Preventive Maintenance at Opportunities of Restricted Duration

Rommert Dekker*
Koninklijke/Shell-Laboratorium, Amsterdam

Eric Smeitink†
Vrije Universiteit, Amsterdam

This article deals with the problem of setting priorities for the execution of maintenance packages at randomly occurring opportunities. These opportunities are of restricted duration, implying that only a limited number of packages can be executed. The main idea proposed is to set up a model for determining the optimal execution time for the individual maintenance packages and to develop cost criteria for deviations from the optimal time. In this article we use the block replacement model, but the approach can be easily extended to include other optimization models as well. Using Monte Carlo simulation the performance of the method is compared with various heuristics, both for a two-package and a multipackage case. © 1994 John Wiley & Sons, Inc.

1. INTRODUCTION

Most preventive maintenance (inspections, component replacements) of production systems require shutdown of the units involved. If these units are used continuously, as is the case in process industry, shutdowns can be very costly, and management will try to minimize their duration and frequency.

It is not uncommon, however, that for a variety of reasons production units have to be shut down for a short time, and in principle, these moments can be used for doing preventive maintenance. In some cases, a major problem in making use of these opportunities is that they cannot be planned in advance (at least not by the maintenance department), as they merely occur at random and are restricted in duration. As a result, traditional maintenance planning and scheduling fails to make effective use of them.

To overcome these problems, a decision support system (DSS) for opportunity-based preventive maintenance has been developed at the Koninklijke/Shell-Laboratorium, Amsterdam. In order to make use of short-lasting opportunities, preventive maintenance work has to be split up into a number of maintenance packages that are small enough to be executed at an opportunity. These packages (typically 40–80 per unit) can be defined by the user in the initialization phase of the DSS.

After all necessary data have been entered, the DSS determines for each maintenance package an optimal control limit. This control limit indicates that the package should

*Present address of Rommert Dekker: Econometrisch Instituut, Erasmus Universiteit, P.O. Box 1738, 3000 DR Rotterdam, the Netherlands.
†Present address of Eric Smeitink: PTT Research, P.O. Box 421, 2260 AK Leidschendam, the Netherlands.

Naval Research Logistics, Vol. 41, pp. 335–353 (1994)

CCC 0894-069X/94/030335-19

be executed at an opportunity if the time since its previous execution exceeds the control limit. At an opportunity, however, more packages may be due than can be executed, and a selection is called for. The DSS supports the user by providing a list of maintenance packages due, which are ranked in order of importance.

In this article we present a method to derive priorities using the opportunistic block replacement model (OBRM) as an underlying long-term optimization model. This method, which has been applied in the DSS, can be easily extended to include minimal-repair and inspection models. We will discuss these extensions in this article. The basic idea is to assign the priorities by looking one opportunity ahead. The higher the cost of deferring execution of a maintenance package until the next opportunity, the higher the position of that package on the list will be. Calculating both the costs of deferring the execution of individual maintenance packages and the optimal control limits using the OBRM as underlying optimization model, one can derive a scheduling priority criterion which is fully consistent with the optimal control limits.

Terminology

In this article we assume that after the execution of a maintenance package the system parts involved are as good as new. For the sake of clearness, therefore, we will use the terminology of replacing individual components rather than executing maintenance packages. We also assume that every maintenance package attends to only one system part. Thus executing a maintenance package is equivalent to preventively replacing a single component. There is no loss of generality because the underlying model (OBRM) that we use for determining the optimal control limits and the cost of deferring preventive replacement can easily be extended to the case where a maintenance package attends to more than one system part (see also Dekker and Smeitink [4]).

Literature

Few articles deal with opportunity maintenance, and for a review we refer to Dekker and Smeitink [4]. An often-used approach for opportunity maintenance (see e.g., Bäckert and Rippin [1], Van der Duyn Schouten and Vanneste [11]) applies Markov decision models in which the states indicate the age for each individual component. This causes a large state space, thereby severely restricting the computational evaluation. Numerical results are therefore presented for up to three nonidentical, or five identical components only. We have not come across any articles so far that also deal with setting priorities for execution of opportunity maintenance.

Comparing our approach to the literature on maintenance scheduling or scheduling in general reveals that our approach sets priorities in a way that is consistent with the determination of the optimum control limits. Pintelon [6] gives an overview of both practical and theoretical priority criteria, but none of them is based on a long-term optimization. Even worse, most scheduling criteria are static, indicating that the priority of a maintenance activity does not increase with the time it is waiting for execution. The examples Pintelon gives on dynamic priority criteria are all based on heuristics. This also holds for a recent example of a maintenance scheduling system given by Ulusoy, Oy, and Suydan [10].

Most of the literature on general maintenance optimization models (see, e.g., the reviews of Pierskalla and Voelker [5], Sherif and Smith [8]) considers a single maintenance activity. Only a few reports exist on multiple activities, and these (see Thomas [9]) merely try to group or combine activities. In both areas one disregards constraints set on the execution of maintenance and consequently one does not deal with priorities.

Outline of the Article

In Section 2 we summarize the results of Dekker and Smeitink [4] and present the general approach. As the procedure for calculating the optimal control limits and the costs of deferring preventive replacements does not take the severeness of the restriction on the opportunity durations into account, one cannot expect our strategy to be optimal. However, achieving optimality for a large number of components is computationally infeasible, and therefore impractical. In that case nearly optimal, well-structured policies are to be preferred.

In order to evaluate critically the value of the priority criterion, we have compared it with other, heuristically derived criteria in two cases, viz., a two-component and a multicomponent model. Section 3 deals with the two-component model, first with identical components, next with nonidentical ones. For two components it is possible to determine better control limits than those obtained from the OBRM and to study the difference. In the case of two identical components, the effect of a restricted opportunity duration can be studied analytically. In all other cases we had to use simulation. In Section 4 we evaluate the performance of the priority criterion in a case with 24 components. In Section 5 we give extensions to our approach. Section 6 concludes the article and indicates further research.

2. GENERAL APPROACH

In this section we will first give a detailed description of the maintenance model that we study. The maintenance strategy that we propose is based on optimality results for the opportunistic block replacement model (OBRM). After briefly summarizing these results we will formulate the maintenance strategy, which we call a ranking strategy.

2.1. The Model and the Strategy

Consider a system consisting of N components, numbered $1, \ldots, N$ and let the lifetime (time to failure) of component i be denoted by the random variable X_i with distribution function $F_i(\cdot)$ with positive support, mean μ_i and variance σ_i^2 (both finite). It is assumed that the lifetimes X_i are independent random variables. If a component fails it is replaced immediately with costs c_i^f for component i.

