Condition-based maintenance for complex systems

Olde Keizer, Minou Catharina Anselma

5

CBM for a redundant, k-out-of-N system with economic dependence

Parts of this chapter have been published as: M.C.A. Olde Keizer, R.H. Teunter, and J. Veldman, *Clustering condition-based maintenance for systems with redundancy and economic dependencies*, European Journal of Operational Research **251** (2) (2016) 531–540. doi:10.1016/j.ejor.2015.11.008.

5

Abstract

In this chapter, we consider a multi-component system subject to both structural dependence (through an active redundant, k-out-of-N setting) and (positive) economic dependence. In case of redundancy, postponing maintenance on some failed components is possible without reducing the availability of the system, while in case of economic dependence, maintaining several components simultaneously can be more cost efficient than performing maintenance on each component separately. No research has been performed yet on clustering CBM tasks for systems with both redundancy and economic dependence. We develop a dynamic programming model to find the optimal maintenance strategy for such systems, and show numerically that it can indeed considerably outperform previously considered policies (failure-based, age-based, block replacement, and more restricted (opportunistic) CBM policies). Moreover, our numerical investigation provides insights into the optimal policy structure.

5.1. Introduction

In practice, many multi-component systems employ redundancy, which is the most common approach to increase availability and prevent downtime of the equipment [1]. Consider for example a gas distribution company that has plants with redundant pumps to distribute gas, to ensure a continuous operation of the system and to prevent that clients will be without gas. A well-known type of redundancy in systems with spares or in so-called fault-tolerant systems is the k-out-of-N system, which has wide applications in both industrial and military systems [2]. A k-out-of-N system is a system consisting of N components which functions as long as at least k components function. Many settings can be viewed as special cases of the k-out-of-N system; the 1-out-of-N system represents a parallel system (fully-redundant), the N-out-of-N system represents a series system (non-redundant), and the k-out-of-N system with $1 < k < N$ is also known as a partially-redundant system [3, 4].

Besides redundancy (i.e., structural dependence), we focus in this chapter on economic dependence, which means that combining maintenance actions can yield a lower total cost than maintaining each component separately [5]. This is for example the case when fixed set-up costs need to be paid, which are independent of the number of components that require maintenance. Economic dependence is very common in most continuous operating systems, such as aircrafts, power-plants, or chemical processing facilities [6]. Combining maintenance on different components is also known as clustering or opportunistic maintenance. There is often a great cost saving potential by implementing an opportunistic maintenance policy for multi-component systems with economic dependence [6].

While scheduling maintenance, systems with redundancy have the unique property that the system could still function even if some components have failed. Hence, failed components do not always require immediate replacement. In case of economic dependence, it might be cheaper to postpone the maintenance until other components require maintenance as well. Of course, doing so does increase the risk of down-time as (some of) the remaining components can also fail unexpectedly. Obviously, condition monitoring could reduce that risk, but CBM has so far not been considered for multi-component systems with redundancy and economic dependence. To the best of our knowledge, we are the first to do so.

The remainder of this chapter is organized as follows. In Section 5.2, a short overview of relevant literature is given. Section 5.3 describes the system under investigation. Section 5.4 explains the dynamic programming model used to optimize the maintenance policy. In Sections 5.5 and 5.6, results of a numerical investigation are presented, describing the optimal policy and comparing it to classical policies. Section 5.7 concludes the chapter.

5.2. Literature review

Many authors have considered one or more of the following elements: CBM, redundancy, or economic dependence. Figure 5.1 gives a schematic overview of the types of systems considered. Although the figure is intended as an illustration, all articles are included that combine at least two of the elements. Interestingly, there is just a single study on CBM for systems with redundancy [7]. Moreover, this study does not consider economic dependence, as we do.

Examples of articles that consider corrective or preventive maintenance policies for systems with redundancy, in the form of k-out-of-N systems, are given by [9–14]. In fact, failure-based maintenance is considered in [9], while block replacement is studied in [10, 11]. Both [12] and [13] study purely corrective maintenance that is initiated as soon as the number of failed components exceeds some critical level, while in [14] both failed and degraded components are replaced at those times. An example of an article that is only focused on CBM is given by [8], in which a preventive replacement threshold is considered for a multi-component system. So far, CBM has only been considered for a system with redundancy in [7], where failure-based maintenance, time-based maintenance, and CBM with a preventive replacement threshold are compared for a k-out-of-N standby system without economic dependence.

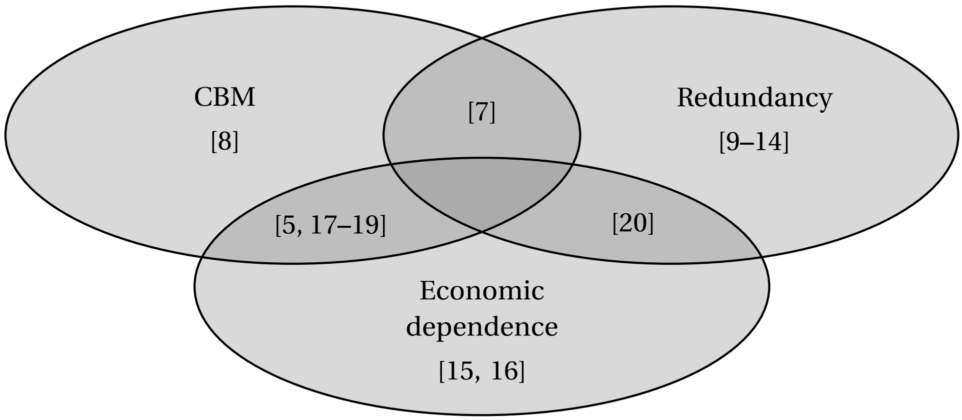

Figure 5.1. Schematic representation of literature on CBM, redundancy, and/or economic dependence.

In [15], a literature review on maintenance policies for multi-component systems with economic dependence is given. An example of a preventive replacement policy developed for a multi-component series system with economic dependence is that by [16], while CBM policies are considered for a two-component series system with economic dependence by [5, 17]. Furthermore, maintenance actions are grouped based on condition-monitoring information for a multi-component system with economic dependence in [18], and a CBM policy with risk-thresholds for preventive and opportunistic replacements is considered for a multi-component system with economic dependence in [19]. Finally, an article that considers systems with different structures (including k-out-of-N) and economic dependence is given by [20], in which periodic age-based maintenance actions are grouped.

