A Review of Multi-Component Maintenance Models with Economic Dependence

Rommert Dekker*, Frank van der Duyn Schouten† and Ralph Wildeman*

Abstract

In this paper we review the literature on multi-component maintenance models with economic dependence. The emphasis is on papers that appeared after 1991, but there is an overlap with Section 2 of the most recent review paper by Cho and Parlar (1991). The classification scheme that we choose is primarily based on the various options of grouping maintenance activities: grouping either corrective or preventive maintenance actions, or combining preventive maintenance activities with corrective actions. As such, this classification links up with the possibilities for grouped maintenance activities that exist in practice.

Keywords: Maintenance, planning, replacement, multiple components

1 Introduction

During the past two decades, there has been a growing interest in the modelling and optimisation of maintenance of systems consisting of multiple components. There are two major reasons for this. First, improvements in analytical techniques and the availability of fast computers have allowed more complex systems to be investigated. Second, people have realised that interactions between components in a system cannot be neglected and should be taken into account in maintenance decisions.

Interactions between components can be classified into three different types (see, e.g., Thomas 1986): economic dependence, structural dependence, and stochastic dependence. Economic dependence implies that costs can be saved when several components are jointly maintained instead of separately, that is, economies of scales can be obtained. Structural dependence applies if components structurally form a part, so that maintenance of a failed component implies maintenance of working components. Finally, stochastic dependence occurs if the state of components influence the lifetime distribution of other components.

Most multi-component maintenance models consider only one of these dependence, since combining more than one makes the models too complicated to solve or analyse. In this review we will focus on economic dependence. As an example of economically dependent components, assume a system in which the maintenance of each component requires preparatory or set-up work that can be shared when several components are maintained simultaneously. The cost of this set-up work is often called the *set-up cost* and may consist of the down-time cost due to production loss if the system cannot be used during maintenance, or of the preparation cost associated with erecting a scaffolding or opening a machine. Set-up costs can be saved when maintenance activities on different components are executed simultaneously, since execution of a group of activities requires only one set-up. This can yield considerable cost savings, and therefore the development of optimisation models for multiple components is an important research issue.

The review article of Cho and Parlar (1991) gives a good overview of the multi-component maintenance literature up to 1991, including machine-interference/repair models, group-replacement models of various types, spare-parts models, and inspection models. In this review we will exclusively deal with multi-component maintenance models that are based on economic dependence. Emphasis will be on new

*Erasmus University Rotterdam
†Tilburg University

modelling and solution techniques that have been published since 1991, but we will also pay attention to that part of the literature before 1991 that exclusively deals with economic dependence. Consequently, there is an overlap with Section 2 of the paper by Cho and Parlar (1991), but our classification scheme differs from theirs. We do not aim to give solely a list of all papers that have appeared. Instead, we want to give insight in the approaches that are followed and explain the implications of certain modelling techniques and assumptions with respect to practical applicability.

Classification

A first classification of the maintenance models with economic dependence is based on the planning aspect: stationary or dynamic. In stationary models, a long-term stable situation is assumed and mostly these models assume an infinite planning horizon. Models of this kind provide static rules for maintenance which do not change over the planning horizon. They generate for example long-term maintenance frequencies for groups of related activities or control limits for carrying out maintenance depending on the state of components.

In dynamic models, short-term information such as a varying deterioration of components or unexpected opportunities can be taken into account. These models generate dynamic decisions that may change over the planning horizon.

2 Stationary Grouping

The stationarity assumption in this kind of models facilitates the mathematical analysis; it is often possible to derive analytical expressions for optimal control parameters and the corresponding optimal costs. A disadvantage of stationary grouping is that short-term information cannot be taken into account.

We divide the stationary-grouping models into three categories: grouping corrective maintenance, grouping preventive/plannable maintenance, and opportunistic maintenance. In the first case, components are only correctively maintained and a failed component can be left in the failed state until its corrective maintenance is carried out simultaneously with other failed components.

In the second case, preventive maintenance is carried out to prevent failures or to decrease operating costs, and this is planned in advance and in such a way that set-up costs can be saved by simultaneous execution.

In the third case, maintenance is not necessarily planned in advance. However, set-up savings can be obtained since (corrective or preventive) maintenance on a component yields an opportunity for maintenance on other components.

Unless stated otherwise, the planning horizon in these models is infinite.

2.1 Grouping Corrective Maintenance

Grouping of corrective maintenance is applicable for systems in which some kind of redundancy is available. In such systems it is possible to leave failed components in a failed condition for some period of time until an appropriate moment for repair arises. Policies of this type are motivated by the existence of economies of scale through simultaneous repair of a number of (identical) components. On the other hand, leaving components in a failed condition for some period of time increases the risk of costly production losses.

Consider n identical components which are operating in parallel and which are subject to random failures. The lifetimes of the individual components are independent and identically distributed random variables with distribution function $F(t)$. When a component is waiting for repair or under repair, this does not affect the functionality of the other components. However, the unavailability of failed components will cause a loss of productivity at a rate $C_1(k)$ per unit of time, when k components are down simultaneously. There is full information about the number of failed components and about the age of non-failed components. At any moment one can decide to start repair or replacement of any number of failed components. Repair of a failed component results in a component which is 'as good as new'. The costs of repair depend on the number of components included in the repair operation, and are denoted by $C_2(k)$.

In Okumoto and Elsayed (1983) it is assumed that repairs are instantaneous, and the cost rate $C_1(k)$ and the cost function $C_2(k)$ are both supposed to be linear in k, that is, $C_1(k) = kC_1$ and $C_2(k) = C_0 + kC_2$, where C_0, C_1, and C_2 are known constants. Okumoto and Elsayed restrict attention to repair rules that are characterised by one single random variable, indicating the time between the start of a fully operational system and the moment of the first simultaneous repair of all failed components. However, without apparently being aware of this, the authors only allow random variables which are independent of the failure process. For example, the policy that only starts a system repair as soon as the number of failed components has passed a given threshold is not taken into consideration. For this rather restricted class, the authors provide sufficient conditions for the existence of a deterministic, unique, and finite optimal repair time. These conditions include the IFR-property of the underlying lifetime distribution $F(t)$, as well as some technical conditions on the model parameters.

Assaf and Shanthikumar (1987) consider the same model as described by Okumoto and Elsayed (1983) and additionally assume that the components have identical negative-exponentially distributed lifetimes. They provide sufficient conditions for the optimality of a control-limit-type policy: start a simultaneous repair on *all* failed components if and only if the number of failed components reaches a threshold value. Also an explicit characterisation of the optimal threshold value is provided. Gertsbakh (1984) essentially obtained the same results using a general optimal stopping approach developed by Aven and Bergman (1986). Note that the fact that in any repair operation *all* failed components are included is according to intuition, but not obvious. Assaf and Shanthikumar extend their model to include the situation of imperfect information, that is, the number of failed components can only be observed by inspections at a fixed cost. It is shown that the optimal policy is characterised by critical numbers $(m, t_0, t_1, ..., t_m)$, such that whenever k $(0 \leq k \leq m)$ failed components are observed on inspection, one waits for t_k units of time and inspects again, and whenever $k > m$ failed components are found, a repair operation is started including all failed components and one inspects again t_0 time units later. Moreover, it is shown that $t_0 \geq t_1 \geq ... \geq t_m$. Finally, Assaf and Shanthikumar show that in both models an optimal policy exists which does not allow operational components to idle.

