

Condition based spare parts supply

Citation for published version (APA):

Lin, X., Basten, R. J. I., Kranenburg, A. A., & van Houtum, G. J. J. A. N. (2017). Condition based spare parts supply. *Reliability Engineering and System Safety*, 168, 240-248. <https://doi.org/10.1016/j.ress.2017.05.035>

DOI:

[10.1016/j.ress.2017.05.035](https://doi.org/10.1016/j.ress.2017.05.035)

Document status and date:

Published: 01/12/2017

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Condition based spare parts supply

X. Lin^a, R.J.I. Basten^{b,*}, A.A. Kranenburg^c, G.J. van Houtum^b

^a*Dow Benelux, P.O. Box 48, 4530 AA, Terneuzen, The Netherlands*

^b*Eindhoven University of Technology, P.O. Box 513, 5600 MB, Eindhoven, The Netherlands*

^c*Consultants in Quantitative Methods, P.O. Box 414, 5600 AK, Eindhoven, The Netherlands*

Abstract

We consider a spare parts stock point that serves an installed base of machines. Each machine contains the same critical component, whose degradation behavior is described by a Markov process. We consider condition based spare parts supply, and show that an optimal, condition based inventory policy is 20% more efficient on average than a standard, state-independent base stock policy. We further propose an efficient and effective heuristic policy.

Keywords: Inventory, spare parts, condition monitoring

1. Introduction

Capital goods, such as lithography equipment used in the semiconductor industry, CT scanners that are used in hospitals, or radar systems on board naval vessels, are expensive, technologically complex systems that are used in the primary processes of their users. As a result, their uptime is of utmost importance; each minute of unavailability may be costly, risky, or both. Spare parts are stocked to prevent downtime: upon failure, a defective component can be replaced quickly by a functioning spare part. It is therefore important to have enough stock on hand. However, spare parts are expensive, which means that stocking too many spare parts is costly. Since making this trade-off poses a challenging problem, there has been a lot of research on spare parts inventory control (see, e.g., Basten and Van Houtum, 2014).

The costs of the spare parts inventories may be reduced by using information on the condition of the components that are installed in the installed base. To this end, we consider a number of machines, each containing the same one critical component that degrades over time. The degradation evolves according to a Markov chain with a finite state space, with at most one state transition per period (see, e.g., Giorgio et al., 2011; Neves et al., 2011; Si et al., 2011, for an

*Corresponding author: T: +31 40 247 2384

Email addresses: xlin2@dow.com (X. Lin), r.j.i.basten@tue.nl (R.J.I. Basten), bram.kranenburg@cqm.nl (A.A. Kranenburg), g.j.v.houtum@tue.nl (G.J. van Houtum)

understanding of how to model degradation using a Markov chain, including the determination of the transition probabilities). The condition is monitored perfectly at the beginning of each period. Since there is at most one transition per period, a component can fail only in a certain period if it is in the last degradation state at the beginning of that period. (This simplification allows us to focus on the key insights; in practice, there may be other failure modes that lead to failure of a component that is in perfect condition.) Upon failure, the component is replaced immediately by a functioning spare part. One stock point is used to stock these spare parts and the base stock level in each period is dependent on the condition of the installed components and on the complete inventory status (stock on hand plus exact position of each outstanding order). If the stock point has no stock on hand when a demand arrives, an emergency procedure is used to obtain the part from a source with ample supply. For this, emergency costs are paid. The other costs that we consider are inventory holding costs.

We model this problem as a discrete-time Markov decision problem (MDP) and we obtain the optimal policy using value iteration (see, e.g., Tijms, 1986). This is very time consuming, especially if the number of degradation states, the lead time, or the number of machines is high. Therefore, we propose three heuristic policies that are easy to compute, and we show that the third policy is close to optimal. In an extensive numerical experiment, we find that the optimal policy, which by definition is a state-dependent base stock policy, achieves average cost savings of 20% compared with a state-independent base stock policy. These savings increase with the precision with which the degradation behavior can be tracked. Interestingly, the possible savings decrease if the size of the installed base increases. This is probably because in a larger installed base, there are virtually always some components that are very new and some components that are close to failure. In other words, effects level out.

Our main contribution is that we study the effect of condition information on spare parts supply *without* changing the maintenance policy. We show that large savings can be obtained, we identify under which circumstances the savings are largest, and we derive an efficient and effective heuristic policy. Our research is especially relevant for situations where preventive replacements are undesired because of the loss of a significant part of the useful lifetime of components, or if preventive replacements are (almost) equally expensive as corrective replacements (e.g., in process manufacturing, operating 24/7). We are aware of only a few papers that are closely related to our work; we explain the differences with our work in Section 2.

The remainder of this paper is organised as follows. We discuss the related literature in Section 2. In Section 3, we introduce our model, and in Section 4 we discuss the resulting

Markov decision process. Next, in Section 5, we discuss the optimal base stock policy, and we discuss the heuristic policies that we propose in Section 6. We then perform an extensive numerical experiment in Section 7. Finally, we draw conclusions in Section 8.

2. Related literature

The relevant literature on spare parts inventory control has started with the paper by Feeney and Sherbrooke (1966). This has led to a huge stream of research on all kinds of spare parts inventory systems. For an overview, we refer to the books by Sherbrooke (2004) and Muckstadt (2005) or the review by Basten and Van Houtum (2014).

In a large part of the literature on spare parts inventory control, one assumes backordering of demands that cannot be fulfilled immediately. In our model, however, we assume that such demands are fulfilled by an emergency source and they can then be seen as lost demands for the stock point under consideration. A recent overview of the literature on inventory control with lost sales, not necessarily considering spare parts, is given by Bijvank and Vis (2011). We discuss three papers in more detail, all considering a discrete-time inventory model with lost sales: Bijvank and Johansen (2012) and Zipkin (2008a,b). Bijvank and Johansen (2012) discuss, among other things, a so-called restricted base stock policy. This is a regular base stock policy, but with a maximum on the order size. The reason to propose this policy is that the authors often find that the “PBSP [pure base stock policy] and the optimal policy coincide in numerous states of the Markov chain” (Bijvank and Johansen, 2012, p.109).

