University of Groningen

Joint condition-based maintenance and inventory optimization for systems with multiple components

Olde Keizer, Minou C.A.; Teunter, Ruud H.; Veldman, Jasper

Published in:
European Journal of Operational Research

DOI:
[10.1016/j.ejor.2016.07.047](https://doi.org/10.1016/j.ejor.2016.07.047)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Olde Keizer, M. C. A., Teunter, R. H., & Veldman, J. (2017). Joint condition-based maintenance and inventory optimization for systems with multiple components. *European Journal of Operational Research*, 257(1), 209-222. <https://doi.org/10.1016/j.ejor.2016.07.047>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Decision Support

Joint condition-based maintenance and inventory optimization for systems with multiple components

Minou C. A. Olde Keizer, Ruud H. Teunter*, Jasper Veldman

Department of Operations, Faculty of Economics and Business, University of Groningen, P.O. Box 800, 9700 AV, Groningen, The Netherlands

ARTICLE INFO

Article history:

Received 16 October 2015

Accepted 22 July 2016

Available online 30 July 2016

Keywords:

Reliability

Maintenance

Inventory

Markov Decision Process

Spare parts

ABSTRACT

Efficient (condition-based) maintenance planning and inventory control of spares for critical components jointly determine the effectiveness of a maintenance strategy and, thereby, balance system uptime and maintenance costs. Duplicating an optimal policy for a single-component system to a multi-component system is not necessarily optimal, while a separate or sequential optimization of the maintenance and inventory decisions is also not guaranteed to yield the lowest costs. We therefore consider the joint optimization of condition-based maintenance and spares planning for multi-component systems. We formulate our model as a Markov Decision Process, and minimize the long-run average cost per time unit. A key insight from our numerical results is that the (s, S) inventory policy, popular in theory as well as practice, can be far from optimal for systems consisting of few components. Significant savings can be obtained by basing both the maintenance decisions and the timing of ordering spare components on the system's condition.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Unexpected failures and resulting downtime account for large losses in the productivity and profitability of a firm (Alsyounf, 2007). Effective maintenance policies can reduce equipment downtime substantially, but rely heavily on the availability of spare components (Jiang, Chen, & Zhou, 2015). Consequently, the joint optimization of maintenance and inventory decisions is an important research area, but most of the existing research, discussed next, considers either maintenance or inventory planning rather than the interface (Van Horenbeek, Scarf, Cavalcante, & Pintelon, 2013). We remark that the (service logistics) inventory literature often uses the term “spare part”, whereas it is common in the maintenance literature to refer to “components” rather than “parts”. To avoid confusion, we will use “spare component” or the shorter “spare” in this paper.

Regarding maintenance decisions, many types of maintenance policies have been both employed in practice and extensively studied under various circumstances, such as corrective, periodic, age-based, and condition-based maintenance (CBM) (Wang, 2002). Compared to other maintenance policies, CBM can be more efficient (Gertsbakh, 1977; 2000), since it bases the maintenance actions on the actual system state. It can reduce the number of fail-

ures (thereby lowering downtime), minimize maintenance costs, and improve operational safety (Rao, 1996). For instance, a CBM policy has been developed for multi-component systems subject to both redundancy and economic dependencies in Olde Keizer, Teunter, and Veldman (2016), without considering inventory decisions. For literature reviews on maintenance policies, we refer to Van der Duyn Schouten (1996), Dekker, Wildeman, and van der Duyn Schouten (1997), Wang (2002). In particular, CBM has been considered by van Noortwijk (2009), Ahmad and Kamaruddin (2012), Bousdekis, Magoutas, Apostolou, and Mentzas (2015), Marseguerra, Zio, and Podofilini (2002), Hong, Zhou, Zhang, and Ye (2014), Li, Deloux, and Dieulle (2016), Rasmekomen and Parlikad (2016). Also inventory strategies have been extensively researched, of which reviews are provided by Kennedy, Wayne Patterson, and Fredendall (2002), Basten and van Houtum (2014), van Houtum and Kranenburg (2015). The spares inventory literature typically treats demand as given, thereby ignoring the underlying maintenance planning, while the majority of research on maintenance assumes an unlimited number of spares. Relatively few contributions exist on the joint optimization of maintenance and inventory. Next, we only discuss those that consider CBM, and refer interested readers to more in-depth reviews in Van Horenbeek, Buré, Cattrysse, Pintelon, and Vansteenwegen (2013), Pierskalla and Voelker (1976), Cho and Parlar (1991).

When considering a joint CBM and inventory policy for a system consisting of a single component, it is proven that a so-called monotonic policy structure is optimal (Kawai, 1983). In such a

* Corresponding author.

E-mail addresses: m.c.a.olde.keizer@rug.nl (M.C.A. Olde Keizer), r.h.teunter@rug.nl (R.H. Teunter), j.veldman@rug.nl (J. Veldman).

Nomenclature

δ_j	binary variable indicating whether or not component j is replaced
μ_j	deterioration parameter of component j
ω	decision variable indicating how many spares to order
c_r^j	cost of a replacement on component j
F	fixed cost per order
H	holding cost per spare per time unit
L_j	fixed failure level of component j
O^j	vector with operating costs for different states of component j
P^j	transition probability matrix of component j
R^j	vector with replacement costs for different states of component j
S	maximum stock level
s	reorder stock level
$\bar{S}$	maximum inventory position
s_h	number of spares on hand
s_l	number of spares ordered l time units ago
T	fixed lead time, $T > 1$
x_j	state of component j

policy, deterioration thresholds are used to determine when to order a spare and, upon arrival of the spare, when to replace the component. Other examples of sequential or joint optimization of CBM and the spares inventory for a single-component system are given by Kawai (1983), Elwany and Gebraeel (2008), Wang, Chu, and Mao (2008a), Rausch and Liao (2010), Louit, Pascual, Banjevic, and Jardine (2011), Zhao and Xu (2012).

In practice, systems often contain multiple components. Applying a single-component policy to such a multi-component system is generally far from optimal for several reasons (Cho & Parlar, 1991). First, different types of dependencies can exist in multi-component systems, which can be economic, structural, or failure-related (Thomas, 1986). In such cases, the optimal maintenance and inventory decisions depend on the complete system state rather than on a single component. Second, multiple components can share a set of identical spares. We remark that identical components may have different failure rates as they can for instance be contained in subsystems that operate under different conditions. Examples are systems consisting of multiple production lines with similar critical components (such as conveyor belts), or gas treatment facilities that use multiple relatively similar pumps to ensure a continuous gas distribution. Obviously, the right number of shared (or pooled) spares should be determined at the system level. It is well known from the inventory pooling literature (see, e.g., Guajardo, Rönnqvist, Halvorsen, & Kallevik, 2015; Karsten & Basten, 2014) that a decomposed approach at the component level leads to much higher inventory levels and costs.

To the best of our knowledge, the integration of CBM and inventory for multi-component systems has only been studied by Wang, Chu, and Mao (2008b), Xie and Wang (2008), Wang, Chu, and Mao (2009), Van Horenbeek and Pintelon (2015). Whereas a shared pool of spares is considered in Wang et al. (2008b), Xie and Wang (2008), Wang et al. (2009), economic and structural dependencies are included in Van Horenbeek and Pintelon (2015). A sequential optimization of maintenance and inventory is considered in Van Horenbeek and Pintelon (2015). However, separate or sequential optimization of maintenance and inventory actions will not necessarily lead to a globally optimal policy (Xie & Wang, 2008). For this reason, maintenance and inventory decisions are jointly optimized in Wang et al. (2008b), Xie and Wang (2008),

