Proceedings of the 35th European Safety and Reliability & the 33rd Society for Risk Analysis Europe Conference
Edited by Eirik Bjorheim Abrahamsen, Terje Aven, Frederic Bouder, Roger Flage, Marja Ylönen

doi: 10.3850/978-981-94-3281-3_ESREL-SRA-E2025-P9997-cd

Balancing downtime and maintenance costs for multi-component systems with economic and structural dependencies

Jussi Leppinen

Department of Mathematics and Systems Analysis, Aalto University, Finland. E-mail: jussi.leppinen@aalto.fi

Ahti Salo

Department of Mathematics and Systems Analysis, Aalto University, Finland. E-mail: ahti.salo@aalto.fi

In the maintenance of multi-component systems, failures can cause downtime costs because immediate maintenance may not be possible. This leads to the need to balance these downtime costs with periodic maintenance costs while ensuring that the system operates with the requisite reliability, given its structural and economic dependencies. Specifically, we formulate a Markov Decision Process model with two novel features. First, the system can have any reliability structure, including structures in which it operates even if some components have failed. Toward this end, the corresponding reliability function is obtained from a binary decision diagram, which significantly increases the range of practical maintenance problems that can be addressed. Second, expected downtime costs are derived from the reliability function for every possible state transition. A modified policy-iteration algorithm is then used to determine the optimal policy to minimise the discounted total costs that consist of maintenance and downtime costs. By varying the unit cost of system downtime, a range of Pareto-optimal policies is derived. Furthermore, we also derive suggestions for changing structural dependencies so that the resulting maintenance costs would be lower.

Keywords: Multi-component system, Markov Decision Process, Downtime, Economic and structural dependencies

1. Introduction

Maintenance costs often constitute a significant portion of the operating expenses of multi-component systems. These costs can be reduced by optimizing the timing of component replacements, as shown in recent studies on determining optimal maintenance policies for multi-component systems (de Jonge and Scarf, 2020; Olde Keizer et al., 2017). However, multiple objectives may need to be considered when choosing policies. Among these, the minimization of maintenance costs is typically prioritized, while other objectives, such as maximization of reliability and reduction of downtime, are treated as constraints. Even in simplified single-objective formulations, the components of the system often have economic, structural, stochastic, and resource dependencies that influence the optimality of the policy (Olde Keizer et al., 2017). Previously, we applied a Markov Decision Process (MDP) to minimize discounted total maintenance costs under economic and structural dependencies, with reliability constraints and periodic maintenance instances (Leppinen et al., 2025).

A system is *available* when it operates as intended. Availability refers to the system's operating probability. Many maintenance models consider system availability (e.g., Geng et al., 2015; Safaei et al., 2020). In contrast, during *downtime*, the system is not operational due to failures or maintenance. Structural dependencies can prolong downtime (Geng et al., 2015). Downtime can be either unplanned due to sudden failures or planned due to scheduled maintenance. Unplanned downtime is generally more costly than planned downtime (e.g., Nguyen et al., 2015) because planned maintenance, such as annual maintenance, can be scheduled during low-demand periods.

System downtime is often modeled using a unit cost of downtime that cumulates costs when the system is not operational. However, when this unit cost is included, total downtime or availability is rarely reported, even though long-term downtime costs can be uncertain, depending on factors such as system utilization. Treating maintenance costs and downtime as separate objectives establishes a trade-off of objectives. This can facilitate a more informed selection among policies calculated with different downtime costs. Thus, we extend Leppinen et al. (2025) by comparing maintenance policies based on total maintenance costs and downtime.

