DEPARTMENT OF OPERATIONS RESEARCH
UNIVERSITY OF AARHUS

Publication no. 99/3

Pure and modified base-stock policies for the lost sales inventory system with negligible set-up costs and constant lead times

Søren Glud Johansen

ISSN 1398-8964

Department of Mathematical Sciences
Telephone: +45 8942 3188

Building 530, Ny Munkegade
DK-8000 Aarhus C, Denmark

Pure and modified base-stock policies for the lost sales inventory system with negligible set-up costs and constant lead times

Søren Glud Johansen
University of Aarhus, Denmark

June 15, 1999

The studied inventory system with continuous review has an easily computed optimal $(S-1, S)$ policy when unsatisfied demands are backlogged. We assume that unsatisfied demands are lost and then it is also easy to compute the best $(S-1, S)$ policy. But, as demonstrated by Roger Hill at the ISIR Symposium in 1996, this pure base-stock policy can never be optimal if $S \geq 2$.

Our focus is on periodic review. We use Erlang's loss formula to derive approximate expressions for the stockout probability and the average cost. These expressions are used to approximate the average cost and to compute a good base-stock. We formulate and implement a Markov decision model to find the optimal replenishment policy. The model is solved by a policy-iteration algorithm. Because the optimal policy is often rather complicated, we introduce modified base-stock policies. They are specified by a pair (S, t) where S is the base-stock and t is a lower bound for the number of review periods between review epochs in which placing a replenishment order is permitted. A simple one has S equal to the base-stock computed from Erlang's formula and fixes t as the largest integer which is less than or equal to the ratio of the number of review periods per delivery period and S. Our numerical examples show that the simple modified base-stock policy provides most of the cost reduction which can be obtained by replacing the best pure base-stock policy by the optimal policy.

Key words: Periodic review, Lost sales, Base-stock, Policy iteration, Simple policy.

1 Introduction

We consider a single-item inventory system with Poisson demand, negligible set-up costs and constant lead times. The system with continuous review has an easily computed optimal pure base-stock policy when unsatisfied demands are backlogged [5]. A pure base-stock policy places a replenishment order to restore the base-stock S whenever the inventory position (stock on hand + on order − backordered) is below S. Hence the reorder level is $S-1$ and the policy is commonly referred to as an $(S-1, S)$ policy.

We assume that unsatisfied demands are lost and, by applying Erlang's loss formula, it is easy to compute the best $(S-1, S)$ policy for the system with continuous review [6]. But, as demonstrated by Hill [3], this pure base-stock policy can never be optimal if $S \geq 2$. We present a model for finding the optimal replenishment policy. Because the optimal policy is often rather complicated, we introduce easily implemented modified base-stock policies. The suggested ones provide most of the cost reduction which can be obtained by replacing the best pure base-stock policy by the optimal policy. Moreover, one of the suggested policies is easy to compute.

Our model assumes that replenishment orders can be placed only at equidistant review epochs and that the constant lead time for one replenishment order is an integer number m of review periods where $m > 1$. (The easy case with $m = 1$ is excluded in order to facilitate the exposition in Section 3.) The demand D in each review period has a Poisson distribution with parameter λ/m, where λ denotes the demand rate per delivery period. We choose the cost per delivery period of holding one unit in stock as the monetary unit. The shortage cost for each lost demand is p. It equals the price at which the item is sold + the penalty cost incurred when a demand is lost − the variable ordering cost per unit. We assume that order set-up costs are negligible and can be ignored. Therefore, the replenishment orders considered are for one unit. The objective is to minimize the long-run average cost AC per delivery period which is our time unit.

Assuming periodic review implies that Erlang's loss formula cannot be applied directly to determine the best $(S-1, S)$ policy as for the system with continuous review. But this formula can be used to derive an approximate expression for AC as a function of the base-stock S. A good S-value is easily computed from this expression. The approximate expression is of interest itself because regularly scheduled shipments, disposed by periodic rather than continuous review, are common practice in order to achieve an efficient utilization of transportation resources and/or to procure transportation services at least cost [2]. In addition to these reasons for assuming periodic review there are two others. The assumption enables us to optimize the model because it is a tractable Markov decision process which can be solved by a policy-iteration algorithm. Moreover, the assumption makes the modified base-stock policies tractable.