Sculli and Wu (1981) earlier dealt with the situation of an n-component system with non-identical Weibull-distributed lifetimes. From the outset it is assumed that the whole system is replaced as soon as one of the components fails. The analysis concentrates on approximation of the mean time between system replacements, that is, the approximation of the mean of the minimum of n independent non-identical Weibull-distributed random variables.

Jansen and Van der Duyn Schouten (1995) show that the non-idling property of the optimal policies in the model of Assaf and Shanthikumar (1987) does not hold in case the repair is not instantaneous. They consider the case in which the costs for production losses far outweigh the actual repair costs of the components, that is, $C_2(k) = 0$. The costs of production losses $C_1(k)$ are non-decreasing and convex in k. The lifetime distributions are again exponentially distributed (like in the model of Assaf and Shanthikumar 1987) and there is full information about the number of failed components. Also the repair time of an individual component is exponentially distributed. There are no economies of scale in repair time, that is, the total time required to repair two components is distributed as the sum of two identical exponentials. There is ample repair capacity, that is, the repair facility can handle all n existing components, if necessary. Due to the assumptions on the cost functions, it follows (in accordance to the result of Assaf and Shanthikumar 1987) that the optimal policy will not allow the repair of a failed component to be postponed until other components have failed (the critical repair limit is equal to one). The control issue that arises in this context is whether it is profitable to put components in cold-standby position by taking them out of operation deliberately (*forced idling*). Jansen and Van der Duyn Schouten show that this phenomenon of forced idling can indeed occur under the optimal policy. Also they show that the optimal policy is monotone in the sense that the number of components put into operation when k operational components are available is a monotone increasing function of k.

2.2 Grouping Preventive/Plannable Maintenance

A big advantage of preventive maintenance is that it is plannable. This is important when work preparation is necessary, so that for example new components can be ordered in time and also enough maintenance crew is available at the planned maintenance execution times.

We distinguish between the models in which all components are in principle always jointly maintained, that is, in which there is only one fixed group (where components that are too young may be left out), and the models in which there can be several groups and an optimal grouping structure has to be determined.

2.2.1 Fixed Group

A number of papers consider the case where groups are made beforehand, either on basis of engineering judgement or as outcome of some optimisation model. It may well be the case that the set-up costs are that high, that without using a formal method it is clear that a certain grouping of activities is optimal. Alternatively, groups may have been created to reduce administrative efforts in registering information on execution (see Dekker and Van Rijn 1996 for a practical situation of this kind).

One important class of papers in this area is based on block replacement. Under this policy, a component is maintained or replaced preventively at fixed intervals, regardless of the events occurring within such an interval. The average costs $g(t)$ when using a block interval of length t is equal to $g(t) = (s + M(t))/t$, where $M(t)$ expresses the expected costs due to failures, minimal repairs and operating costs over the interval, and s denotes the cost of preventive replacement. It will be immediately clear that block replacement may also be applied to a group of components while using the same interval, and then we have $g(t) = (s_G + \sum_{i=1}^{n} M_i(t))/t$ for maintaining a group of n components preventively at costs s_G each time. One such an example, considering minimal repairs and failure replacements within the interval, and dealing with multiple components, is Sheu (1991). Notice that the block policy allows a coordination of the replacement of components, but has the disadvantage that it does not react to events within the cycle, such as failure replacements.

Archibald and Dekker (1994) consider therefore an extension of block replacement, which originates from Berg and Epstein (1976). If at a scheduled moment for replacement the age of a component is below a threshold value, its replacement is delayed to the next scheduled moment (so-called modified block replacement). Upon failure the components get a replacement. This policy captures both the flexibility of age replacement and the coordination of block replacement, which allows it to be used for multiple components (using the same interval for all components). While its loss in average costs compared to an optimal age-based policy is only marginal, the policy is much easier evaluated and optimised; it requires only a number of steps that is linear in the number of components.

To conclude this section on preventive maintenance in a fixed group, we mention three papers in which the assumption to do all maintenance at a fixed interval is relaxed, but where the only preventive maintenance option is still to replace the whole system. However, as opposed to block replacement, the decision to replace the whole system may be based on the actual age information of the individual components. As such, these models have strong resemblance with the models treated in Section 2.3 (including the curse of dimensionality due to explosion of the state space in case of many components with non-exponential lifetime distributions). However, the typical aspect of Section 2.3, the simultaneous execution of preventive and corrective maintenance activities, is lacking here.

In Dekker and Roelvink (1995) the group is only fixed for preventive maintenance and not for corrective maintenance, in which case only the failed component is replaced. The reason for this is that preventive maintenance is done while the machine is not needed for production, and saving set-up work is interesting, while upon failures the machine should be brought back to a running condition as soon as possible and a set-up saving is relatively low. Dekker and Roelvink investigate how the decision to execute the group can be based on the actual ages in case the underlying components have generally distributed times to failure. As the optimal policy is likely to have a complex structure (Özekici 1988 shows this for a related problem, see Section 2.3), they consider a simple approximative replacement criterion based on marginal-cost considerations and evaluate it with discrete-time Markov decision chains for two components and continuous-time simulation for multiple. Although this approach yields a better policy than block replacement, it does not provide an expression for the associated average costs.

Sivazlian and Mahoney (1978) consider a related problem in continuous time under periodic review, where the components are subject to stochastic deterioration (but no failure). When the n-component vector, which expresses the state of deterioration, pierces a certain surface, the entire multi-component system is replaced. In the absence of this situation nothing is replaced. The conditions for minimising the long-term cost of maintaining a system which operates under the stated policy, were found through solution of a problem in variational analysis. They do not provide a general solution approach, but give two simplified examples with identical components.

Ohashi (1982) considers a coherent system where the components have exponentially distributed up- and downtimes. Both economic and stochastic dependence between the components is taken into account. He sets up a functional equation for the optimal policy, which can only be solved numerically by applying policy improvement. Next he derives some properties of the optimal policy, like monotonicity in the number of components failed.

2.2.2 Optimisation over Groups

While in the previous subsection we have considered a fixed group, we consider here the optimisation over the possible groupings. This makes the problems more difficult, because of the combinatorial aspects that are involved.

There are two solution approaches distinguished here: direct and indirect grouping. In the former approach, the components are partitioned into a number of fixed groups and are then always maintained in these groups. In the latter approach we distinguish two cases, which will be referred to as *standard indirect grouping* and the *joint-overhaul problem.* In the first case (standard indirect grouping), an occasion for preventive maintenance is created at a basis maintenance interval, and each component is maintained at an integer multiple of this interval. In that case the groups are not fixed over time, but are formed indirectly when the maintenance of different components coincides. Notice that the basis interval is the minimum maintenance interval in this case. In the second case of indirect grouping (the joint-overhaul problem), the basis interval is the maximum interval, and in that case all components are overhauled at this interval, whereas on each component a *minor* overhaul is performed an integer number of times in between.

Direct Grouping

According to a direct-grouping strategy, all components are partitioned into a number of groups. The components in a group are always jointly maintained with a maintenance interval that is optimal for that group. The objective is to find the partition which minimises the long-run average costs. In general, the problem of finding an optimal partition is a set-partitioning problem, and as such NP-complete (see Garey and Johnson 1979). Consequently, an optimal direct-grouping strategy can be found for only a small number of components, unless a special structure is assumed.