However, there are randomly occurring maintenance opportunities at which times one can decide to replace one or more components preventively, i.e., before failure, with costs $c_i^p < c_i^f$ for component i. A new component has the same characteristics, i.e., lifetime distribution function $F_i(\cdot)$ and costs c_i^f and c_i^p, as the one it replaces (preventively or upon failure).

We assume that the opportunities for preventive maintenance occur according to a renewal process, independently of the lifetime processes. The generic random variable Y denotes the time between two successive opportunities (abbreviated to TBO) and we assume that its distribution function $G(\cdot)$ has finite mean and variance.

The opportunities for preventive maintenance have a restricted duration. We assume that the durations of the opportunities are independent, identically distributed random variables and denote the generic variable by L. For simplicity we further assume that replacing a component preventively takes one time unit for each component, so that at each opportunity at most L components can be replaced. The extension to different replacement durations is considered in Section 5.

Furthermore we assume that the times needed for replacements are so short compared with the lifetimes of the components that they can be neglected when considering the failure processes. The duration of replacements is typically on the order of hours, whereas mean times to failure are on the order of months or years. Hence, this assumption is no restriction in practice.

The only information available at an opportunity is the time elapsed since the last preventive replacement, t_i, for each component i. Thus the decision of which components to replace preventively must be based on the values $t_1, \ldots, t_N$. This situation, which often occurs in practice, is referred to as *block replacement* in the literature. Note that since decisions can only be taken at opportunities which occur according to a renewal process, there is no need to consider the actual time, or the past evolution of the opportunity process.

Objective

The objective is to generate at each opportunity a ranked list of maintenance packages (components to be replaced). Using this list and more detailed information on, e.g., manpower available, maintenance management can schedule the activities to be executed at an opportunity. In practice there can be various reasons to deviate from the list, such as a lack of spare parts.

In our model, however, we assume that components are always replaced in order of decreasing priority. The more formal objective in our mathematical model is then to decide at every opportunity which components to replace preventively so as to minimize the long-term average costs. The main problem in the cost minimization is that the opportunities have a restricted duration. Apart from deciding for each component if it should be replaced at all at a given opportunity, one must also decide *which* components should actually be replaced if there are more than L candidates. Giving priority to some components automatically implies that the preventive replacement of one or more other components has to be postponed. This interaction between the components renders the search for an optimal strategy computationally infeasible.

Outline of the Strategy

The strategy we propose is based on optimality results for the case without restrictions on the opportunity durations. In this case there is no interaction between the components so that the problem decomposes into solving N independent opportunistic block replacement models (OBRM).

In the OBRM a single component, say number i, is replaced preventively if at an opportunity the time since its last preventive replacement, t_i, exceeds a certain control limit. Under some mild conditions (see Section 2.2 below) it can be shown that there exists a finite optimal *control limit*, t_i^*, that minimizes the expected long-term average cost.

The optimal control limit strategy with control limit t_i^* can be characterized in a different way. In Section 2.3 we will introduce a cost function $R_i(t_i)$ so that the control limit strategy is equivalent to the following *one-opportunity-look-ahead strategy*. If at an opportunity the time since the last preventive replacement of component i is t_i, then the component i is replaced preventively if the expected cost of deferring its replacement, $R_i(t_i)$, is positive. In the case of restricted opportunity durations the components are assigned priorities according to the values $R_i(t_i)$. The component for which deferring replacement is most expensive gets the highest priority, etc.

2.2. Summary of Results for the OBRM

In this section we summarize the results from Dekker and Smeitink [4] concerning the equivalence of the optimal control limit strategy and a one-opportunity-look-ahead strategy. These results are an extension of the results obtained by Berg [3] for the same model without opportunities, i.e. when preventive maintenance can be carried out at any time instant.

Consider component i and suppose that a finite control limit t is used. Thus component i is preventively replaced at an opportunity if $t_i > t$. Preventive replacements at opportunities constitute renewals since both the opportunity process and the lifetime process of component i have a renewal. Let the random variable Z_t denote the time between t and the first opportunity after t (forward recurrence time of the opportunity process) and let $M_i(\cdot)$ denote the renewal function associated with the lifetime distribution function $F_i(\cdot)$ of component i. Thus $M_i(x)$ is the expected number of failures of component i in the interval $[0, x]$ if, starting with a new component at time 0, component i is only replaced upon failure. It then follows from renewal theory that the expected long-term average costs $\Phi_i(t)$ for component i, using control limit t, is given by

$$\Phi_i(t) = \frac{c_i^p + c_i^f \int_0^\infty M_i(t + z)\, dP(Z_t \leq z)}{t + E[Z_t]}, \tag{1}$$

where the integral in the numerator represents the expected number of failures during a renewal cycle (time between two successive preventive replacements). The denominator represents the expected cycle length.

Under the conditions stated below it can be proven that there exists a finite optimal control limit, t_i^*, minimizing (1).

Conditions

The main condition that suffices for the existence of a finite optimal control limit is that

$$\frac{c_i^p}{c_i^f} < \frac{1}{2}\left(1 - \frac{\sigma_i^2}{\mu_i^2}\right). \tag{2}$$

The other, more technical conditions are that the distribution functions $F_i(\cdot)$ and $G(\cdot)$ must be continuously differentiable with finite first and second moment and that the renewal density function $m_i(\cdot)$, defined as $m_i(x) = d/dx\, M_i(x)$, is continuously increasing in x. These conditions are assumed to be satisfied throughout.

The One-Opportunity-Look-Ahead Strategy

Now we formulate a one-opportunity-look-ahead strategy that compares the following two possibilities: preventive replacement at this opportunity or at the next opportunity. We denote the expected time between two successive opportunities by ν and define for every component i the function $\eta_i(\cdot)$ by

$$\eta_i(t_i) = \frac{c_i^f}{\nu} \int_0^\infty \{M_i(t_i + y) - M_i(t_i)\}\, dG(y). \tag{3}$$

The interpretation of $\eta_i(t_i)$ is the following. If at an opportunity the time elapsed since the last preventive replacement of component i is t_i, then $\eta_i(t_i)$ is the expected average cost due to failures of component i between this opportunity and the next one if the component is not preventively replaced.