To our knowledge, despite the widespread presence in industry of systems with both redundancy and economic dependence, no research has been performed on CBM for such systems. We will cover this gap by developing a dynamic programming model to optimize the CBM policy for a k-out-of-N system with economic dependence. An advantage of applying dynamic programming for this explorative study is that the structure of the optimal policy can be explored (numerically) and compared to commonly used strategies. Previously, dynamic programming has also been applied in, for example, [16, 17].

5.3. System description

5.3.1. Deterioration model

We consider a k-out-of-N system, defined here as a system consisting of N (non-identical) components that functions as long as at least k components function ($1 \leq k \leq N$). All functioning components are fully operational and subject to failure, that is, we consider an active redundant system. The deterioration process that we consider shows similarities to that of the two-component series system considered by [17]. Let S_N denote the set of all component labels, i.e., $S_N := \{1, 2, \ldots, N\}$. The condition, or deterioration level, of component i can be described by a random variable X_t^i, for $i \in S_N$ and $t \in \mathbb{N}$. After a replacement (say at time t_r), component i is assumed to be as-good-as-new, i.e., $X_{t_r}^i = 0$, while a deterioration level exceeding the fixed failure level L_i implies that component i has failed. Immediately following possible replacements, which are assumed to be instantaneous, component i is subject to a random increase in deterioration Y_t^i, for $i \in S_N$ and $t \in \mathbb{N}$. Hence, if component i is not replaced at the start of time unit t, its deterioration level at time unit $t+1$ is equal to $X_{t+1}^i = X_t^i + Y_t^i$. If, on the other

Nomenclature

δ_i	Binary variable indicating whether or not a replacement (preventive or corrective) should be performed on component i
μ_i	Deterioration parameter of component i
c_c^i	Cost of a corrective replacement on component i
c_p^i	Cost of a preventive replacement on component i, $c_p^i < c_c^i$
c_s	Fixed set-up cost for maintenance
C_t	Optimum cumulative cost from period t on
f_i	Pdf of the deterioration increments of component i
k	Number of components in the system that need to function for the system to function
L_i	Fixed failure level of component i
N	Number of components in the system
p	Penalty for a system failure
$R(x_1, x_2, \ldots, x_N)$	System reliability, given deterioration levels of $x_1, x_2, \ldots, x_N$ for components $1, 2, \ldots, N$, respectively
$R_i(x_i)$	Reliability of component i, given a deterioration level of x_i
S_N	Set of all component labels, $S_N := \{1, 2, \ldots, N\}$
X_t^i	Condition of component i at time t, $X_0^i = 0$
$\bar{X}_t^i$	Condition of component i after possible maintenance has been performed
Y_t^i	Increase in deterioration on component i during t, $X_{t+1}^i = \bar{X}_t^i + Y_t^i$

hand, component i is replaced, its deterioration level becomes $X_{t+1}^i = Y_t^i$. The probability density function of the i.i.d. non-negative deterioration increments Y_t^i of component i is denoted by f_i, for $i \in S_N$ and $t \in \mathbb{N}$, which can be either continuous or discrete.

5.3.2. Maintenance actions and corresponding costs

At the start of each time unit, the condition of each component is known. If there are less than k functioning components, the system has failed, and a penalty cost p for a system failure is incurred. Next, a decision is needed on what components to replace. If component i is replaced, either the preventive replacement cost c_p^i or the corrective replacement cost c_c^i is incurred, depending on whether the component has failed. Typically, a corrective replacement is more expensive than a preventive replacement (i.e., $c_c^i > c_p^i$), due to for example the unexpected nature of a failure or the fact that a preventive replacement could be an easy fix, while a failure can cause additional damage on other parts of the equipment, implying a more expensive replacement. This distinction is also made in, for example, [5, 7, 11, 16, 18, 21]. In addition, as soon as at least one of the components is replaced, the shared set-up cost c_s is incurred. In practice, this shared set-up cost can arise from e.g. ordering spare parts, traveling to the plant, or shutting down the system.

5.4. Maintenance cost model

The dynamic programming model that we will use for optimizing the maintenance policy is partially based on that in [17]. We introduce the function $R_i(\cdot)$, which equals 1 if component i is functioning (i.e., its deterioration level x_i is below the fixed failure level), and 0 if it has failed. Hence,

$$R_i(x_i) = \begin{cases} 1, & \text{if } x_i < L_i, \\ 0, & \text{if } x_i \geq L_i. \end{cases}$$

We further introduce the function $R(x_1, x_2, \ldots, x_N)$, which equals 1 if the system functions, and 0 if it has failed, given that components $1, 2, \ldots, N$ have deterioration levels $x_1, x_2, \ldots, x_N$, respectively. This function can be expressed as follows:

$$R(x_1, x_2, \ldots, x_N) = \begin{cases} 1, & \text{if } \sum_{i=1}^N R_i(x_i) \geq k, \\ 0, & \text{if } \sum_{i=1}^N R_i(x_i) < k. \end{cases}$$

Maintenance costs are determined at the start of a time unit t. For this, the binary variable δ_i is used to indicate whether any replacements (preventive or corrective) are performed on component i, while $\bar{x}_t^i$ denotes the condition of component i after possible replacements have been performed, i.e.,

$$\delta_i = \begin{cases} 1, & \text{if a replacement (preventive or corrective) is performed on component } i, \\ 0, & \text{if no replacement is performed on component } i, \end{cases}$$

$$\bar{x}_t^i = (1 - \delta_i) \cdot x_t^i.$$

The penalty for a system failure at the start of time unit t (before possible replacements are performed) is given by

$$p \cdot (1 - R(x_t^1, x_t^2, \ldots, x_t^N)). \tag{5.1}$$

If at least one component is replaced, the set-up cost c_s is incurred, which can be expressed as

$$\left(1-\prod_{i=1}^{N}(1-\delta_i)\right)\cdot c_s. \tag{5.2}$$

Furthermore, if component i is replaced while it is still functioning, the preventive replacement cost c_p^i is incurred. Hence, the total preventive replacement costs at time t for all components are equal to

$$\sum_{i=1}^{N}\delta_i\cdot R_i(x_t^i)\cdot c_p^i. \tag{5.3}$$

If component i has failed, it could be replaced correctively, meaning that the total corrective replacement costs (depending on whether or not the component is replaced) are given by