The first heuristic policy that we propose is also a base stock policy with a maximum, although we use a different way of deriving this maximum (see Section 6.2). Zipkin (2008a) discusses various heuristic policies, one of which is the myopic policy (based on Morton, 1971). Our second heuristic policy is also a myopic policy, but it is different since we have to make some approximations to cope with our complex demand process (see Section 6.3).

In both papers, Zipkin assumes that demands in consecutive periods are independent. However, in Zipkin (2008b) he mentions an extension to Markov-modulated demands, resulting in a state-dependent inventory policy. He explains that the state space of possible supply orders is bounded. In our paper, the demand process follows a specific Markov-chain-driven counting process (with demands connected to transitions) whose structure is explicitly exploited in the analysis and the derivation of the heuristic policies.

There is also literature on varying demand rates and state-dependent inventory policies in models with backlogging. For example, Song and Zipkin (1993) consider a single stock point

that faces demand that follows a Markov-modulated Poisson process. Considering continuous review, holding costs for inventory on hand and penalty costs for backorders, the authors show that the optimal policy is a base stock policy. Although the demand process at each point in time is dependent on an underlying Markov chain, there is no direct link with the state of the components in the installed base.

Another stream of research on state-dependent inventory policies uses advance demand information (ADI). In most of the literature, ADI means that customers place orders that will lead to an actual demand only after a certain demand lead time. The seminal paper in this stream of research is the paper of Hariharan and Zipkin (1995). The authors consider both a single location system and a serial system. In both cases they assume a continuous review, base stock policy with full backordering. Replenishment orders are triggered by the customers' orders, which result in actual demands after a certain demand lead time, thus making perfect ADI. ADI may also be imperfect. For example, Topan et al. (2016) consider three types of imperfectness: The demand lead time may be stochastic, a demand that is preceded by ADI may not materialize, and a demand may materialize that is not preceded by ADI. Topan et al. assume a single stock point with periodic review and lost sales if a demand cannot be fulfilled from stock. They give the setting of spare parts inventory control and condition monitoring as one example where their model applies.

A key difference with the work of Topan et al. (2016) is that in our model, we explicitly model the degradation behavior of the critical components and we derive the imperfect ADI from that behavior. A related paper is that of Deshpande et al. (2006). The authors assume that a part-age signal can be observed each period, which is then compared with a certain threshold value. Depending on the number of parts that have a signal above the threshold value, the authors calculate a conditional mean and variance of a normally distributed lead time demand. These are used to set the base stock level, assuming holding costs per unit on hand and backorder costs per backorder.

Finally, as mentioned in the introduction (Section 1), our paper is related to the stream of literature on condition based maintenance (CBM). In particular, the delay time model is of interest, as introduced by Christer (1982). In this model, if a component becomes defective (which is not self-announcing) there is a certain delay time after which the component actually fails. This makes it worthwhile to perform inspections to observe the condition of the component. For an overview of the literature on CBM, including a review of diagnostics and prognostics techniques, see Jardine et al. (2006). Two more recent overviews of the literature on

CBM are those by Alaswad and Xiang (2017) and Olde Keizer et al. (2017). Within the stream of literature on CBM, also the reducing effect of CBM on spare parts supply costs has been studied; see, for example, Bjarnason et al. (2014), Elwany and Gebraeel (2008), Van Horenbeek and Pintelon (2015), Rausch and Liao (2010), Wang et al. (2008), Wang et al. (2009), Wang (2011), Wang et al. (2015a) or Wang et al. (2015b). Van Horenbeek et al. (2013) review the literature on joint optimization of spare parts inventory control and maintenance (not necessarily CBM).

As already stated, we distinguish ourselves from the latter studies by considering the effect of conditioning monitoring on spare parts supply *without* changing the maintenance policy. We are aware of two papers with the same focus. The first paper is that of Lout et al. (2011), who assume a single system for which at most one spare part is kept on stock. The authors further assume backordering when a spare part is demanded but not available. In contrast, we consider an arbitrary number of systems, allow any spare parts inventory level, and assume that an emergency shipment is executed in an out-of-stock situation. This also implies that we have a different cost structure. The second paper is that of Li and Ryan (2011). They model deterioration of each part as a Wiener process and use that to estimate the distribution of the remaining useful life of each part. This estimate is updated each period using Bayesian updating and it is used to estimate the distribution of the demand for spare parts in the upcoming periods. This type of modelling of degradation is a key difference with our work. Another difference is that Li and Ryan assume a zero replenishment lead time.

3. Model description

We consider a group of N ($\in \mathbb{N}$) identical machines, each containing one critical component. The component is subject to a degradation process on a finite state space $\mathcal{I}' = \{0, \dots, I\}$, with state 0 representing the perfect working condition and state I representing failure. In practice, the states may have meaningful interpretations, such as ‘Good’, ‘Minor defects only’, ‘Maintenance required’, and ‘Down’. Alternatively, they may correspond to consecutive intervals of a continuous degradation parameter (see also Giorgio et al., 2011; Neves et al., 2011; Si et al., 2011). Time is divided into periods of unit length and we assume an infinite horizon. We assume that the length of a period is short compared to the average lifetime of the critical component (e.g., the period length may be one week, while the lifetime of the critical component is in the range of 1 to 10 years). It is then reasonable to assume that a component degrades at most one state per period. The state transition can occur at any time during the period, and a transition

to state I (failure) is self-announcing. A failed component is replaced in negligible time by a functioning spare part.

This spare part is demanded from one local stock point. If this local stock point has no stock on hand when a demand arrives, an emergency procedure is used and the local stock point faces a lost sale. Using the emergency procedure leads to additional costs of c_e (> 0), which may include some downtime costs and the higher costs (compared to a normal replenishment) that have to be made to achieve a short emergency lead time. We assume that the failed part is still replaced in the same period in which it failed. The stock point can order new components that arrive after a deterministic replenishment lead time of L ($\in \mathbb{N}$) periods. Notice that adding variable ordering costs per part ordered does not influence the optimal ordering policy, since the total number of parts to order (through the regular plus emergency procedures) is not influenced by this policy. There are no fixed ordering costs.

At the beginning of each period t , we consider the following sequence of events:

1. Spare parts in the pipeline come one period closer; items that were ordered L periods earlier arrive at the stock point.
2. The state of each critical component is observed. Since a failed component is replaced before the beginning of the next period, a component will never be in state I at the beginning of a period. We therefore introduce the state space $\mathcal{I} = \mathcal{I}' \setminus I$ of states that can be observed at the beginning of a period.
3. A replenishment order may be placed.
4. Holding costs c_h (> 0) are paid for the complete inventory position, so for components on hand and in the replenishment pipeline.