It can be shown that the functions $\Phi_i(\cdot)$ and $\eta_i(\cdot)$ intersect exactly once, in the minimum point t_i^* *of* $\Phi_i(\cdot)$. Thus

$$\eta_i(t_i) = \Phi_i(t_i) \Longleftrightarrow t_i = t_i^*. \tag{4}$$

Moreover, we have that

$$\eta_i(t_i) \geq \Phi_i^* \Longleftrightarrow t_i \geq t_i^*, \tag{5}$$

where $\Phi_i^* = \Phi_i(t_i^*)$. It follows directly from (5) that the optimal control limit strategy is equivalent with the following *one-opportunity-look-ahead strategy*. Replace component i preventively at an opportunity if $\eta_i(t_i) \geq \Phi_i^*$. Thus at each opportunity the expected average cost, $\eta_i(t_i)$, of deferring preventive replacement of component i is compared with the minimum expected long-term average cost Φ_i^*. In Section 2.3 we will see that this interpretation provides a very useful basis for assigning priorities in case of restricted opportunity durations.

Another advantage of the one-opportunity-look-ahead strategy is that using (4) and (5) the optimal control limit t_i^* and associated cost Φ_i^* can be calculated in an efficient way. In Dekker and Smeitink [4] numerical procedures are given for general lifetime distributions and Coxian-2 distributions for the TBO. A random variable Y has a Coxian-2 distribution if it can be represented as

$$Y = \begin{cases} E_1, \text{ with probability } 1 - b, \\ E_1 + E_2, \text{ with probability } b, \end{cases}$$

where E_1 and E_2 are independent, exponentially distributed random variables.

2.3. The Ranking Strategy

The strategy we propose in case of restricted opportunity lengths is based on the one-opportunity-look-ahead strategy for the OBRM. Define for each component i, $i = 1, \ldots, N$, its *ranking criterion* $R_i(t_i)$ by

$$R_i(t_i) = \eta_i(t_i) - \Phi_i^*. \tag{6}$$

The cost function $R_i(t_i)$ reflects the expected average cost of deferring the preventive replacement of component i at an opportunity as a function of the time elapsed since its last preventive replacement, t_i. If at an opportunity there are two components that have exceeded their control limit then it makes good sense to give priority to the component with the highest ranking, since deferring its replacement is more expensive. Therefore, the components are placed on an ordered ranking list at each opportunity, with the component with the highest ranking on top, the one with the second highest ranking in the second place, etc.

We denote by OBRC 1 the ranking strategy that prescribes the preventive replacement at each opportunity of the first L components on the ranking list, provided that their ranking criterion is positive. From (5) it immediately follows that

$$R_i(t_i) \geq 0 \Longleftrightarrow t_i \geq t_i^*, \tag{7}$$

so that the ranking strategy OBRC 1 is fully consistent with the optimal control limits t_i^* obtained from the OBRM.

In order to compare strategy OBRC 1 with other strategies of the same type we define the class of *ranking strategies* as all those strategies based on individual control limits $\hat{t}_i$ and ranking criteria $\hat{R}_i(t_i)$. under such a strategy component i is preventively replaced at an opportunity if $\hat{t}_i \geq t_i$ whenever possible. Priorities are assigned in decreasing order of the values $\hat{R}_i(t_i)$. Thus OBRC 1 is the ranking strategy based on the control limits t_i^* and the ranking criteria $R_i(t_i)$ obtained from the OBRM.

In general, the optimal strategy will not be in the class of ranking strategies. Ranking strategies are very appealing, however, because of their simple structure and the fact that the computational effort is only linear in the number of components, N. The optimal strategy, by contrast, has a complex structure and will be very difficult, if not impossible, to obtain.

Now the question arises whether OBRC 1 is the best possible strategy in the class of ranking strategies. We expect not, since the control limits t_i^* and the cost functions $R_i(t_i)$ themselves do not account for the interaction between the components. As the only consequence of the restricted opportunity durations is that preventive replacements must sometimes be delayed, we conjecture that the control limit for component i in the optimal ranking strategy is smaller than t_i^*. A proof of this conjecture for the special case with two identical components and exponential times between opportunities has been given by Smeitink [7].

3. TWO-COMPONENT MODEL

We now investigate how well the ranking strategy OBRC 1 performs in the simplest case, i.e., the two-component model with exponentially distributed times between opportunities. (In Section 4 nonexponential times between opportunities are considered.) In this case the restricted duration of the opportunities can be represented in the following way: with probability p only one component can be replaced at an opportunity, and with probability $1 - p$ both components can be replaced.

3.1. Two Identical Components

For two identical components and exponentially distributed times between opportunities the optimal ranking strategy for block replacement, referred to as OBRC 2, was obtained by Smeitink [7]. We use this result to compare the suboptimal ranking strategy OBRC 1 with the optimal ranking strategy OBRC 2. As we assume in this section that the two components are identical, we suppress unnecessary subscripts, i.e. we write $\eta(\cdot)$ instead of $\eta_i(\cdot)$, etc.

The fact that the two components are identical implies that they have the same optimal control limit, t^*, and minimum expected long-term average costs, Φ^*, in case of unrestricted opportunity durations. Hence it follows from (5) and the definition of the ranking criterion (6) that

$$R(t_1) \geq R(t_2) \Longleftrightarrow t_1 \geq t_2. \tag{8}$$

Thus the ranking strategy OBRC 1 prescribes that one should preventively replace component i at an opportunity if $t_i \geq t^*$, where t_i is the time elapsed since the last replacement of component i, $i = 1, 2$. Further, if at an opportunity with duration for only one component both components should be preventively replaced, i.e., if $t_{i1} \geq t_{i2} \geq t^*$, then component i_1 gets priority if $t_{i1} > t_{i2}$. In case $t_{i1} = t_{i2}$ either component 1 or component 2 is replaced, each with probability $\frac{1}{2}$.

Now there are two possibilities. A component is replaced preventively either at the first or at the second opportunity after it has reached the control limit t^*. In the latter case the preventive replacement of this component was *blocked* at the first opportunity after it had reached its control limit; i.e., that opportunity had a duration allowing replacement of only one component and the other component gained priority.

Due to blocking the times between two successive preventive replacements of the same component resulting from using control limit t^* are now stochastically larger than in the case of unrestricted opportunity durations. Thus the control limit t^* will in general not be optimal. It will be clear, however, that the optimal ranking strategy OBRC 2 assigns priorities in the same way as OBRC 1 and that both components have the same optimal control limit, to be denoted by t_p^*. The subscript p refers to the situation that with probability p only one component can be preventively replaced at an opportunity. In order to compensate for the blocking phenomenon we expect that $t_p^* \leq t^*$. That this is indeed the case is a direct consequence of the first part of the following result.

$$p_a > p_b \Rightarrow t_{p_a}^* < t_{p_b}^* \tag{9}$$

and

$$p_a > p_b \Rightarrow \Phi_{p_a}^* > \Phi_{p_b}^*, \tag{10}$$

where Φ_p^* denotes the minimum expected long-term average costs associated with the optimal control limit t_p^*. The analytical results for this model require the blocking probability $b(t)$, which is defined as the limiting probability that the preventive replacement of component i is blocked at the first opportunity after it has reached its control limit, if the same control limit t is used for both components.