Van Dijkhuizen and Van Harten (1995) assume the following structure: each component has an individual optimal maintenance frequency and the optimal frequency of a group of components is equal to the *highest* frequency of the individual components in that group. Consequently, the maintenance of components can only be advanced. The authors show that in this special case the problem is polynomially solvable with a dynamic-programming algorithm. The authors furthermore extend their model to deal with multiple levels of set-up costs. Because of the special structure of restricted frequencies, the authors are able to derive certain dominance results, which enable them to solve the problem with multiple set-ups to optimality for a reasonable number of components by applying a branch-and-bound algorithm. When the number of components is much larger than fifty, the running time of this algorithm becomes too long and in that case a greedy heuristic can be applied.

The direct-grouping approach mostly performs worse than the indirect-grouping approach described below. Van Eijs, Heuts and Kleijnen (1992) show that in an inventory environment (where the same grouping approaches apply for items that are jointly replenished), the indirect-grouping strategy always outperforms the direct-grouping approach. Only when the set-up cost is low, the difference between the two strategies becomes smaller. However, in that case grouping is not very economical at all.

Indirect Grouping

An indirect grouping strategy is determined by $n+1$ parameters: a real value $T > 0$ and for each component i an integer k_i, $i = 1, \dots, n$.

In the *standard indirect-grouping* models, these parameters have the following meaning: every T time units an occasion for preventive maintenance is created, and component i is preventively maintained at the integer multiple k_iT of T. For example, let T be equal to one month, and $k_1 = 1$ and $k_2 = 3$, then component 1 is preventively maintained every month, and component 2 every three months. After preventive maintenance a component is considered as good as new.

In the *Joint-Overhaul Problem* (JOP), the parameters T and k_i, $i = 1, \ldots, n$, have the following meaning: every T time units all components are jointly overhauled, called a major overhaul, and component i receives a minor overhaul every T/k_i time units, that is, $k_i - 1$ times in between. For example, let T be equal to one month, and $k_1 = 1$ and $k_2 = 3$, then components 1 and 2 are jointly overhauled every month, component 1 does not receive extra minor overhauls in between, whereas component 2 receives two extra overhauls in between, viz. every 1/3 of a month. After a major overhaul a component is considered as good as new, whereas after a minor overhaul this need not necessarily be the case.

Standard indirect grouping and the joint-overhaul problem are very much related to each other and the same solution techniques can be applied in either case. Many of the techniques originate from inventory theory, where a similar problem, viz. the joint-replenishment problem, is analysed. Many heuristics for the joint-replenishment problem have appeared in the literature (see for overview up to 1989 the review article of Goyal and Satir 1989), and several of them have been applied to the indirect-grouping models in maintenance. This shows how the results in one area of operational research can be fruitful for another. However, this is not a one-way process. Although the joint-replenishment problem has received much more attention in the literature than the related models in maintenance, a new efficient *optimal* solution method for the problem was found by following an approach that was first developed for the indirect-grouping models in maintenance.

To discuss the models below in more detail, we introduce the function $M_i(x)$, denoting the expected cumulative deterioration costs of component i (due to failures, repairs, operating costs, etc.), x time units after its latest preventive maintenance.

We will consider first the standard indirect-grouping models. Let S be the set-up cost incurred each time an occasion for preventive maintenance is created (that is, every T time units), and let s_i be the extra cost of maintaining component i preventively on an occasion. Under the standard indirect-grouping strategy, the total average costs are equal to

$$\frac{S}{T} + \sum_{i=1}^{n} \frac{s_i + M_i(k_iT)}{k_iT}. \tag{1}$$

The problem of finding optimal values for T and k_i is a mixed continuous-integer programming problem and in general such problems are difficult to solve. However, since the function (1) is separable in the vector $k = (k_1, \ldots, k_n)$, several fast (heuristic) solution methods can be developed.

The idea of standard indirect grouping originates from inventory theory (see Goyal 1973) and was introduced in maintenance by Goyal and Kusy (1985) and further developed by Goyal and Gunasekaran (1992). Goyal and Kusy (1985) assume that the deterioration-cost functions $M_i(\cdot)$ are of the form $M_i(x) = a_ix + b_ix^e$, with $a_i, b_i \geq 0$ constants for component i, and with $e \geq 1$ the same for all components. Goyal and Gunasekaran (1992) take $e = 2$, but they also take individual down-time and utilisation factors into account. In both cases the authors propose a heuristic similar to that of Goyal (1973) for the joint-replenishment problem. The heuristic iteratively determines values for T and k_i. First, it initialises the k_i to one and then finds a corresponding optimal value for T, with which new values for the k_i are determined, and so on, until no further improvement is found. The heuristic appears to be fast, which is partly due to the special (polynomial) deterioration functions $M_i(\cdot)$, which enables the authors to derive an analytic expression for T and to find values for k_i in little time as well.

The iterative algorithm of Goyal and Kusy (1985) is improved by Vos de Wael (1995), who moreover extends the standard direct-grouping approach into a model where the components are maintained according to a block-replacement model. In that case, the cost functions $M_i(\cdot)$ include the renewal function and hence loose their simple form, so that the T has to be determined numerically. The algorithm is applied to the maintenance of traffic-control systems.

Dekker, Van Egmond, Frenk and Wildeman (1995) extend the standard indirect-grouping approach of Goyal et al. (1985,1992) to deal with general deterioration-cost functions $M_i(\cdot)$. They show that many well-known maintenance models can be incorporated, such as minimal repair, inspection and block replacement. Furthermore, the authors propose an alternative solution approach using global-optimisation theory. First, the authors solve a relaxed problem, where the constraints $k_i \in \mathbb{N}$ are replaced by $k_i \geq 1$. The solution of this relaxed problem is subsequently used in a fast solution approach. The authors show that the problem can in most cases be solved to optimality in very little (linear) time by using Lipschitz optimisation. This is possible for the special deterioration cost functions of Goyal et al. (1985,1992), but also for the minimal-repair and inspection models. For the cases where this is not possible (for example for block replacement), the authors present heuristics that perform better than previously published ones. The optimal solution approach is also applied to the joint-replenishment problem, see Dekker, Frenk and Wildeman (1995), where the method can be made even faster (though still with a linear time complexity) since the cost functions in this problem have a very simple form.

In all cases above, the possibility of so-called *empty* occasions is ignored, which occur when the smallest integer k_i is larger than one. For example, if there are two components and $k_1 = 2$ and $k_2 = 3$, then two out of six occasions are not used for maintenance and so the average set-up cost equals $4S/6T$ instead of S/T. This implies that a correction factor $\Delta(k)$, $k = (k_1, \ldots, k_n)$, has to be applied (see Dagpunar 1982 in the context of the joint-replenishment problem). For example, if $k = (2,3)$ then $\Delta(k) = 4/6$. However, incorporation of this correction factor is computationally prohibitive, since the correction factor is hard to evaluate for large values of n and the corresponding objective function is not separable any more in the vector k. Dekker, Frenk and Wildeman (1996) present some results for the problem with correction factor and they show that the optimal solution for the function (1) is a good feasible solution for the problem with correction factor. The authors show by numerical experiments that the gap between this feasible solution and an optimal solution is quite small (less than one per cent on average).