5

$$\sum_{i=1}^{N}\delta_i\cdot\left(1-R_i(x_t^i)\right)\cdot c_c^i. \tag{5.4}$$

In the dynamic programming model, let C_t denote the optimum cumulative cost from period t onwards. These costs consist of the direct costs at time t (consisting of the penalty cost for a system failure (5.1), the set-up costs (5.2), the preventive replacement costs (5.3), and the corrective replacement costs (5.4)) and the expected costs from period $t+1$ onwards. Hence, C_t can be expressed as follows:

$$\begin{aligned}C_t(x_t^1,x_t^2,\ldots,x_t^N) = \min_{\delta_1,\delta_2,\ldots,\delta_N}\Bigg\{&p\cdot(1-R(x_t^1,x_t^2,\ldots,x_t^N))\\&+\left(1-\prod_{i=1}^{N}(1-\delta_i)\right)\cdot c_s+\sum_{i=1}^{N}\delta_i\cdot R_i(x_t^i)\cdot c_p^i+\sum_{i=1}^{N}\delta_i\cdot\left(1-R_i(x_t^i)\right)\cdot c_c^i\\&+\int_0^\infty\cdots\int_0^\infty C_{t+1}\left(\bar{x}_t^1+y_1,\bar{x}_t^2+y_2,\ldots,\bar{x}_t^N+y_N\right)\cdot\prod_{i=1}^{N}\left[f_i(y_i)\right]\\&\mathrm{d}y_N\cdots\mathrm{d}y_1\Bigg\}.\end{aligned} \tag{5.5}$$

We are interested in minimizing the long-run average cost per time unit. This can be achieved by starting with (for example) $C_T(x_T^1,x_T^2,\ldots,x_T^N)=0$ for all x_T^i and sufficiently high T, and recursively applying (5.5) until both the average cost per time unit, $C_t(x_t^1, x_t^2, \ldots, x_t^N) - C_{t+1}(x_{t+1}^1, x_{t+1}^2, \ldots, x_{t+1}^N)$, and the corresponding maintenance policy have converged. Subsequently, the long-run average cost will become independent of the initial system state. Furthermore, the resulting optimal maintenance policy will only depend on the system state, and thus becomes independent of time.

In the next section, we numerically investigate the performances of our dynamic program, and compare these with previously considered maintenance policies.

5.5. Numerical investigation

5.5.1. Base case

We first consider a 1-out-of-2 system consisting of two identical components, with a discrete state space. In this setting with two components, the optimal policy can easily be presented in two-dimensional tables, showing the optimal maintenance action for each combination of deterioration states of both components. This allows easy interpretation of the optimal policy and comparison to other policies. By considering identical components, we can omit the superscripts denoting to which component a certain cost corresponds.

We next describe a base case scenario. We remark that the specific parameter settings are somewhat arbitrary and indeed selected as they allow us to highlight important characteristics of the optimal policy. Effects of parameter variations will later be explored in a sensitivity analysis. The preventive replacement costs are $c_p = 5$ per component, whilst the corrective replacement costs are $c_c = 11$ per component. In case at least one component is replaced, the set-up costs of $c_s = 4$ are incurred, and if the system fails the penalty costs $p = 300$ are incurred. The deterioration increments follow a Poisson distribution with mean $\mu = 0.7$, and a component fails once its deterioration level reaches the fixed failure limit $L = 5$. The choice of a Poisson distribution guarantees a discrete state space, which enables us to consider a finite number of states in the dynamic programming model.

The optimal maintenance policy and corresponding average cost obtained with the dynamic program converge relatively quickly for this case and other scenarios; at most 30 iterations are needed. Our computations are made using Python 3.4.3 on a computer with a 3.30 GHz quad core processor and 8.00 GB of RAM. For this example, it takes 0.05 seconds to find the optimal maintenance policy and corresponding average cost.

		X_2					
		0	1	2	3	4	5
X_1	0	00	00	00	00	01	00
	1	00	00	00	00	01	00
	2	00	00	00	00	01	01
	3	00	00	00	11	11	11
	4	10	10	10	11	11	11
	5	00	00	10	11	11	11

Figure 5.2. Optimal maintenance policy for a 1-out-of-2 system with identical components and $c_p = 5$, $c_c = 11$, $c_s = 4$, $p = 300$, $\mu = 0.7$, and $L = 5$.

5

Figure 5.2 shows the long-run optimal maintenance policy for this example, where for any combination of the deterioration levels X_1 and X_2 the optimal maintenance actions are represented. These maintenance actions are independent of time. Note that '00' corresponds to no replacements at all, '10' means that only component 1 is replaced, '01' means that only component 2 is replaced, and '11' corresponds to a complete system replacement. The long-run average costs corresponding to this optimal maintenance policy are equal to 3.42 per time unit. Note the non-monotonic behavior of the optimal policy, for instance by considering the '$X_1 = 0$' row. For deterioration levels of up to 3 for component 2, it is not replaced as component 1 is in good condition, and there is redundancy. However, if component 2 further deteriorates into state 4 (implying imminent failure), then it is preventively replaced in order to avoid higher corrective replacement costs. This motivation disappears if component 2 has already failed (state 5). In that case, the corrective replacement is postponed until component 1 has a deterioration level of at least 2, thus utilizing the redundancy. Of course, this behavior is driven by the (realistic) assumption that a corrective replacement is more costly than a preventive replacement. None of the more common maintenance policies allows this sort of behavior, as we will discuss after listing them.

Failure-based maintenance: This strategy can be interpreted as a simple inspection policy, where only failed components are replaced.

Multi-threshold maintenance without OM: Components are correctively replaced upon failure, and preventively as soon as their deterioration level reaches a certain threshold value ξ ($0 \le \xi \le L$), which needs to be optimized. No opportunistic maintenance (OM) is considered.