If one would have a modified problem in which holding costs are only paid for the stock on hand, then there is a one-to-one coupling with our problem; see Remark 3.1. A similar coupling exists when holding costs would be incurred at the end of each period; see Remark 3.2.

We define $q_{i,j}$ to be the transition probability of one component's state $i \in \mathcal{I}$ to $j \in \mathcal{I}$. For ease of notation, we define $q_{i,I} \equiv q_{i,0}$ and we refer to state 0 as the subsequent state of state $I - 1$. Then, for all $i \in \mathcal{I}$ it holds that $q_{i,i} \geq 0$, $q_{i,i+1} > 0$, $q_{i,i} + q_{i,i+1} = 1$, and $q_{i,j} = 0$ for all $j \notin \{i, i+1\}$ (notice that $q_{I-1,0} = q_{I-1,I} > 0$). In order to avoid technical complications (related to possible periodicity of the Markov chain per component), we assume that for at least one $i \in \mathcal{I}$, it holds that $q_{i,i} > 0$ (*strict inequality*). Notice that $1/(1 - q_{i,i}) = 1/q_{i,i+1}$ represents the average number of periods that a component stays in state i and that the average lifetime of the critical component is equal to $\sum_{i \in \mathcal{I}} 1/q_{i,i+1}$.

We are interested in finding the inventory policy that minimizes the (undiscounted) average costs per period for an infinite planning horizon. Since the inventory policy does not influence the total number of spare parts that will be used (remember that c_e is defined to be *additional* costs), we may ignore the variable ordering costs per component.

Remark 3.1. Suppose that one has the same problem as our problem, except that holding costs are only paid for stock on hand. Let \hat{c}_h and \hat{c}_e be the holding costs and emergency costs for the modified problem. Then this modified problem is coupled to our problem formulation with $c_h = \hat{c}_h$ and $c_e = \hat{c}_e + c_h L$. Under each policy, in our problem formulation, one then pays exactly $c_h L$ extra for each demanded part, both when satisfied from stock and when satisfied via the emergency procedure. Hence, both problems have the same optimal policy and the optimal costs differ by a constant factor $c_h L N / [\sum_{i \in \mathcal{I}} 1/q_{i,i+1}]$.

Remark 3.2. Suppose that one has a modified problem in which holding costs are incurred at the end of a period. Then the coupling with our problem is as follows. Let \hat{c}_h and \hat{c}_e be the holding costs and emergency costs for the modified problem. This modified problem is coupled to our problem formulation with $c_h = \hat{c}_h$ and $c_e = \hat{c}_e + c_h$. Under each policy, in our problem formulation, one then pays exactly c_h extra for each demanded part, both when satisfied from stock and when satisfied via the emergency procedure. Hence, both problems have the same optimal policy and the optimal costs differ by a constant factor $c_h N / [\sum_{i \in \mathcal{I}} 1/q_{i,i+1}]$. Notice that the modification here and the one of Remark 3.1 can also be combined.

Remark 3.3. Our model may alternatively be used in the following situation. In practice, there is often a certain degradation threshold defined that, when reached, triggers preventive replacement of a component. In our model, this means reaching state I . If a spare part is available, the preventive maintenance can be performed immediately. Otherwise, there is no failure, but still an emergency shipment may be requested to preventively replace the part as soon as possible. This emergency shipment will not be as expensive as the ‘normal’ emergency shipment that we consider in our model.

4. Markov decision process

The behavior of the system that we presented in Section 3 can be described by a discrete-time Markov decision process (MDP). We define its states, describe the system transitions between states, and finally show the resulting costs.

We model an MDP with state space $\mathcal{S} = \{(\mathbf{m}, \mathbf{s})\}$, with:

- The degradation vector $\mathbf{m} = (m_0, \dots, m_{I-1})$ and m_i denotes the number of parts in degradation state $i \in \mathcal{I}$. It holds that $\sum_{i \in \mathcal{I}} m_i = N$.
- The inventory vector $\mathbf{s} = (s_0, \dots, s_{L-1})$, s_0 denotes the stock on hand, and s_l (for $l \in \{1, \dots, L-1\}$) denotes the number of parts that arrive in l periods.

Notice that these states describe the situation before a new replenishment order is placed. Therefore, there is no s_L in our state description.

The action space of possible actions that can be taken in state (\mathbf{m}, \mathbf{s}) is denoted by $\mathcal{A}'(\mathbf{m}, \mathbf{s}) = \mathbb{N}_0 = \mathbb{N} \cup 0$. Notice that $a \in \mathcal{A}'(\mathbf{m}, \mathbf{s})$ represents the number of spare parts to order in the current period. In Section 5, we introduce the action space $\mathcal{A}(\mathbf{m}, \mathbf{s})$ of actions that can be taken in an optimal policy, and we show that $\mathcal{A}(\mathbf{m}, \mathbf{s})$ has a finite number of elements.

We define the transition vector $\mathbf{d} = (d_0, \dots, d_{I-1})$ with d_i denoting the number of parts degrading in a certain period from state $i \in \mathcal{I}$ to its subsequent state. We define the set $\mathcal{D}(\mathbf{m}) = \{\mathbf{d} \mid 0 \leq d_i \leq m_i, \forall i \in \mathcal{I}\}$ as the set of all possible transition vectors, given the current degradation vector \mathbf{m} . We use $\tilde{q}(\mathbf{m}, \mathbf{d})$ to denote the probability of observing transition vector $\mathbf{d} \in \mathcal{D}(\mathbf{m})$ in a period if the system state with respect to the status of the installed components at the beginning of that period is \mathbf{m} :

$$\tilde{q}(\mathbf{m}, \mathbf{d}) = \prod_{i \in \mathcal{I}} \binom{m_i}{d_i} (q_{i,i+1})^{d_i} (1 - q_{i,i+1})^{m_i - d_i}.$$