We will now consider the joint-overhaul problem (JOP). In this case, S is defined as the cost of a joint overhaul (carried out every T time units), and s_i is the cost of carrying out a minor overhaul on component i. We first assume that both a major and a minor overhaul restores a component into an as-good-as-new state. The JOP is to find values for T and k_i that minimise the total average costs, which are equal to

$$\frac{S + \sum_{i=1}^{n} \{(k_i - 1)s_i + k_i M_i(T/k_i)\}}{T}. \tag{2}$$

Again this is a mixed continuous-integer programming problem, which is separable in the vector $k = (k_1, \ldots, k_n)$. This separability is used in all solution techniques described below.

The problem was first formulated as such by Sule and Harmon (1979). Like Goyal and Kusy (1985) these authors assume that the deterioration functions $M_i(\cdot)$ are of the form $M_i(x) = a_i x + b_i x^e$, with $a_i, b_i \geq 0$ constants for component i, and with $e \geq 2$ the same for all components. The authors use an iterative heuristic that was proposed by Silver (1976) in the inventory context. The algorithm is fast, partly because of the easy expressions for $M_i(\cdot)$.

Hariga (1994a) was the first who gave the JOP its name. The author uses the same functions $M_i(\cdot)$ as Sule and Harmon (1979) and applies a relaxation of the JOP which is similar to the relaxation of Dekker, Van Egmond, Frenk and Wildeman (1995) for standard indirect grouping, that is, with the constraints $k_i \in \mathbb{N}$ replaced by $k_i \geq 1$. However, the author does not subsequently solve the problem to optimality, but instead applies two heuristics that have been developed previously in the context of inventory theory. The first heuristic is similar to that of Hariga and Jackson (1990), the second is similar to the one of Maxwell and Muckstadt (1985) and of Jackson, Maxwell and Muckstadt (1985). In the second heuristic, the values of k_i are restricted to a power of two, and the interval T is restricted to a power-of-two multiple of some basic planning period. The restriction enables the author to derive a worst-case bound for the performance of the heuristic. This worst-case bound depends on the exponent e, but the costs according to the heuristic are always smaller than a factor two of the optimal costs.

Duffuaa and Ben-Daya (1994) add an extra term in the cost function so as to compensate for so-called coordination costs due to the fact that components are not overhauled according to individually optimal times. However, we do not agree with their problem formulation, since the original formulation already takes this into account. Besides, the authors use the same functions $M_i(\cdot)$ as Sule and Harmon (1979), and in that case the extra cost term can be incorporated in the constant a_i, so that there is no extra contribution in this respect. The solution approach applied by Duffuaa and Ben-Daya (1994) is the one of Kaspi and Rosenblatt (1991) for the joint-replenishment problem.

In all cases above it is assumed that the functions $M_i(\cdot)$ are polynomial as in the case of Sule and Harmon (1979). In that case, but also for more complex functions $M_i(\cdot)$, the JOP can be efficiently solved to optimality following the approach of Dekker, Van Egmond, Frenk and Wildeman (1995) for standard indirect grouping. Indeed, notice that the function (2) can be written as

$$\frac{S-\sum_{i=1}^{n} s_i}{T}+\sum_{i=1}^{n}\frac{s_i+M_i(T/k_i)}{T/k_i}=\frac{\tilde{S}}{T}+\sum_{i=1}^{n}\frac{s_i+M_i(T/k_i)}{T/k_i},$$

which resembles (1), implying that a similar optimal solution approach as that of Dekker, Van Egmond, Frenk and Wildeman (1995) can be followed to find an optimal solution to JOP in little time by applying Lipschitz optimisation.

In all cases above it is furthermore assumed that a component is as good as new both after a major and a minor overhaul. Hariga (1994b) relaxes this assumption by assuming that only after a major overhaul a component is renewed. This is reflected by a reduction in the operating age of a component at a minor overhaul, instead of resetting the age to zero. On the other hand, the author restricts the analysis to functions $M_i(\cdot)$ that are equal to those of Sule and Harmon (1979), with the exponent $e = 2$ for all components. The author then solves a relaxation similar to Hariga (1994a) and subsequently applies an easy iterative algorithm to find near-optimal values for T and k_i.

A special case of the JOP is presented by Epstein and Wilamowsky (1985) for a deterministic two-component system. In this system the components 1 and 2 have lifelimits T_1 and T_2, respectively, at which time replacement is mandatory and before which time the probability of failure is assumed to be zero. A cycle starts with two new components, and each component is replaced individually after T_1 and T_2 time units, respectively (it is assumed that T_1 and T_2 are relatively prime). The objective is to determine after how many T_1's *or* T_2's the two components are jointly replaced. This problem is similar to the JOP in the sense that the two components are jointly replaced after T time units, and each component is replaced a number of times in between. However, in the case of Epstein and Wilamowsky the value T is a multiple of either T_1 or T_2, but not of both as is the case in the JOP. This implies that if T is a multiple of T_1 (and not of T_2), then the last replacement interval of component 2 in a cycle is smaller than the usual T_2 time units, and if T is a multiple of T_2 (and not of T_1), then this holds for component 1 and T_1. The authors present theorems that reduce the total number of possible options for T, and subsequently they calculate the average costs of the remaining possibilities. The optimal value of T corresponds with the option with minimal costs. It is difficult to extend this approach for three components or more, since then the computational effort is too big.

Notice that a solution to the problem of Epstein and Wilamowsky (1985) is equivalent with a solution to the JOP. For example, let $T = m_1T_1$ and $T = m_2T_2$, with $m_1 \in \mathbb{N}$ and $m_2 \notin \mathbb{N}$, then the solution with m_2 replaced by the value of m_2 rounded above to the nearest integer has the same number of individual replacements and consequently the same costs, but in that case it is also a solution to the JOP. Consequently, the problem of Epstein and Wilamowsky can be written as the JOP, with the extra restriction that the individual replacement intervals (T/k_1 and T/k_2) are not larger than T_1 and T_2, respectively. The latter restriction can be realised by defining the functions $M_i(\cdot)$ infinite for values beyond T_1 and T_2, respectively, so that this does not complicate the problem.

2.3 Opportunistic Maintenance

In this subsection we consider models in which the combination of preventive and corrective maintenance is allowed. An advantage of this opportunistic maintenance is that also corrective maintenance can be used to save set-up costs. This is not possible in the models of Section 2.2, since there only preventive/plannable maintenance is grouped. However, a disadvantage of opportunistic maintenance is that it is often not known in advance which actions are taken, in which case no planning and no work preparation is possible.

The combination of preventive and corrective maintenance has been systematically studied for the first time by researchers from the RAND corporation in the early sixties (see, for example, Radner and Jorgenson 1963, McCall 1963, and Jorgenson and McCall 1963). In the last two decades, many extensions on the basic model have been studied. In this subsection this development will be reviewed.