The paper is organized as follows. Section 2 describes and illustrates with numerical examples how a pure base-stock policy is evaluated by Erlang's loss formula with three different specifications of the traffic intensity. Section 3 introduces useful concepts to describe and evaluate an arbitrary replenishment policy. Our policy-iteration algorithm for computing the optimal replenishment policy is presented and illustrated in Section 4. Section 5 introduces modified base-stock policies. Numerical examples illustrate two such policies: a simple one and the best one. Section 6 has concluding remarks.

2 Approximate evaluations of a pure base-stock policy

Recall that the constant delivery period is our time unit and that our monetary unit is the cost per unit time of holding one unit in stock. The long-run average cost per unit time for an $(S-1, S)$ policy is

$$AC(S) = p \cdot \lambda \cdot B(S) + AS(S) \tag{1}$$

where $B(S)$ is the long-run fraction of demand lost (hereafter referred to as the 'stockout probability') and $AS(S)$ is the average stock.

For the system with continuous review, the stockout probability is specified by Erlang's loss formula

$$B_\rho(S) = \frac{\frac{\rho^S}{S!}}{\sum_{i=0}^{S} \frac{\rho^i}{i!}} \tag{2}$$

and the average stock is

$$AS_\rho(S) = S - (1 - B_\rho(S)) \cdot \rho. \tag{3}$$

Here the traffic intensity ρ equals λ. These results hold for any distribution of the delivery times (when they are independent) because the number of single-unit orders outstanding is the same as the number of busy servers in an $M/G/S/S$ queue [6]. For this queue, the equilibrium distribution is independent of the form of the distribution for the service times (which equal the delivery times).

We shall use the above results to get approximate expressions for the system with periodic review. Periodic review implies that there is a delay from the epoch in which a demand is satisfied to the review epoch in which the system places the replenishment order triggered by this demand. Note that this delay refers to *satisfied* demand only. Therefore, we cannot conclude (as for the backlogging system considered by Axsäter [1]) that the delay is uniformly distributed between 0 and $1/m$, and it is not easy to compute the average delay exactly. Furthermore, because the delays for consecutive satisfied demands are not independent, we cannot conclude from the above insensitivity result that the equilibrium distribution for the number of single-unit orders outstanding does not depend on how the delay is distributed. But we find it interesting to investigate how close the stockout probability and the average stock are approximated by Eqs. 2 and 3 with the following three specifications of the traffic

		m=2		m=5		m=10	
$S=1$	a	33.3333%	0.6667	33.3333%	0.6667	33.3333%	0.6667
	b	38.4615%	0.6154	35.4839%	0.6452	34.4262%	0.6557
	c	38.6579%	0.6134	35.5185%	0.6448	34.4352%	0.6556
	e	38.6579%	0.6134	35.5185%	0.6448	34.4352%	0.6556
$S=2$	a	7.6923%	1.5385	7.6923%	1.5385	7.6923%	1.5385
	b	10.7296%	1.4421	8.8905%	1.4989	8.2879%	1.5185
	c	10.8585%	1.4382	8.9107%	1.4983	8.2929%	1.5183
	e	10.8877%	1.4413	8.9172%	1.4988	8.2948%	1.5185
$S=3$	a	1.2658%	2.5063	1.2658%	2.5063	1.2658%	2.5063
	b	2.1865%	2.3887	1.6038%	2.4588	1.4297%	2.4825
	c	2.2301%	2.3839	1.6098%	2.4580	1.4311%	2.4823
	e	2.2703%	2.3886	1.6167%	2.4588	1.4329%	2.4825
$S=4$	a	0.1580%	3.5008	0.1580%	3.5008	0.1580%	3.5008
	b	0.3405%	3.3771	0.2200%	3.4512	0.1873%	3.4760
	c	0.3501%	3.3720	0.2212%	3.4504	0.1875%	3.4758
	e	0.3666%	3.3772	0.2236%	3.4512	0.1883%	3.4760

