$ and suppose that t is a full production period in the optimal solution of $P(t)'$. Then the optimal solution of $P(t)$ differs from the optimal solution of $P(t)'$ only in the fact that the full production in t is reallocated.*

3.2 Iterative algorithm

In our algorithm we only perform the greedy algorithm explicitly to obtain the optimal solution of $P(l)'$ and $P(l)$. After this initialization, we determine the optimal solutions of $P(t)'$ and $P(t)$ for all $t \in \{1, \dots, l-1\}$, by starting at period $l-1$ and moving downward to period 1.

An iteration of the algorithm consists of two steps. For a certain period $t \in \{1, \dots, l-1\}$, we are given the optimal solution of $P(t+1)'$. First, we will determine the optimal solution of $P(t)'$, and then, if this solution is not feasible for $P(t)$, i.e., if t is a full production period, it will be changed into an optimal one of $P(t)$, i.e., one without full production in t . We will first sketch an iteration and then give the details.

2-step iteration for period t

Given: The optimal solution of $P(t+1)'$.

Step 1: Determine the optimal solution of $P(t)'$.

Move in the optimal solution of $P(t+1)'$ the fraction f from $t+1$ to t . As we have seen in Lemma 3, it is only necessary to check now whether there exists a feasible and profitable move of full production from a period up to t to a period later than t .

Step 2: Convert the optimal solution of $P(t)'$ to an optimal solution of $P(t)$.

If t is a full production period in $P(t)'$, then the cheapest feasible alternative for

production of the C units is to be found. According to Lemma 4 no other changes have to be considered.

We will now explain both steps in greater detail.

Details of Step 1

Let I_t denote the inventory level of period t in the optimal solution of $P(t+1)'$ and let I_t' denote the inventory level after the fractional production has been moved from $t+1$ to t , i.e., $I_t' = I_t + f$. All other inventory levels do not change.

- (a) If $I_t' \geq C > I_t$ then the inventory at the end of period t is increased from a value less than C to a value greater than or equal to C . Let $s \leq t < u$ be such that $I_{s-1} < C$, $I_u < C$ and all ending inventories of the periods $s, s+1, \dots, u-1$ are larger than or equal to C . Note that it is not feasible to move production from a period that is earlier than s , or to a period that is later than u . Therefore, we only have to check whether it is advantageous to move production from the most expensive production period in $\{s, \dots, t\}$ to the (earliest) cheapest available period in $\{t+1, \dots, u\}$. If this is the case, the move is made.
- (b) If $I_t' > I_t \geq C$ or $C > I_t' > I_t$ then the full production periods remain the same. If $I_t' < C$, then moving production from a period up to t to a period after t will always lead to a negative inventory level in period t . On the other hand, if $I_t \geq C$, then every profitable move that we can make now would also have been feasible with respect to the optimal solution of $P(t+1)'$. That would contradict the optimality of the latter solution.

The following lemma is now obvious.

Lemma 5 *Only if $I_t' \geq C > I_t$ it may be necessary to move a full production period in Step 1. If the move is made as specified under (a), the resulting solution is optimal for $P(t)'$.*

Note that after step 1 we have the schedule that is the starting point for the next iteration.

Details of Step 2

The starting point is the optimal solution of $P(t)'$ created in step 1. We transform this schedule into a feasible solution for $P(t)$, if necessary. Thus, nothing remains to be done if t is not a full production period. Otherwise, the C units produced in t should be transferred to another period. We determine the choices that we have as follows. Let $u \geq t$ be the period such that the ending inventory of the periods $\{t, \dots, u-1\}$ is at least C and $I_u < C$. Moving the production to a period after u will result in an infeasible solution. Therefore, the production of C units in t is moved to the (earliest) cheapest available period in the set $\{1, \dots, t-1, t+1, \dots, u\}$. Clearly, this solution is feasible and optimality follows from Lemma 4.

Lemma 6 *If step 2 is performed as described above then the optimal schedule of $P(t)$ results.*

3.3 Running time of the algorithm

In the following we will add some details to steps 1 and 2 that are necessary to derive the $O(T^3)$ running time of the overall algorithm. We will do so for each part of the algorithm separately, i.e., for the initialization (performing the greedy algorithm explicitly), step 1 and step 2.

Complexity of the initialization

We will use the following result.