Multi-threshold maintenance with OM: This policy is similar to the policy above, but each component has an additional (opportunistic) threshold value. A component is replaced preventively if its deterioration level exceeds the preventive maintenance threshold ξ $(0 \leq \xi \leq L)$, independent of the states of other components, and also if its deterioration level exceeds the opportunistic maintenance (OM) threshold ζ $(0 \leq \zeta \leq \xi)$, provided that at least one other component requires a replacement. Both ξ and ζ need to be optimized.

Age-based maintenance: A component is replaced correctively upon failure or preventively as soon as it reaches age A. This replacement age A needs to be optimized.

Block replacement without intermediate CM: The complete system is replaced periodically with periodicity P, which needs to be optimized. If a component fails, it remains in the failed state until the next scheduled system replacement, so no corrective maintenance (CM) is performed between two system replacements.

Block replacement with intermediate CM: This policy is similar to the policy above, but in addition components are replaced correctively upon failure. This does not affect the moment of the next system replacement. Also here the periodicity P needs to be optimized.

The costs corresponding to the different maintenance policies are obtained by using simulation, for which we consider 10,500 periods (of which the first 500 time units are used as the warm-up time) and 100 replications. The resulting minimal long-run average costs per time unit along with the corresponding standard error are given in Table 5.1. Also the optimal parameters of the different maintenance policies are included. From Table 5.1, we observe that the standard errors are small, indicating significant cost differences between the various policies. Hence, we can conclude that failure-based maintenance performs worst, which makes sense as this is the least advanced policy. Also the preventive maintenance policies, such as age-based maintenance and block replacement, perform much worse than the optimal case, with costs that are about 50% higher. An interesting observation is the fact that replacing failed components between two consecutive system replacements in the block replacement strategy reduces costs substantially compared to leaving failed components in the failed state until the next planned replacement.

Furthermore, even compared to the most advanced CBM policies which use threshold values to schedule replacements, a large cost saving can apparently

Table 5.1. Minimal long-run average costs per time unit and optimal parameters for different maintenance policies.

Maintenance policy	Minimal average costs per time unit	Standard error	Optimal parameter(s)
Optimal policy	3.42	-	-
Failure-based maintenance	9.01 (+163.5%)	0.038	-
Multi-threshold maintenance without OM	3.77 (+ 10.2%)	0.005	$\xi = 3$
Multi-threshold maintenance with OM	3.72 (+ 8.8%)	0.005	$\xi = 3, \zeta = 2$
Age-based maintenance	5.24 (+ 53.2%)	0.011	$A = 4$
Block replacement without intermediate CM	5.31 (+ 55.3%)	0.013	$P = 3$
Block replacement with intermediate CM	5.01 (+ 46.5%)	0.018	$P = 4$

be obtained by allowing for a different structure of the policy. Figures 5.3 and 5.4 show the optimal multi-threshold maintenance policies without and with opportunistic maintenance, respectively. It turns out that these policies are too restrictive, by scheduling replacements more often than actually required, leading to costs that are about 10% higher than necessary.

X_1 \ X_2	0	1	2	3	4	5
0	00	00	00	01	01	01
1	00	00	00	01	01	01
2	00	00	00	01	01	01
3	10	10	10	11	11	11
4	10	10	10	11	11	11
5	10	10	10	11	11	11

Figure 5.3. Multi-threshold maintenance without OM, $\xi = 3$.

X_1 \ X_2	0	1	2	3	4	5
0	00	00	00	01	01	01
1	00	00	00	01	01	01
2	00	00	00	11	11	11
3	10	10	11	11	11	11
4	10	10	11	11	11	11
5	10	10	11	11	11	11

Figure 5.4. Multi-threshold maintenance with OM, $\xi = 3$ and $\zeta = 2$.

5.5.2. Sensitivity analysis with respect to corrective replacement cost

The non-monotonic behavior of the optimal policy, in that it could be optimal to preventively replace a component that is close to failure (in order to avoid a more expensive corrective replacement), and to postpone a corrective replacement on a failed component until the other component has deteriorated a bit more, is caused by the assumption that the costs of a corrective replacement c_c exceed those of a preventive replacement c_p. To illustrate this, Figure 5.5 shows the optimal policy for corrective replacement costs of $c_c = 9$ and $c_c = 7$ (whereas we considered $c_c = 11$ in the original example). Indeed, as the corrective replacement costs approach the preventive replacement costs $c_p = 5$, performing a preventive replacement to avoid a corrective replacement becomes less rewarding. In fact, for $c_c = 7$, this effect has disappeared completely. In line with this, preventive replacements are performed less often as the costs of a corrective replacement decrease.

Furthermore, to give an indication of the performances of our model for these lower corrective replacement costs, Table 5.2 shows the minimal long-run average costs per time unit for different maintenance policies and for different values of the corrective replacement costs c_c along with the percentage increase in costs compared to the optimal policy. For ease of comparison, we also included the original case with $c_c = 11$.

Although the non-monotonic behavior of the optimal policy may disappear for values of c_c close to c_p, we observe that the cost difference compared to for example the multi-threshold policies increases for decreasing values of c_c. Indeed,

X_1 \ X_2	0	1	2	3	4	5
0	00	00	00	00	00	00
1	00	00	00	00	01	00
2	00	00	00	00	01	01
3	00	00	00	11	11	11
4	00	10	10	11	11	11
5	00	00	10	11	11	11

(a) $c_c = 9$.

X_1 \ X_2	0	1	2	3	4	5
0	00	00	00	00	00	00
1	00	00	00	00	00	00
2	00	00	00	00	01	01
3	00	00	00	11	11	11
4	00	00	10	11	11	11
5	00	00	10	11	11	11

(b) $c_c = 7$.

Figure 5.5. Optimal policy for $c_p = 5$, $c_s = 4$, $p = 300$, $\mu = 0.7$, and $L = 5$ for different values of the corrective replacement costs c_c.

Table 5.2. Minimal costs per time unit for different maintenance policies for $c_p = 5$, $c_s = 4$, $p = 300$, $\mu = 0.7$, and $L = 5$ for different values of the corrective replacement costs c_c.