Using an imbedded Markov chain technique it can be shown that

$$b(t) = \frac{p}{pt/\nu + 1}, \qquad t \geq 0. \tag{11}$$

Table 1. Comparison of the restricted and the unrestricted case.

			Restricted ($p = 1$)				
	Unrestricted		OBRC 2			OBRC 1	
ν	t^*	Φ^*	t_p^*	Φ_p^*	$b(t_p^*)$	$\Phi_p(t^*)$	$b(t^*)$
0	2.60	0.7820					
0.5	2.18	0.7948	2.12	0.7989	0.191	0.7991	0.187
1.0	1.85	0.8276	1.70	0.8524	0.370	0.8531	0.351
2.0	1.41	0.9278	1.19	1.0241	0.627	1.0248	0.587
3.0	1.17	1.0396	0.96	1.1947	0.758	1.1948	0.719
5.0	0.92	1.2320	0.76	1.4342	0.868	1.4342	0.845

Due to symmetry, both components experience the same blocking probability. Thus, in the long run, the random time between two successive preventive replacements of the same component is given by $t + Y_{t,p}$, where

$$Y_{t,p} = \begin{cases} Y_1, & \text{with probability } 1 - b(t), \\ Y_1 + Y_2, & \text{with probability } b(t), \end{cases} \tag{12}$$

with Y_1 and Y_2 independent, exponentially distributed random variables, both with mean ν. Notice that $Y_{t,p}$ has a Coxian-2 distribution so that the numerical procedures from Dekker and Smeitink [4] can be used to calculate t_p^* and Φ_p^* (see also Section 2.2).

In Table 1 we compare the ranking strategy with control limit t^* (OBRC 1) and the optimal ranking strategy with the control limit t_p^* (OBRC 2) for various expected times between opportunities ν. The lifetimes of the components have a Weibull distribution with mean $\mu = 10$ and shape $\beta = 2$. The costs of a failure are $c^f = 20$ and a preventive replacement costs $c^p = 1$. In the case of restricted opportunity durations we use $p = 1$, so that at every opportunity at most one component can be replaced. From (9)–(11) it follows that $p = 1$ represents an extreme case, because it yields the smallest optimal control limit and the highest cost and blocking probability. Thus the results for the unrestricted case and the restricted case with $p = 1$ bound the results for $0 < p < 1$.

Conclusions

Using the control limit t^* instead of the optimal control limit t_p^* results in only slightly higher costs. Thus the ranking strategy OBRC 1 with the easily obtained control limit t^* is nearly optimal in this case. The relative difference between the control limits can be much larger, due to the flatness of the cost curve $\Phi_p(t)$ around its minimum. Further we notice that for moderate values of ν the restricted opportunity durations already cause a substantial increase of the expected long-term average costs as compared with the unrestricted case.

3.2. Two Nonidentical Components

For nonidentical components we did not obtain the optimal ranking strategy. As we also wanted to gain an idea in this case of how far from optimal OBRC 1 is within the class of ranking strategies, we first derived approximations for the optimal control limits $t_{p,i}^*$ within the subclass of ranking strategies with fixed ranking criteria $R_i(t_i)$, $i = 1, 2$

defined in (6). We then compared OBRC 1 with the ranking strategy based on these new control limits $t^*_{p,i}$ and the same ranking criteria $R_i(t_i)$ that are used in OBRC 1. We will refer to this strategy as OBRC 3. The expected long-term average costs for component i resulting from OBRC 3 are denoted by $\Phi^*_{p,i}$, $i = 1, 2$.

Note that OBRC 3 is not the optimal ranking strategy, as it still uses the ranking criteria $R_i(t_i)$, which in general are not optimal in combination with control limits different from t^*_i. This is only the case for identical components. But comparing OBRC 1 with OBRC 3 provides an estimate of the improvement resulting from using the optimal control limits instead of the control limits t^*_i.

The approximations for the optimal control limits $t^*_{p,i}$ were obtained in the following way. For a given pair of control limits $(\bar{t}_{p,1}, \bar{t}_{p,2})$ we simulated the model and obtained a point estimate for the associated expected average costs, to be denoted by $\bar{\Phi}_{p,i}(\bar{t}_{p,i})$. The priorities were assigned according to (6), irrespective of the pair of control limits under consideration. Thus if $t_1 \geq \bar{t}_{p,1}$ and $t_2 \geq \bar{t}_{p,2}$, then the component with the highest value $R_i(t_i) = \eta_i(t_i) - \Phi^*_i$ was assigned the highest priority. Starting with the pair of optimal control limits for the case of unrestricted opportunity durations, (t^*_1, t^*_2), we approximately obtained the pair of control limits $(t^*_{p,1}, t^*_{p,2})$ that minimize $\bar{\Phi}_{p,1}(\bar{t}_{p,1}) + \bar{\Phi}_{p,2}(\bar{t}_{p,2})$ by using a trial-and-error procedure and conjecturing that $t^*_{p,i} \leq t^*_i$; i.e., we only considered values $\bar{t}_{p,i} \leq t^*_i$.

In order to investigate the value of ranking we also considered two randomized strategies. These strategies use the same control limits t^*_i as OBRC 1 to decide whether or not component i should be preventively replaced. However, if both components exceed their control limit at an opportunity with duration for only one component then strategy RANDOM 1 selects one of the components at random, that is, with equal probabilities $\frac{1}{2}$. RANDOM 2 is a modification of RANDOM 1 as it gives priority to the component whose preventive replacement was delayed at the previous opportunity. Thus RANDOM 2 precludes that the preventive replacement of the same component is postponed twice.

In Table 2 we list simulation results for four different combinations of components. In all cases considered the lifetimes follow a Weibull distribution with mean 10. The shape parameter β of the Weibull lifetime distribution and the failure costs c^f of the components are varied with fixed cost of preventive replacement $c^p = 1$. Φ^* is the minimum expected long-term average cost in case of unrestricted opportunity durations. The TBO are exponentially distributed with fixed mean $\nu = 1$ and, just as in the previous section, we consider the extreme case; i.e., we assume that at an opportunity at most one component can be preventively replaced ($p = 1$). We list the point estimations for the average cost and blocking probabilities for the individual components. The last column contains the sum of the individual costs and, in parentheses, the half-width of the corresponding 95% confidence interval.