Opportunistic replacement basically refers to the situation in which preventive maintenance (or replacement) is carried out at opportunities, either by choice or by restriction. In some papers it is assumed that the opportunity-generating process is completely independent of the failure process (see Vanneste 1992, Dekker and Dijkstra 1992, and Dekker and Smeitink 1991,1994). In this subsection, however, we will restrict attention to the situation in which the opportunities are by definition equal to failure epochs of individual components. Due to economies of scale in maintenance cost functions, the unpleasant event of a failing component is at the same time considered as an opportunity for preventive maintenance on other components. Note that the model with failure-independent arrival process of opportunities can be used as a (mathematically tractable) approximation of the situation where the opportunities are generated by the (many) components themselves. One should realise that in many situations combination of corrective and preventive repair is not realistic. The need for corrective maintenance arises unexpectedly, while preventive maintenance can be planned. Hence, by combining both types of repair one looses the plannable character of preventive maintenance or one has to leave failed equipment unrepaired for some period of time. However, there are situations in which both disadvantages are acceptable. In particular, when corrective repair on one single component requires dismantling of the whole system, combining a corrective repair on one component with preventive repair on neighbouring components might be worthwhile. There are two options for combination. On one hand, preventive maintenance can be advanced on occurrence of a failure, whose repair cannot be postponed. On the other hand, when failed components can be kept idling for some limited amount of time, one has the option to delay the corrective-maintenance action until the first scheduled preventive-maintenance action.

An example of the latter situation is described by Ritchken and Wilson (1990). They consider the situation as described in Section 2.1, with general (but identical) lifetime distributions for the individual components. The repair is instantaneous. Attention is restricted to a class of policies characterised by two critical numbers m and T, while the only possible action is a complete overhaul of the whole system, that is, repair of all failed components and preventive maintenance on all non-failed components. Such a maintenance operation is carried out if and only if the number of failed components has reached the level m, or T time units have passed since the previous system overhaul. The analysis is based on the observation that the moments of system overhaul generate renewal epochs for the stochastic process describing the ages of individual components. Ritchken and Wilson derive explicit expressions for the average costs per unit of time as a function of the control parameters m and T. Using some properties of this function, they present an algorithm to compute optimal values for m and T from a finite number of function evaluations.

The rest of this subsection will be devoted to the situation in which immediate repair of monitored failures is compulsory. Monitored failures of individual components are considered as opportunities for preventive maintenance of non-failed (but deteriorated) other components, giving rise to advancing previously planned preventive-maintenance actions. Consider a system consisting of $n+1$ non-identical independently operating components, labelled 0 to n, with lifetime distribution $F_i(t)$ for component i. Corrective maintenance is assumed to start on occurrence of a failure. The cost and reward structure generally is as follows. When the system is operating a reward is earned, which depends on the condition of each of the components. In the maintenance costs, economies of scale are incorporated in case maintenance activities are combined. A typical example of this is the situation in which a fixed cost S is incurred for any maintenance operation (either preventive, corrective or combined), while additionally an individual cost s_i or c_i is incurred, whenever component i is subject to preventive or corrective maintenance, respectively. All repairs are supposed to be instantaneous and to result in components which are 'as good as new'. When a corrective maintenance action on component i is combined with preventive maintenance of component k, the associated costs equal $S + s_i + a_k$, which is S cheaper than in case of separate actions. Attention is restricted to stationary dynamic policies, that is, policies which prescribe actions based on the actual condition of all components. As condition variable usually the 'age' of a component is used, which is defined as the time since the last maintenance action (although also different condition descriptions are used, see for example Fard and Zheng 1991, who use the present failure rate as description of the condition of a component). In this description the symbol '∞' is used to indicate that a component has failed. A stationary policy can now be described as a function from $\mathbf{R}^{n+1}$ to the class of all subsets of $\{0, 1, \ldots, n\}$. Let π denote a stationary policy. Then $\pi(x_0, \ldots, x_n) = \{i_1, \ldots, i_k\}$ has the following interpretation: when at a decision epoch the age configuration of the components is given by $(x_0,\ldots,x_n)$, then the components $i_1,\ldots,i_k$ are maintained simultaneously. A policy π is called *monotone* when $\pi(\mathbf{x}) \subset \pi(\mathbf{y})$ whenever $\mathbf{x} \leq \mathbf{y}$.

Radner and Jorgenson (1963) consider a series system in which the role of component 0 is different from that of the other components, in the sense that failures of component 0 are not monitored, while failures of other components are. The reward is either zero or one, depending on the non-monitored condition of component 0 (failed or non-failed). All other (monitored) components have non-identical exponentially distributed lifetimes. The maintenance costs have a general economy-of-scale structure: s_i when component i is maintained individually, and s_{0i} when components i and 0 are maintained simultaneously, with $s_i \leq s_{0i} \leq s_i + s_0$ for $1 \leq i \leq n$. It is shown that the optimal policy has the following structure: there exist critical numbers $(m_1,\ldots,m_n)$ and M, such that component 0 is replaced individually as soon as its age reaches the level M, component i is replaced alone upon failure when the age of component 0 is less than m_i, while component i is replaced upon failure together with component 0 in case the age of component 0 is between m_i and M. An explicit expression for the average costs per unit time is derived (see also McCall 1963) and brute-force optimisation of this explicit expression is proposed to find the optimal values of the control parameters $(m_1,\ldots,m_n)$ and M.

A natural extension of the model of Radner and Jorgenson (1963) is obtained, when the exclusive role of component 0 is dropped. Berg (1976) analyses a system with two identical components with exponentially distributed lifetimes, for which preventive maintenance of one of the components can be combined with a corrective-maintenance action on the other. In spite of the exponentiality assumption, there is an incentive for preventive replacement due to rewards which are decreasing with age. However, preventive replacement of a component is only allowed in combination with corrective replacement of the other component. It is shown that the optimal policy is of control limit type: replace the working component on failure of the other component if and only if its age exceeds a critical level. The method of proof is by semi-Markov decision theory.

Liang (1985) also elaborates on the situation in which preventive maintenance of a component can only be performed in combination with corrective maintenance of another component (*piggybacking*). He remarks that this type of policy can be far from optimal, in particular when the failure rate of the piggybackee is low. In this situation individual age replacement will perform better.

Berg (1978) generalises the model of Berg (1976) to the situation where preventive replacement is also allowed at non-failure epochs. Moreover, the components are no longer identical nor exponentially distributed, as in Berg (1976). Attention is restricted to the class of control-limit policies, characterised by four parameters (m_1, m_2) and (M_1, M_2): replace component i correctively when it fails and preventively when its age reaches level M_i; at replacement (corrective or preventive) of component i, replace component j preventively as well if its age exceeds level $m_j, 1 \leq i \neq j \leq 2$. This class does not necessarily contain the overall optimal policy. The stationary joint probability distribution of the ages of both components is characterised as the solution of an integral equation, which is explicitly solved for the case of two identical components with exponentially distributed lifetimes.

A generalisation of Berg (1976) in another direction is proposed by Tatsuno, Ohi and Nishida (1983). They deal with a two-component model in which one of the components with an exponentially distributed lifetime is continuously monitored, while the other component has a general lifetime distribution and is not monitored. When the latter component fails it remains unobserved, until a costly inspection is carried out. In the meantime an undetected failure will cause production losses. At failure epochs of the observed component, one has the option to replace only this failed component, or to replace this component and simultaneously do an inspection on the unobserved component (followed by a minimal repair when a failure is detected), or to replace the failed and the unobserved component simultaneously (without inspection). Inspections or replacement of the unobserved component separate from failures of the observed component are not allowed. The relevant state information is two dimensional: the age of the unobserved component and the time since the last inspection, at which the unobserved component was found in operational condition. The authors show that, under the IFR-assumption, the optimal policy has a control-limit-type structure.