Maintenance policy	$c_c = 11$	$c_c = 9$	$c_c = 7$
Optimal policy	3.42	3.29	3.13
Failure-based maintenance	9.01 (+163.5%)	8.48 (+157.8%)	7.96 (+154.3%)
Multi-threshold maintenance without OM	3.77 (+ 10.2%)	3.72 (+ 13.1%)	3.60 (+ 15.0%)
Multi-threshold maintenance with OM	3.72 (+ 8.8%)	3.60 (+ 9.4%)	3.45 (+ 10.2%)
Age-based maintenance	5.24 (+ 53.2%)	5.03 (+ 52.9%)	4.80 (+ 53.4%)
Block replacement without intermediate CM	5.31 (+ 55.3%)	5.23 (+ 59.0%)	5.14 (+ 64.2%)
Block replacement with intermediate CM	5.01 (+ 46.5%)	4.86 (+ 47.7%)	4.70 (+ 50.2%)

the optimal policy has a completely different structure than considered in other maintenance policies. Hence, our model still significantly outperforms previously considered maintenance policies, indicating that it is of value for varying cost scenarios.

5.5.3. Sensitivity analysis with respect to maintenance set-up cost

Another cost parameter that significantly influences the resulting optimal policy is the set-up cost c_s. If chosen too low, no need for clustering exists, while if chosen too high, all replacements are clustered. In the latter case, the non-monotonic behavior of the optimal policy also disappears. To illustrate this, Figure 5.6 shows

X_1 \ X_2	0	1	2	3	4	5
0	00	00	00	00	00	00
1	00	00	00	00	01	00
2	00	00	00	00	11	11
3	00	00	00	11	11	11
4	00	10	11	11	11	11
5	00	00	11	11	11	11

(a) $c_s = 6$.

X_1 \ X_2	0	1	2	3	4	5
0	00	00	00	00	00	00
1	00	00	00	00	00	00
2	00	00	00	00	11	11
3	00	00	00	11	11	11
4	00	00	11	11	11	11
5	00	00	11	11	11	11

(b) $c_s = 8$.

Figure 5.6. Optimal policy for $c_p = 5$, $c_c = 11$, $p = 300$, $\mu = 0.7$, and $L = 5$ for different values of the set-up costs c_s.

Table 5.3. Minimal costs per time unit for different maintenance policies for $c_p = 5$, $c_c = 11$, $p = 300$, $\mu = 0.7$, and $L = 5$ for different values of the set-up costs c_s.

Maintenance policy	$c_s = 4$	$c_s = 6$	$c_s = 8$
Optimal policy	3.42	3.87	4.29
Failure-based maintenance	9.01 (+163.5%)	9.49 (+145.2%)	9.98 (+132.6%)
Multi-threshold maintenance without OM	3.77 (+ 10.2%)	4.52 (+ 16.8%)	5.15 (+ 20.0%)
Multi-threshold maintenance with OM	3.72 (+ 8.8%)	4.18 (+ 8.0%)	4.63 (+ 7.9%)
Age-based maintenance	5.24 (+ 53.2%)	6.01 (+ 55.3%)	6.76 (+ 57.6%)
Block replacement without intermediate CM	5.31 (+ 55.3%)	5.98 (+ 54.5%)	6.64 (+ 54.8%)
Block replacement with intermediate CM	5.01 (+ 46.5%)	5.57 (+ 43.9%)	6.13 (+ 42.9%)

the optimal policy for set-up costs of $c_s = 6$ and $c_s = 8$ (whereas we considered $c_s = 4$ in the original example). Indeed, as the set-up cost increases, replacements are clustered more often, while for $c_s = 8$, all replacements are clustered. In line with this, we observe that performing a preventive replacement to avoid a corrective replacement becomes less rewarding as the set-up cost increases.

Table 5.3 shows the minimal long-run average costs per time unit for different maintenance policies and for different values of the set-up cost c_s along with the percentage increase in costs compared to the optimal policy. Even though performing a preventive replacement to avoid a corrective replacement becomes less rewarding for increasing set-up costs, we observe that our dynamic program still significantly outperforms previously considered maintenance policies. Interestingly, the cost difference with the multi-threshold maintenance policy without opportunistic maintenance increases as the set-up costs increase, while the cost difference for the same policy with opportunistic maintenance decreases. Hence, clustering becomes more rewarding for higher set-up costs.

5.5.4. Sensitivity analysis with respect to number of components

The k-out-of-N system is a system in which k components need to function, but redundancy is incorporated by installing $N > k$ components. When deciding on the number of redundant components to install, one is trading off the maintenance costs (for additional components) and the system availability. First, to gain insights into this tradeoff, we will determine the long-run average cost per time unit for different values of k and N. We continue to consider identical

Table 5.4. Average cost per time unit for different values of k and N, for $c_p = 5$, $c_c = 11$, $c_s = 4$, $p = 300$, $\mu = 0.7$, and $L = 5$.

		N				
		1	2	3	4	5
k	1	3.54	3.42	3.42	3.42	3.42
	2	-	6.69	5.19	5.18	5.18
	3	-	-	9.66	6.89	6.89
	4	-	-	-	12.41	8.57

5

components as in the base case, and set $c_p = 5$, $c_c = 11$, $c_s = 4$, $p = 300$, $\mu = 0.7$, and $L = 5$. In all cases, 30 iterations of the dynamic program proved to be sufficient for the cost to converge. Computation time was not much affected by the value of k, but exponentially increasing in N. It took approximately 0.01, 0.05, 0.5, 7.6, and 180 seconds to analyze systems with $N = 1, 2, 3, 4$, and 5 components, respectively. Table 5.4 shows the resulting average costs. It follows that $N = 2$ is sufficient for a system with $k = 1$. Adding more components will not reduce costs, as the additional maintenance costs do not outweigh the reduced unavailability cost. Note that the cost does not increase in N either, even for relatively large values of N. Such large systems retain efficiency by keeping some components continuously in the failed state. For all k, it holds true that by installing one redundant component (i.e., $N = k + 1$), costs are reduced substantially compared with $N = k$. This reduction in costs increases to 30.9 percent as k increases to four. Furthermore, adding another redundant component offers little benefit for any value of k considered.