Conclusions

The first observation to be made is that both OBRC strategies are better than the strategies that assign priorities at random. As expected, OBRC 3 is slightly better than OBRC 1 since it uses better control limits and RANDOM 2 is better than RANDOM 1 since it precludes that the preventive replacement of a component is deferred at two consecutive opportunities.

The OBRC strategies are better than the RANDOM strategies because they give priority to the most important component. A measure for the importance of a component

Table 2. Effect of the restricted opportunity duration.

	Combination 1			Combination 2		
	Comp. 1 $\beta = 2.0$ $c^f = 5$ $\Phi^* = 0.380$	Comp. 2 $\beta = 2.0$ $c^f = 10$ $\Phi^* = 0.560$	Total 0.940	Comp. 1 $\beta = 2.0$ $c^f = 20$ $\Phi^* = 0.828$	Comp. 2 $\beta = 4.0$ $c^f = 20$ $\Phi^* = 0.375$	Total 1.203
OBRC 1						
E[av. cost]	0.391	0.564	0.955	0.868	0.386	1.254
p_{block}	0.407	0.057	(0.008)	0.312	0.278	(0.015)
OBRC 3						
E[av. cost]	0.389	0.560	0.949	0.865	0.385	1.250
p_{block}	0.374	0.101	(0.008)	0.285	0.336	(0.013)
RANDOM 1						
E[av. cost]	0.390	0.572	0.962	0.872	0.419	1.291
p_{block}	0.239	0.161	(0.007)	0.242	0.294	(0.014)
RANDOM 2						
E[av. cost]	0.387	0.573	0.960	0.867	0.406	1.273
p_{block}	0.242	0.167	(0.008)	0.269	0.332	(0.014)

	Combination 3			Combination 4		
	Comp. 1 $\beta = 2.0$ $c^f = 20$ $\Phi^* = 0.828$	Comp. 2 $\beta = 4.0$ $c^f = 50$ $\Phi^* = 0.513$	Total 1.341	Comp. 1 $\beta = 2.0$ $c^f = 50$ $\Phi^* = 1.428$	Comp. 2 $\beta = 4.0$ $c^f = 50$ $\Phi^* = 0.509$	Total 1.937
OBRC 1						
E[av. cost]	0.897	0.532	1.429	1.543	0.585	2.128
p_{block}	0.536	0.151	(0.016)	0.304	0.555	(0.064)
OBRC 3						
E[av. cost]	0.891	0.511	1.402	1.521	0.577	2.098
p_{block}	0.507	0.182	(0.015)	0.288	0.621	(0.033)
RANDOM 1						
E[av. cost]	0.880	0.622	1.502	1.594	0.694	2.288
p_{block}	0.285	0.300	(0.017)	0.285	0.335	(0.034)
RANDOM 2						
E[av. cost]	0.874	0.593	1.467	1.578	0.605	2.183
p_{block}	0.328	0.346	(0.018)	0.365	0.453	(0.034)

with respect to preventive maintenance is the difference $\rho = c^f/\mu - \Phi^*$, which indicates the cost per unit time that can be saved by optimally executing preventive maintenance at opportunities (of unrestricted duration). Note that c^f/μ is the long-term average cost resulting from replacement upon failure only. In general ρ increases with increasing failure cost c^f and Weibull shape parameter β. For higher values of β the Weibull probability density function is more peaked, so that we have better information about the lifetime of the component. This in turn implies that preventive replacement can be more effective.

Consider for example Combination 3 of Table 2. From the difference between the blocking probabilities for both components it follows that most times the OBRC strategies give priority to the more important Component 2. In contrast, the blocking probabilities following from the random strategies are roughly the same.

Comparing Combination 1 of Table 2 with relatively unimportant components and Combination 4 with much more important components we see that the difference between the OBRC strategies and the random strategies increases with the importance of the components. Also the difference between OBRC 1 and OBRC 3 increases with increasing importance of the components. This is due to the fact that the value of preventive maintenance for important components is more sensitive with respect to the preventive replacement interval, i.e., the control limit used. Strategy OBRC 1 does not account for blocking in calculating the control limits, whereas OBRC 3 does.

A last remark is in order. Although OBRC 1 clearly outperforms RANDOM 1 (and RANDOM 2) we also see from Table 2 that the differences are not very large. This is due to the fact that in the case of two components RANDOM 1 assigns priority to the same component as OBRC 1 with probability $\frac{1}{2}$. Remember that the random strategies only assign the priorities at random, but that they use the same control limits t_i^* as OBRC 1. In Section 4 below we will see that the value of ranking can be much higher in a multicomponent case.

4. A MULTICOMPONENT CASE

In this section we evaluate the performance of the ranking strategy OBRC 1 based on the control limits t_i^* and the priority criterion $R_i(t_i)$ defined in (6) in a multicomponent case, again by using simulation. We did not try to obtain better control limits as we did in Section 3.2. The two-component case is a simple example in this respect. However, we did consider alternative priority criteria and compared the performance in various cases. All other criteria were also used in combination with the control limits t_i^* from the OBRM, so that only the selection from the components due for replacement, i.e., the assignment of priorities, would differ. Below we will specify the unit to be maintained, the alternative criteria, and the type of restrictions on the opportunity durations.

4.1. Outline of the Unit to be Maintained

We considered a 24-component unit. Component lifetimes were assumed to be independent and to be following a Weibull distribution. Both lifetime and cost parameters were varied widely. Mean lifetimes were taken from the range of 5, 10, and 20 time units, and Weibull shape factors were either 2 or 4. Component failure costs were either 5, 10, 20, or 50. Full data are given in Table 3. The last four columns of this table give the optimal control limits and associated average costs for exponentially and Coxian-2 distributed TBO, respectively.

4.2. Alternative Priority Criteria

As in Section 3 we also considered two randomized selection strategies. RANDOM 1 selects at random (i.e., with equal probabilities) from all the components due. RANDOM 2 first gives priority to those components that were already due at the previous opportunity, but could not be replaced. It selects randomly from this group, and if more components can be replaced, a random selection is made from the newly due components. These strategies can be regarded as base cases, as they give us an idea about the value of setting priorities. Next to these criteria, we considered a heuristic criterion, called

Table 3. Description of the components.