Discrete-time MDP models have addressed system downtime using a constant failure cost (Andersen et al., 2022; Olde Keizer et al., 2018) or a unit cost of downtime (Xu et al., 2021; Zheng et al., 2023). A constant failure cost is a conservative estimate if it is paid in full regardless of the realized failure time between maintenance instances (Zheng et al., 2023). Using a unit cost of downtime requires calculating the expected failure time between decisions. This has been calculated for systems where k-out-of-n components must function (Xu et al., 2021; Zheng et al., 2023). In this paper, we derive downtime costs with the help of binary decision diagrams (BDDs), which allow the modeling of complex reliability structures in discrete-time MDPs and thus broaden the scope of practical maintenance applications.

Recent studies on MDP-based maintenance scheduling models for multi-component systems assume fixed structural dependencies and focus on the analysis of optimal policy structure or solution times (Olde Keizer et al., 2018; Xu et al., 2021; Andersen et al., 2022; Zheng et al., 2023). Few studies compare optimal policies with varying structural dependencies (Nguyen et al., 2022). To broaden the scope of MDP models, we compare different maintenance policies across structural updates, addressing both technical and performance-related dependencies. This approach aims to guide investments by identifying the most effective structural changes to reduce downtime or maintenance costs, particularly when resources are limited.

Section 2 describes the MDP model with extensions. Section 3 demonstrates the use of the model. Section 4 concludes the paper and suggests future research avenues.

2. Maintenance model

2.1. *System states, dependencies, costs*

The multi-component system has m components denoted by the set $M = \{1, ..., m\}$. The components can only be replaced at predetermined periodic *maintenance instances*, denoted by $t^k = k\Delta$, $k \in \mathbb{N}$, where the constant Δ is *the length of a maintenance interval*. We denote *replacement decisions* at instance t^k with a binary vector $z^k \in \{0, 1\}^{1\times m}$ where $z_i^k = 1$ if component i is replaced. The superscript k is omitted in the notation if the index of instance t^k does not need to be specified. A finite number of replacement decisions (at most 2^m) is indicated by set Z: $z \in Z$.

We assume that the ages of the components are known. The age of the component i at time t^k is a_i^k, and the combination of components ages with $a^k = (a_i^k)_{i\in M}$. The components are always replaced with new ones with zero age. Consequently, the ages of the components are discrete multiples of the maintenance interval.

In contrast to Leppinen et al. (2025), we allow multiple components to fail between maintenance instances. The failure probabilities of the components depend on their age, according to some continuous distributions. We assume that component failures occur independently and are detected instantly. Each component i has a binary *failure status* f_i^k at maintenance instance t^k, where $f_i^k = 1$ ($f_i^k = 0$) indicates the component has failed (is operational). The *failure time* t_i^f of component i satisfies $t_i^f \leq t^k$ if $f_i^k = 1$. A *failure state* of the system at instance t^k is a binary vector $f^k = (f_i^k)_{i\in M}$ that combines the failure statuses of the components. There are 2^m failure states. A failure state where exactly the components $X \subseteq M$ are failed is denoted by f_X.

The *failure status* of the system at instance t^k is denoted by the binary variable $f_s^k \in \{0, 1\}$. The system operates when $f_s^k = 0$. This value depends on the reliability structure and the failure state of the system: $f_s^k = BR(f^k)$ where $BR(\cdot)$ is a binary function that describes the reliability structure of the system. We model the function using a binary decision diagram (BDD) (for details, see Section 2.3). This is an extension of the model of Leppinen et al. (2025) where only a series reliability structure is considered. The component i is *critical* if its failure causes system failure when all other components operate.

At instance t^k, the state of the component i is a combination of its age and failure status: $s_i^k = (a_i^k, f_i^k)$, which contains relevant information for maintenance decisions. Thus, the state of the system is $s^k = (a^k, f^k)$, which is denoted for brevity as $s = (a, f)$. The state space of possible states is denoted by S.