The subsequent degradation vector, given the current degradation vector \mathbf{m} and the transition vector \mathbf{d} is given by $f(\mathbf{m}, \mathbf{d}) = (f_0(\mathbf{m}, \mathbf{d}), \dots, f_{I-1}(\mathbf{m}, \mathbf{d}))$, with:

$$f_i(\mathbf{m}, \mathbf{d}) = \begin{cases} m_0 - d_0 + d_{I-1} & \text{if } i = 0; \\ m_i - d_i + d_{i-1} & \text{if } i \in \mathcal{I} \setminus 0. \end{cases}$$

We are now ready to define the subsequent inventory vector, given the current inventory vector \mathbf{s} , the transition vector \mathbf{d} , and the action $a \in \mathcal{A}'(\mathbf{m}, \mathbf{s})$ (notice that the decision on a is taken before \mathbf{d} is observed) as $g^a(\mathbf{s}, \mathbf{d}) = (g_0^a(\mathbf{s}, \mathbf{d}), \dots, g_{L-1}^a(\mathbf{s}, \mathbf{d}))$, with:

$$g_l^a(\mathbf{s}, \mathbf{d}) = \begin{cases} (s_0 - d_{I-1})^+ & \text{if } l = 0; \\ s_{l+1} & \text{if } l \in \{1, 2, \dots, L-2\}; \\ a & \text{if } l = L-1. \end{cases}$$

Note that $(x)^+ = \max\{0, x\}$.

We next define the transition probability from one system state (\mathbf{m}, \mathbf{s}) to the next $(\mathbf{m}', \mathbf{s}')$, given that action $a \in \mathcal{A}'(\mathbf{m}, \mathbf{s})$ is taken as:

$$p_a((\mathbf{m}, \mathbf{s}), (\mathbf{m}', \mathbf{s}')) = \sum_{\{\mathbf{d} \in \mathcal{D}(\mathbf{m}) \mid f(\mathbf{m}, \mathbf{d}) = \mathbf{m}'; g^a(\mathbf{s}, \mathbf{d}) = \mathbf{s}'\}} \tilde{q}(\mathbf{m}, \mathbf{d}).$$

Now that the transition probabilities are defined, we are ready to focus on the costs. The expected one-step costs in the current period, depending on the current system state (\mathbf{m}, \mathbf{s}) and the action $a \in \mathcal{A}'(\mathbf{m}, \mathbf{d})$ that is taken, are defined as:

$$c_a(\mathbf{m}, \mathbf{s}) = c_h \left(a + \sum_{l \in \{0, \dots, L-1\}} s_l \right) + c_e \sum_{\mathbf{d} \in \mathcal{D}(\mathbf{m})} \tilde{q}(\mathbf{m}, \mathbf{d}) (d_{I-1} - s_0)^+. \quad (1)$$

The first term shows that holding costs are paid for the complete inventory position, while the second term shows that emergency costs are paid for all demands, d_{I-1} , that cannot be fulfilled from stock on hand, s_0 , where the number of demands depends on the exact degradation that is observed out of all possible degradations or transition vectors, $\mathcal{D}(\mathbf{m})$.

We denote with $V_n^\pi(\mathbf{m}, \mathbf{s})$ the total (undiscounted) expected costs with n periods left to the time horizon, when the current system state is (\mathbf{m}, \mathbf{s}) and the policy π is used. This policy is denoted by $\pi = \{\pi(\mathbf{m}, \mathbf{s}) \mid (\mathbf{m}, \mathbf{s}) \in \mathcal{S}\}$, with $\pi(\mathbf{m}, \mathbf{s})$ being the function that gives the ordering action $a \in \mathcal{A}'(\mathbf{m}, \mathbf{s})$ given a system state (\mathbf{m}, \mathbf{s}) . $V_0^\pi(\mathbf{m}, \mathbf{s}) = c_\pi(\mathbf{m}, \mathbf{s})$ and $V_n^\pi(\mathbf{m}, \mathbf{s})$ for $n \in \mathbb{N}$ is recursively calculated as follows:

$$V_n^\pi(\mathbf{m}, \mathbf{s}) = c_\pi(\mathbf{m}, \mathbf{s}) + \sum_{(\mathbf{m}', \mathbf{s}') \in \mathcal{S}} p_\pi((\mathbf{m}, \mathbf{s}), (\mathbf{m}', \mathbf{s}')) V_{n-1}^\pi(\mathbf{m}', \mathbf{s}').$$

Since we are interested in the long-run average costs in an infinite horizon setting, we define:

$$g(\pi) = \lim_{n \rightarrow \infty} \frac{V_n^\pi(\mathbf{m}, \mathbf{s})}{n}.$$

Our goal is to find the optimal ordering policy π^* that minimizes the long-run average costs:

$$\pi^* = \arg \min_{\pi} g(\pi).$$

5. Optimal policy

In this section, we first give some properties that an optimal ordering policy satisfies. We use these properties to reduce the action space and we then use value iteration to find the optimal policy.

Lemma 1. *It is never optimal to order more than N spare parts in one period.*

Proof. Components may degrade at most one state per period. Since there are N components installed, we may observe at most N failures per period. If we then compare policy π_1 ordering $N' > N$ components in a certain period with policy π_2 ordering N components in that period and $N' - N$ components extra in the next period, we see that applying policy π_2 cannot lead to incurring more emergency costs, but it does lead to incurring less inventory holding costs than applying policy π_1 . Therefore, a policy that orders more than N spare parts in one period, may never be optimal. \square

As a result, we know that the action space of actions that may be taken in an optimal policy is finite. We introduce such an action space formally after Corollary 1.