Comp.	μ	β	c^f	c^p	TBO exp. t^*	TBO exp. Φ^*	TBO Coxian2 t^*	TBO Coxian2 Φ^*
1	5	2	5	1	2.252	0.774	2.336	0.765
2	5	2	10	1	1.258	1.181	1.322	1.142
3	5	2	20	1	0.710	1.856	0.745	1.733
4	5	2	50	1	0.318	3.645	0.328	3.222
5	5	4	5	1	2.020	0.539	2.127	0.509
6	5	4	10	1	1.468	0.715	1.599	0.642
7	5	4	20	1	1.037	0.988	1.178	0.825
8	5	4	50	1	0.609	1.653	0.740	1.212
9	10	2	5	1	5.092	0.380	5.234	0.378
10	10	2	10	1	3.041	0.560	3.169	0.553
11	10	2	20	1	1.848	0.828	1.955	0.806
12	10	2	50	1	0.930	1.439	0.999	1.349
13	10	4	5	1	4.761	0.246	4.949	0.241
14	10	4	10	1	3.748	0.302	3.960	0.291
15	10	4	20	1	2.933	0.375	3.164	0.353
16	10	4	50	1	2.072	0.513	2.318	0.461
17	20	2	5	1	10.955	0.189	11.133	0.188
18	20	2	10	1	6.812	0.275	6.982	0.274
19	20	2	20	1	4.362	0.397	4.521	0.294
20	20	2	50	1	2.407	0.650	2.542	0.637
21	20	4	5	1	10.454	0.119	10.670	0.118
22	20	4	10	1	8.501	0.143	8.729	0.142
23	20	4	20	1	6.925	0.172	7.164	0.170
24	20	4	50	1	5.253	0.221	5.503	0.216

CORF (combination of relevant factors). The CORF ranking criterion $R_i^{\mathrm{corf}}(t_i)$ was defined as

$$R_i^{\mathrm{corf}}(t_i) = \frac{c_i^f t_i \beta_i \nu}{c_i^p \mu_i 2 \mu_i}. \tag{13}$$

The idea behind this criterion is that the greatest priority should be given to those components which have a high cost of failure, for which a relatively long time has elapsed since the last preventive replacement, which have a peaked lifetime distribution and finally, for which opportunities occur relatively infrequently.

4.3. Description of the Opportunities

Apart from the exponential distribution we also considered a Coxian-2 distribution with squared coefficient of variation $c_Y^2 = 0.75$ for the TBO. The mean time between opportunities was set to one time unit in both cases.

Setting priorities is only needed when there are restrictions on the number of components that can be executed. As the outcome of the comparison may depend on the type of restriction we considered two types of restriction. In the first one, the number L of components which can be replaced at an opportunity is constant. We varied this number from 0 (when no components can be replaced at all) to 24 (when there is no restriction at all). Next to this we considered a variable restriction. In case STOCH 1

L was drawn with equal probabilities from the values 3, 6, 9, 12, and 15 ($E[L] = 9$), while in case STOCH 2 L was drawn in an equal fashion from the values, 6, 9, 12, 15, and 18 ($E[L] = 12$).

4.4. Results and Discussion

In Table 4 the total expected long-term average costs are given for the unit as a whole. For all cases we used the same random number seeds. Half-lengths of the 95% confidence intervals are given in parentheses. From the table the following observations can be made.

- The effect of the restricted opportunity duration on the average costs is only substantial for $L \leq 12$ and increases rapidly if L goes to 0.

Table 4. Effect of the opportunity restriction for several MPs.

L	TBO distr	OBRC 1	CORF	RANDOM 1	RANDOM 2
0		59.50	59.500	59.500	59.500
1	EXP	41.31 (0.12)	45.27 (0.13)	51.68 (0.08)	51.44 (0.07)
	C2	40.79 (0.07)	44.99 (0.08)	51.62 (0.07)	51.47 (0.05)
2	EXP	31.81 (0.13)	34.80 (0.15)	44.06 (0.14)	42.35 (0.15)
	C2	30.86 (0.11)	34.13 (0.12)	43.78 (0.12)	42.19 (0.13)
3	EXP	26.57 (0.14)	28.62 (0.15)	37.21 (0.18)	34.30 (0.21)
	C2	25.37 (0.08)	27.68 (0.08)	36.71 (0.12)	33.50 (0.14)
4	EXP	23.05 (0.13)	24.42 (0.10)	31.36 (0.18)	28.45 (0.18)
	C2	21.80 (0.09)	23.40 (0.09)	30.47 (0.13)	26.92 (0.15)
5	EXP	21.09 (0.13)	21.97 (0.13)	27.47 (0.20)	25.22 (0.17)
	C2	19.64 (0.07)	20.67 (0.08)	26.09 (0.12)	23.21 (0.13)
6	EXP	19.97 (0.09)	20.51 (0.09)	24.88 (0.15)	23.27 (0.14)
	C2	18.36 (0.06)	18.95 (0.07)	22.99 (0.11)	21.10 (0.10)
9	EXP	18.59 (0.08)	18.73 (0.08)	20.88 (0.10)	20.45 (0.11)
	C2	16.96 (0.05)	17.07 (0.05)	18.79 (0.08)	18.39 (0.08)
12	EXP	18.35 (0.08)	18.40 (0.08)	19.52 (0.10)	19.46 (0.09)
	C2	16.70 (0.05)	16.73 (0.05)	17.44 (0.07)	17.42 (0.07)
15	EXP	18.23 (0.11)	18.25 (0.11)	18.76 (0.13)	18.77 (0.13)
	C2	16.63 (0.05)	16.63 (0.05)	16.91 (0.06)	16.89 (0.06)
18	EXP	18.16 (0.09)	18.16 (0.09)	18.31 (0.10)	18.33 (0.10)
	C2	16.62 (0.04)	16.62 (0.04)	16.69 (0.04)	16.68 (0.04)
21	EXP	18.22 (0.10)	18.22 (0.10)	18.24 (0.10)	18.24 (0.10)
	C2	16.62 (0.05)	16.62 (0.05)	16.62 (0.05)	16.62 (0.05)
24	EXP	18.18 (0.09)	18.18 (0.09)	18.18 (0.09)	18.18 (0.09)
	C2	16.61 (0.04)	16.61 (0.04)	16.61 (0.04)	16.61 (0.04)
STO	EXP	19.11 (0.12)	19.29 (0.12)	21.86 (0.17)	21.26 (0.16)
I	C2	17.46 (0.06)	17.64 (0.07)	19.89 (0.11)	19.24 (0.10)
STO	EXP	18.47 (0.10)	18.56 (0.10)	19.93 (0.14)	19.76 (0.13)
II	C2	16.83 (0.04)	16.90 (0.04)	17.93 (0.06)	17.77 (0.05)