The structural and economic dependencies are modeled using a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where $\mathcal{V} = \{0\} \cup \mathcal{D} \cup \mathcal{R}$ represents the nodes, and $\mathcal{A}$ represents directed arcs (x, y) with associated costs (Leppinen et al., 2025). The root node 0 represents a fixed set-up cost c_0 whenever maintenance is performed. The nodes in $\mathcal{D}$ represent the disassembly actions, while the nodes in $\mathcal{R}$ correspond to the replacement actions of components, with $|\mathcal{R}| = m$. Components may require disassembly of other components prior to replacement, and arcs define these dependencies. An arc (x, y) with weight c_{xy} indicates that the action corresponding to node y can be performed at cost c_{xy} if x is also performed. If x is the root node, y can be done independently. Thus, maintenance costs depend on both structural dependencies and the choice of simultaneous actions.

We assume that corrective replacement is more expensive than preventive replacement. First, a *system failure surplus* $r_s \geq 0$ is paid if the system fails before replacements. The cost r_s represents a corrective surplus for the setup costs. Second, a *component-specific corrective replacement surplus* $r_i \geq 0$ is charged for every failed component i that is replaced. Third, *the unit cost of downtime* $c_d \geq 0$ determines the downtime cost if the system is failed during the maintenance interval such that the downtime cost is scaled from c_d according to the actual *downtime*, the duration between the system failure and the next available maintenance instance. Section 2.4 describes the calculation of the expected system downtime.

2.2. *System dynamics*

The failure time t_i^f of the component i is assumed to follow a cumulative distribution function $F_i(t)$ such that the component fails before age $a_i \geq 0$ with probability $F_i(a_i)$. If component i of age a_i^k is operational at maintenance instance t^k, it stays operational until the next instance with conditional probability

$$R_i(a_i^k) := \frac{1 - F_i(a_i^k + \Delta)}{1 - F_i(a_i^k)}. \tag{1}$$

This is the *reliability of component* i at age a_i^k. We assume that the reliability of every component decreases with age.

We assume that the components age regardless of whether they have failed or not. When a component i has state (a_i^k, f_i^k) at a maintenance instance t^k, its state changes during the maintenance interval (t^k, t^{k+1}) to state (a_i^{k+1}, f_i^{k+1}) as follows:

- If the component is replaced, it transitions to state $(\Delta, 0)$ with probability $R_i(0)$ and to state $(\Delta, 1)$ with probability $1 - R_i(0)$.
- If the component is operational and is not replaced, it transitions to $(a_i^k + \Delta, 0)$ with probability $R_i(a_i^k)$ and to $(a_i^k + \Delta, 1)$ with probability $1 - R_i(a_i^k)$.
- If the component has failed and is not replaced, it transitions to $(a_i^k + \Delta, 1)$ with probability 1.

Thus, the state transition depends on which components are replaced at t^k and which fail during the interval (t^k, t^{k+1}).

After replacements, the system state $s = (a, f)$ is given by the post-decision state $s^z = (a^z, f^z)$, where $a_i^z = 0$ and $f_i^z = 0$ if component i is replaced, and $a_i^z = a_i$ and $f_i^z = f_i$ if component i is not replaced. Let $M_z \subseteq M$ denote components that operate in the post-decision state ($f_i^z = 0$). From state s^z the system transitions to state $\tilde{s} = (\tilde{a}, \tilde{f})$. During the transition, every component $i \in M$ ages deterministically: $\tilde{a}_i = a_i^z + \Delta$. However, state $\tilde{s}$ can have $2^{|M_z|}$ different failure states, as each operational component can fail. The component $i \in M_z$ fails with probability $1 - R_i(a_i^z)$.