Wang et al. (2009). All three papers consider an (s, S) inventory policy, which means that an order is placed to refill the inventory position to S units once it drops below s . It is well-known that the order level, order-up-to level (s, S) policy is optimal under quite general conditions for inventory systems (Iglehart, 1963; Sahin, 1990; Scarf, 1959). This policy has also been considered by many authors for controlling spare part inventories (e.g., Cohen, Kleindorfer, Lee, & Pyke, 1992; van Jaarsveld, Dollevoet, & Dekker, 2015; Kennedy et al., 2002; Kranenburg & van Houtum, 2009; Strijbosch, Heuts, & van der Schoot, 2000; Svoronos & Zipkin, 1991; Williams, 1984; Zohrul Kabir & Al-Olayan, 1996). In practice, this policy is often referred to as the min-max policy, where an order is placed up to the maximum if the inventory position drops to (or below) the minimum. This policy is available in all major software packages for stock control (e.g., Slimstock) or Enterprise Resource Planning (e.g., SAP). Intuitively, however, components only need to be replaced and thus require a spare when they are close to failure. The condition information that is used for scheduling maintenance can thus also be used for deciding when to order spares. No research has been performed yet on this (just-in-time) condition-based ordering for multi-component systems, despite the obvious cost savings potential. To cover this gap, we are the first to consider the joint optimization of the condition-based maintenance and inventory decisions for a multi-component system with a shared pool of spares. We benchmark the performance of our condition-based inventory policy against that of an (s, S) policy and the optimal policy for a single component. Another contribution is that we are the first to provide an exact method for a multi-component system by formulating the problem as a Markov Decision Process, whereas Wang et al. (2008b), Xie and Wang (2008), Wang et al. (2009) use a simulation-based approach. In this way, we are able to obtain structural insights through a numerical study. Many Markovian maintenance models have been developed for deteriorating single-component systems, e.g., Byon, Ntaimo, and Ding (2010), Elwany, Gebraeel, and Maillart (2011), Ulukus, Kharoufeh, and Maillart (2012), Borrero and Akhavan-Tabatabaei (2013).

The remainder of this paper is organized as follows. Section 2 gives a description of the system, while we formulate the model as a Markov Decision Process in Section 3. Next, we present a numerical study and comparison to the (s, S) policy for a two-component system in Section 4, followed by a sensitivity analysis in Section 5. In Section 6, we consider a system with more than two components. Section 7 concludes the paper.

2. System description

2.1. Model description

To obtain structural insights, we consider a discrete-time system consisting of N components, which function and deteriorate independently. We model the condition of a component j , $j = 1, 2, \dots, N$, discretely using $L_j + 1$ different states, $0, 1, \dots, L_j$, where state 0 means that the component is as-good-as-new, and state L_j means that the component has failed. The components share a pool of spares. Although the components are identical, they may be contained in different subsystems that may operate under different conditions, possibly leading to different failure rates. If a component is replaced, the old component is discarded, and the new component is in the as-good-as-new state 0. Since repair times are typically small (days) compared to the expected lifetime of a component (years) and lead times of spares (months), replacements are assumed to be instantaneous, but can only be scheduled if the required spares are on hand. Spares can be ordered in any amount, and arrive after a fixed lead time of T time units (typically in the order of months). After possible maintenance and inventory actions have been performed, component j is subject to deterioration. We assume that deterioration worsens, rather than improves,

the state of a component. To that extent, we assume that the deterioration increments of component j follow a Poisson distribution with parameter μ_j .

2.2. Order of events and costs

The order of events is as follows. At the start of each time unit, an order for spares can arrive (provided that an order has been placed T time units ago). Next, a so-called operating cost is incurred for each component, which depends on the state of the component. Let $O^j = (O_0^j, O_1^j, \dots, O_{L_j}^j)$ denote the vector of operating costs. If component j is in state u , it incurs an operating cost of O_u^j per time unit. In this way, it is not only possible to include a downtime cost for a failed component (by choosing a high value for $O_{L_j}^j$), but also to include costs for losses in revenue due to deterioration of the component. After operating costs have been incurred, components can be replaced, provided that sufficient spares are on hand. In practice, a corrective replacement is sometimes more expensive than a preventive replacement, because a component failure can cause damage on other elements of the system as well. This distinction has been considered by Kawai (1983), Bouvard, Artus, Bérenguer, and Cocquempot (2011), Park and Pham (2012), de Jonge, Klingenberg, Teunter, and Tinga (2015). Similar to Kawai (1983), we therefore define $R^j = (R_0^j, R_1^j, \dots, R_{L_j}^j)$ as the vector of state-dependent replacement costs for component j including the purchase price of the new component, where $R_u^j \geq R_v^j$ for $u > v$. After possible maintenance actions have been performed, spares can be ordered, for which a fixed cost per order F is incurred. This cost is independent of the number of spares that are ordered. For each spare that is still on hand at the end of a time unit, a holding cost H is incurred.

3. Markov Decision Process formulation

The Markov Decision Process (MDP) model consists of a set of possible states $\mathcal{I}$ at the start of each time unit, a set of possible actions $\mathcal{A}_{\{i\}}$ that can be performed at each state $i \in \mathcal{I}$, transition probabilities $p^a(i; \bar{i})$ that the system moves from state $i \in \mathcal{I}$ to state $\bar{i} \in \mathcal{I}$ if action $a \in \mathcal{A}_{\{i\}}$ is chosen, and the cost function $c^a(i)$ which returns the expected costs of performing action $a \in \mathcal{A}_{\{i\}}$ at state $i \in \mathcal{I}$. To limit the size of the state space, we assume that the inventory position (spares on hand plus spares on order) cannot exceed some maximum $\bar{S}$. This maximum is only needed to ensure boundedness of the state space, and will be set sufficiently large in our numerical experiments to ensure that it does not affect the results.

State space. The state space keeps track of the state of each component $(x_1, x_2, \dots, x_N)$, the status of each order $(s_1, s_2, \dots, s_{T-1})$, where s_l denotes the number of spares ordered l time units ago, $l = 1, 2, \dots, T-1$, and the number of spares on hand (s_h). This results in the following $(N+T)$ -dimensional state space:

$$\mathcal{I} = \{(x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h)\},$$

where, for $j = 1, 2, \dots, N$ and $l = 1, 2, \dots, T-1$:

$$x_j \in \{0, 1, \dots, L_j\} \quad (\text{state of component } j),$$

$$s_l \in \{0, 1, \dots, \bar{S}\} \quad (\text{number of spares ordered } l \text{ time units ago}),$$

$$s_h \in \{0, 1, \dots, \bar{S}\} \quad (\text{number of spares on hand}).$$

The size of the state space is restricted by imposing a logical constraint. As we assumed that the inventory position cannot exceed the maximum of $\bar{S}$, we have $s_1 + s_2 + \dots + s_{T-1} + s_h \leq \bar{S}$.

Action space. At the start of each time unit, a decision is needed on whether or not to replace some components and on the number of spares to order. Hence, the $(N+1)$ -dimensional action space can be represented as follows:

$$\mathcal{A} = \{(\delta_1, \delta_2, \dots, \delta_N, \omega)\},$$

where, for $j = 1, 2, \dots, N$:

$$\delta_j = \begin{cases} 1, & \text{if component } j \text{ is replaced,} \\ 0, & \text{otherwise,} \end{cases}$$

$$\omega \in \{0, 1, \dots, \bar{S}\} \quad (\text{number of spares to order}).$$

The following sets of actions are allowed for different states of the state space, for any possible values of x_j ($j = 1, 2, \dots, N$), s_l ($l = 1, 2, \dots, T-1$), and s_h :

$$\mathcal{A}_{\{(x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h)\}} = \left\{ (\delta_1, \delta_2, \dots, \delta_N, \omega) : \sum_{j=1}^N \delta_j \leq s_h, \right. \\ \left. \omega \leq \bar{S} - s_h - \sum_{l=1}^{T-1} s_l + \sum_{j=1}^N \delta_j \right\}.$$

Transition probabilities. If no maintenance is performed, component j will move from state x_j to state $\bar{x}_j$ with probability $p_{x_j \bar{x}_j}^j$. If component j is replaced, however, its condition becomes as-good-as-new, and it moves from state 0 to state $\bar{x}_j$ with probability $p_{0 \bar{x}_j}^j$. Hence, we define the probability of moving from state x_j to state $\bar{x}_j$ under maintenance action δ_j (denoted by $p_{x_j \bar{x}_j}^j(\delta_j)$) as follows, for $j = 1, 2, \dots, N$:

$$p_{x_j \bar{x}_j}^j(\delta_j) = \begin{cases} p_{x_j \bar{x}_j}^j, & \text{if } \delta_j = 0, \\ p_{0 \bar{x}_j}^j, & \text{if } \delta_j = 1. \end{cases}$$