Lemma 7 *Let $1 \leq t_1 \leq t_2 \leq T - 1$. Consider the optimal schedules for subplans (t_1, t_2) and $(t_1, t_2 + 1)$ under the restriction that the fractional periods are the last production periods. Then the set of full production periods in the schedule for subplan (t_1, t_2) is a subset of the set of full production periods in the schedule for subplan $(t_1, t_2 + 1)$.*

Proof. This follows from the fact that the choice periods for the smaller subplan are a subset of the set of choice periods of the larger subplan (cf. Lemma 1). If $\sum_{j=t_1}^{t_2} d_j < kC \leq \sum_{j=t_1}^{t_2+1} d_j$ for some k then one extra choice has to be made in the larger subplan. \square

It follows immediately from Lemma 7 that performing the initialization for all subplans with first period t_1 has a total running time that is of the same order as the running time of the initialization for subplan (t_1, T) only. The latter can easily be implemented in $O(T^2)$ time. Hence, the overall algorithm takes $O(T^3)$ in the initialization step.

Complexity of step 1

We proceed with the iterations of the actual algorithm for a specific subplan. Again, we consider for convenience subplan $(1, T)$. After the initialization, we determine for each period t in this subplan the period $\tau(t)$ which is the cheapest available period in the set $\{1, \dots, t\}$. This takes $O(T)$ time. The periods $\tau(t)$ may change during the algorithm.

The following lemma restricts the number of possible changes performed by step 1.

Lemma 8 *Suppose that in step 1 of the iteration for period $t \in \{1, \dots, l\}$, it is necessary to check whether production should be moved from a period up to t to a period after t . Let $s \leq t$ be such that $I_{s-1} < C$, whereas all inventory levels from periods s to t are at least C . Then it is not necessary to check whether production should be moved in step 1 of the iterations for periods $\{s, \dots, t - 1\}$.*

Proof. If in step 1 of the iteration for t no profitable move of full production is detected, then the statement follows directly from Lemma 5 and the fact that for all periods in $\{s, \dots, t - 1\}$ the ending inventory is at least C .

Now suppose that a full production period is moved in step 1 of the iteration for t from a period in $\{s, \dots, t\}$ to a period after t . Note that this move reduces the inventory of t to a

level below C . Suppose that the lemma is false and there are periods in $\{s, \dots, t-1\}$ for which it is profitable to move a full production period in step 1. Consider the first iteration for which this happens and let v be the corresponding period. A profitable move with respect to v consists of moving full production from a period in $\{s, \dots, v\}$ to a period after v , but not later than t . It is easy to see that this move would also have been a feasible and profitable one with respect to the solution given at the start of the iteration for t . As this was the optimal solution to $P(t+1)'$, we have derived a contradiction. Hence, the lemma holds. \square

If we want to check whether it is profitable to move a full production period in step 1 of iteration t , then we need to determine the first period after t that has an inventory level below C . Denote this period by $u(t)$.

Lemma 9 *Suppose that in step 1 of the iteration for period $t \in \{1, \dots, l\}$, it is checked whether full production needs to be moved from a period up to t to a period after t . Let $v < t$ be the next period in which such a check is made. Then $u(v) = u(t)$ or $u(v) \leq t$.*

Proof. This follows from the observation already made in the proof of Lemma 8, that if a full production period is moved in step 1 of the iteration for t , the inventory of t is reduced to a level below C . \square

Step 1 of the iteration of period t is implemented as follows. We are given an optimal solution of $P(t+1)'$. First, $u(t)$ is determined. The cheapest available period $\tau(u)$ in the set $\{1, \dots, u(t)\}$ is given. Furthermore, we can determine s and the most expensive production period in $\{s, \dots, t\}$, say τ' , in $t - s + 1$ comparisons. By Lemma 8 the overall complexity of this operation is $O(T)$, since this part of step 1 never needs to be performed in a period after s anymore, if it is performed in t . Using Lemma 9, it follows that also the determination of the relevant $u(t)$ values can be implemented in linear time. Suppose full production is moved from a period $\tau' \geq s$ to $\tau(u)$. Note that $\tau(u) \leq t$ is not possible, because then the solution of $P(t+1)'$ would not have been optimal. Because τ' becomes available, we may have to update $\tau(t')$ for some periods t' in $\{\tau', \dots, t\}$. Because $\tau' \geq s$, this operation takes $O(T)$ in total as well.