If the state $\tilde{s} = (\tilde{a}, \tilde{f})$ is such that $\tilde{a}_i = a_i^z + \Delta$ for every $i \in M$ and $\tilde{f}_i = 1$ for every $i \notin M_z$, this state $\tilde{s}$ is one of the *reachable* states of s^z, denoted by $RE(s^z)$. Each state s^z has $2^{|M_z|}$ reachable states. The reachable states are mutually exclusive and collectively exhaustive. The transition probability from state s^z to state $\tilde{s} \in RE(s^z)$ is

$$p(\tilde{s}|s^z) = \prod_{i \in M_z} \left[(1-\tilde{f}_i)R_i(a_i^z) + \tilde{f}_i(1-R_i(a_i^z))\right]. \quad (2)$$

The probability of failure of a component depends on its age, which is zero if the component is replaced. Components that do not operate after the replacement decision z do not impact the probabilities of state transition because they evolve deterministically: for all $i \notin M_z$ $\tilde{a}_i = a_i^z + \Delta$ and $\tilde{f}_i = f_i^z = 1$.

The reliability of the system in state s^z is the sum of all transition probabilities from state s^z to those reachable states $\tilde{s} \in RE(s^z)$ in which the system is operational

$$R(s^z) = \sum_{\tilde{s} \in RE(s^z),\, BR(\tilde{f})=0} p(\tilde{s}|s^z). \quad (3)$$

The system failure can be constrained by setting a reliability threshold $\rho \in (0,1)$ that limits the failure probability of the system

$$R(s^z) \geq \rho \quad \forall s^z \in S. \quad (4)$$

As components age, their reliability decreases. Thus, for every component-specific reliability threshold ρ_i there is a maximum age $a_i^{\max}$ such that $R_i(a_i) < \rho_i$ when $a_i \geq a_i^{\max}$. Due to the reliability threshold (4), the reliability of critical components i must be $R_i(a_i) \geq \rho$. This determines the maximum ages for critical components. In contrast, non-critical components j can age and fail without causing the system to fail. However, we also set a maximum age for these components as they are required to operate with a high enough reliability. This also serves to uphold the salvage value of the system, for example. Furthermore, we assume that failed components must be replaced before they exceed their maximum age, even when the system is operational. Thus, the age of every component is bounded, and the system has 2^m failure states, resulting in a finite state space S.

The objective is to develop a *feasible* replacement policy that minimizes expected overall costs in the long term. At any given maintenance instance, replacement decision z is feasible if

(i) post-decision state s^z fulfills the reliability threshold (4),
(ii) the system operates in the post-decision state: $BR(f^z) = 0$,
(iii) components are replaced before they reach their maximum age,
(iv) structural dependencies are satisfied (as in Leppinen et al., 2025).

The set of feasible maintenance actions for the state s is denoted by $Z_s \subset Z$.

2.3. *Binary decision diagram*

Binary decision diagrams (BDD) model failure states of the system as paths between components from the top of the diagram to the bottom. An operational (failed) component corresponds to the path to the left (right) of the component. If a path ends in branch "1", the system has failed; and, conversely, if it ends in branch "0", it is still operational. We illustrate a BDD with an example in Figure 1. The system has four components: engine 1 (E1), engine 2 (E2), chassis (C), and wheels (W). The reliability structure is shown on the left of the figure, and the corresponding BDD is on the right. For example, the system operates if E1 is failed and other components operate. This failure state corresponds to the path highlighted in orange in Figure 1. We discuss the example in more detail in Section 3.2.

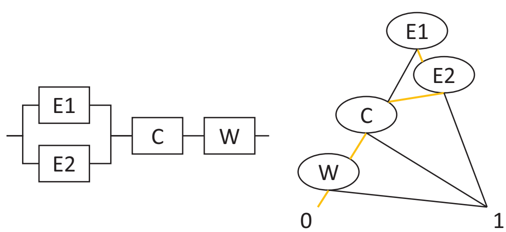

Fig. 1. Reliability structure and BDD of the example system.