Second, we observed non-monotonic behavior of the optimal maintenance policy for the 1-out-of-2 system in Section 5.5.1. To investigate whether this result carries over to systems with more components, we now consider the 3-out-of-4 system in some more detail. We choose 3-out-of-4 as it is the largest system solvable in a few seconds. The corresponding optimal policy describes for each component when it should be replaced, depending on the complete system state. Such a policy can become rather complex for large N, but for the 3-out-of-4 system it is summarized in Table 5.5. We observe that this maintenance policy exhibits non-monotonic behavior as well. Compare for instance the actions in states $(4, 0, 0, 0)$ and $(5, 0, 0, 0)$, i.e., states where component 1 is in state 4 or 5 and components 2, 3, and 4 are as-good-as-new. If $X_1 = 4$, it is optimal to only replace component 1, while if $X_1 = 5$, no replacements are performed at all. Instead, the corrective replacement on component 1 is postponed until components 2, 3, and 4 have deteriorated somewhat more, thus exploiting the redundancy and

Table 5.5. Optimal replacement decision for component i, $i = 1,2,3,4$, for the 3-out-of-4 system.

State of component i	Optimal action for component i
$X_i = 0,1,2$	Do not replace component i
$X_i = 3$	Replace component i if $X_j > 2$ for some $j \neq i$ and/or $\sum_{j=1,j\neq i}^{N} X_j > 5$
$X_i = 4$	Replace component i
$X_i = 5$	Replace component i if $\sum_{j=1,j\neq i}^{N} X_j > 1$

economic dependence. This implies that the obtained insights do indeed carry over to systems with more than two components.

Third, to explore the performances of the dynamic program for systems with more than two components, Table 5.6 shows a cost comparison for the 3-out-of-4 system with other maintenance policies. From comparing Table 5.6 for a 3-out-of-4 system with Table 5.1 for a 1-out-of-2 system, we observe first that the cost difference with the multi-threshold maintenance policy decreases to 3.8 percent for $N = 4$. This implies that a multi-threshold policy might be suitable for systems with many components, although it could still be rewarding to postpone corrective replacements and thus utilize the redundancy. Second, we observe that the optimal multi-threshold maintenance policy does not include opportunistic replacements. Indeed, as the number of components increases, an opportunistic replacement threshold of $\zeta = 2$ (as for the 1-out-of-2 system) would imply

Table 5.6. Minimal long-run average costs per time unit and optimal parameters for different maintenance policies for the 3-out-of-4 system.

Maintenance policy	Minimal costs per time unit	Standard error	Optimal parameter(s)
Optimal policy	6.89	-	-
Failure-based maintenance	33.15 (+381.1%)	0.063	-
Multi-threshold maintenance without OM	7.15 (+ 3.8%)	0.010	$\xi = 3$
Multi-threshold maintenance with OM	7.15 (+ 3.8%)	0.010	$\xi = 3$, $\zeta = 3$
Age-based maintenance	11.16 (+ 62.0%)	0.018	$A = 3$
Block replacement without intermediate CM	10.70 (+ 55.3%)	0.026	$P = 3$
Block replacement with intermediate CM	10.06 (+ 46.0%)	0.020	$P = 3$

much more frequent replacements in case of a 3-out-of-4 system, which does not outweigh the savings in set-up costs.

5.5.5. Sensitivity analysis with respect to monitoring accuracy

To gain somewhat more insight into the behavior of the optimal policy, we now consider an example with better condition information. To achieve that, we return our attention to the 1-out-of-2 system and assume a fixed failure level of $L = 10$. In this way, each component can be in 11 different states rather than 6 as in the previous example. Furthermore, we consider preventive replacement costs $c_p = 3$, corrective replacement costs $c_c = 8$, set-up costs $c_s = 5$, and a penalty cost $p = 500$ for two identical components with a mean increase in deterioration per time unit of $\mu = 2.5$. It takes 0.13 seconds to find the optimal policy in this example. Figure 5.7 presents the resulting optimal maintenance policy. The costs corresponding to this policy are equal to 4.85 per time unit. Similarly to the previous case, it turns out that it is optimal to perform a preventive replacement in case of an imminent failure, while a corrective replacement is postponed until the other component has deteriorated to a certain extent. Furthermore, we observe an almost linear diagonal structure in the optimal policy, indicating that the multi-threshold policies are too restricted as they only allow (graphically) for horizontally

X_1 \ X_2	0	1	2	3	4	5	6	7	8	9	10
0	00	00	00	00	00	00	00	00	00	01	00
1	00	00	00	00	00	00	00	00	01	01	00
2	00	00	00	00	00	00	00	00	01	01	00
3	00	00	00	00	00	00	00	01	01	01	01
4	00	00	00	00	00	00	11	11	11	11	11
5	00	00	00	00	00	11	11	11	11	11	11
6	00	00	00	00	11	11	11	11	11	11	11
7	00	00	00	10	11	11	11	11	11	11	11
8	00	10	10	10	11	11	11	11	11	11	11
9	10	10	10	10	11	11	11	11	11	11	11
10	00	00	00	10	11	11	11	11	11	11	11

Figure 5.7. Optimal maintenance policy for a 1-out-of-2 system with identical components and $c_p = 3$, $c_c = 8$, $c_s = 5$, $p = 500$, $\mu = 2.5$, and $L = 10$.

Table 5.7. Minimal costs and optimal parameters for different maintenance policies.

Maintenance policy	Minimal costs per time unit	Standard error	Optimal parameter(s)
Optimal policy	4.85	-	-
Failure-based maintenance	30.04 (+519.4%)	0.098	-
Multi-threshold maintenance without OM	5.43 (+ 12.0%)	0.013	$\xi = 6$
Multi-threshold maintenance with OM	5.01 (+ 3.3%)	0.013	$\xi = 6, \zeta = 3$
Age-based maintenance	6.91 (+ 42.5%)	0.053	$A = 2$
Block replacement without intermediate CM	5.91 (+ 21.9%)	0.011	$P = 2$
Block replacement with intermediate CM	5.90 (+ 21.6%)	0.012	$P = 2$

linear structures. The structure of our optimal policy shows similarities to that obtained for a two-component series system by [17], although in this case with redundancy, we observe that preventive replacements are scheduled less often, and that corrective replacements are sometimes postponed. Also our structure is not completely linear, so allowing for a more flexible structure could reduce costs.