For the case of two identical components (in which $m_1 = m_2 = m$ and $M_1 = M_2 = M$) and generally distributed lifetimes, Van der Duyn Schouten and Vanneste (1990) provide an efficient numerical procedure of policy-iteration type to find the optimal control parameters m and M. This special case was also numerically investigated by Vergin and Scriabin (1977) (see also Bäckert and Rippin 1985 for an extension to three components). Van der Duyn Schouten and Vanneste (1990) also show by extensive numerical experiments that the best (m, M)-policy is almost always less than 1% off from the overall optimal policy. The simple structure of (m, M)-policies and their near optimality, combined with the existence of an efficient computational procedure, make this class of (m, M)-policies of special interest for systems with two identical components. However, the question arises whether this also holds true for general multi-component systems. For a more complete overview of two-component systems we refer to Section 2.4 of Cho and Parlar (1991).

Haurie and l'Ecuyer (1982) consider a system with n identical components (from now on we consider n components in stead of $n+1$, since there is no longer a single special component) and a very simple cost structure: $S + s\nu$ for a maintenance action including ν components, irrespective whether it concerns preventive or corrective maintenance. Rewards are not taken into account. In this situation a preventive replacement will never be carried out in isolation but only in combination with a corrective replacement (recall that in this subsection corrective replacement on monitored components is compulsory upon failure). Haurie and l'Ecuyer (1982) show by counterexample that the optimal policy is not necessarily monotone.

Özekici (1988) considers the multi-component model with non-identical components, very general assumptions on the aging process and cost functions, and no a-priori restrictions as far as preventive-maintenance actions are concerned. Under these general conditions, the optimal policy can have a rather complex structure. Özekici shows, however, that the optimal policy π^* is *component-monotone*: if $i \in \pi^*(x_1, \ldots, x_n)$ then $i \in \pi^*(x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$ for all $y \geq x_i$. Also it is shown that π^* is *system-monotone*: if $\pi^*(\mathbf{x}) = \{1, \ldots, n\}$ then $\pi^*(\mathbf{y}) = \{1, \ldots, n\}$ for all $\mathbf{y} \geq \mathbf{x}$. System monotonicity refers to the property that whenever a replacement of the whole system is prescribed this will also happen under every worse age configuration. Worth mentioning is that Özekici also provides some early references for the opportunity-based replacement model dating back to 1956.

The results of Haurie and l'Ecuyer (1982) and Özekici (1988) reveal that the detailed age information on component level should be used in making maintenance decisions, but in general the resulting optimal decision rules are too complex to be numerically computed, let alone to be applied. For the case of identical components, Van der Duyn Schouten and Vanneste (1993) propose a model with sublimated age information, where for a single component only four possible states are distinguished: good, doubtful, bad, and failed. Under a given lifetime distribution, these four states can be seen as representing certain age intervals. The uncontrolled evolution of a component through these states occurs according to a Markov chain. When a component enters the bad state, a preventive-maintenance action on component level is required, while a corrective-maintenance action on component level is compulsory on entrance in the failed state. The moments at which corrective or preventive maintenance on component level is carried out, provide opportunities to replace the whole system simultaneously. Attention is restricted to decision rules of control-limit type: replace the whole system at a component-replacement epoch when the number of doubtful components exceeds a given threshold or replace the whole system at the first component-replacement epoch after the moment at which the number of doubtful components has reached a given threshold.

A recent contribution to the analysis of the general multi-component system comes from Wijnmalen and Hontelez (1996). They restrict attention to the class of $(m_1, \ldots, m_n), (M_1, \ldots, M_n)$-policies, which is the natural generalisation of the class of policies introduced by Berg (1978) for two-component systems. The exact computational optimisation procedure proposed by Van der Duyn Schouten and Vanneste (1990) for the system with two identical components cannot be easily generalised to multi-component systems. Wijnmalen and Hontelez (1996) propose an approximating optimisation method by decomposing the n-dimensional problem into n one-dimensional problems. Each one-dimensional system is modelled as a one-dimensional Markov decision process with discount opportunities. These discount opportunities in the single-component problem correspond to occurrences of compulsory maintenance of other components. The approximation is that the processes generating these discount opportunities are supposed to be independent Poisson processes, which are also supposed to be independent of the failure processes. The rates of these Poisson processes are determined by the thresholds m_i and M_i. As such, the one-dimensional problem correspond to a discretisation of the model of Radner and Jorgenson (1963). This gives rise to an iterative procedure in which alternatively the maintenance limits and the opportunity-arrival rates are adjusted until a sufficient degree of convergence is achieved. This procedure has a strong resemblance with the computational procedure developed by Federgruen, Groenevelt, and Tijms (1984) to compute 'optimal' can-order policies for multi-item inventory systems. For comments on the limitation of this procedure we refer to Van Eijs (1994).

3 Dynamic Grouping

The advantage of dynamic-grouping models is that they take short-term information into account. Situations that are not stationary, such as a varying use of components, and unexpected events that may create an opportunity to do maintenance at lower costs, can now be incorporated.

The dynamic-grouping models are divided into two categories: those with a finite horizon and those with a rolling horizon. Finite-horizon models consider the system in this horizon only, and hence assume implicitly that the system is not used afterwards, unless a so-called *residual value* is incorporated to estimate the industrial value of the system at the end of the horizon.

Rolling-horizon models also use a finite horizon, but they do so repeatedly and based on a long-term (infinite-horizon) plan. That is, once decisions of the finite horizon are implemented or when new information becomes available, a new horizon is considered, and a tentative plan based on the long-term is adapted according to short-term information.

3.1 Finite Horizon

The replacement decision variables are parametrised with time and to keep the models simple, time is usually discretised. In the latter case, quite general specifications for the costs are possible (changes over time and relations between components). Stinson and Khumawala (1987) consider such a problem. They formulate it as a mixed-integer nonlinear programming problem for which some heuristics are developed, which performance is compared to the optimal policy obtained by enumeration. Such an approach does not give any structure results and hence little insight. The optimal policy is often quite complex and time dependent. Moreover, it is usually not robust against changes in the horizon length. Some important applications of this type of modelling can be found in the scheduling of power system outages, for which Kralj and Petrovic (1987) give an overview.

Another example can be found in the model-based approach for road maintenance by Worm and Van Harten (1996). Roads can be split up into lanes, and lanes into segments of a certain length (say hundred meter). Doing the same action on adjacent lanes or segments yields set-up savings. The savings per segment tend to be a concave increasing function that is constant when more than a certain number of adjacent segments undergo the same action. The authors first apply an infinite-horizon single-component model to determine the best optimal action for each lane segment, and subsequently they apply a very simple rule for the grouping of maintenance. For given actions in a certain year, the authors consider whether it is profitable to execute the same action on other lane segments.

3.2 Rolling Horizon

Rolling-horizon models aim to bridge the gap between finite- and infinite-horizon models and to combine the advantages of both. Though the horizon itself is finite, the decisions are based on a long-term plan (infinite horizon) and are adapted according to short-term information. This yields more stable solutions than finite-horizon models, and also provides more insight into the actions that are taken. Another advantage is that interactive planning with the decision maker is possible.