With BDD, the reliability of the system $R(s^z)$

in the post-decision state s^z is a function of the reliabilities of the components $R(s^z) = RS(R_1(a_1^z), ..., R_m(a_m^z))$ by combining the different paths, i.e., failure states, from the top of the diagram to the bottom "0". In a single path, the component i operates with probability $R_i(a_i^z)$ or fails with probability $1-R_i(a_i^z)$. The probabilities of the events that correspond to the failure statuses along each path are multiplied, and the path probabilities are added together. This is an alternative formulation for equation (3). For example, for the system in Figure 1 we get

$$R(s^z) = R_{\text{E1}}(a_{\text{E1}}^z)R_{\text{C}}(a_{\text{C}}^z)R_{\text{W}}(a_{\text{W}}^z) + (1 - R_{\text{E1}}(a_{\text{E1}}^z))R_{\text{E2}}(a_{\text{E2}}^z)R_{\text{C}}(a_{\text{C}}^z)R_{\text{W}}(a_{\text{W}}^z). \tag{5}$$

Furthermore, a closed-form representation of the function $BR(\cdot)$, which maps the failure state to the failure status of the system, is obtained by replacing $R(s^z)$ with $1-f_s$, replacing every $R_i(a_i^z)$ with $1-f_i$, and solving the equation with respect to f_s.

2.4. *Expected system downtime*

For component i whose nonnegative failure time t_i^f follows a cumulative distribution function (CDF) $F_i(x)$, the expected failure time is

$$E(t_i^f) = \int_0^\infty (1 - F_i(x))\, dx. \tag{6}$$

Moreover, the failure times of components can be conditioned on their failure statuses at consecutive maintenance instances. If component i fails between maintenance instances t^k and t^{k+1}, we can scale the CDF between 0 and 1 on the interval (t^k, t^{k+1}). In this case, the expected failure time is

$$E(t_i^f) := \int_0^\Delta (1 - F_{a_i^k}^\Delta(x))\, dx. \tag{7}$$

where $F_{a_i^k}^\Delta(x) := \frac{F_i(a_i^k+x)-F_i(a_i^k)}{F_i(a_i^k+\Delta)-F_i(a_i^k)}$ is the scaled CDF of component i over the maintenance interval it has failed.

System failure depends on the reliability structure. Using the BDD, the expected failure time of the system can be calculated for every component failure combination $X \subseteq M_z$ that can occur during a maintenance interval (t^k, t^{k+1}). We derive the reliability of the system as a function of time, $R^{\tilde{s}|s^z}(t)$, for $t \in (t^k, t^k + \Delta)$ when the system transitions from the post-decision state s^z to $\tilde{s}$ where $\tilde{f}_i = 1$ for all $i \in X \cup M_z^{\complement}$ and $\tilde{f}_i = 0$ otherwise. Since we know the component failure statuses before and after the state transition, the reliability of each component during the state transition can be expressed as

$$R_i^{\tilde{s}|s^z}(t) = \begin{cases} 1 - F_{a_i^z}^\Delta(t), & \text{if } 0 = f_i^z < \tilde{f}_i = 1 \\ 1, & \text{if } f_i^z = \tilde{f}_i = 0 \\ 0, & \text{if } f_i^z = 1. \end{cases} \tag{8}$$

If a component fails during the state transition ($i \in X$), its reliability is $1 - F_{a_i^z}^\Delta(t)$. If component i remains operational during the state transition, its reliability is 1. If it has failed and has not been replaced, its reliability is 0. We define

$$R^{\tilde{s}|s^z}(t) := RS(R_1^{\tilde{s}|s^z}(t), ..., R_m^{\tilde{s}|s^z}(t)). \tag{9}$$

When definition (9) is applied in equation (6), we can calculate the expected system failure time during state transition from state s^z to $\tilde{s}$ as

$$E(t_{\tilde{s}|s^z}^f) = \int_0^\Delta R^{\tilde{s}|s^z}(t)\, dt. \tag{10}$$

The expected system failure time $E(t_{s^z}^f)$ for post-decision state s^z is the expected value over state transitions to each reachable state $\tilde{s} \in RE(s^z)$:

$$E(t_{s^z}^f) = \sum_{\tilde{s} \in RE(s^z)} p(\tilde{s}|s^z)\, E(t_{\tilde{s}|s^z}^f). \tag{11}$$

Finally, the expected system downtime $ED(s^z)$ of the post-decision state s^z is the time from the expected system failure time to the next available maintenance instance $ED(s^z) = \Delta - E(t_{s^z}^f)$.