We conclude that step 1 takes $O(T)$ for each subplan.

Complexity of step 2

In this step we need to know the cheapest available period in $\{1, \dots, u\}$, which is exactly $\tau(u)$. Therefore, this step is trivially seen to take only $O(T)$ time for each subplan. Note that the rest of the data is not updated, since the schedule at the beginning of step 2, is also the schedule at the beginning of the next iteration.

To summarize, we have shown the following complexities.

Initialization: $O(T^2)$ running time for all subplans (t_1, t_2) with t_1 fixed.

Step 1: $O(T)$ running time for each subplan (t_1, t_2) .

Step 2: $O(T)$ running time for each subplan (t_1, t_2) .

Hence, the total complexity of our algorithm is $O(T^3)$.

4 Concluding remarks

We have presented an $O(T^3)$ dynamic programming algorithm for solving the economic lot-sizing problem with constant capacities, concave production costs and linear holding costs. Our algorithm is an improvement upon the algorithm by Florian and Klein [6] by a factor T . However, the latter algorithm also solves the more general problem in which the holding costs are concave. For our approach the linearity of the holding costs seems essential. It allows us to formulate an equivalent problem without holding costs, for which it is easy to calculate the change in costs when a full production period is moved.

The improvement in running time of our algorithm is based on the idea that many similar subproblems have to be solved, and that it is worthwhile to exploit the fact that the optimal solutions to these problems are partially the same. The only possible way in which a further improvement could be achieved, seems to be the existence of a relationship between the optimal fractional periods of similar subplans. Until now, we have not been able to identify such a relationship.

References

- [1] A. Aggarwal and J.K. Park. Improved Algorithms for Economic Lot-Size Problems. Working paper, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1990.
- [2] G.B. Bitran and H.H. Yanasse. Computational Complexity of the Capacitated Lot Size Problem. *Management Science*, 28(10):1174–1185, 1982.
- [3] H.-D. Chen, D.W. Hearn, and C.-Y. Lee. A New Dynamic Programming Algorithm for the Single Item Capacitated Dynamic Lot Size Model. Research Report 92-7, Department of Industrial and Systems Engineering, University of Florida, Gainesville, Florida, 1990.
- [4] C.-S. Chung and C.-H.M. Lin. An $O(T^2)$ Algorithm for the $NI/G/NI/ND$ Capacitated Lot Size Problem. *Management Science*, 34:420–426, 1988.
- [5] A. Federgruen and M. Tzur. A Simple Forward Algorithm to Solve General Dynamic Lot Sizing Models with n Periods in $O(n \log n)$ or $O(n)$ Time. *Management Science*, 37:909–925, 1991.
- [6] M. Florian and M. Klein. Deterministic Production Planning with Concave Costs and Capacity Constraints. *Management Science*, 18:12–20, 1971.
- [7] M. Florian, J.K. Lenstra, and A.H.G. Rinnooy Kan. Deterministic Production Planning: Algorithms and Complexity. *Management Science*, 26:669–679, 1980.

- [8] S. Van Hoesel, A. Wagelmans, and B. Moerman. Using Geometric Techniques to Improve Dynamic Programming Algorithms for the Economic Lot-Sizing Problem and Extensions. Report 9212/A, Econometric Institute, Erasmus University Rotterdam, Rotterdam, The Netherlands, 1992. (*to appear in European Journal of Operational Research*).
- [9] A.F. Veinott, Jr. Unpublished class notes. Program in Operations Research, Stanford University, Stanford, California, 1963.
- [10] A. Wagelmans, S. Van Hoesel, and A. Kolen. Economic Lot-Sizing: An $O(n \log n)$ Algorithm that Runs in Linear Time in the Wagner–Whitin Case. *Operations Research*, 40:S145–156, 1992.