Table 1: Stockout probabilities (in %) and average stocks computed approximately (a, b and c) and exactly (e) for various base-stock policies when $\lambda = 0.5$.

intensity ρ.
a. *Neglecting the delays*
The first specification neglects the delays and sets ρ equal to λ.
b. *Setting the average delay equal to* $\frac{1}{2m}$
The second specification assumes that the average delay is $\frac{1}{2m}$ and sets ρ equal to $\lambda \cdot (1 + \frac{1}{2m})$.
c. *Setting the average delay equal to* $\frac{1}{m\cdot(1-e^{-\lambda/m})} - \frac{1}{\lambda}$
The third specification is exact for $S = 1$ because it restricts attention to demands which occur as the first in a review period. The average delay for such demands is equal to the expected delay for the first demand occuring after an arbitrary review epoch. Because the time from any epoch to the epoch in which the next demand occurs is exponential with rate λ, we conclude that the desired average delay is

$$
\begin{aligned}
&\sum_{i=1}^{\infty} \int_{(i-1)/m}^{i/m} (i/m - t) \cdot \lambda e^{-\lambda t} dt \\
&= \sum_{i=1}^{\infty} e^{-\lambda \cdot (i-1)/m} \cdot \int_0^{1/m} (1/m - \tau) \cdot \lambda e^{-\lambda \tau} d\tau \\
&= \frac{\frac{1}{m} - \frac{1}{\lambda}\cdot(1-e^{-\lambda/m})}{1-e^{-\lambda/m}} \\
&= \frac{1}{m\cdot(1-e^{-\lambda/m})} - \frac{1}{\lambda}.
\end{aligned}
$$

The third specification sets ρ equal to λ multiplied by 1 plus this expression for the average delay.

For various values of λ and S, Tables 1, 2 and 3 illustrate the results obtained by the three specifications of ρ compared to the exact values (e) found as

		m=2		m=5		m=10	
$S = 1$	a	50.0000%	0.5000	50.0000%	0.5000	50.0000%	0.5000
	b	55.5556%	0.4444	52.3810%	0.4762	51.2195%	0.4878
	c	55.9616%	0.4404	52.4564%	0.4754	51.2393%	0.4876
	e	55.9616%	0.4404	52.4564%	0.4754	51.2393%	0.4876
$S = 2$	a	20.0000%	1.2000	20.0000%	1.2000	20.0000%	1.2000
	b	25.7732%	1.0722	22.3660%	1.1460	21.1917%	1.1725
	c	26.2300%	1.0626	22.4436%	1.1443	21.2115%	1.1721
	e	26.2019%	1.0690	22.4413%	1.1455	21.2111%	1.1724
$S = 3$	a	6.2500%	2.0625	6.2500%	2.0625	6.2500%	2.0625
	b	9.6974%	1.8712	7.5793%	1.9834	6.9050%	2.0225
	c	9.9996%	1.8563	7.6249%	1.9808	6.9160%	2.0218
	e	10.0511%	1.8709	7.6375%	1.9833	6.9196%	2.0225
$S = 4$	a	1.5385%	3.0154	1.5385%	3.0154	1.5385%	3.0154
	b	2.9413%	2.7868	2.0418%	2.9225	1.7803%	2.9687
	c	3.0789%	2.7684	2.0599%	2.9194	1.7845%	2.9679
	e	3.1493%	2.7874	2.0728%	2.9226	1.7878%	2.9687

Table 2: Stockout probabilities (in %) and average stocks computed approximately (a, b and c) and exactly (e) for various base-stock policies when $\lambda = 1.0$.