Table 5.7 shows the different (minimal) costs for the various possible maintenance strategies along with the standard errors and, if applicable, the corresponding optimal parameter(s). Again, failure-based maintenance is the most expensive maintenance strategy, while also the preventive maintenance policies age-based maintenance and block replacement perform much worse than CBM. The difference between the multi-threshold policy with opportunistic maintenance and the optimal maintenance policy is equal to 3.3% in this case (versus 12.0% if no opportunistic maintenance is included).

5.6. Continuous degradation

In the numerical investigation so far, we assumed that the deterioration increments of all components follow a Poisson distribution. This guarantees a discrete state space, and hence a finite number of states in the dynamic program. This allowed easy representation and interpretation of the results. Also, there are many practical cases where condition monitoring is not perfect and therefore not done on a continuous scale. We did already investigate the effect of having better condition information (i.e., of having a larger state space), and in this section, we explore whether the resulting insights carry over to the case of continuous degradation. Since the gamma process is the most appropriate process to model continuous degradation over time [22], we consider exponentially distributed deterioration increments. We consider the same costs as in the first example in the numeric investigation, i.e., $c_p = 5$, $c_c = 11$, $c_s = 4$, and $p = 300$, set the failure level equal to $L = 2$, and scale the deterioration parameter to $\alpha = \frac{1}{0.7} \cdot \frac{5}{2} \approx 3.6$. Due to the large number of nested integrals required in the dynamic program, we approximate the integrals numerically by applying the extended midpoint rule [23], dividing the area into a 45x45 grid. Furthermore, we assume that the maximum deterioration level of a component never exceeds $1.5 \cdot L$, because the extended midpoint rule cannot cope with infinite integral bounds. Initial testing reveals that this does not affect the results. The maintenance policy that we obtain by applying the dynamic program is an approximation of the optimal policy, since we still discretize the state space for numerical reasons. The resulting policy is shown in Figure 5.8. The computing time was approximately equal to 3.8 seconds for 20 iterations. As for the discrete state space, the property that a preventive

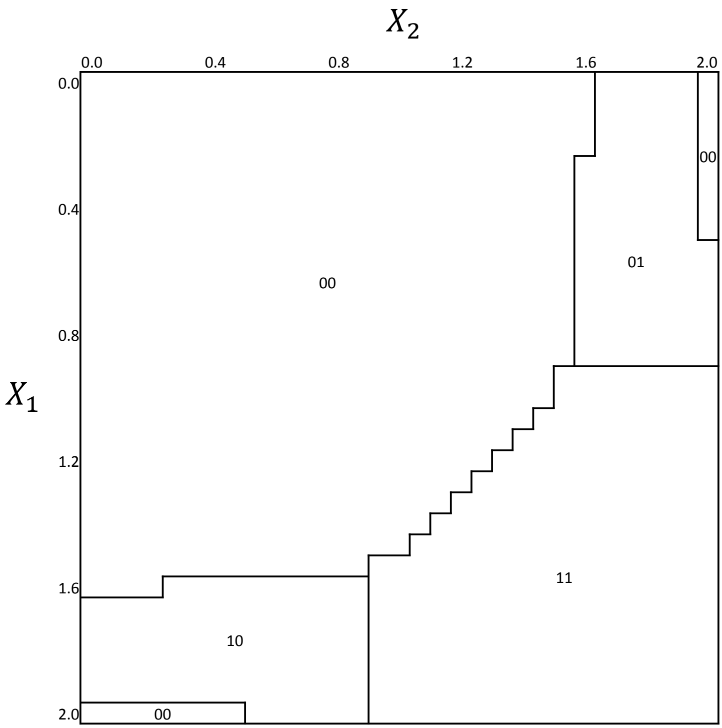

Figure 5.8. Approximation of the optimal maintenance policy for a continuously deteriorating 1-out-of-2 system with identical components and $c_p = 5$, $c_c = 11$, $c_s = 4$, $p = 300$, $\alpha \approx 3.6$, and $L = 2$.

Table 5.8. Minimal costs and optimal parameters for different maintenance policies.

Maintenance policy	Minimal costs per time unit	Standard Error	Optimal parameter(s)
Approximated optimal policy	3.02	0.006	-
Failure-based maintenance	8.23 (+172.7%)	0.040	-
Multi-threshold maintenance without OM	3.24 (+ 7.4%)	0.008	$\xi = 1.43$
Multi-threshold maintenance with OM	3.06 (+ 1.5%)	0.007	$\xi = 1.45$, $\zeta = 0.85$
Age-based maintenance	4.23 (+ 40.3%)	0.010	$A = 5$
Block replacement without intermediate CM	4.20 (+ 39.3%)	0.015	$P = 4$
Block replacement with intermediate CM	4.06 (+ 34.6%)	0.010	$P = 4$

replacement is performed to avoid a (more expensive) corrective replacement is clearly visible.

Table 5.8 shows the minimal costs obtained with the other maintenance policies along with the standard errors and the corresponding optimal parameter(s). Since we discretized the state space in the dynamic program, the resulting average cost is not accurate. For a fair comparison with the other maintenance policies, we therefore use simulation to find the average cost corresponding to the (approximated) optimal policy. Note that although we approximated the optimal dynamic program, it still outperforms the other maintenance policies.

We can conclude that our policy is applicable to systems with various types of degradation, but also that performing a preventive replacement to avoid a corrective replacement can be profitable for a wide class of systems with redundancy and economic dependence, where a corrective replacement is more costly than a preventive replacement.

5.7. Conclusion

We are the first to consider CBM for a multi-component system with both (active) redundancy, through a k-out-of-N structure, and economic dependence. We developed a dynamic programming model to find the optimal maintenance strategy for such systems. For various systems, either with a discrete or a continuous state space, we observed some interesting characteristics of the optimal policies. First, we observed a rather different structure of the optimal policy than those possible by using more common CBM policies with degradation threshold values corresponding to different maintenance actions. In particular, we observed non-monotonic behavior in the optimal policy in that for certain states of the system, it could be optimal to preventively replace a component that is close to failure, while a corrective replacement on that same component would be postponed until the other component has deteriorated a bit more. Furthermore, a numerical comparative cost study revealed that the optimal maintenance policy considerably outperforms all previously considered policies (failure-based, age-based, block replacement, and more restricted CBM policies).