Lemma 2. *Let $D_{L+1}^{max}(\mathbf{m})$ be the maximum demand over $L + 1$ periods, given the current degradation vector \mathbf{m} . It can be calculated as follows:*

$$D_{L+1}^{max}(\mathbf{m}) = N \cdot \left\lfloor \frac{L+1}{I} \right\rfloor + \sum_{i=I+I\lfloor(L+1)/I\rfloor-(L+1)}^{I-1} m_i.$$

Proof. Notice that if $I > L + 1$, the expression reduces to $\sum_{i=I-(L+1)}^{I-1} m_i$ and we are simply counting the number of components that are in the last $L + 1$ degradation states $i \in \mathcal{I}$. If $I = L + 1$, then the second part of the expression reduces to the summation over an empty set and the first part of the expression reduces to N . If $I < L + 1$, then machines may experience multiple failures in the next $L + 1$ periods. $\lfloor (L + 1)/I \rfloor$ counts the number of times that a component in one machine may pass through all degradation states in the next $L + 1$ periods. Each time that all states are passed through, the machine experiences one failure. This number is multiplied by the number of machines N . In addition, more failures may be experienced if components are already in the last few degradation states. These additional failures are counted in the second term of the expression. \square

Corollary 1. *Under an optimal policy, the inventory position, given by $a + \sum_{l=0}^{L-1} s_l$ ($a \in \mathcal{A}'(\mathbf{m}, \mathbf{s})$), will never be increased to a higher level than $D_{L+1}^{max}(\mathbf{m})$.*

Proof. The ordering decision is taken at the beginning of the period, before demand in that period has realized. As a result, components that we order at the beginning of a period t , may be used from period $t + L$ on. In the periods t up to and including $t + L$, we may observe at most $D_{L+1}^{max}(\mathbf{m})$ failures, so we require at most $D_{L+1}^{max}(\mathbf{m})$ spare parts. If our inventory position would be increased to a higher level than $D_{L+1}^{max}(\mathbf{m})$, then the additional spare parts can only be needed from period $t + L + 1$ on. With a similar argument as we used in Lemma 1, we see that it is better to order such spare parts at least one period later. \square

As a result, we can define the action space of actions that may be taken under an optimal policy to be $\mathcal{A}(\mathbf{m}, \mathbf{s}) = \left\{ a \in \mathbb{N}_0 \mid a \leq \min \left\{ N, \left(D_{L+1}^{\max}(\mathbf{m}) - \sum_{l=0}^{L-1} s_l \right)^+ \right\} \right\}$.

Using value iteration (see, e.g., Tijms, 1986), the optimal policy can be found. The value function $V_n(\mathbf{m}, \mathbf{s})$ can be determined recursively as follows:

$$V_n(\mathbf{m}, \mathbf{s}) = \min_{a \in \mathcal{A}(\mathbf{m}, \mathbf{s})} \left\{ c_a(\mathbf{m}, \mathbf{s}) + \sum_{(\mathbf{m}', \mathbf{s}') \in \mathcal{S}} p_a((\mathbf{m}, \mathbf{s}), (\mathbf{m}', \mathbf{s}')) V_{n-1}(\mathbf{m}', \mathbf{s}') \right\}. \quad (2)$$

$V_0(\mathbf{m}, \mathbf{s}) = 0$ for all $(\mathbf{m}, \mathbf{s}) \in \mathcal{S}$ and computation is stopped if:

$$\begin{aligned} \max_{(\mathbf{m}, \mathbf{s}) \in \mathcal{S}} \{V_n(\mathbf{m}, \mathbf{s}) - V_{n-1}(\mathbf{m}, \mathbf{s})\} - \min_{(\mathbf{m}, \mathbf{s}) \in \mathcal{S}} \{V_n(\mathbf{m}, \mathbf{s}) - V_{n-1}(\mathbf{m}, \mathbf{s})\} \\ \leq e \left(\min_{(\mathbf{m}, \mathbf{s}) \in \mathcal{S}} \{V_n(\mathbf{m}, \mathbf{s}) - V_{n-1}(\mathbf{m}, \mathbf{s})\} \right), \end{aligned}$$

with $e = 10^{-6}$. The stationary policy whose actions minimize the right hand side of Equation (2) for all $(\mathbf{m}, \mathbf{s}) \in \mathcal{S}$ will be negligibly close in costs to the minimal average costs policy (Tijms, 1986).

Computation of the optimal policy using value iteration is computationally inefficient: the size of the state space increases exponentially with the number of machines, the number of degradation states, and the length of the replenishment lead time. With respect to the structure of the optimal policy; in examples, we see a certain monotonicity in the optimal decision $a^*(m, s)$ as a function of the degradation state \mathbf{m} and the inventory status \mathbf{s} . See, e.g., Table 1, where the optimal policy is given for one example. It may be even possible to prove a formal result; see Remark 5.1. But in any case, such monotonicity results would have a limited advantage for the computation time of the optimal policy, and it therefore remains relevant to derive efficient and effective heuristic policies.

Remark 5.1. Zipkin (2008b) shows via so-called $L^\# - convexity$ that a standard lost sales problem with Markov-modulated demand has the following monotonicity property (as a clear structure result, in total even more can be shown): for a given \mathbf{m} , the optimal decision $a^*(\mathbf{m}, \mathbf{s})$ is non-increasing as a function of the inventory status \mathbf{s} , where $\mathbf{s} \leq \mathbf{s}'$ if and only if $\sum_{j=i}^{I-1} s_j \leq \sum_{j=i}^{I-1} s'_j$ for all $i = 0, \dots, I-1$. Given Remarks 3.1 and 3.2, our problem is equivalent to the standard lost sales problem, but then with a Markov-chain-driven counting process (see Zipkin, 2000), where demands are coupled to transitions instead of states. In other words, we do not have Markov-modulated demand and cannot directly apply the results of Zipkin (2008b). We do see, however, that the above monotonicity property holds for our problem in all instances that we have checked, but we have not been able to prove this.

		m					
		(2, 0, 0)	(1, 1, 0)	(1, 0, 1)	(0, 2, 0)	(0, 1, 1)	(0, 0, 2)
s	(0, 0)	0	1	1	1	1	2
	(1, 0)	0	0	0	0	1	1
	(0, 1)	0	0	0	0	0	1
	(2, 0)	0	0	0	0	0	0
	(1, 1)	0	0	0	0	0	0
	(0, 2)	0	0	0	0	0	0

Table 1: Optimal ordering decisions ($a \in \mathcal{A}(\mathbf{m}, \mathbf{s})$) given \mathbf{m} and \mathbf{s} , when $N = 2$, $L = 2$, $I = 3$, $c_e = 10^5$, $c_h = 1$, and $(q_{0,1}, q_{1,2}, q_{2,3}) = (1/50, 1/35, 1/15)$

6. Heuristic policies

In this section, we propose three heuristic policies. First, however, we present a reference policy, in Section 6.1, that we use to compare our heuristic policies with and that we use as the basis of our first heuristic policy in Section 6.2. We next propose the so-called myopic policy in Section 6.3 and the best-of-two policy in Section 6.4.