		m=2		m=5		m=10	
$S = 1$	a	60.0000%	0.4000	60.0000%	0.4000	60.0000%	0.4000
	b	65.2174%	0.3478	62.2641%	0.3774	61.1651%	0.3883
	c	65.7703%	0.3423	62.3705%	0.3763	61.1933%	0.3881
	e	65.7703%	0.3423	62.3705%	0.3763	61.1933%	0.3881
$S = 2$	a	31.0345%	0.9655	31.0345%	0.9655	31.0345%	0.9655
	b	37.9427%	0.8364	33.9358%	0.9099	32.5088%	0.9370
	c	38.7206%	0.8226	34.0757%	0.9073	32.5450%	0.9363
	e	38.6289%	0.8307	34.0607%	0.9089	32.5414%	0.9368
$S = 3$	a	13.4328%	1.7015	13.4328%	1.7015	13.4328%	1.7015
	b	19.1685%	1.4844	15.7289%	1.6095	14.5789%	1.6546
	c	19.8716%	1.4604	15.8438%	1.6051	14.6076%	1.6535
	e	19.8449%	1.4831	15.8462%	1.6093	14.6089%	1.6546
$S = 4$	a	4.7957%	2.5719	4.7957%	2.5719	4.7957%	2.5719
	b	8.2445%	2.2796	6.0929%	2.4505	5.4288%	2.5105
	c	8.7138%	2.2460	6.1608%	2.4446	5.4450%	2.5090
	e	8.7784%	2.2816	6.1790%	2.4509	5.4502%	2.5106

Table 3: Stockout probabilities (in %) and average stocks computed approximately (a, b and c) and exactly (e) for various base-stock policies when $\lambda = 1.5$.

	m=2		m=5		m=10	
a	-56.910%	16.858%	-29.359%	6.300%	-16.115%	3.075%
	(0.5,4)	(1.5,1)	(0.5,4)	(1.5, 1)	(0.5,4)	(1.5,1)
b	-7.133%	1.615%	-1.609%	0.283%	-0.550%	0.073%
	(0.5,4)	(1.5,1)	(0.5,4)	(1.5,1)	(0.5,4)	(1.5,1)
c	-4.499%	-1.562%	-1.094%	-0.261%	-0.413%	-0.067%
	(0.5,4)	(1.5,4)	(0.5,4)	(1.5,3)	(0.5,4)	(1.5,3)

Note. Each pair in parenthesis specifies the values of λ and S for which the difference in the line above is computed.

Table 4: The largest positive or negative differences (in %) between the approximate and exact values of the stockout probabilities and the average stocks computed in Tables 1, 2 and 3.

described in the next section. We have computed how much the approximate values differ (in %) from their exact values for each of the three specifications and for each considered m-value. Table 4 shows the largest positive or negative percentage computed (first line) and the pair (λ, S) for which it is computed (second line). We are not able to explain why the largest differences occur for the pairs reported. But we see that the specifications b and c perform reasonable well (especially for $m = 5$ and $m = 10$) whereas the a-specification has a poor performance. For modest S-values, our favorite is the c-specification. We use it in Section 5 to specify the simple modified base-stock policy.

Let $AC_\rho(S)$ denote the approximate average cost computed from Eq. 1 when $B(S)$ and $AS(S)$ are specified by Eqs. 2 and 3 with the traffic intensity ρ. Note that

$$AC_\rho(S+1) - AC_\rho(S) = (p+1)\cdot\lambda\cdot\Delta B_\rho(S) + 1$$

where $\Delta B_\rho(S) = B_\rho(S+1) - B_\rho(S)$ is negative and strictly increasing in S. Therefore, the largest S which minimizes $AC_\rho(S)$ is

$$S_\rho = \min\left\{ S \,\middle|\, \Delta B_\rho(S) > -\frac{1}{(p+1)\cdot\lambda} \right\}. \tag{4}$$

3 Exact evaluation of a replenishment policy

We consider an arbitrary stationary replenishment policy for which the inventory position is at most S. A Markov chain is used to evaluate this policy exactly. The concepts introduced in this section to describe policy evaluation are useful in Sections 4 and 5 to describe policy improvement and modified base-stock policies, respectively.