Dekker, Smit and Losekoot (1991) were the first to consider this kind of maintenance planning. The authors derive for each component a penalty function, which expresses the extra costs for deviating from individually optimal execution times. These execution times are determined by individual long-term maintenance rules such as block replacement. Within a given finite planning horizon, the maintenance on each component is individually planned according to these long-term maintenance rules. Subsequently, it is possible to shift the execution times within the planning horizon to enable the grouping of maintenance activities, which saves set-up costs. The penalty functions indicate how much one has to pay for this deviation, and whether it is cost-effective, that is, whether the reduction in set-up costs outweighs the penalty costs. The objective is to find the partition of maintenance activities that minimises the total costs over the horizon. The authors formulate this grouping problem as a set-partitioning problem, analogously to the direct-grouping approach described in Section 2.2. Furthermore, they prove some lemmas that reduce the computational effort. However, this reduction depends on the data and does not decrease the time complexity, so that for more than ten activities the method becomes computationally prohibitive.

Wildeman, Dekker and Smit (1995) improve this method and derive some general conditions under which the problem reduces to a dynamic-programming problem that can be solved to optimality with a quadratic time complexity. This enables them to handle many components. The authors show furthermore that even if the conditions are not satisfied, their dynamic-programming algorithm can be used as a good heuristic. They prove that in the worst case the costs according to the dynamic-programming algorithm are less than a factor two of the optimal costs. The method allows the incorporation of unexpected activities, such as corrective maintenance which provides an opportunity for other maintenance. Furthermore, the authors show how a varying use or a changing deterioration of components can be incorporated. The method serves as a basis for a decision support system and enables interactive planning with the maintenance manager, since new decisions are easily determined and one can see how these decisions influence the final maintenance plan. The above method is included in a five-phase rolling-horizon approach, where new decisions can be taken when new information becomes available or when the decisions of the current finite planning horizon have been carried out.

Dekker, Wildeman and Van Egmond (1996) extend the above approach to age-replacement policies for a discrete-time Markov decision chain, both for an infinite and a finite horizon. Notice that for an infinite horizon the approach is used as a heuristic for stationary grouping. The approach is numerically validated and compared with an optimal solution. The performance of the approach turns out to be very good: the deviation from the optimal costs is less than one per cent provided that certain harmonisation and (finite-) horizon effects are incorporated. This result is important for two reasons. First, the approach can be considered as an extension of both finite- and infinite-horizon models, with the extra advantage that short-term information can be incorporated on a day-to-day basis. Fortunately, this extension does not affect the consistency with these models, as the numerical performance shows. The other reason is that the current Markov decision models allow only few components. Optimal solutions can be calculated for less than five components only, whereas the approach of Dekker, Wildeman and Van Egmond (1996) allows many components, since the time complexity of the dynamic-programming algorithm is only quadratic.

References

Archibald, T.W. and R. Dekker. 1994. Modified block replacement for multiple component systems. Technical Report 9412/A, Econometric Institute, Erasmus University Rotterdam (to appear in IEEE Transactions on Reliability).

Assaf, D. and J.G. Shanthikumar. 1987. Optimal group maintenance policies with continuous and periodic inspections. *Management Science*, **33**, 1440–1452.

Aven, T. and B. Bergman. 1986. Optimal replacement times — a general set-up. *Journal of Applied Probability*, **23**, 432–442.

Bäckert, W. and D.W.T. Rippin. 1985. The determination of maintenance strategies for plants subject to breakdown. *Computers and Chemical Engineering*, **9**, 113–126.

Berg, M. 1976. Optimal replacement policies for two-unit machines with increasing running costs. *Stochastic Processes and their Applications*, **4**, 89–106.

Berg, M. 1978. General trigger-off replacement procedures for two-unit systems. *Naval Research Logistics Quarterly*, **25**, 15–29.

Berg, M. and B. Epstein. 1976. A modified block replacement policy. *Naval Research Logistics Quarterly*, **23**, 15–24.

CHO, D.I. AND M. PARLAR. 1991. A survey of maintenance models for multi-unit systems. *European Journal of Operational Research*, **51**, 1–23.

DAGPUNAR, J.S. 1982. Formulation of a multi item single supplier inventory problem. *Journal of the Operational Research Society*, **33**, 285–286.

DEKKER, R. AND M.C. DIJKSTRA. 1992. Opportunity-based age replacement: Exponentially distributed times between opportunities. *Naval Research Logistics*, **39**, 175–190.

DEKKER, R., J.B.G. FRENK, AND R.E. WILDEMAN. 1995. An efficient optimal solution method for the joint replenishment problem. Technical Report 9542, Econometric Institute, Erasmus University Rotterdam (submitted for publication).

DEKKER, R., J.B.G. FRENK, AND R.E. WILDEMAN. 1996. How to determine maintenance frequencies for multi-component systems? A general approach. In *Reliability and Maintenance of Complex Systems*, pages ???–??? Springer-Verlag, Berlin. Proceedings of the NATO-ASI, Antalya, 1995.

DEKKER, R. AND I.F.K. ROELVINK. 1995. Marginal cost criteria for preventive replacement of a group of components. *European Journal of Operational Research*, **84**, 467–480.

DEKKER, R. AND E. SMEITINK. 1991. Opportunity-based block replacement: The single component case. *European Journal of Operational Research*, **53**, 46–63.

DEKKER, R. AND E. SMEITINK. 1994. Preventive maintenance at opportunities of restricted duration. *Naval Research Logistics*, **41**, 335–353.

DEKKER, R., A.C.J.M. SMIT, AND J.A. LOSEKOOT. 1992. Combining maintenance activities in an operational planning phase: a set-partitioning approach. *IMA Journal of Mathematics Applied in Business and Industry*, **3**, 315–331.

DEKKER, R., R.P. VAN EGMOND, J.B.G. FRENK, AND R.E. WILDEMAN. 1995. A general approach for the coordination of maintenance frequencies. Technical Report 9539/A, Econometric Institute, Erasmus University Rotterdam (submitted for publication).

DEKKER, R. AND C. VAN RIJN. 1996. PROMPT, a decision support system for opportunity-based preventive maintenance. In *Reliability and Maintenance of Complex Systems*, pages ???–??? Springer-Verlag. Proceedings of the NATO-ASI, Antalya, 1995.

DEKKER, R., R.E. WILDEMAN, AND R. VAN EGMOND. 1996. Joint replacement in an operational planning phase. Technical Report 9438/A, Econometric Institute, Erasmus University Rotterdam (to appear in the *European Journal of Operational Research*).

DUFFUAA, S.O. AND M. BEN-DAYA. 1994. An extended model for the joint overhaul scheduling problem. *International Journal of Operations & Production Management*, **14**, 37–43.

EPSTEIN, S. AND Y. WILAMOWSKY. 1985. Opportunistic replacement in a deterministic environment. *Computers & Operations Research*, **12**, 311–322.

FARD, N. AND X. ZHENG. 1991. An approximate method for non-repairable systems based on opportunistic replacement policy. *Reliability Engineering and System Safety*, **33**, 277–288.

FEDERGRUEN, A., H. GROENEVELT, AND H.C. TIJMS. 1984. Coordinated replenishments in a multi-item inventory system with compound poisson demands. *Management Science*, **30**, 344–357.