- OBRC 1 always results in the lowest expected long-term average costs. However, if $L > 9$, the difference with CORF is small compared with the effect of ranking itself, which is expressed by the differences with the two RANDOM strategies. For $L \leq 9$, the difference with CORF becomes larger if L decreases (up to 10% for $L = 1$). Note that $L = 0$ implies that no preventive replacements can be carried out at all and that $L = 24$ implies that there is no restriction at all, so all four policies produce the same average costs in these two extreme cases.
- Comparing OBRC 1 and CORF with the two RANDOM strategies reveals that the value of ranking depends substantially on the severeness of the restriction (i.e., the value of L). For $2 < L < 21$ these differences decrease with L. For $L = 1$, however, the differences are smaller than for $L = 2$. This can be explained from the following example.

EXAMPLE: Suppose we have to pick L ($L \leq 4$) numbers out of the four numbers 10, 10, 3, 1. In Table 5 below we compare the value of the sum of the numbers in case we have a ranked list at our disposal with the expected value of the sum if we have to choose at random. This example shows that the savings induced by ranking first increase and then decrease with L.

- As expected, RANDOM 2 (with priority for those replacements that have been blocked before) produces lower average costs than RANDOM 1. The difference between the two strategies is only significant, however, for $L \leq 9$. The difference increases with L for $L < 4$, whereas for $4 \leq L \leq 9$, it decreases with L. This can also be explained by the above-mentioned reasoning about the value of ranking.
- The average costs in case of Coxian-2 distributed TBO with $c_Y^2 = 0.75$ are smaller than those for exponentially distributed TBO ($c_Y^2 = 1$), indicating that the average costs increase with the coefficient of variation of the TBO distribution (a result which has also been reported by Dekker and Smeitink [4]). The (relative) effect of shortening the opportunity durations is for $L \leq 6$ larger and for $L > 9$ somewhat smaller for a smaller coefficient of variation.
- The results of the cases with L random show that the variation in the opportunity restriction increases the average costs significantly. Apparently, the effects of a more severe and (equally) lighter restriction are not offset by each other. This can also be inferred from Table 4, where for each strategy the effect of being able to replace an extra component is decreasing for increasing L.

The overall conclusion is that the simulation study certainly indicates the value of the priority scheduling criterion based on (6) that is used in OBRC 1. Subjects for further research are indicated in Section 6.

5. EXTENSIONS

In this section we will first discuss two maintenance models other than the OBRM for which a priority ranking criterion can be derived in an analogous way, viz., a minimal repair model and an inspection model (see Barlow and Proschan [2]). These maintenance models have in common that an optimal control limit strategy exists in the case of unrestricted opportunity durations and that component failures do not influence the

Table 5. Example: The effect of ranking.

L	Value (ranked list)	E[value(random)]	Difference
1	10	6	4
2	20	12	8
3	23	18	5
4	24	24	0

times between successive preventive maintenance actions. Models with this property can be analyzed in the same way as the OBRM and are easily extended to the case where more than one system part is addressed by an individual maintenance package. The important thing to note is that for all models the priority criterion has the same meaning and can therefore be used to *set priorities between activities of various types.*

Further we will show how the ranking criteria can be used to include nonidentical replacement times for the components (or equivalently, nonidentical execution times for the maintenance packages). The main point that we make is that it is not a good idea to use normalized ranking values, i.e. the original ranking values $R_i(t_i)$ divided by the execution times, as priority scheduling criterion in that case. Instead, the *additivity* of the ranking criterion should be used to formulate a knapsack problem. The additivity of the ranking criterion can also be fruitfully exploited for other optimization purposes.

5.1. Minimal-Repair Model

As in the OBRM, we assume that component i can be preventively replaced at an opportunity against costs c_i^p. After a preventive replacement the component is as good as new. However, if component i fails then it only gets a minimal repair, which costs c_i^f. After a minimal repair the component is assumed to have been brought back to the state it was in just before failure. Thus t time units after its last preventive replacement the component has age t, irrespective of the number of component failures. It then follows that the expected number of failures of component i in an interval of length t, starting with a new component, equals

$$Q_i(t) = \int_0^t q_i(u)\, du, \tag{14}$$

where $q_i(\cdot)$ denotes the failure rate of component i, given by

$$q_i(u) = \frac{f_i(u)}{1 - F_i(u)}. \tag{15}$$

From this point the analysis is analogous to the analysis of the OBRM with $M_i(\cdot)$ and $m_i(\cdot)$ replaced by $Q_i(\cdot)$ and $q_i(\cdot)$, respectively.

5.2. Inspection Model

In this model we assume that failures have no direct consequences and hence do not reveal themselves. However, if component i has failed a (virtual) cost c_i^f per unit of time is incurred, for example, due to a decreased safety level. At an opportunity a component can be inspected against costs c_i^p for component i. If a failed component is found upon inspection then it is repaired without any additional costs. It is assumed that a component is always as good as new after inspection. The expected time that component i is down in an interval of length t, starting with a new component, equals

$$S_i(t) = \int_0^t (t - u)\, dF_i(u). \tag{16}$$

Defining $s_i(t) = d/dt\, S_i(t) = F_i(t)$, the analysis is analogous to the analysis of the OBRM, now with $S_i(\cdot)$ and $s_i(\cdot)$ replacing $M_i(\cdot)$ and $m_i(\cdot)$, respectively.

5.3. Additivity of the Ranking Criterion

The interpretation of the ranking criterion $R_i(t_i) = \eta_i(t_i) - \Phi_i^*$ is that it indicates the expected average cost of deferring the execution of maintenance package MP i until the next opportunity. With this interpretation in mind it is easily seen that the cost of deferring the execution of a collection S of MPs, all with a positive ranking criterion, is given by the sum of the individual ranking criteria of the MPs in S.

Suppose that, due to some side constraints, one wants to deviate from the priorities indicated by the ranking list. This situation may occur if, for example, the necessary spare parts or manpower to execute some of the maintenance packages with a high ranking criterion is not available. If one has the option to execute either the subset S_1 or S_2 of maintenance packages with a positive ranking criterion, then it will be advantageous to execute the collection S_1 if and only if

$$\sum_{i \in S_1} R_i(t_i) \geq \sum_{i \in S_2} R_i(t_i). \tag{17}$$

Thus, apart from indicating the scheduling priority of individual MPs, the ranking criteria can also be used as cost figures in additional optimization routines. This is in contrast with, e.g., the criterion R_i^{corr} defined in (13), which can only be used as a relative value to indicate the scheduling priority of MP i.