The total costs depend on replacement decisions, the failure statuses of the system and components, and the expected system downtime costs. Specifically, they consist of expected downtime costs

$$c_D(s, z) = c_d ED(s^z) \tag{12}$$

and maintenance costs

$$c_M(s, z) = c(z) + r_s BR(f) + \sum_{i \in M} (r_i f_i), \tag{13}$$

where $c(z)$ is the cost of the replacement decision obtained from the directed graph with minimum cost arborescence (Leppinen et al., 2025) and costs r_s and r_i are corrective replacement surpluses. The state and action-specific cost is $c(s,z) = c_M(s,z) + c_D(s,z)$.

2.5. *Solution algorithm*

A policy $U\colon S \rightarrow Z$ is a function that prescribes a feasible maintenance action portfolio $z \in Z_s$ for each state s, that is, $U(s) = z$. We focus on *stationary policies* that do not depend on the current maintenance instance t^k. The optimal policy minimizes the net present value of expected maintenance and downtime costs in the long run. The long-term costs of a discounted MDP are represented by a value vector $v^U \in \mathbb{R}^{|S|}$ where $v^U(s)$ represents the expected long-term cost of the policy U when the system is currently in state s.

The model is solved with the Anderson-accelerated Gauss-Seidel Modified Policy Iteration Algorithm (AAGSMPI) (Leppinen et al., 2025). Due to our extensions, this algorithm is more time-consuming than Leppinen et al. (2025), given that for every post-decision state, the number of reachable states is larger and the expected system failure time (11) must be calculated.

3. Model illustration

3.1. *Comparison of policies*

We apply the model to a four-component system with two engines (E1 and E2), a chassis (C), and wheels (W) (Leppinen et al., 2025). In the first instance, we assume that all components need to operate for the system to operate. Thus, the reliability of the system in any post-decision state s^z is $R(s^z) = R_{\text{E1}}(a^z_{\text{E1}})R_{\text{E2}}(a^z_{\text{E2}})R_{\text{C}}(a^z_{\text{C}})R_{\text{W}}(a^z_{\text{W}})$. Table 1 shows the maintenance costs for the components and the Weibull distribution parameters.

The system has a positive economic dependence with a fixed set-up cost $c_0 = 388$. Its technical structural dependencies are represented in Figure 2 as a directed graph (Leppinen et al., 2025). The components must be disassembled before replacements. Disassembly costs are added to the weights of the arcs that lead to component

Table 1. Maintenance costs and failure distribution parameters of system components

	E1	E2	C	W
Disassembly cost	23	28	167	0
Replacement cost	393	403	413	1000
r_i	190	190	50	503
Weibull shape param.	5.1	5.1	5.5	4.0
Weibull scale param.	10.8	10.8	9.9	9.0

replacement nodes (E1, E2, C, W). Additionally, both engines must be disassembled (node DE12) before the chassis can be replaced. The chassis and engines must be disassembled before replacing the wheels. The system failure surplus is $r_s = 110$.

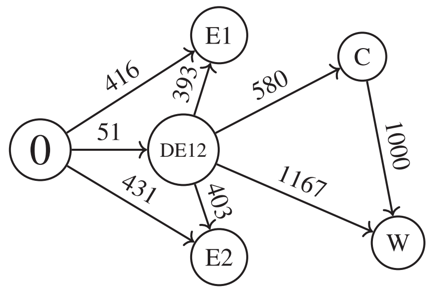

Fig. 2. Directed graph of the 4-component system

We set the reliability threshold at 0.50, the discount factor at 0.99, and the maintenance interval length at 1. We calculate six optimal policies by varying only the unit cost of downtime, $c_d \in \{200, 2000, 4000, 10000, 20000, 65000\}$. We simulate the policies over 150 consecutive maintenance instances. In the initial state, the components are new and operational with zero age. After replacements, failure times are drawn from component-specific Weibull distributions. A component fails if its age exceeds the drawn failure time.