In our numerical experiments, we assume that the deterioration increments (i.e., the increases in deterioration between two time units) of component j follow a Poisson distribution with parameter μ_j , $j = 1, 2, \dots, N$. To illustrate this, let X_j be Poisson distributed with mean μ_j . Then, the transition probabilities of component j are given by

$$p_{x_j \bar{x}_j}^j = \begin{cases} P(X_j = \bar{x}_j - x_j), & \text{if } \bar{x}_j < L_j, \\ P(X_j \geq \bar{x}_j - x_j), & \text{if } \bar{x}_j = L_j. \end{cases}$$

The system now moves from state $(x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h)$ to state $(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_N, \bar{s}_1, \bar{s}_2, \dots, \bar{s}_{T-1}, \bar{s}_h)$, depending on the actions $(\delta_1, \delta_2, \dots, \delta_N, \omega)$ performed, with probability $\prod_{j=1}^N p_{x_j \bar{x}_j}^j(\delta_j)$ if the transition is possible, and with probability 0 otherwise:

$$p^{(\delta_1, \delta_2, \dots, \delta_N, \omega)}((x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h); \\ (\bar{x}_1, \bar{x}_2, \dots, \bar{x}_N, \bar{s}_1, \bar{s}_2, \dots, \bar{s}_{T-1}, \bar{s}_h)) \\ = \begin{cases} \prod_{j=1}^N p_{x_j \bar{x}_j}^j(\delta_j), & \text{if } \bar{s}_h = s_h - \sum_{j=1}^N \delta_j + s_{T-1}, \bar{s}_1 = \omega, \text{ and} \\ & \bar{s}_l = s_{l-1} \text{ for } l \in \{2, 3, \dots, T-1\}, \\ 0, & \text{otherwise.} \end{cases}$$

Expected costs. During each time unit, costs are incurred. The costs of performing action $(\delta_1, \delta_2, \dots, \delta_N, \omega)$ in state $(x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h)$ are given by the sum of the operating costs, the replacement costs (if replacements are performed), the order cost (if an order for spares is placed), and the holdings costs (for the spares that remain on hand), i.e.:

$$c^{(\delta_1, \delta_2, \dots, \delta_N, \omega)}(x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h) \\ = \sum_{j=1}^N O_{x_j}^j + \sum_{j=1}^N \delta_j \cdot R_{x_j}^j + I_{\{\omega > 0\}} \cdot F + \left(s_h - \sum_{j=1}^N \delta_j \right) \cdot H,$$

where $I_{\{\cdot\}}$ denotes the indicator function that equals one if the expression between brackets is true and zero otherwise. Note that this MDP formulation is constructed for the case with a lead time of $T > 1$. In case $T = 1$, there is no need to keep track of when orders are placed; an order simply arrives before the next decision is made. Therefore, the state space reduces to an $(N + 1)$ -dimensional one that only keeps track of the state of each component and the number of spares on hand. The action space and transition probabilities also simplify for $T = 1$.

3.1. (s, S) inventory policy

The inventory position (denoted by IP) includes both the spares on hand ($s_h - \sum_{j=1}^N \delta_j$) and on order ($\sum_{l=1}^{T-1} s_l$). The (s, S) inventory policy places an order for $S - IP$ spares as soon as the inventory position is smaller than or equal to s . It can be seen as a special case of our MDP, by restricting the set of actions that are allowed in each state as follows, for any x_j ($j = 1, 2, \dots, N$), s_l ($l = 1, 2, \dots, T - 1$), and s_h :

$$\mathcal{A}_{\{(x_1, x_2, \dots, x_N, s_1, s_2, \dots, s_{T-1}, s_h)\}} \\ = \left\{ (\delta_1, \delta_2, \dots, \delta_N, \omega) : \sum_{j=1}^N \delta_j \leq s_h, \omega = I_{\{s_h + \sum_{l=1}^{T-1} s_l - \sum_{j=1}^N \delta_j \leq s\}} \right. \\ \left. \cdot \left(S - s_h - \sum_{l=1}^{T-1} s_l + \sum_{j=1}^N \delta_j \right) \right\}.$$

The remainder of the MDP formulation above remains unchanged.

3.2. Value iteration

As a performance criterion, we are interested in the long-run average cost per time unit. At the start of Section 3, we assumed that the inventory position cannot exceed the maximum of $\bar{S}$. This assumption guarantees that both the state space and the action space are finite. Furthermore, the cost function is bounded by definition. If, in addition, the model would be *unichain*¹, then a stationary average optimal policy exists, and the value iteration algorithm can be applied to find such a policy (Puterman, 1994). However, our model can contain multiple recurrent states. Consider for example the case where $N = 2$ and $T = 3$, and consider the stationary policy $f(i) = (0, 0, 0)$ for all $i \in \mathcal{I}$ (i.e., never perform maintenance and never order spares). Under this policy, in the long-run, each component will be in the failed state (L_j), and the initial number of spares on hand will remain unchanged. The resulting transition matrix therefore contains $\bar{S} + 1$ recurrent states: $(L_1, L_2, 0, 0, 0)$, $(L_1, L_2, 0, 0, 1)$, $\dots$, $(L_1, L_2, 0, 0, \bar{S})$. Hence, we are dealing with a model that is multichain rather than unichain. Nevertheless, it is realistic to assume that any optimal policy orders a new spare (at some point in time) for every component that is replaced, and that a failed component is replaced as soon as a spare is available, as long as the operating cost for a failed component (i.e., the downtime cost) is chosen sufficiently high. This implies that the model does satisfy the *Weak Unichain Assumption*²

¹ An MDP is unichain if the transition matrix corresponding to every deterministic stationary policy consists of a single recurrent class and a (possibly empty) set of transient states (Puterman, 1994).

² For each average cost optimal stationary policy, the associated Markov chain has no two disjoint closed sets (Tijms, 1994).

as defined in Tijms (1994). Therefore, the minimal average cost per time unit is independent of the initial state, and can be efficiently determined by applying the value iteration algorithm (Tijms, 1994), which is shown below. Here v_n denotes the value function obtained with the n th iteration.

Value iteration algorithm

Step 0: Initialization

Choose an $\varepsilon > 0$ and $v_0(i)$ with $0 \leq v_0(i) \leq \min_{a \in \mathcal{A}(i)} c^a(i)$ for all $i \in \mathcal{I}$. Set $n := 1$.

Step 1:

For each $i \in \mathcal{I}$, compute the value function $v_n(i)$ as

$$v_n(i) := \min_{a \in \mathcal{A}(i)} \left\{ c^a(i) + \sum_{\tilde{i} \in \mathcal{I}} p^a(i; \tilde{i}) v_{n-1}(\tilde{i}) \right\}, \quad (1)$$

and select a stationary policy f_n which minimizes the right-hand side of (1) as

$$f_n(i) \in \arg \min_{a \in \mathcal{A}(i)} \left\{ c^a(i) + \sum_{\tilde{i} \in \mathcal{I}} p^a(i; \tilde{i}) v_{n-1}(\tilde{i}) \right\}. \quad (2)$$

Step 2:

Let $M_n := \max_{i \in \mathcal{I}} \{v_n(i) - v_{n-1}(i)\}$ and $m_n := \min_{i \in \mathcal{I}} \{v_n(i) - v_{n-1}(i)\}$. Stop the algorithm with policy f_n if $0 \leq M_n - m_n \leq \varepsilon \cdot m_n$. Otherwise, set $n := n + 1$ and repeat Steps 1 and 2.

Let g^* denote the minimal average cost per time unit, let f_n be a stationary policy which satisfies (2) for $n \geq 1$, and let $g_i(f_n)$ denote the corresponding one-step difference $v_n(i) - v_{n-1}(i)$. Since the latter is independent of the state i , we denote it by $g(f_n)$ instead. Then it holds that $m_n \leq g^* \leq g(f_n) \leq M_n$, for $i \in \mathcal{I}$, where the sequences $\{m_n, n \geq 1\}$ and $\{M_n, n \geq 1\}$ are non-decreasing and non-increasing, respectively (Tijms, 1994). Furthermore, $g(f_n)$, the average cost resulting from policy f_n , deviates at most 100ε percent from g^* .