5.4. Different Preventive Replacement Times

Until now we have assumed that replacing a component, or equivalently, executing a maintenance package (MP) takes one time unit. If the execution times are different, then instead of executing the packages in decreasing order of their ranking value $R_i(t_i)$, we can use the additivity of the ranking criterion to formulate the following knapsack problem. Let E_i denote the execution time of maintenance package i. We assume that at the beginning of an opportunity, maintenance management knows its duration l(a realization of the random variable L). Given l, the execution times E_i and the ranking values $R_i(t_i)$, we have to choose a subset S of MPs to be executed in order to

$$\begin{aligned} &\text{maximize} \sum_{i \in S} R_i(t_i), \\ &\text{subject to} \sum_{i \in S} E_i \leq l. \end{aligned} \tag{18}$$

As only those activities with a positive ranking value $R_i(t_i)$ need to be considered, the above knapsack problem will be of small to moderate size and can be easily solved.

Another way to incorporate the different execution times would be to define new ranking values $\hat{R}_i(t_i) = R_i(t_i)/E_i$ and to execute the MPs in decreasing order of the new ranking values $\hat{R}_i(t_i)$. The next example illustrates the problem resulting from this approach, namely, that MPs with large execution times will not be high on the new ranking list, although executing them results in high savings.

EXAMPLE: Consider the following situation at an opportunity with duration l = 10 (Table 6). If the MPs are executed in decreasing order of the values $\hat{R}_i(t_i)$ then the net return is 15 + 10 + 24 = 49 with a total execution time 6. Although there are four

Table 6. Example: Discrimination of MPs with large execution times

MP	$\hat{R}_i(t_i)$	E_i	$R_i(t_i)$
1	15	1	15
2	10	1	10
3	6	4	24
4	5	9	45

time units left, MP 4 cannot be executed, as its execution requires nine time units. The optimal solution is to execute MP 1 and MP 4, yielding a return of 15 + 45 = 60 with a total execution time 10.

6. CONCLUSIONS AND SUBJECTS FOR FURTHER RESEARCH

In this article we have presented a priority scheduling criterion with a sound theoretical justification. The resulting maintenance strategy performed better than other, more heuristically derived criteria we considered. The strategy has a simple structure and can be used on line, due to the modest computational requirement, which is only linear in the number of maintenance packages.

The fact that the priorities are assigned on the basis of a cost comparison makes the approach very flexible. The ranking criteria can, e.g., be used as inputs to more advanced scheduling routines that account for manpower requirements. This can also be done in a planned environment, where maintenance packages can in principle be executed every weekend, say, but where sometimes priorities must be set due to capacity restrictions. Further, the method presented in this article can be easily extended to include other underlying long-term optimization models.

The two main subjects for research that we want to indicate are closely related. In the multicomponent case of Section 4 we did not obtain better control limits as we did in the two-component case. However, if the restrictions on the opportunity durations are severe, then preventive maintenance actions must often be deferred. The random time between two successive preventive replacements of component i using control limit t_i^* will be substantially larger than $t_i^* + Z_{t_i^*}$ in that case. Thus the long-term optimization with the OBRM as underlying model (unrestricted opportunity durations) will produce suboptimal control limits.

One option for a strategy improvement would be to do a simulation run in order to estimate the distribution of the number of opportunities O_i, that the preventive replacement of component i is deferred, given that the control limits $t_1^*, \ldots, t_N^*$ and the ranking criteria $R_i(t_i)$ defined in (6) are used. The random time between two successive preventive replacements of component i, using control limit t, is then approximately given by $t + Z_t + O_i Y$ for values of t close to t_i^*. Using these adjusted random times in the OBRM for all components separately, possibly better control limits could be obtained.

It will be clear that such a complicated strategy improvement procedure cannot be implemented as an automatic procedure in a DSS. In practice it will be more important to know, for the maintenance strategy suggested, whether the restrictions on the opportunity durations are so severe that certain components will (almost) never be preventively replaced. A simple criterion that would indicate those components without requiring simulation would be very useful.

ACKNOWLEDGMENT

The authors like to thank Marcel van der Lee and Theo Mandos for carrying out the simulation studies reported in this article.

REFERENCES

[1] Bäckert, W., and Rippin, D.W.T., "The Determination of Maintenance Strategies for Plants Subject to Breakdown," *Computers & Chemical Engineering,* **9**, 113–126 (1985).
[2] Barlow, R.E., and Proschan, F., *Mathematical Theory of Reliability,* Wiley, New York, 1965.
[3] Berg, M., "A Marginal Cost Analysis for Preventive Maintenance Policies," *European Journal of Operational Research,* **4**, 136–142 (1980).
[4] Dekker, R., and Smeitink, E., "Opportunity Based Block Replacement," *European Journal of Operational Research,* **53**, 46–63 (1991).
[5] Pierskalla, W.P., and Voelker, J.A., "A Survey of Maintenance Models: The Control and Surveillance of Deterioriating Systems," *Naval Research Logistics Quarterly,* **23**, 353–388 (1976).
[6] Pintelon, L., "Performance Reporting and Decision Tools for Maintenance Management," Ph.D. thesis, Catholic University of Leuven, Belgium, 1990.
[7] Smeitink, E., "Stochastic Models for Repairable Systems," Ph.D. thesis, Vrije Universiteit Amsterdam, 1992.
[8] Sherif, Y.S., and Smith, M.L., "Optimal Maintenance Models for Systems Subject to Failure—A Review," *Naval Research Logistics Quarterly,* **28**, 47–74 (1981).
[9] Thomas, L.C., A Survey of Maintenance and Replacement Models for Maintainability and Reliability of Multi-Item Systems, *Reliability Engineering,* **16**, 297–309 (1986).
[10] Ulusoy, G., Or, I., and Soydan, N., "Design and Implementation of a Maintenance Planning System," *International Journal of Production Economics,* **24**, 263–272 (1992).
[11] Van der Duyn Schouten, F.A., and Vanneste, S., Analysis and Computation of (n, N) Strategies for Maintenance of a Two-Component System, *European Journal of Operational Research,* **48**, 260–274 (1990).

Manuscript received June 19, 1991
Revised manuscript received March 15, 1993
Accepted August 24, 1993