6.1. Reference policy

Our *reference policy* is the *optimal state-independent (SID) base stock policy*. Comparing with this policy, we can find the value of incorporating degradation information in the inventory control policy. The SID base stock policy has one parameter, the base stock level S_{SID} . If the inventory position is lower than S_{SID} when an ordering decision is taken, then a number of spare parts is ordered such that the inventory position is raised to S_{SID} : $a = (S_{\text{SID}} - \sum_{l=0}^{L-1} s_l)^+$. The optimal base stock level, S_{SID}^* , can be found using bisection, since the cost function is convex in the base stock level (see Downs et al., 2001, and combine with Remarks 3.1 and 3.2). An obvious lower bound on the optimal base stock level is 0. Using the best base stock level found thus far and the corresponding average costs, it is easy to find an upper bound on the optimal base stock level, since a lower bound on the average costs for a given base stock level is given by $c_h S_{\text{SID}}$ (see Equation 1).

6.2. Modified SID policy

We call our first heuristic policy the *modified SID policy* (MOD), since it is based on the optimal SID base stock policy. We follow the ordering decisions of the optimal SID base stock

policy, except when this would lead to an inventory position that is higher than $D_{L+1}^{\max}(\mathbf{m})$; the order size is then adjusted such that the inventory position will be exactly $D_{L+1}^{\max}(\mathbf{m})$: $a = (\min\{S_{\text{SID}}, D_{L+1}^{\max}(\mathbf{m})\} - \sum_{l=0}^{L-1} s_l)^+$. The reason to propose this policy is that we have shown in Corollary 1 that having an inventory position that is higher than $D_{L+1}^{\max}(\mathbf{m})$ can never be optimal. As a result, this adjustment means that we sometimes have a lower inventory position than the SID base stock policy, leading to lower inventory holding costs, while not increasing the emergency costs that we face. This policy is thus at least as good as the optimal SID base stock policy. However, the simple procedure that we use to lower the base stock level means that it is very probable that there are better ways to set the base stock level. This is what we aim for with the policy in the next section.

6.3. Myopic policy

We propose a second heuristic policy, which we call the *myopic policy* (MYO). The basic idea of the policy is to make two approximate steps to estimate the number of parts that we need one lead time from now, i.e., in period $t + L$ if we are now in period t :

- When determining the total demand distribution over the periods t up to and including $t + L$, we consider at most one failure during those periods for each machine. If the number of degradation states exceeds the replenishment leadtime ($L + 1 \leq I$), there can be at most one failure for each machine. Otherwise, there is generally a very small probability that one component fails twice during $L + 1$ periods, since the average lifetime of the component is typically much larger than the replenishment leadtime: the lifetime of the critical component is in the order of 1 to 10 years, while the replenishment leadtime is in the order of 1 to 10 weeks. However, in some instances of our numerical experiment, we are faced with a serious violation of this assumption; see Section 7.2.
- We ignore that some demands in the periods t up to and including $t + L - 1$ are satisfied by an emergency supply, i.e., that they cannot be fulfilled from stock. Since emergency costs are typically high, stock levels are high and few demands need to be satisfied by an emergency supply. However, if the emergency costs are relatively low compared with the holding costs, then more demands are satisfied by an emergency supply (see also Section 7.2).

The heuristic policy next makes an explicit trade-off between holding costs and emergency costs, but it does so myopically by considering the demand distribution over the next $L + 1$ periods only, thus ignoring the effect of the decision on the next periods.

In more detail, the policy works as follows. Let us assume that we are at the beginning of period t and that we have to decide how much to order. We need to determine the demand in the periods t up to and including $t + L$. We approximate this demand by ignoring that a machine may face more than one component failure in these periods. This effectively means that we change the degradation state space back to $\mathcal{I}' = \mathcal{I} \cup I$, set $q_{I,0} = 0$, and set $q_{I,I} = 1$. We thus make degradation state I an absorbing state. We then define $P_{i,t'}$ to be the probability that a machine that is in state $i \in \mathcal{I}'$ at the beginning of period t is in state I (has failed) at the beginning of period $t + t'$, $t' \in \mathbb{N}_0$. Obviously, $P_{i,0} = 0$ for all $i \in \mathcal{I}$ and $P_{I,t'} = 1$ for all $t' \in \mathbb{N}_0$. We can then calculate the probabilities for all $i \in \mathcal{I}$ and $t' \geq 1$ recursively as follows:

$$P_{i,t'} = q_{i,i+1}P_{i+1,t'-1} + q_{i,i}P_{i,t'-1}.$$

We next define $Q(\mathbf{m}, J)$ to be the probability of exactly J (out of N) components failing in periods t up to and including $t + L$, given the current degradation vector \mathbf{m} . To calculate this probability, we need to consider only those components that are in the last $L + 1$ degradation states at the beginning of period t . We therefore distinguish the cases $L + 1 \leq I$ and $L + 1 > I$.

Let us first consider the case that $L + 1 \leq I$. From each of the states $i = I - (L + 1), \dots, I - 1$, at most m_i components can fail in the periods t up to and including $t + L$. As an example, assume that $L = 1$ and $\mathbf{m} = (2, 2)$. If we want to calculate the probability of observing exactly two failures in periods t and $t + 1$, we need to consider (i) the probability of two failing components that are currently in degradation state 1, while no component fails that is in state 0, (ii) the probability of exactly one failing component that is in state 1 and also exactly one failing component that is in state 0, and (iii) the probability of two failing components that are currently in state 0, while no component fails that is in state 1. There are thus three different ways in which we may observe two component failures, and to denote those three ways, we introduce the set $\mathcal{J}(\mathbf{m}, J)$. In our example, $\mathcal{J}((2, 2), 2) = \{(0, 2), (1, 1), (2, 0)\}$, and in general:

$$\mathcal{J}(\mathbf{m}, J) = \{\mathbf{j} \in \mathbb{N}_0^{L+1} \mid j_i \leq m_{i+I-(L+1)}; \sum_{i=0}^L j_i = J\},$$

where j_i indicates the number of components that is in degradation state i at the beginning of period t and that has failed after $L + 1$ periods (at the beginning of period $t + L + 1$). We can now calculate $Q(\mathbf{m}, J)$ as follows:

$$Q(\mathbf{m}, J) = \sum_{\mathbf{j} \in \mathcal{J}} \prod_{i \in \{0, \dots, L\}} \binom{m_{i+I-(L+1)}}{j_i} (P_{i+I-(L+1), L+1})^{j_i} (1 - P_{i+I-(L+1), L+1})^{m_{i+I-(L+1)} - j_i}.$$