4. Numerical investigation

In this section, we construct a base case and obtain some numerical results that illustrate the structure of the optimal policy. We compare its performances to those obtained by (s, S) inventory policies and the case where the decisions for each component are optimized separately, and analyze the different cost components. In Section 5, we will perform a sensitivity analysis on this base case, while we vary the number of components in Section 6.

4.1. Base case

In our base case, we consider a system consisting of two components with identical failure rates that deteriorate independently. By considering two components, the resulting optimal maintenance and inventory decisions can be represented two-dimensionally, which allows easy interpretation. Considering components with identical failure rates further aids in interpreting the structure of the optimal policy. It also allows us to drop the superscripts denoting to which component a certain parameter corresponds. This (realistic) base case serves to illustrate some characteristics of the optimal policy and the effects of limiting the possible inventory actions to an (s, S) inventory policy.

We consider a fixed failure level of $L = 4$, so each component can be in one of five different states. Furthermore, we assume that

Table 1
Model parameters for the base case.

Parameter	Interpretation	Value
L	Failure level	4
O	Operating cost for states 0,1,2,3,4	(0, 0, 0, 0, 100)
R	Replacement cost for states 0,1,2,3,4	(5, 5, 5, 5, 5)
F	Fixed cost per order	0
H	Holding cost per spare per time unit	0.5
T	Fixed lead time	3
μ	Deterioration parameter	0.2

inspections are performed four times a year, so one time unit has a length of three months. Although current technology allows systems to be monitored continuously by placing sensors, we observe in practice that a mixture of periodic and continuous review is applied. In the process industries, the use of (Wi-Fi) sensors can be prohibited due to safety reasons, while placing sensors may not yet be cost-effective for other systems (e.g., large systems such as a railway network). Furthermore, effective condition-based maintenance requires a measurable parameter that correlates strongly with the onset of failure (Tsang, 1995), which often appears to be problematic. Gas distribution companies, for instance, lack proper models that can link internal fouling of heat exchangers to available process data (e.g., temperatures, flows) (Veldman, Wortmann, & Klingenberg, 2011), making periodic visual inspection the only suitable way to appropriately determine the asset's condition. Periodic and continuous review will thus continue to co-exist, certainly for the near future. An expected lifetime of E years is equivalent to selecting the deterioration parameter as $\mu = \frac{1}{E}$ per time unit. We assume that the deterioration increments per time unit of each component follow a Poisson distribution with parameter $\mu = 0.2$, i.e., we assume an expected lifetime of 5 years (or 20 time units). This implies that, when no maintenance is performed, each component is subject to transition probabilities p_{uv} of moving from state u to state v during one time unit, summarized in the following upper-triangular transition matrix P .

$$P = \begin{bmatrix} 0.82 & 0.16 & 0.02 & 0.00 & 0.00 \\ 0 & 0.82 & 0.16 & 0.02 & 0.00 \\ 0 & 0 & 0.82 & 0.16 & 0.02 \\ 0 & 0 & 0 & 0.82 & 0.18 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

Each order arrives after three time units, which equals nine months and is realistic for spares. Indeed, a study on forecasting intermittent demand using a large dataset from the UK Royal Air Force with lead times of spares such as valves, cables, and diodes shows an average lead time of nine months for these spares (Teunter & Duncan, 2009). In this base case, we do not include a fixed cost per order, i.e., we select $F = 0$, but we will investigate the influence of this cost in the sensitivity analysis in Section 5.3. Although our model is capable of handling state-dependent replacement costs, we choose $R = (5, 5, 5, 5, 5)$ (i.e., the replacement cost is independent of the state of the component), because this allows us to focus on the key issues in the paper. The parameter settings are listed in Table 1.

4.2. Optimal maintenance and inventory policy for the base case

We apply the value iteration algorithm to find the long-run average cost per time unit for the joint condition-based maintenance and inventory policy. We use $\varepsilon = 0.0005$ for the stopping criterion, which implies that the resulting average cost deviates at most 0.05 percent from the actual average cost. Furthermore, we use $v_0(i) = 0$ for all $i \in \mathcal{I}$ as our initial value function. Results indicate that a maximum inventory position of $\bar{S} = 2$ is sufficiently

Fig. 1. Convergence of the sequences $\{M_n, n \geq 1\}$ and $\{m_n, n \geq 1\}$ in the value iteration algorithm.

high for this example, as increasing this value does not alter the resulting optimal policy and corresponding average cost. All experiments are performed using Python 3.4.3 on a computer with a 3.30 gigahertz quad core processor and 16.0 gigabytes of RAM. For this example, the state space consists of 250 states and the value iteration algorithm converges after 24 iterations, which takes about 0.04 seconds. Fig. 1 shows the convergence of the sequences $\{M_n, n \geq 1\}$ and $\{m_n, n \geq 1\}$. It appears that these sequences behave nicely, and converge relatively quickly.

Fig. 2 shows the resulting optimal maintenance and inventory policy for the base case. In this figure, the optimal maintenance and inventory actions are shown for any combination of the states x_1 and x_2 of components 1 and 2, respectively, given a realization of the number of spares ordered l time units ago ($s_l, l = 1, 2$) and the number of spares on hand (s_h). Note that the choice of $\bar{S}$ determines the number of realizations that (s_1, s_2, s_h) can take. The average cost per time unit for the optimal policy is equal to 1.57 per time unit. From Fig. 2, we observe that the optimal maintenance decisions for one component depend on the states of both components. This is caused by the fact that the components share a pool of spares. Indeed, when sufficient spares are on hand, i.e., when $s_h = 2$, the optimal replacement decision of component j depends solely on its own deterioration level x_j (replace if $x_j \geq 2$), $j = 1, 2$. However, when only one spare is on hand, this decision depends on the complete system state. Consider for example the case where $(s_1, s_2, s_h) = (0, 0, 1)$, i.e., one spare on hand and no outstanding orders. In principle, the component with the highest degradation level will be replaced, provided that the component is at least in state 2. However, when components 1 and 2 are equally close to failure, both in state 2 or 3, no replacement is performed. Instead, the spare is saved for the component that gains most deterioration in the near future, and an additional spare is ordered. On the other hand, if a new spare has already been ordered two time units ago, i.e., $(s_1, s_2, s_h) = (0, 1, 1)$, no spare will be reserved, as a new spare will arrive before the next inspection. The optimal replacement decisions for a component depend thus on the state of each component, as well as the number of spares on hand and the status of each order.

4.3. Optimal maintenance and inventory policy for a single component

As stated in Cho and Parlar (1991), the optimal policy for a single component is not necessarily optimal for the complete (multi-component) system. To investigate this, we now treat each component separately in the base case. To that extent, we apply the MDP to find the optimal ordering and replacement decisions for a

Fig. 2. Matrices with optimal maintenance and inventory actions for x_1 and x_2 , for different realizations of (s_1, s_2, s_h) .

single component. It is proven that a monotonic policy structure is optimal for a single component, in which deterioration thresholds are used to decide on when to order a spare and perform a replacement (Kawai, 1983). Indeed, results indicate that a spare is ordered when the component is in state 0 or higher (so immediately), and that a replacement is performed if the component is in state 2 or higher, provided that a spare is available. The corresponding minimal average costs are equal to 0.92 per time unit, which, for two components, is about 17 percent more expensive than our optimal policy. This implies that a significant cost reduction can be obtained by jointly optimizing the decisions for both components. In the remainder of this paper, we refer to the separate optimization of the decisions for each component as the “single component” policy.

4.4. Condition-based maintenance with an (s, S) inventory policy for the base case

Similarly to the optimal condition-based ordering policy assessed in Section 4.2, we use $\varepsilon = 0.0005$ and $v_0(i) = 0$ for all $i \in \mathcal{I}$ as our initial value function for assessing the CBM model with an (s, S) inventory policy. Because we do not consider a fixed cost per order in this base case, an $(S - 1, S)$ policy structure will be optimal, so we only consider the cases where $s = S - 1$. In the sensitivity analysis, we will also evaluate (s, S) policies for other values of s , when we do include a fixed cost per order. Table 2 shows the average cost per time unit for different values of S , along with the number of iterations required in the value iteration algorithm. Costs rapidly increase for $S > 2$, and are therefore not reported. From Table 2, it follows that the lowest costs are obtained for the $(S - 1, S)$ policy with $S = 2$. These costs are about 14 percent

Table 2

Performance of the $(S - 1, S)$ policy for different values of S .