The Markov chain focuses on the inventory system at the review epochs. The state of the system is described by a vector $\mathbf{x}$ with components $x_j,\ j = 0, 1, .., S$. The component x_0 specifies the stock on hand. The other components are either positive or zero and they are ordered so that

$$m \geq x_j \geq x_{j+1} \geq 0,\ j = 1, 2, .., S-1. \tag{5}$$

A positive component x_j identifies an outstanding order for one unit and specifies that $x_j - 1$ review periods are gone since this order was placed. Hence $\max\{j|x_j > 0\}$ is the number of single-unit orders outstanding. If $x_1 = 0$ then this number is 0. The number of orders to be delivered at the end of the coming review period is 0 if $x_1 < m$ and $\max\{j|x_j = m\}$ else. Our motivation for letting x_j be one plus the number of review periods elapsed since the jth order outstanding was placed is that then we can distinguish between when the jth order has just been placed and when less than j orders are outstanding.

The inventory position for the state described by the vector $\mathbf{x}$ equals $x_0 + \max\{j|x_j > 0\}$. It is easy to verify, by induction on k, that the number of states with inventory position k is

$$n_k = \frac{(m+k)!}{m!k!}$$

and that the number of states for which the inventory position is at most k is

$$N_k = \sum_{j=0}^{k} n_j = \frac{(m+1+k)!}{(m+1)!k!}.$$

In particular N_S is the total number of states. We identify the states with the integers from 1 to N_S. The state space I_S is the set of these integers.

For each state $i \in I_S$, let $\mathbf{y}(i)$ denote the $1+S$ dimensional vector describing the state. Note that it specifies that the number of orders outstanding is

$$NOO(i) = \begin{cases} 0 & \text{if } y_1(i) = 0 \\ \max\{j|y_j(i) > 0\} & \text{else} \end{cases}$$

and that the number of orders to be delivered at the end of the coming review period is

$$NOD(i) = \begin{cases} 0 & \text{if } y_1(i) < m \\ \max\{j|y_j(i) = m\} & \text{else.} \end{cases}$$

The numbers of single-unit orders which can be placed in state i when the inventory position is bound by S are the elements of the set

$$A_S(i) = \{0, 1, .., S - y_0(i) - NOO(i)\}.$$

If the number of orders placed in state i is a, say, then this state is changed immediately into another one which is denoted by $o^a(i)$. The vector $\mathbf{y}(o^a(i))$ describing the new state has the components

$$y_j(o^a(i)) = \begin{cases} y_j(i), & j \leq NOO(i) \\ 1, & NOO(i) < j \leq NOO(i) + a \\ 0, & j > NOO(i) + a. \end{cases}$$

Note that $NOD(o^a(i)) = NOD(i)$ because the constant delivery times are assumed to be at least two review periods.

Let $NOP(i)$ denote the number of single-unit orders prescribed to be placed in state i by the considered policy. Suppose that the demand satisfied during the review period following a review epoch in which the state is i equals d. Then the components of the vector describing the state $s(i, d)$ at the next review epoch are specified as follows. The component for the stock on hand is $y_0(s(i, d)) = y_0(i) - d + NOD(i)$. If $1 \leq j \leq NOO(i) - NOD(i)$ then j identifies an old order for which the component is $y_j(s(i, d)) = y_{NOD(i)+j}(i) + 1$. Otherwise j either identifies an order placed in state i or no order, and then the components are

$$y_j(s(i,d)) = \begin{cases} 2, & NOO(i) - NOD(i) < j \leq NOO(i) - NOD(i) + NOP(i) \\ 0, & j > NOO(i) - NOD(i) + NOP(i). \end{cases}$$

Recall that the demand D during each review period has a Poisson distribution with parameter λ/m. The positive probabilities $P_{i,j}$ of a transition from state i at one review epoch to state j at the next one are

$$P_{i,s(i,d)} = \begin{cases} \Pr\{D = d\}, & d = 0, 1, .., y_0(i) - 1, \\ \Pr\{D \geq y_0(i)\}, & d = y_0(i). \end{cases} \tag{6}$$

The expected cost incurred during the review period is [4, page 183]

$$c_i = p \cdot U(y_0(i)) + y_0(i)/m - \sum_{k=1}^{y_0(i)} U(k)/\lambda \tag{7}$$

where

$$U(k) = E[D - k]^+ = \sum_{d=k}^{\infty} \Pr\{D > k\},\ k = 0, 1, ...$$

Let g denote the long-run average cost incurred per review period and let v_i denote the expected incremental cost incurred over an infinite horizon when the inventory system is started in state i rather than in state 1. If the Markov chain has no two disjoint closed sets (this condition is satisfied when the considered policy is a reasonable one), then g and the state values v_i satisfy the following system of linear equations [7, Theorem 3.1.1]

$$v_i = c_i - g + \sum_{j=1}^{N_S} P_{i,j} \cdot v_j,\ i = 1, 2, .., N_S. \tag{8}$$