If $L + 1 > I$, we ignore that a machine may face more than one component failure in the periods t up to and including $t + L$. The calculations simplify:

$$\mathcal{J}(\mathbf{m}, J) = \{\mathbf{j} \in \mathbb{N}_0^I \mid j_i \leq m_i; \sum_{i=0}^{I-1} j_i = J\},$$

and

$$Q(\mathbf{m}, J) = \sum_{\mathbf{j} \in \mathcal{J}} \prod_{i \in \{0, \dots, I-1\}} \binom{m_i}{j_i} (P_{i,L+1})^{j_i} (1 - P_{i,L+1})^{m_i - j_i}.$$

Now that we have an approximate demand distribution, we ignore that some demands are satisfied by an emergency supply in the periods t up to and including $t + L - 1$. We can then calculate the (approximate) marginal $(L + 1)$ -period costs if the inventory position at the beginning of period t is increased to base stock level $S(\mathbf{m}) + 1$ instead of $S(\mathbf{m})$, given degradation vector \mathbf{m} at the beginning of the first period that we consider:

$$\Delta C(S(\mathbf{m})) = c_h(L + 1) - c_e \left(1 - \sum_{J=0}^{S(\mathbf{m})} Q(\mathbf{m}, J) \right). \quad (3)$$

We have to pay additional holding costs for periods t up to and including $t + L$ (first term). The emergency costs decrease if there are more than $S(\mathbf{m})$ demands (second term). Since this second term at the right hand side of Equation (3) increases in $S(\mathbf{m})$, we see that $\Delta C(S(\mathbf{m}))$ is increasing in $S(\mathbf{m})$, which makes the (approximate) optimal costs convex in $S(\mathbf{m})$. Hence, the (approximate) optimal base stock level $S^*(\mathbf{m})$ is the smallest base stock level $S(\mathbf{m})$ for which $\Delta C(S(\mathbf{m})) \geq 0$, i.e., for which the following inequality holds:

$$\sum_{J=0}^{S(\mathbf{m})} Q(\mathbf{m}, J) \geq 1 - \frac{c_h(L + 1)}{c_e},$$

and $a = (S^*(\mathbf{m}) - \sum_{l=0}^{L-1} s_l)^+$.

6.4. Best-of-two policy

Evaluation of the reference policy (the SID policy), the modified SID policy (MOD), and the myopic policy (MYO) can be done using value iteration for smaller problem instances, and using simulation for larger problem instances (see also Section 7.2). The myopic policy is expected to lead to lower costs than the modified SID policy in most cases. On the other hand, the modified SID policy cannot perform worse than the reference policy. Therefore, we propose the *best-of-two policy* (BO2), which is the superior of the other two policies: we take the solution (resulting policy) of both MOD and MYO, evaluate both using value iteration or simulation (as explained above), and choose out of those two policies the policy that leads to the lowest costs. By definition, this policy is thus the most time consuming one.

Parameter		Values used in all test beds	Additional in test bed		
			1	2	3
N	(Number of machines)	1; 5	—	10	—
L	(Replenishment lead time, in weeks)	—	1; 2	2; 5	1
I	(Number of degradation states)	2	3	5	3; 5
DPV	(Degradation probability vector: $q_{i,i+1}$ for $i \in \mathcal{I}$)	100v1; 100v2; 250	—	—	—
c_e	(Emergency costs, in Euros)	10^4 ; 10^5	—	—	—
c_h	(Inventory holding costs, in Euros/unit/week)	1; 200; 10^3	—	—	—

Table 2: Parameters used in each test bed (full factorial); for DPV , see Table 3

Abbreviation	I (Number of degradation states)		
	2	3	5
100v1	(1/50, 1/50)	(1/50, 1/35, 1/15)	(1/50, 1/20, 1/15, 1/10, 1/5)
100v2	(1/50, 1/50)	(1/50, 1/25, 1/25)	(1/50, 2/25, 2/25, 2/25, 2/25)
250	(1/125, 1/125)	(1/125, 2/125, 2/125)	(1/125, 4/125, 4/125, 4/125, 4/125)

Table 3: Degradation probability vectors, DPV : $q_{i,i+1}$ for $i \in \mathcal{I}$

7. Numerical experiment

We perform an extensive numerical experiment, using three test beds of problem instances. We explain the design of our experiment in Section 7.1 and we discuss the results in Section 7.2.

7.1. Design

The parameters that we use are given in Tables 2 and 3; we explain below how the problem instances are generated using these parameters. Notice that the test beds differ only for N , L , and I . Per test bed we use a full factorial design; test beds 1, 2, and 3 consist of 144, 216, and 108 problem instances, respectively. All parameter values are chosen such that we get a wide range of problem instances that are realistic in practice, e.g., for manufacturers in the high-tech industry, where equipment is expensive and often has high downtime costs (in the order of 1 to 100 k-Euros per hour). Therefore, the additional costs for an emergency supply (c_e), which include costs for a longer downtime when an emergency supply is used, are much higher than the inventory holding costs (c_h) in practice. The values for those parameters are chosen accordingly in our problem instances.

The aim of test bed 1 is twofold. First, we aim to see how much costs can be saved when using the optimal state-dependent (SD) base stock policy (Section 5) instead of the optimal SID base stock policy (Section 6.1), and second, we aim to see how much of these cost savings

are captured by our three heuristic policies (Sections 6.2, 6.3, and 6.4). Test bed 2 is then used to explore the performance of the best heuristic policy, the best-of-two policy (Section 6.4), on a wider range of parameter settings. We also use this test bed to investigate in which cases the myopic policy (Section 6.3) does not perform well. Finally, we use test bed 3 to understand what is lost if partial degradation information is used while more degradation information is available.