$(S - 1, S)$ policy with	Average cost per time unit	Iterations
$S = 1$	1.92	28
$S = 2$	1.79	23

higher than with condition-based ordering of spares. The corresponding maintenance (and inventory) policy is shown in Fig. 3.

Comparing Figs. 2 and 3, we observe that the optimal replacement actions corresponding to the condition-based inventory policy and the best $(S - 1, S)$ inventory policy, with $S = 2$, are almost identical.

The optimal inventory actions do differ, however, between the two policies. The difference in costs hence seems to be mainly caused by the different inventory decisions. The long-run average cost per time unit corresponding to a certain policy consists of costs corresponding to component downtime (i.e., the operating cost), replacements, ordering, and holding costs. For the different policies, Fig. 4 shows the corresponding minimal average cost per time unit, divided into the different cost components. As we did not include a fixed cost per order in this example, the ordering costs equal zero and can thus be left out. Fig. 4 shows that among the different policies, the replacement cost remains approximately unchanged. This is explained by the fact that all (nearly) failed components are ultimately replaced by any sensible policy. However, due to the different timing of placing orders among the different policies, the holding cost and downtime cost do vary. The single component policy has low operating costs, but very high holding costs. For different values of S , the $(S - 1, S)$ policy also has either low operating costs or low holding costs, but is unable to simultaneously minimize these. Even the best $(S - 1, S)$ policy,

Fig. 3. Matrices with optimal maintenance and best $(S - 1, S)$ inventory actions, for $S = 2$, for x_1 and x_2 , for different realizations of (s_1, s_2, s_h) .

Fig. 4. Average cost per time unit divided per cost component, for different strategies.

with $S = 2$, increases the total inventory and holding costs with about 31 percent compared to the condition-based inventory policy, as spares are ordered too soon for relatively new components. Indeed, it has been shown in Syntetos, Teunter, Babai, and Tranchel (2016) that under a continuous (s, S) policy spares are often ordered too early, indicating that delayed ordering could be profitable. Basing the decisions on when to order spares on the complete system state rather than applying a basic inventory policy (which only considers the number of spares on hand and for which the demand process is memoryless) is an effective way to delay orders, and thus to reduce costs significantly.

5. Sensitivity analysis

In this section, we will vary several parameters of the base case, divided into different categories. We distinguish between the parameters that are related to the replacement cost, R , the holding

cost, H , and the order cost, F . Furthermore, we will vary the deterioration parameter μ and the lead time T , and consider two components that deteriorate at different rates.

5.1. Influence of replacement cost

In the original example, we chose the replacement cost vector as $R = (5, 5, 5, 5, 5)$. Hence, the replacement costs are independent of the state of the component. In the sequel, we denote the replacement cost by c_r . This means that the replacement cost vector can be obtained as $R = (c_r, c_r, c_r, c_r, c_r)$. In Fig. 5, we vary the preventive replacement cost c_r between 2 and 15. Naturally, a higher value for c_r will increase the total average cost per time unit. The optimal policy, the best $(S - 1, S)$ policy, and the single component policy are almost not affected by the value of c_r , indicating that the increases in total cost are mainly caused by the more expensive replacements. Indeed, the absolute difference in costs between the different policies remains approximately equal as c_r increases, thus decreasing the relative difference. Nevertheless, even for relatively high values of c_r , the optimal (condition-based) inventory policy significantly outperforms both the best $(S - 1, S)$ policy and the single component policy.

5.2. Influence of holding cost

In Fig. 6, we vary the holding cost H , which was equal to 0.5 for the base case, between 0 and 2.2 per time unit. For small values of H , it is relatively cheap to hold spares on hand. The optimal policy thus approaches the $(1, 2)$ policy for small H . For high holding costs, however, having spares on hand becomes expensive. For this reason, the $(0, 1)$ policy is the best performing $(S - 1, S)$

Fig. 5. Average cost per time unit for different policies for different values of c_r , along with the percentage increase in cost from applying the best $(S-1, S)$ inventory policy or the single component policy rather than the optimal policy.

Fig. 6. Average cost per time unit for different policies for different values of H , along with the percentage increase in cost from applying the best $(S-1, S)$ policy or the single component policy rather than the optimal policy.

policy for high H , while the $(1, 2)$ policy results in much higher costs. Nevertheless, significant savings can be obtained by applying the condition-based inventory policy, which better coordinates the optimal ordering and replacement decisions. Furthermore, we observe in Fig. 6 that the single component policy performs quite well for small values of H , which is in line with Fig. 4 (where we found that the holding costs mainly cause the cost difference with the optimal policy). For values of H larger than 1.1, however, we observe that the costs corresponding to the single component policy no longer increase, indicating that spares are no longer held on hand. The cost difference with the optimal policy thus decreases for $H > 1.1$.

To get some insights into the optimal policy structure, Fig. 7 shows the different cost components (operating, replacement, and holding costs) corresponding to the optimal policy for different values of H .

The jumps in the average holding cost per time unit are caused by changes in the corresponding optimal policy; as H increases, replacements are performed more frequently to avoid high holding costs. Between these jumps, the optimal policy remains (almost) unchanged, which means that the costs corresponding to replacements and the operating costs remain unchanged. As H increases, the average holding cost will increase with the same rate, thus gradually increasing total costs, until a new jump occurs. As H increases further, replacements will be performed as soon as a spare arrives to avoid paying the holding cost. Consider for example $H = 10$ (not shown in the figure). The corresponding minimal total average costs are equal to 2.26 per time unit, of which 0.46, 1.80, and 0 are the average operating cost, replacement cost, and holding cost, respectively. Thus, as H increases, holding costs will be avoided completely by postponing the ordering of spares and performing replacements as soon as an order for spares

Fig. 7. Average cost per time unit for the optimal policy divided into different cost components for different values of the holding cost H .

arrives. This can also be observed from the corresponding optimal maintenance and inventory policy, which is shown in Fig. A.1 of Appendix A.

5.3. Including a fixed cost per order

In the base case, we did not include a fixed cost per order, i.e., we chose $F = 0$. In Fig. 8, we do include a fixed cost per order, and vary it between 0 and 4.

For a non-zero fixed cost per order F , the $(S-1, S)$ policy is no longer the best (s, S) policy, but instead the $(0, 2)$ policy becomes best. Also in the optimal strategy, spares are ordered in sets of two more frequently as F increases, thus reducing the cost difference with the best (s, S) policy. With no fixed cost per order, the optimal policy benefits most from basing the inventory decisions on the system's condition in such a way that spares arrive just in time. This benefit is reduced when a fixed order cost incites to order in (larger) batches. When we optimize the decisions for each component separately (i.e., when we apply the single component

strategy), spares are ordered individually for each component, independent of F . The cost difference with this policy thus remains substantial for any fixed cost per order.

5.4. Influence of the deterioration parameter and lead time

In the original example, we considered a deterioration parameter of $\mu = 0.2$, which can be interpreted as an expected lifetime of 5 years. In Fig. 9, we consider expected lifetimes between 3 and 15 years.

We observe that for relatively small expected lifetimes of the components (of less than six years), spares are needed quite often. Indeed, the optimal policy approaches the $(1, 2)$ policy. As the expected lifetime increases, fewer spares are needed, and a condition-based inventory policy becomes more rewarding. For an expected lifetime of 7 years, the $(0, 1)$ policy is quite efficient, but for higher expected lifetimes, this policy is not suitable as spares are ordered too soon. The optimal policy significantly outperforms the best $(S-1, S)$ policy by only ordering spares once the components have deteriorated considerably. For expected lifetimes exceeding 9 years, the single component policy even outperforms the best $(S-1, S)$ policy, as it is able to postpone the ordering of spares for relatively new components.