Furthermore, these equations together with the normalizing equation

$$v_1 = 0 \tag{9}$$

has a unique solution. We have used Gauss elimination to find the unique solution numerically. The exact average cost AC per unit time equals m multiplied by the g-value found.

The exact stockout probabilities (in %) reported in Tables 1, 2 and 3 are computed as $100 \cdot \bar{g} \cdot m/\lambda$ where $\bar{g}$ is the solution to Eqs. 8 and 9 with c_i set equal to $U(y_0(i))$. The exact average stocks reported are the average costs AC computed for $p = 0$.

4 Computing the optimal replenishment policy

We shall now present a policy-iteration algorithm [7, Section 3.2] for finding the replenishment policy which is optimal subject to the condition that the inventory position is at most S. How this S can be fixed is discussed after the presentation of the algorithm.

The iterations of the policy-iteration algorithm start with a replenishment policy which prescribes the number $NOP(i)$ of single-unit orders to be placed in each state $i \in I_S$ so that the inventory position is at most S. This policy is available from the previous iteration or - in the first iteration - from the outset. Each iteration consists of the following three steps.

Step 1 (value determination). For the current policy specified by $NOP(i),\ i \in I_S$, use Eqs. 8 and 9 to compute the average cost g per review period and the state values $v_i,\ i \in I_S$.

Step 2 (policy improvement). For each state $i \in I_S$, determine an action $a(i)$ yielding the minimum in $\min\left\{v_{o^a(i)} | a \in A_S(i)\right\}$. The new policy is obtained by choosing $NOP_{\text{new}}(i) = a(i)$ for all $i \in I_S$ with the convention that $NOP_{\text{new}}(i)$ is chosen equal to the old action $NOP(i)$ when this action minimizes the policy-improvement quantity.

Step 3 (convergence test). If the new policy equals the old one, the algorithm is stopped with the policy found. Otherwise, go to step 1 with the old policy replaced by the new one.

It is our experience from numerical examples that the algorithm converges after few iterations, typically at most 5. The algorithm requires per iteration the solving of N_S equations. The number N_S of states increases rapidly in S. For example, if $m = 10$, then we have $N_0 = 1$, $N_1 = 12$, $N_2 = 78$, $N_3 = 364$, $N_4 = 1365$ and $N_5 = 4368$. Therefore, even for this modest value of m, the computational burden of the algorithm becomes overwhelming unless S is small. Tijms [7] suggests a value-iteration algorithm to solve large-scale Markov decision problems. But we desist from discussing such problems in this paper.

For $m = 10$ and various values of the demand rate λ and the shortage cost p, Table 5 reports the average cost AC and the stockout probability B for four policies: the best pure base-stock policy (denoted by $(S, 0)$), the simple modified base-stock policy (SIMP) introduced in the next section, the best modified base-stock policy (denoted by (S, t) as explained in the next section) and the optimal policy (OPT). For each reported case, Eq. 4 provides the same S_ρ-value with the three specifications of ρ mentioned in Section 2 and this value turns out to be the best base-stock. OPT is computed in the following way. First, the upper bound S for the inventory position is fixed as the best pure base-stock and the policy-iteration algorithm, initialized with the best pure base-stock policy, is run. Next, starting with the policy found, the algorithm is run again with S increased by one to search for a better policy. For all cases reported, no better policy was found by this search.