In practice, companies generally replace components immediately upon failure, despite the redundancy in their equipment. The insights we obtained in the structure of the optimal policy can, for example, help companies decide on when to postpone corrective replacements to save costs, thus utilizing their redundancy.

The structure of the optimal policy is, however, rather complex, especially when considering more than two components. For that reason, future research should be performed on investigating whether the (numerically) observed optimal policy structure can be (approximately) expressed using policy parameters. This would facilitate a much faster optimization. Furthermore, in the current model the system is inspected every time unit, after which possible replacements are performed. Including aperiodic inspection could be profitable [5], and is certainly possible in this setting. Although in practice minimizing the total maintenance costs might be one of the main goals, maximizing availability is also deemed important. In the current model, availability is guaranteed by imposing a large penalty for system unavailability, but no information is available on the actual availability. This could be an interesting field of future research as well.

References

[1] R. Moghaddass, M. Zuo, M. Pandey, Optimal design and maintenance of a repairable multi-state system with standby components, Journal of Statistical Planning and Inference 142 (8) (2012) 2409–2420. doi:10.1016/j.jspi.2012.02.045.

[2] W. Kuo, M. Zuo, Optimal reliability modeling: principles and applications, John Wiley & Sons, 2003.

[3] P. Boland, F. Proschan, The reliability of K out of N systems, The Annals of Probability 11 (3) (1983) 760–764. doi:10.1214/aop/1176993520.

[4] K. Misra, New trends in system reliability evaluation, Elsevier Science, 2012.

[5] B. Castanier, A. Grall, C. Bérenguer, A condition-based maintenance policy with non-periodic inspections for a two-unit series system, Reliability Engineering & System Safety 87 (1) (2005) 109–120. doi:10.1016/j.ress.2004.04.013.

[6] H. Wang, A survey of maintenance policies of deteriorating systems, European Journal of Operational Research 139 (3) (2002) 469–489. doi:10.1016/S0377-2217(01)00197-7.

[7] L. Lu, J. Jiang, Analysis of on-line maintenance strategies for k-out-of-n standby safety systems, Reliability Engineering & System Safety 92 (2) (2007) 144–155. doi:10.1016/j.ress.2005.11.012.

[8] L. Wang, J. Chu, W. Mao, A condition-based replacement and spare provisioning policy for deteriorating systems with uncertain deterioration to failure, European Journal of Operational Research 194 (1) (2009) 184–205. doi:10.1016/j.ejor.2007.12.012.

[9] E. Bjarnason, S. Taghipour, D. Banjevic, Joint optimal inspection and inventory for a k-out-of-n system, Reliability Engineering & System Safety 131 (2014) 203–215. doi:10.1016/j.ress.2014.06.018.

[10] K. de Smidt-Destombes, M. van der Heijden, A. van Harten, Availability of k-out-of-N systems under block replacement sharing limited spares and repair capacity, International Journal of Production Economics 107 (2) (2007) 404–421. doi:10.1016/j.ijpe.2006.08.013.

[11] M. Park, H. Pham, A generalized block replacement policy for a k-out-of-n system with respect to threshold number of failed components and risk costs, IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans 42 (2) (2012) 453–463. doi:10.1109/TSMCA.2011.2162499.

[12] K. de Smidt-Destombes, M. van der Heijden, A. van Harten, Joint optimisation of spare part inventory, maintenance frequency and repair capacity for k-out-of-N systems, International Journal of Production Economics 118 (1) (2009) 260–268. doi:10.1016/j.ijpe.2008.08.058.

[13] K. de Smidt-Destombes, M. van der Heijden, A. van Harten, On the availability of a k-out-of-N system given limited spares and repair capacity under a condition based maintenance strategy, Reliability Engineering & System Safety 83 (3) (2004) 287–300. doi:10.1016/j.ress.2003.10.004.

[14] K. de Smidt-Destombes, M. van der Heijden, A. van Harten, On the interaction between maintenance, spare part inventories and repair capacity for a k-out-of-N system with wear-out, European Journal of Operational Research 174 (1) (2006) 182–200. doi:10.1016/j.ejor.2005.02.043.

[15] R. Dekker, R. Wildeman, F. van der Duyn Schouten, A review of multi-component maintenance models with economic dependence, Mathematical Methods of Operations Research 45 (3) (1997) 411–435. doi:10.1007/BF01194788.

[16] X. Zhou, L. Xi, J. Lee, Opportunistic preventive maintenance scheduling for a multi-unit series system based on dynamic programming, International Journal of Production Economics 118 (2) (2009) 361–366. doi:10.1016/j.ijpe.2008.09.012.

[17] F. Barbera, H. Schneider, E. Watson, A condition based maintenance model for a two-unit series system, European Journal of Operational Research 116 (2) (1999) 281–290. doi:10.1016/S0377-2217(98)00189-1.

[18] K. Bouvard, S. Artus, C. Bérenguer, V. Cocquempot, Condition-based dynamic maintenance operations planning & grouping. Application to commercial heavy vehicles, Reliability Engineering & System Safety 96 (6) (2011) 601–610. doi:10.1016/j.ress.2010.11.009.

[19] Z. Tian, H. Liao, Condition based maintenance optimization for multi-component systems using proportional hazards model, Reliability Engineering & System Safety 96 (5) (2011) 581–589. doi:10.1016/j.ress.2010.12.023.

[20] H. Vu, P. Do, A. Barros, C. Bérenguer, Maintenance grouping strategy for multi-component systems with dynamic contexts, Reliability Engineering & System Safety 132 (2014) 233–249. doi:10.1016/j.ress.2014.08.002.

[21] B. de Jonge, W. Klingenberg, R. Teunter, T. Tinga, Optimum maintenance strategy under uncertainty in the lifetime distribution, Reliability Engineering & System Safety 133 (0) (2015) 59 – 67. doi:10.1016/j.ress.2014.09.013.

[22] J. van Noortwijk, A survey of the application of gamma processes in maintenance, Reliability Engineering & System Safety 94 (1) (2009) 2–21. doi:10.1016/j.ress.2007.03.019.

[23] W. Press, S. Teukolsky, W. Vetterling, B. Flannery, Numerical recipes in FORTRAN: the art of scientific computing, Cambridge University Press, 1992.