GAREY, M.R. AND D.S. JOHNSON. 1979. *Computers and Intractability, A Guide to the Theory of NP-Completeness*. Freeman.

GERTSBAKH, I.B. 1984. Optimal group preventive maintenance of a system with observable state parameter. *Advances in Applied Probability*, **16**, 923–925.

GOYAL, S.K. 1973. Determination of economic packaging frequency for items jointly replenished. *Management Science*, **20**, 293–298.

GOYAL, S.K. AND A. GUNASEKARAN. 1992. Determining economic maintenance frequency of a transport fleet. *International Journal of Systems Science*, **4**, 655–659.

GOYAL, S.K. AND M.I. KUSY. 1985. Determining economic maintenance frequency for a family of machines. *Journal of the Operational Research Society*, **36**, 1125–1128.

GOYAL, S.K. AND A.T. SATIR. 1989. Joint replenishment inventory control: Deterministic and stochastic models. *European Journal of Operational Research*, **38**, 2–13.

HARIGA, M. 1994a. Coordinated overhaul scheduling of production units. *Naval Research Logistics*, **41**, 565–577.

HARIGA, M. 1994b. A deterministic maintenance-scheduling problem for a group of non-identical machines. *International Journal of Operations & Production Management*, **14**, 27–36.

HARIGA, M.A. AND P.L. JACKSON. 1990. The warehouse scheduling problem: Formulation and algorithm. Technical Report 932, School of OR & IE, Cornell University, Ithaca, NY.

HAURIE, A. AND P. L'ECUYER. 1982. A stochastic control approach to group preventive replacement in a multicomponent system. *IEEE Transactions on Automatic Control*, **27**, 387–393.

JACKSON, P.L., W.L. MAXWELL, AND J.A. MUCKSTADT. 1985. The joint replenishment problem with a power-of-two restriction. *IIE Transactions*, **17**, 25–32.

JANSEN, J. AND F.A. VAN DER DUYN SCHOUTEN. 1995. Maintenance optimization on parallel production units. *IMA Journal of Mathematics Applied in Business & Industry*, **6**, 113–134.

JORGENSON, D.W. AND J.J. MCCALL. 1963. Optimal replacement policies for a ballistic missile. *Management Science*, **9**, 358–379.

KASPI, M. AND M.J. ROSENBLATT. 1991. On the economic ordering quantity for jointly replenished items. *International Journal of Production Research*, **29**, 107–114.

KRALJ, B.L. AND R. PETROVIC. 1987. Optimal preventive maintenance scheduling of thermal generating units in power systems — a survey of problem formulations and solution methods. *European Journal of Operational Research*, **35**, 1–15.

LIANG, T.Y. 1985. Optimum piggyback preventive maintenance policies. *IEEE Transactions on Reliability*, **34**, 529–538.

MAXWELL, W.L. AND J.A. MUCKSTADT. 1985. Establishing consistent and realistic reorder intervals in production-distribution systems. *Operations Research*, **33**, 1316–1341.

MCCALL, J.J. 1963. Operating characteristics of opportunistic replacement and inspection policies. *Management Science*, **10**, 85–97.

OHASHI, M. 1982. Group replacement policy for a maintained coherent system. *Journal of the Operations Research Society of Japan*, **25**, 228–237.

OKUMOTO, K. AND E.A. ELSAYED. 1983. An optimum group maintenance policy. *Naval Research Logistics Quarterly*, **30**, 667–674.

ÖZEKICI, S. 1988. Optimal periodic replacement of multi-component reliability systems. *Operations Research*, **36**, 542–552.

RADNER, R. AND D.W. JORGENSON. 1963. Opportunistic replacement of a single part in the presence of several monitored parts. *Management Science*, **10**, 70–84.

RITCHKEN, P. AND J.G. WILSON. 1990. (m, T) group maintenance policies. *Management Science*, **36**, 632–639.

SCULLI, D. AND S.Y. WU. 1981. Block replacement for units whose components have weibull life distributions. *Computers and Operations Research*, **8**, 295–302.

SHEU, S.H. 1991. A generalized block replacement policy with minimal repair and general random repair costs for a multi-unit system. *Journal of the Operational Research Society*, **42**, 331–341.

SILVER, E.A. 1976. A simple method of determining order quantities in joint replenishments under deterministic demand. *Management Science*, **22**, 1351–1361.

SIVAZLIAN, B.D. AND J.F. MAHONEY. 1978. Group replacement of a multicomponent system which is subject to deterioration only. *Advances in Applied Probability*, **10**, 867–885.

STINSON, J.P. AND B.M. KHUMAWALA. 1987. The replacement of machines in a serially dependent multi-machine production system. *International Journal of Production Research*, **25**, 677–688.

SULE, D.R. AND B. HARMON. 1979. Determination of coordinated maintenance scheduling frequencies for a group of machines. *AIIE Transactions*, **11**, 48–53.

TATSUNO, K., F. OHI, AND T. NISHIDA. 1983. Opportunistic maintenance policy with minimal repair. *Mathematica Japonica*, **28**, 327–335.

THOMAS, L.C. 1986. A survey of maintenance and replacement models for maintainability and reliability of multi-item systems. *Reliability Engineering*, **16**, 297–309.

VAN DER DUYN SCHOUTEN, F.A. AND S.G. VANNESTE. 1990. Analysis and computation of (n, N)-strategies for maintenance of a two-component system. *European Journal of Operational Research*, **48**, 260–274.

VAN DER DUYN SCHOUTEN, F.A. AND S.G. VANNESTE. 1993. Two simple control policies for a multicomponent maintenance systems. *Operations Research*, **41**, 1125–1136.

VAN DIJKHUIZEN, G. AND A. VAN HARTEN. 1995. Optimal clustering of repetitive frequency-constrained maintenance jobs with shared setups. Technical report, University of Twente, The Netherlands.

VAN EIJS, M.J.G. 1994. On the determination of the control parameters of the optimal can-order strategy. *Zeitschrift für Operations Research*, **39**, 289–304.

VAN EIJS, M.J.G., R.M.J. HEUTS, AND J.P.C. KLEIJNEN. 1992. Analysis and comparison of two strategies for multi-item inventory systems with joint replenishment costs. *European Journal of Operational Research*, **59**, 405–412.

VANNESTE, S.G. 1992. A generalized age-replacement model. *Probability in the Engineering and Informational Sciences*, **6**, 525–541.

VERGIN, R.C. AND M. SCRIABIN. 1977. Maintenance scheduling for multicomponent equipment. *AIIE Transactions*, **9**, 297–305.

VOS DE WAEL, S. 1995. A decision support system for LampRemplace (in Dutch: Strategieën voor LampRemplace). Master's thesis, Tilburg University, The Netherlands.

WIJNMALEN, D.J.D. AND J.A.M. HONTELEZ. 1996. Coordinated condition-based repair strategies for components of a multi-component maintenance system with discounts. *(To appear in the European Journal of Operational Research)*, **?**, ??–??

WILDEMAN, R.E., R. DEKKER, AND A.C.J.M. SMIT. 1995. A dynamic policy for grouping maintenance activities. Technical Report 9537/A, Econometric Institute, Erasmus University Rotterdam (submitted for publication).

WORM, J.M. AND A. VAN HARTEN. 1996. Model based decision support for planning of road maintenance. To appear in Reliability Engineering & Systems Safety.