Furthermore, we considered a fixed lead time of $T = 3$ (nine months) in the original example. In Fig. 10, we vary T between 2 and 9 time units, i.e., we consider lead times between 6 and 27 months. Since our state space is $(N+T)$ -dimensional, the choice of T influences the size of the state space and the computation time. A maximum inventory position of $\bar{S} = 2, 2, 2, 2, 3, 3, 3$, and 4 is sufficiently large for lead times of $T = 2, 3, 4, 5, 6, 7, 8$, and 9, for which the state space consists of 150, 250, 375, 525, 2100, 3000, 4125, and 17,875 states, and the optimal policy can be found within 0.03, 0.04, 0.05, 0.08, 0.42, 0.99, 2.99, and 82.60 seconds, respectively. For a lead time of just 2 time units, it suffices to order one spare at a time, and thus the $(0, 1)$ policy performs well, while for lead times of 5 and 9 time units the $(1, 2)$ and the $(2, 3)$ policy performs well, respectively. In between, however, the best $(S-1, S)$ policy does not result in a close-to-optimal solution. Instead, significant cost savings can be obtained by basing both the maintenance decisions as well as the inventory decisions on the actual system state. Furthermore, the single component policy is

Fig. 8. Average cost per time unit for different policies for different values of F , along with the percentage increase in cost from applying the best (s, S) policy or the single component policy rather than the optimal policy.

Fig. 9. Average cost per time unit for different policies for different values of μ , along with the percentage increase in cost from applying the best $(S-1, S)$ policy or the single component policy rather than the optimal policy.

Fig. 10. Average cost per time unit for different policies for different values of T , along with the percentage increase in cost from applying the best $(S-1, S)$ policy or the single component policy rather than the optimal policy.

significantly more expensive than the optimal policy for all values of T considered. For a large lead time, the sharing of spares becomes more rewarding.

5.5. Non-identical failure rates

So far, we assumed that the two components are identical. Even though both components require identical spares, one of the components could deteriorate faster than the other, because they can operate in different environments or at different intensities. For that reason, we now consider two components that are deteriorating at different rates. To be able to compare our results to the base case, we decided to keep the demand for spares approximately unchanged. We vary the deterioration parameter of component 1 (μ_1) between 0.05 and 0.2, while we set the deterioration parameter of component 2 to $\mu_2 = 0.4 - \mu_1$. The results are shown in Fig. 11.

First, we observe that the single component policy outperforms the best $(S-1, S)$ policy if the components have different deterioration rates. For relatively low values of μ_1 (and thus for large differences between the components), the single component policy performs quite well, indicating that an optimization on a component level makes sense for different components. The $(S-1, S)$ policy results in much higher costs than our optimal policy for all values of μ_1 and μ_2 considered. The $(1, 2)$ policy is the best performing $(S-1, S)$ policy in case $\mu_1 < \mu_2$, but results in costs that are at least 10 percent higher than those of our optimal policy.

6. Varying the number of components

In Sections 4 and 5, we considered a system consisting of two components. Analyzing larger systems is more complex and computationally time consuming, as the condition of each component

Fig. 11. Average cost per time unit for different policies for different values of μ_1 (while $\mu_2 = 0.4 - \mu_1$), along with the percentage increase in cost from applying the best $(S-1, S)$ policy or the single component policy rather than the optimal policy.

Fig. 12. Average cost per time unit for different policies for different values of N , along with the percentage increase in cost from applying the best $(S-1, S)$ policy or the single component policy rather than the optimal policy.

needs to be tracked. Also the optimal policy is hard to represent graphically. Nevertheless, we are interested in extending our system to more than two components. For that reason, we again consider the system from the base case, with N components (between one and six) rather than two. The cost and deterioration parameters remain unchanged, as well as the lead time for spares. A maximum inventory position of $\bar{S} = 1, 2, 2, 3, 3$, and 4 is sufficiently large for $N = 1, 2, 3, 4, 5$, and 6, respectively, for which the state space consists of 20, 250, 1250, 12,500, 62,500, and 546,875 states, and our MDP formulation can be used to find the optimal policy in 0.0, 0.0, 0.2, 2.1, 13.3, and 220.3 seconds, respectively. The results are shown in Fig. 12. Compared with the single component policy, the optimal policy benefits more from the pooling of spares as N increases. The cost difference with this policy thus increases up

to 39 percent for $N = 6$. Furthermore, for $N = 1$, we observe that the $(S-1, S)$ policy performs best for $S = 1$. If S is chosen larger than 1, then too many spares are held on hand, resulting in higher holding costs. The cost difference between the $(1, 2)$ and the $(2, 3)$ policy, for example, is thus exactly equal to 0.5: the holding cost for one extra component. As N increases, the $(0, 1)$ policy is no longer the best performing $(S-1, S)$ policy, as too few spares are ordered, resulting in high downtime costs. The $(1, 2)$ policy thus performs best for N between 2 and 4. For N larger than 4, however, the $(1, 2)$ policy provides too few spares, and the $(2, 3)$ policy performs better. Intuitively, Fig. 12 illustrates that the cost difference between the optimal policy and the best performing $(S-1, S)$ policy will always be between 0 and 0.5 (the holding cost for one extra spare). Indeed, we observed in Section 4.4 that the cost

Fig. 13. Average cost gap between the optimal policy and the best $(S-1, S)$ policy and the single component policy from 20 uniform problem instances.

difference is mainly caused by the different inventory actions (rather than the replacement actions). Naturally, as N increases, the corresponding minimal cost will also increase, thus reducing the relative cost difference with the $(S-1, S)$ policy. The cost gap between the optimal policy and the best $(S-1, S)$ policy depicted in Fig. 12b shows non-monotonic behavior, which is somewhat unexpected but can be explained by the integer decision variable S that may lead to non-monotonic results for specific parameter settings. To show that this behavior is indeed instance specific, we next present the average cost gap over 20 randomly selected instances, drawing key parameters with equal probabilities from the following sets: replacement cost c_r : $(2, 3, \dots, 8)$, fixed cost per order F : $(0.0, 0.5, 1.0, 1.5)$, holding cost H : $(0.0, 0.1, \dots, 1.0)$, fixed lead time T : $(2, 3, 4)$, and expected lifetime E : $(3, 4, \dots, 7)$. The resulting parameters are shown in Table B.1 in Appendix B. For each instance, we determine the cost corresponding to the optimal policy, the single component policy, and the $(S-1, S)$ policy for different values of S . Fig. 13 shows the resulting average cost increase (in percentage) of the 20 problem instances from applying either the single component policy or the best $(S-1, S)$ policy rather than the optimal policy. Indeed, we observe that the average cost gap with the best $(S-1, S)$ policy decreases as the number of components increases, although it does not drop below 5 percent. In addition, we again observe that the pooling of spares becomes more beneficial as the number of components increases. The average cost gap with the single component policy increases to 33 percent.

7. Conclusion

This paper is the first to consider the joint optimization of condition-based maintenance and a condition-based spares inventory for a multi-component system. For a single-component system, it is optimal to use deterioration thresholds to determine when a spare should be ordered and when a replacement should be performed. For a multi-component system, however, such a monotonic policy structure is not necessarily optimal. The components share a pool of spares, so optimal decisions at the component level need not be optimal at the system level.

Through an exact numerical analysis, we first obtained several structural insights for a system consisting of two components. We found that a monotonic policy structure is indeed not necessarily optimal for a multi-component system. Instead, the maintenance and inventory decisions should both be based on all available infor-

mation, i.e., on the state of each component, the number of spares on hand, and the size and expected arrival time of each outstanding order. We observe that basing the inventory decisions on this complete system state rather than applying a more standard inventory policy, such as the (s, S) strategy, can reduce costs significantly. In particular, it can be beneficial to reserve the last spare if both components are close to failure, and to delay an order for spares if both components are in good condition. In a sensitivity analysis, we varied several cost parameters (such as the replacement cost, the holding cost, and the fixed cost per order), the expected lifetime, and the lead time for spares. The main results are robust for a variety of parameter settings. In addition, we considered two components that deteriorate at different rates, and we varied the number of components between one and six. Results indicate, for example, that the differences in costs between different policies (the optimal policy, the (s, S) inventory policy, and a separate optimization for each component) are mainly caused by the different ordering decisions rather than the replacement decisions. A separate optimization thus makes more sense than an (s, S) policy for slowly deteriorating components, as the latter is not able to postpone the ordering of spares. Also for systems with more than two components significant cost savings can be obtained by basing both the replacement and ordering decisions on the complete system state, as the (s, S) policy will result in higher inventory-related costs by ordering spares too early.