Each case in Table 5 is identified by a pair (λ, p). OPT equals the base-stock policy with $S = 1$ for the two cases (0.5,2.5) and (0.5,5). For the four cases

	$p = 2.5$			$p = 5$			$p = 10$		
	policy	AC	B	policy	AC	B	policy	AC	B
$\lambda = 0.5$	(1,0)	1.086	34.44	(1,0)	1.517	34.44	(2,0)	1.933	8.29
	SIMP	1.086	34.44	SIMP	1.517	34.44	SIMP	1.925	8.48
	(1,0)	1.086	34.44	(1,0)	1.517	34.44	(2,6)	1.924	8.62
	OPT	1.086	34.44	OPT	1.517	34.44	OPT	1.924	8.62
$\lambda = 1.0$	(2,0)	1.703	21.21	(2,0)	2.233	21.21	(3,0)	2.714	6.92
	SIMP	1.678	21.78	SIMP	2.223	21.78	SIMP	2.698	7.11
	(2,8)	1.668	24.13	(2,4)	2.222	21.51	(3,3)	2.698	7.11
	OPT	1.668	24.13	OPT	2.222	21.51	OPT	2.695	7.20
$\lambda = 1.5$	(2,0)	2.157	32.54	(3,0)	2.750	14.61	(4,0)	3.328	5.45
	SIMP	2.137	33.38	SIMP	2.725	15.05	SIMP	3.306	5.58
	(2,5)	2.137	33.38	(3,3)	2.725	15.05	(4,3)	3.303	5.95
	OPT	2.137	33.38	OPT	2.721	15.14	OPT	3.296	5.78

Note. $(S,0)$ denotes the best pure base-stock policy, SIMP denotes the simple modified base-stock policy, (S,t) denotes the best modified base-stock policy and OPT denotes the optimal policy.

Table 5: Average cost AC and stockout probability B (in %) for four policies when $m = 10$.

(0.5,10), (1,2.5), (1,5) and (1.5,2.5), OPT has $S = 2$ and it prescribes to place one order for one unit in the states described by the vectors (0,0,0), (1,0,0) and $(0, x_1, 0)$ for $x_1 = t+1, t+2, .., 10$, where t is specified by the best modified base-stock policy. For the two cases (1,10) and (1.5,5), OPT has $S = 3$ and it prescribes to place one order for one unit in the states described by the vectors $\mathbf{x}$ with the components listed in Table 6. We desist from listing the vectors describing the states in which OPT with $S = 4$ prescribes to place orders for the last case (1.5,10). But it is noteworthy that OPT never prescribes to place more than one single-unit order in any state.

The approach used for computing OPT is applicable for any parameter setting. We suggest to run the policy-iteration algorithm first with the upper bound S for the inventory position set equal to S_ρ where the traffic intensity ρ is determined e.g. by the c-specification mentioned in Section 2. The initial policy can be the pure base-stock policy specified by this S. But the simple modified base-stock policy presented in the next section is a better choice. When the algorithm has found the policy which is optimal subject to the condition that the inventory position is at most S, then increase S by one and run the algorithm again to investigate whether a better policy exists.

inventory position	x_0	x_1	x_2	x_3
0	0	0	0	0
1	1	0	0	0
	0	≥ 3	0	0
2	2	0	0	0
	1	≥ 4	0	0
	0	10	≥ 4	0
	0	9	$\geq z_9$	0
	0	8	$\geq z_8$	0

Note. If $(\lambda, p) = (1, 10)$ then $z_8 = 6$ and $z_9 = 5$. If $(\lambda, p) = (1.5, 5)$ then $z_8 = 5$ and $z_9 = 4$.

Table 6: Components for the vectors $\mathbf{x}$ describing states in which OPT prescribes to place an order for one unit when (λ, p) equals (1,10) or (1.5,5).

5 Modified base-stock policies

For the continuous-review $(S-1, S)$ system with $S = 2$, Hill [3] has shown that the long-run average cost can be decreased by imposing a lower limit on the time between epochs in which placing a replenishment order is permitted. The cost reduction arises because with a short time between two consecutive ordering epochs it is most likely to have no demand during the interval between the epochs in which the two orders are delivered. The expected incremental shortage cost caused by extending this interval shall in optimum balance the holding cost rate (our monetary unit) for the unit coming from the second order.