Downtime occurs whenever the system fails. In this case, it will not be operational until the next maintenance instance. For any maintenance policy, the total downtime is calculated as the cumulative downtime over 150 maintenance instances, and the corresponding downtime cost is computed by multiplying this downtime by the unit cost of downtime. The total cost is the discounted sum of the maintenance and downtime costs incurred over 150 maintenance instances.

The results in Figure 3 are averages of 10 000 simulations per policy. For example, with $c_d = 10000$, the average total downtime, expressed as unavailability, is 1.27%. The average total costs are 53 729, of which 82% are maintenance costs and 18% downtime costs. For comparatively high unit costs of downtime $c_d > 2000$, downtime costs remain moderate due to frequent maintenance, leading to reduced total downtime. Thus, total costs increase primarily due to higher maintenance expenses.

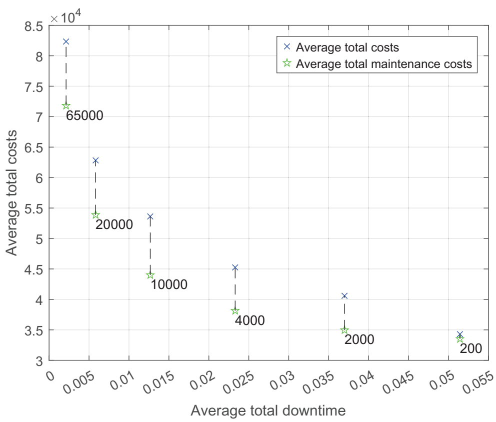

Fig. 3. Average total costs separated into maintenance and downtime costs and average total downtimes of six policies calculated with six unit costs of downtime labeled in the figure.

3.2. *Updates in structural dependencies*

We explore improving system maintainability by relaxing disassembly requirements in three scenarios that alter the directed graph in Figure 2.

(1) Independent chassis replacement: The chassis is directly connected to the root node with an arc of weight 580. The arc between nodes C and W is removed, and the updated weight of the arc between nodes DE12 and W is 1000.
(2) Independent wheel replacement: Node W is connected to the root node with an arc of weight 1000, and the arc between nodes C and W is removed.
(3) Independent replacement of chassis and wheels: Nodes E1, E2, C, and W are directly connected to the root node with arcs of weights 416, 431, 580, and 1000, respectively.

Optimal policies are calculated for the original system and three scenarios with four unit costs of downtime $c_d \in \{200, 4000, 10000, 30000\}$ and the reliability threshold 0.83. Total maintenance costs are assessed using the simulation described in Section 3.1. The results are shown in rows 3–6 of Table 2.

Table 2. Average total maintenance costs are shown for the original system, and the percentage decrease with 0.1% precision in average total maintenance costs are shown for systems with reduced disassembly for the chassis, wheels, or both, a system with parallel engines, and a system with parallel engines and reduced disassembly.

c_d	200	4000	10000	30000
System	Average total maintenance costs			
O: Original	41992	43924	44010	53883
C: Chassis	1.6%	0.3%	0.0%	0.0%
W: Wheels	2.2%	0.3%	0.0%	0.0%
C+W	2.6%	0.3%	0.0%	0.0%
P: Parallel engines	18.1%	21.8%	11.2%	13.2%
P+C+W	19.0%	22.5%	12.1%	14.2%

Compared to Figure 3, maintenance costs increase when $c_d \leq 4000$ due to a stricter reliability threshold. Reduced disassembly leads to lower maintenance costs when $c_d = 200$ because engine disassembly costs can be saved if the chassis or wheels are replaced without engines. However, the changes in maintenance costs are negligible when $c_d \geq 10000$, indicating that the optimal policy suggests grouping replacements. Thus, it can be difficult to justify implementing the changes to the disassembly requirements, especially if downtime costs are high.