If a company decides to invest in monitoring equipment to increase the efficiency of their maintenance strategy, we believe that the obtained condition information should also be used to apply a more cost efficient inventory policy. Standard inventory policies order spares although the system may still be in a very good state, leading to unnecessary holding costs. Implementing a condition-based inventory policy can reduce holding costs substantially, whilst maintaining a high level of availability, by only ordering spares for somewhat deteriorated components.

In this paper, we assumed that the system is monitored through periodic inspections. In order to investigate the benefits of continuous, real-time monitoring versus periodic inspections, future research could consider condition-based, or just-in-time, ordering of spares for situations with continuous monitoring. Furthermore, with the cost structure proposed in this paper, it is relatively easy to include economic dependencies, in the form of a fixed set-up cost for maintenance. Other possible directions for future research consider uncertain lead times, which often occur in practice, or imperfect information. The latter occurs when an inspection does not reveal the actual state of each component, but merely reveals the probability of being in a certain state.

Acknowledgments

This project was co-funded by the Dutch Institute for Advanced Logistics (Dinalog). We also thank Gerlach van der Heide for his valuable contribution in speeding up the MDP algorithm.

Appendix A. Optimal maintenance and inventory policy for $H = 10$

Fig. A.1 depicts the optimal condition-based maintenance and inventory policy for the case with holding costs equal to $H = 10$ per spare per time unit (see Section 5.2).

Appendix B. Randomly selected problem instances

Table B.1 shows the realizations of the replacement cost c_r , fixed cost per order F , holding cost H , lead time T , and expected lifetime E for the problem instances analyzed in Section 6.

Fig. A.1. Optimal maintenance and inventory policy for $N = 2$, $L = 4$, $O = (0, 0, 0, 0, 100)$, $R = (5, 5, 5, 5, 5)$, $F = 0$, $H = 10$, $T = 3$, and $\mu = 0.2$.

Table B1

Parameters of the randomly selected problem instances.

	Instance number																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
c_r	2	5	4	3	6	3	5	8	8	2	4	2	5	8	5	5	7	6	5	5
F	0.0	2.0	0.5	0.5	0.5	0.0	1.0	2.0	2.0	0.5	0.0	1.5	0.0	1.0	0.0	0.0	1.5	1.5	0.0	1.5
H	0.2	0.8	0.6	0.3	0.4	0.7	0.1	0.4	0.9	0.8	0.5	0.4	1.0	0.2	0.4	0.2	0.7	0.2	0.7	0.6
T	2	4	4	2	3	2	2	3	3	2	3	3	2	3	2	3	3	2	3	4
E	4	5	4	4	7	6	3	5	6	3	4	4	7	5	4	3	5	7	3	4

References

- Ahmad, R., & Kamaruddin, S. (2012). An overview of time-based and condition-based maintenance in industrial application. *Computers and Industrial Engineering*, 63(1), 135–149. doi:10.1016/j.cie.2012.02.002.
- Alsyouf, I. (2007). The role of maintenance in improving companies productivity and profitability. *International Journal of Production Economics*, 105(1), 70–78. doi:10.1016/j.ijpe.2004.06.057.
- Basten, R., & van Houtum, G. (2014). System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science*, 19(1), 34–55. doi:10.1016/j.sorms.2014.05.002.
- Borrero, J., & Akhavan-Tabatabaei, R. (2013). Time and inventory dependent optimal maintenance policies for single machine workstations: An MDP approach. *European Journal of Operational Research*, 228(3), 545–555. doi:10.1016/j.ejor.2013.02.011.
- Bousdekis, A., Magoutas, B., Apostolou, D., & Mentzas, G. (2015). A proactive decision making framework for condition-based maintenance. *Industrial Management and Data Systems*, 115(7), 1225–1250. doi:10.1108/IMDS-03-2015-0071.
- Bouvard, K., Artus, S., Béranger, C., & Cocquempot, V. (2011). Condition-based dynamic maintenance operations planning & grouping. Application to commercial heavy vehicles. *Reliability Engineering and System Safety*, 96(6), 601–610. doi:10.1016/j.res.2010.11.009.
- Byon, E., Ntamo, L., & Ding, Y. (2010). Optimal maintenance strategies for wind turbine systems under stochastic weather conditions. *IEEE Transactions on Reliability*, 59(2), 393–404. doi:10.1109/TR.2010.2046804.
- Cho, D., & Parlar, M. (1991). A survey of maintenance models for multi-unit systems. *European Journal of Operational Research*, 51(1), 1–23. doi:10.1016/0377-2217(91)90141-H.
- Cohen, M., Kleindorfer, P., Lee, H., & Pyke, D. (1992). Multi-item service constrained (s, S) policies for spare parts logistics systems. *Naval Research Logistics*, 39(4), 561–577. doi:10.1002/1520-6750(199206)39:4<561::AID-NAV3220390409>3.0.CO;2-5.
- Dekker, R., Wildeman, R., & van der Duyn Schouten, F. (1997). A review of multi-component maintenance models with economic dependence. *Mathematical Methods of Operations Research*, 45(3), 411–435. doi:10.1007/BF01194788.
- Van der Duyn Schouten, F. (1996). Maintenance policies for multicomponent systems: An overview. In *Reliability and maintenance of complex systems*. In NATO ASI series: Vol. 154 (pp. 117–136). Berlin Heidelberg: Springer. doi:10.1007/978-3-662-03274-9_8.
- Elwany, A., & Gebraeel, N. (2008). Sensor-driven prognostic models for equipment replacement and spare parts inventory. *IEEE Transactions*, 40(7), 629–639. doi:10.1080/07408170701730818.
- Elwany, A., Gebraeel, N., & Maillart, L. (2011). Structured replacement policies for components with complex degradation processes and dedicated sensors. *Operations Research*, 59(3), 684–695. doi:10.1287/opre.1110.0912.
- Gertsbakh, I. (1977). *Models of preventive maintenance*. North-Holland.
- Gertsbakh, I. (2000). *Reliability theory: With applications to preventive maintenance*. Springer.
- Guajardo, M., Rönnqvist, M., Halvorsen, A., & Kallevik, S. (2015). Inventory management of spare parts in an energy company. *Journal of the Operational Research Society*, 66(2), 331–341. doi:10.1057/jors.2014.8.