For the periodic-review inventory system considered in this paper, we shall now introduce modified base-stock policies which are easy to implement. Each modified base-stock policy is specified by a pair (S, t) where S is the base-stock and t is a lower bound for the number of review periods between review epochs in which placing a replenishment order is permitted. We refer to this t as the minimal time between ordering epochs. The $1 + S$ dimensional vectors $\mathbf{x}$ describing the states of the (S, t) system have components satisfying Condition 5 supplemented with

$$x_j \geq t + x_{j+1} \text{ if } x_{j+1} > 0,\ j = 1, 2, .., S-1.$$

Therefore, the number $\tilde{N}_{S,t}$ of states is at most N_S and equals N_S only when $t = 0$ or $S \leq 2$. The state space $\tilde{I}_{S,t}$ is a subset of I_S. The exact average cost $AC(S, t)$ for the modified base-stock policy equals m multiplied by the g-value found by solving an adapted version of Eqs. 8 and 9. We delete in Eq. 8 the linear equations for the superfluous states $i \in I_S \backslash \tilde{I}_{S,t}$.

A simple modified base-stock policy is specified in the following way. The base-stock S is set equal to S_ρ where the traffic intensity ρ is determined e.g. by the c-specification mentioned in Section 2 and the minimal time t between ordering epochs is fixed as the ratio m/S truncated to an integer. A best modified base-stock policy is specified by a pair (S,t) which minimizes $AC(S,t)$. If the minimum is attained for more than one pair, then our tie-breaking rule prescribes to select the pair with the smallest S and the largest t. But ties have never occured in the numerical examples investigated by us.

The cases reported in Table 5 (and numerous others) illustrate that the simple and best modified base-stock policies provide most of the cost reduction which can be obtained by replacing the best pure base-stock policy by the optimal policy. Both modified base-stock policies are easy to implement. The simple one is easy to compute whereas extensive computations are needed to find the best one. We suggest to compute the best modified base-stock policy by a neighborhood search procedure which is initialized with the simple modified base-stock policy. For each investigated (S,t) system, the procedure requires the the solution of $\tilde{N}_{S,t}$ equations to evaluate the system. This evaluation is computational demanding if $\tilde{N}_{S,t}$ is large and it is avoided by restricting attention to the simple modified base-stock policy. This policy is often a good choice and we recommend to implement it. It outperforms the best pure base-stock policy when $S \geq 2$ and its simplicity makes it worthy for implementation.

6 Conclusions

We have used Erlang's loss formula to derive an approximate expression for the long-run average cost of a periodic-review inventory system controlled by a pure base-stock policy. A good base-stock S is easily computed from this expression. If $S \geq 2$ then we recommend to modify the base-stock policy by never placing more than one single-unit order at any review epoch and by imposing a lower bound t for the number of review periods between review epochs in which placing an order is permitted. We suggest to fix t as the largest integer which is less than or equal to the ratio of the number m of review periods per delivery period and S. Our suggestion implies for increasing m that the minimal time between ordering epochs approaches the ratio of the delivery time for one order and S. Therefore, we suggest to control the continuous-review inventory system by a modified base-stock policy with the minimal time between ordering epochs specified by this ratio.

References

[1] Axsäter, S., "Optimization of order-up-to-S policies in two-echelon inventory systems with periodic review", *Naval Research Logistics* 40 (1993), 245-253.

[2] Graves, S. C., "A multiechelon inventory model with fixed replenishment intervals", *Management Science* 42:1 (1996), 1-18.

[3] Hill, R. M., "On the suboptimality of $(S-1,S)$ lost sales inventory policies", *International Journal of Production Economics* 59 (1999), 387-393.

[4] Johansen, S. G., and Thorstenson, A., "Optimal and approximate (Q, r) inventory policies with lost sales and gamma-distributed lead time", *International Journal of Production Economics* 30-31 (1993), 179-194.

[5] Schultz, C. T., "On the optimality of the $(S-1, S)$ policy", *Naval Research Logistics* 37 (1990), 715-723.

[6] Smith, S. A., "Optimal inventories for an $(S-1, S)$ system with no backorders", *Management Science* 23:5 (1977), 522-528.

[7] Tijms, H. C., *Stochastic Models*, Wiley, Chisester, 1994.