We further improve the system's maintainability with a parallel engine configuration. Now, the system operates even if one of the engines has failed. The reliability of this system in any post-decision state s^z is given by equation (5). Optimal policies are calculated for the system with parallel engines and for the system combining parallel engines with reduced disassembly requirements using the four unit costs of downtime. The reliability threshold is 0.83 for the system and 0.50 for non-critical engines. Percentage decrease in cumulative maintenance costs are shown in rows 7–8 of Table 2.

Maintenance costs are lower when a single engine failure does not cause the system to fail. After engine parallelization, reduced disassembly also provides minor cost savings. In particular, increasing c_d from 200 to 4000 results in similar total maintenance costs (as $0.819 \cdot 41992 \approx 0.782 \cdot 43924$ when comparing the values from rows 3 and 7 of Table 2). This implies that the two policies are similar among the states that the system is most likely to reach and that the other parameters have a more significant impact on the structure of the optimal policy than the unit cost of downtime. However, maintenance costs increase with $c_d \geq 10000$.

4. Conclusions

We have developed an optimization model for multi-component systems that accounts for maintenance and downtime costs as well as economic and structural dependencies. In keeping with intuition, more preventive maintenance is carried out when the unit cost of downtime is high, but exceptions may occur, as observed with the system with parallel engines. Our model is novel in that it guides optimal system decisions in view of maintenance and downtime costs in the presence of structural dependencies and reliability constraints.

References

Andersen, J. F., A. R. Andersen, M. Kulahci, and B. F. Nielsen (2022). A numerical study of Markov Decision Process algorithms for multi-component replacement problems. European Journal of Operational Research 299(3), 898–909.

de Jonge, B. and P. A. Scarf (2020). A review on maintenance optimization. European Journal of Operational Research 285(3), 805–824.

Geng, J., M. Azarian, and M. Pecht (2015). Opportunistic maintenance for multi-component systems considering structural dependence and economic dependence. Journal of Systems Engineering and Electronics 26(3), 493–501.

Leppinen, J., A. Punkka, T. Ekholm, and A. Salo (2025). An optimization model for determining cost-efficient maintenance policies for multi-component systems with economic and structural dependencies. Omega 130, 103162.

Nguyen, K.-A., P. Do, and A. Grall (2015). Multi-level predictive maintenance for multi-component systems. Reliability Engineering & System Safety 144, 83–94.

Nguyen, V.-T., P. Do, A. Vosin, and B. Iung (2022). Artificial-intelligence-based maintenance decision-making and optimization for multi-state component systems. Reliability Engineering & System Safety 228, 108757.

Olde Keizer, M. C., S. D. P. Flapper, and R. H. Teunter (2017). Condition-based maintenance policies for systems with multiple dependent components: A review. European Journal of Operational Research 261(2), 405–420.

Olde Keizer, M. C., R. H. Teunter, J. Veldman, and M. Z. Babai (2018). Condition-based maintenance for systems with economic dependence and load sharing. International Journal of Production Economics 195, 319–327.

Safaei, F., E. Châtelet, and J. Ahmadi (2020). Optimal age replacement policy for parallel and series systems with dependent components. Reliability Engineering & System Safety 197, 106798.

Xu, J., Z. Liang, Y.-F. Li, and K. Wang (2021). Generalized condition-based maintenance optimization for multi-component systems considering stochastic dependency and imperfect maintenance. Reliability Engineering & System Safety 211, 107592.

Zheng, M., J. Lin, T. Xia, Y. Liu, and E. Pan (2023). Joint condition-based maintenance and spare provisioning policy for a k-out-of-N system with failures during inspection intervals. European Journal of Operational Research 308(3), 1220–1232.