- Hong, H., Zhou, W., Zhang, S., & Ye, W. (2014). Optimal condition-based maintenance decisions for systems with dependent stochastic degradation of components. *Reliability Engineering and System Safety*, 121, 276–288. doi:10.1016/j.res.2013.09.004.
- van Houtum, G., & Kranenburg, B. (2015). Spare parts inventory control under system availability constraints. *International series in operations research and management science*: Vol. 227. US: Springer. doi:10.1007/978-1-4899-7609-3.
- Iglehart, D. (1963). Optimality of (s,S) policies in the infinite horizon dynamic inventory problem. *Management Science*, 9(2), 259–267. doi:10.1287/mnsc.9.2.259.
- van Jaarsveld, W., Dollevoet, T., & Dekker, R. (2015). Improving spare parts inventory control at a repair shop. *Omega*, 57(B), 217–229. doi:10.1016/j.omega.2015.05.002.
- Jiang, Y., Chen, M., & Zhou, D. (2015). Joint optimization of preventive maintenance and inventory policies for multi-unit systems subject to deteriorating spare part inventory. *Journal of Manufacturing Systems*, 35, 191–205. doi:10.1016/j.jmsy.2015.01.002.
- de Jonge, B., Klingenberg, W., Teunter, R., & Tinga, T. (2015). Optimum maintenance strategy under uncertainty in the lifetime distribution. *Reliability Engineering and System Safety*, 133(0), 59–67. doi:10.1016/j.res.2014.09.013.
- Karsten, F., & Basten, R. (2014). Pooling of spare parts between multiple users: How to share the benefits? *European Journal of Operational Research*, 233(1), 94–104. doi:10.1016/j.ejor.2013.08.029.
- Kawai, H. (1983). An optimal ordering and replacement policy of a Markovian degradation system under complete observation: Part I. *Journal of the Operations Research Society of Japan*, 26(4), 279–291.
- Kawai, H. (1983). An optimal ordering and replacement policy of a Markovian deterioration system under incomplete observation: Part II. *Journal of the Operations Research Society of Japan*, 26(4), 293–307.
- Kennedy, W., Wayne Patterson, J., & Fredendall, L. (2002). An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 76(2), 201–215. doi:10.1016/S0925-5273(01)00174-8.
- Kranenburg, A., & van Houtum, G. (2009). A new partial pooling structure for spare parts networks. *European Journal of Operational Research*, 199(3), 908–921. doi:10.1016/j.ejor.2009.01.057.
- Li, H., Deloux, E., & Dieulle, L. (2016). A condition-based maintenance policy for multi-component systems with Lévy copulas dependence. *Reliability Engineering and System Safety*, 149, 44–55. doi:10.1016/j.res.2015.12.011.
- Louit, D., Pascual, R., Banjevic, D., & Jardine, A. (2011). Condition-based spares ordering for critical components. *Mechanical Systems and Signal Processing*, 25(5), 1837–1848. doi:10.1016/j.ymssp.2011.01.004.
- Marseguerra, M., Zio, E., & Podofillini, L. (2002). Condition-based maintenance optimization by means of genetic algorithms and Monte Carlo simulation. *Reliability Engineering and System Safety*, 77(2), 151–165. doi:10.1016/S0951-8320(02)00043-1.
- van Noortwijk, J. (2009). A survey of the application of gamma processes in maintenance. *Reliability Engineering and System Safety*, 94(1), 2–21. doi:10.1016/j.res.2007.03.019.
- Olde Keizer, M., Teunter, R., & Veldman, J. (2016). Clustering condition-based maintenance for systems with redundancy and economic dependencies. *European Journal of Operational Research*, 251(2), 531–540. doi:10.1016/j.ejor.2015.11.008.
- Park, M., & Pham, H. (2012). 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), 453–463. doi:10.1109/TSMCA.2011.2162499.
- Pierskalla, W., & Voelker, J. (1976). A survey of maintenance models: the control and surveillance of deterioration systems. *Naval Research Logistics Quarterly*, 23, 353–388. doi:10.1002/nav.3800230302.
- Puterman, M. (1994). *Markov decision processes: Discrete stochastic dynamic programming*. New York: John Wiley & Sons.
- Rao, B. (1996). *Handbook of condition monitoring*. Elsevier Advanced Technology.
- Rasmekomen, N., & Parlikad, A. (2016). Condition-based maintenance of multi-component systems with degradation state-rate interactions. *Reliability Engineering and System Safety*, 148, 1–10. doi:10.1016/j.res.2015.11.010.
- Rausch, M., & Liao, H. (2010). Joint production and spare part inventory control strategy driven by condition based maintenance. *IEEE Transactions on Reliability*, 59(3), 507–516. doi:10.1109/TR.2010.2055917.
- Sahin, I. (1990). *Regenerative inventory systems*. New York, NY: Springer-Verlag.
- Scarf, H. (1959). The optimality of (s,S) policies in the dynamic inventory problem. *Mathematical methods in the social sciences* (pp. 196–202). Stanford University Press. *Proceedings of the first Stanford symposium*
- Strijbosch, L., Heuts, R., & van der Schoot, E. (2000). A combined forecast-inventory control procedure for spare parts. *Journal of the Operational Research Society*, 51(10), 1184–1192. doi:10.1057/palgrave.jors.2601013.
- Svoronos, A., & Zipkin, P. (1991). Evaluation of one-for-one replenishment policies for multiechelon inventory systems. *Management Science*, 37(1), 68–83. doi:10.1287/mnsc.37.1.68.
- Syntetos, A., Teunter, R., Babai, M., & Transchel, S. (2016). On the benefits of delayed ordering. *European Journal of Operational Research*, 248(3), 963–970. doi:10.1016/j.ejor.2015.08.003.
- Teunter, R., & Duncan, L. (2009). Forecasting intermittent demand: A comparative study. *Journal of the Operational Research Society*, 60(3), 321–329. doi:10.1057/palgrave.jors.2602569.
- Thomas, L. (1986). A survey of maintenance and replacement models for maintainability and reliability of multi-item systems. *Reliability Engineering*, 16(4), 297–309. doi:10.1016/0143-8174(86)90099-5.
- Tijms, H. (1994). *Stochastic models: An algorithmic approach*. Chichester: John Wiley & Sons.
- Tsang, A. (1995). Condition based maintenance: tools and decision making. *Journal of Quality in Maintenance Engineering*, 1(3), 3–17. doi:10.1108/13552519510096350.
- Uluks, M., Kharoufeh, J., & Maillart, L. (2012). Optimal replacement policies under environment-driven degradation. *Probability in the Engineering and Informational Sciences*, 26, 405–424. doi:10.1017/S0269964812000083.
- Van Horenbeek, A., Buré, J., Cattrysse, D., Pintelon, L., & Vansteenwegen, P. (2013). Joint maintenance and inventory optimization systems: A review. *International Journal of Production Economics*, 143(2), 499–508. doi:10.1016/j.ijpe.2012.04.001.
- Van Horenbeek, A., & Pintelon, L. (2015). A joint predictive maintenance and inventory policy. In P. Tse, J. Mathew, K. Wong, R. Lam, & C. Ko (Eds.), *Engineering asset management - systems, professional practices and certification*. In Lecture notes in mechanical engineering (pp. 387–399). Springer International Publishing. doi:10.1007/978-3-319-09507-3_34.
- Van Horenbeek, A., Scarf, P., Cavalcante, C., & Pintelon, L. (2013). The effect of maintenance quality on spare parts inventory for a fleet of assets. *IEEE Transactions on Reliability*, 62(3), 596–607. doi:10.1109/TR.2013.2270409.
- Veldman, J., Wortmann, H., & Klingenberg, W. (2011). Typology of condition based maintenance. *Journal of Quality in Maintenance Engineering*, 17(2), 183–202. doi:10.1108/13552511111134600.
- Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3), 469–489. doi:10.1016/S0377-2217(01)00197-7.
- Wang, L., Chu, J., & Mao, W. (2008a). A condition-based order-replacement policy for a single-unit system. *Applied Mathematical Modelling*, 32(11), 2274–2289. doi:10.1016/j.apm.2007.07.016.
- Wang, L., Chu, J., & Mao, W. (2008b). An optimum condition-based replacement and spare provisioning policy based on Markov chains. *Journal of Quality in Maintenance Engineering*, 14(4), 387–401. doi:10.1108/13552510810909984.
- Wang, L., Chu, J., & Mao, W. (2009). A condition-based replacement and spare provisioning policy for deteriorating systems with uncertain deterioration to failure. *European Journal of Operational Research*, 194(1), 184–205. doi:10.1016/j.ejor.2007.12.012.
- Williams, T. (1984). Stock control with sporadic and slow-moving demand. *Journal of the Operational Research Society*, 35(10), 939–948. doi:10.1057/jors.1984.185.
- Xie, J., & Wang, H. (2008). Joint optimization of condition-based preventive maintenance and spare ordering policy. In *Proceedings of the fourth international conference on wireless communications, networking and mobile computing* (pp. 1–5). IEEE. doi:10.1109/WiCom.2008.1468.
- Zhao, J., & Xu, C. (2012). A joint policy for condition-based maintenance and spare provisioning using simulation. In *Proceedings of the IEEE conference on prognostics and system health management (PHM)* (pp. 1–7). doi:10.1109/PHM.2012.6228941.
- Zohrul Kabir, A., & Al-Olayan, A. (1996). A stocking policy for spare part provisioning under age based preventive replacement. *European Journal of Operational Research*, 90(1), 171–181. doi:10.1016/0377-2217(94)00246-0.