List of COSOR-memoranda - 1993

Number	Month	Author	Title
93-01	January	P. v.d. Laan C. v. Eeden	Subset selection for the best of two populations: Tables of the expected subset size
93-02	January	R.J.G. Wilms J.G.F. Thiemann	Characterizations of shift-invariant distributions based on summation modulo one.
93-03	February	Jan Beirlant John H.J. Einmahl	Asymptotic confidence intervals for the length of the shortt under random censoring.
93-04	February	E. Balas J. K. Lenstra A. Vazacopoulos	One machine scheduling with delayed precedence constraints
93-05	March	A.A. Stoorvogel J.H.A. Ludlage	The discrete time minimum entropy H_∞ control problem
93-06	March	H.J.C. Huijberts C.H. Moog	Controlled invariance of nonlinear systems: nonexact forms speak louder than exact forms
93-07	March	Marinus Veldhorst	A linear time algorithm to schedule trees with communication delays optimally on two machines
93-08	March	Stan van Hoesel Antoon Kolen	A class of strong valid inequalities for the discrete lot-sizing and scheduling problem
93-09	March	F.P.A. Coolen	Bayesian decision theory with imprecise prior probabilities applied to replacement problems
93-10	March	A.W.J. Kolen A.H.G. Rinnooy Kan C.P.M. van Hoesel A.P.M. Wagelmans	Sensitivity analysis of list scheduling heuristics
93-11	March	A.A. Stoorvogel J.H.A. Ludlage	Squaring-down and the problems of almost-zeros for continuous-time systems
93-12	April	Paul van der Laan	The efficiency of subset selection of an ε -best uniform population relative to selection of the best one
93-13	April	R.J.G. Wilms	On the limiting distribution of fractional parts of extreme order statistics

Number	Month	Author	Title
93-14	May	L.C.G.J.M. Habets	On the Genericity of Stabilizability for Time-Day Systems
93-15	June	P. van der Laan C. van Eeden	Subset selection with a generalized selection goal based on a loss function
93-16	June	A.A. Stoorvogel A. Saberi B.M. Chen	The Discrete-time H_∞ Control Problem with Strictly Proper Measurement Feedback
93-17	June	J. Beirlant J.H.J. Einmahl	Maximal type test statistics based on conditional processes
93-18	July	F.P.A. Coolen	Decision making with imprecise probabilities
93-19	July	J.A. Hoogeveen J.K. Lenstra B. Veltman	Three, four, five, six or the Complexity of Scheduling with Communication Delays
93-20	July	J.A. Hoogeveen J.K. Lenstra B. Veltman	Preemptive scheduling in a two-stage multiprocessor flow shop is NP-hard
93-21	July	P. van der Laan C. van Eeden	Some generalized subset selection procedures
93-22	July	R.J.G. Wilms	Infinite divisible and stable distributions modulo 1
93-23	July	J.H.J. Einmahl F.H. Ruymgaart	Tail processes under heavy random censorship with applications
93-24	August	F.W. Steutel	Probabilistic methods in problems of applied analysis
93-25	August	A.A. Stoorvogel	Stabilizing solutions of the H_∞ algebraic Riccati equation
93-26	August	R. Perelaer J.K. Lenstra M. Savelsbergh F. Soumis	The Bus Driver Scheduling Problem of the Amsterdam Transit Company
93-27	August	J.M. van den Akker C.P.M. van Hoesel M.W.P. Savelsbergh	Facet Inducing Inequalities for Single-Machine Scheduling Problems
93-28	September	H.J.C. Huijberts H. Nijmeijer	Dynamic disturbance decoupling of nonlinear systems and linearization

Number	Month	Author	Title
93-29	September	L.C.G.J.M. Habets	Testing Reachability and Stabilizability of Systems over Polynomial Rings using Gröbner Bases
93-30	September	Z. Liu J.A.C. Resing	Duality and Equivalencies in Closed Tandem Queueing Networks
93-31	November	A.G. de Kok J. van der Wal	Assigning identical operators to different machines
93-32	November	M.J.A. van Eenige J.A.C. Resing J. van der Wal	A Matrix-Geometric Analysis of Queueing Systems with Periodic Service Interruptions
93-33	November	A.G. de Kok	Backorder lead time behaviour in (b, Q) -inventory models with compound renewal demand
93-34	November	J. van der Wal	The lead time shift theorem for the continuous review (s, Q) and (s, S) inventory systems
93-35	November	J.H.C.M. Verrijdt A.G. de Kok	Distribution Planning for a Divergent 2-echelon Network without Intermediate Stocks under Service Restrictions
93-36	November	A.G. Lagodimos A.G. de Kok J.H.C.M. Verrijdt	The Robustness of Multi-Echelon Service Models under Autocorrelated Demands
93-37	November	C.P.M. van Hoesel A.P.M. Wagelmans	An $O(T^3)$ algorithm for the economic lot-sizing problem with constant capacities