The degradation probability vectors (DPV ; $q_{i,i+1}$ for $i \in \mathcal{I}$) are based on a period length of 1 week. The mean lifetime is equal to 100, 100, and 250 weeks for the vectors 100v1, 100v2, and 250, respectively. Further, they are chosen such that the expected duration in the perfect state is the same as the total expected duration in the other states. The vector ‘100v1’ leads to an increasing degradation probability with an increasing degradation state (degradation keeps going faster). The other two vectors have constant degradation probabilities, except when they are in the perfect state (state 0). In fact, all values in the ‘250’ vectors can be found by taking the values in the corresponding ‘100v2’ vector and dividing them by 2.5. Notice that when comparing, for example, ‘100v1’ for $I = 2$ and $I = 5$, then the former vector can be seen as an aggregated version of the latter vector: there is less information on the exact degradation state. We use this fact in test bed 3 to understand what is lost if only partial degradation information is used when more is available.

7.2. Results

The results for test beds 1 and 3 are obtained using value iteration. Some problem instances in test bed 2 are too large to use value iteration, and therefore we use simulation for all problem instances in that test bed, except for those that are also part of test bed 1. To be more precise, we use the *batch means* method (see, e.g., Law, 2007, pp. 520–521), as follows:

- Perform a simulation run of length m periods (our choice of m is explained below), resulting in m observations (costs in a period): Y_i for $i \in \{0, \dots, m - 1\}$.
- Divide the run into n batches (we choose $n = 10$) and calculate the mean value for each batch: $\bar{Y}_j = \frac{1}{k} \sum_{i=0}^{k-1} Y_{jk+i}$ for $j \in \{0, \dots, n - 1\}$, with $k = m/n$.
- Calculate the grand sample mean: $\bar{\bar{Y}} = \frac{1}{n} \sum_{j=0}^{n-1} \bar{Y}_j$.
- Calculate the $100(1 - \alpha)$ percent confidence interval (we choose $\alpha = 0.1$) for $\bar{\bar{Y}}$: $\bar{\bar{Y}} \pm t_{n-1, 1-\alpha/2} \sqrt{S^2(n)/n}$, with $t_{n-1, 1-\alpha/2}$ being the upper $1 - \alpha/2$ critical point for the t -

distribution with $n-1$ degrees of freedom ($t_{9,0.95} \approx 1.833$), and $S^2(n) = \sum_{j=0}^{n-1} (\bar{Y}_j - \bar{Y})^2 / (n-1)$.

We have chosen m for each problem instance such that the width of the confidence interval divided by the grand sample mean is less than 1%.

Determining MOD and evaluating its costs for test beds 1 and 3, takes a few minutes per instance; the parameter values do not greatly influence this time. For MYO, this takes from a few minutes up to a few hours, depending mainly on the number of degradation states (I), the replenishment lead time (L), and the number of machines (N). The computation times for BO2 thus mainly depend on those times for MYO. The length of each simulation run m , as used for the additional instances of test bed 2, is usually a few million periods, with exceptions of up to 500 million. Computation times of a few hours are quite high, but remember that the optimal policy cannot be determined at all for many problem instances (especially also those with high I , L , and N).

The cost savings that we show in Tables 4 and 5 are calculated as: $\frac{1}{P} \sum_{p=1}^P \frac{\text{Costs}_{\text{SID}}(p) - \text{Costs}_{\text{ALT}}(p)}{\text{Costs}_{\text{SID}}(p)}$, with P being the number of problem instances in the test bed with the parameter as indicated in the table (e.g., $P = 72$ in test bed 1 if $N \in \{1, 5\}$), the problem instances being numbered $1, \dots, P$, ‘SID’ referring to the optimal SID base stock policy, and ‘ALT’ referring to either OPT for the optimal policy, MOD (modified SID policy), MYO (myopic policy), or BO2 (best-of-two policy).

Table 4 gives the results for test bed 1. Using the optimal degradation state-dependent policy instead of the optimal state-independent policy leads to drastic cost savings of 19.6% on average and 73.4% at maximum (the latter value is not shown in the table). Most of these savings are also achieved by BO2, and it is clear that the performance of that policy depends heavily on the performance of MYO. We come back to this in our discussion of the results on test bed 2.

Other interesting things to notice are that:

- The cost savings reduce with an increasing number of machines (N). One reason is that the SID base stock policy improves due to pooling effects. Another reason is that the demand for spare parts gets more stable, since chances increase that if some machines have a high degradation level, others have a low degradation level; there are fewer ‘extreme’ states for which it really helps to adapt the base stock policy.
- The cost savings reduce with an increasing lead time (L). If the lead time is 1, the SD

Parameter	Value	# Problem instances	Costs		Cost savings		
			SID	OPT	MOD	MYO	BO2
N	1	72	193.7	23.9%	7.6%	23.0%	23.2%
	5	72	377.5	15.2%	1.7%	14.0%	14.0%
L	1	72	278.9	21.7%	9.3%	21.3%	21.3%
	2	72	292.2	17.5%	0.0%	15.6%	15.9%
I	2	72	285.6	9.6%	0.0%	8.9%	9.0%
	3	72	285.6	29.5%	9.3%	28.1%	28.2%
DPV	100v1	48	327.9	21.6%	5.1%	20.0%	20.0%
	100v2	48	327.9	19.5%	5.1%	18.4%	18.5%
	250	48	201.0	17.5%	3.6%	17.0%	17.3%
(c_e, c_h)	$(10^4, 10^3)$	24	240.0	0.3%	0.0%	0.1%	0.3%
	$(10^4, 200)$	24	152.5	14.2%	0.2%	14.1%	14.1%
	$(10^4, 1)$	24	1.8	23.4%	7.4%	21.5%	22.1%
	$(10^5, 10^3)$	24	1035.9	27.2%	4.5%	26.8%	26.8%
	$(10^5, 200)$	24	281.3	32.6%	7.2%	29.6%	29.6%
	$(10^5, 1)$	24	2.1	19.6%	8.6%	18.8%	18.8%
Average		144	285.6	19.6%	4.6%	18.5%	18.6%

Table 4: Results for test bed 1

policies can stock spare parts as soon as any machine reaches the last degradation state. If the lead time is higher, the SD policies cannot wait so long and will look more like the SID policy.

- The possible cost savings increase drastically with an increasing number of degradation states (I). We come back to this when we discuss the results for test bed 3.
- The degradation probability vectors (DPV) have no clear influence on the results.
- The emergency costs (c_e) and holding costs (c_h) have a huge influence on the cost savings. It seems that the potential cost savings are minor if the ratio of emergency costs over holding costs is small, probably because preventing an emergency saves only little money, relative to the cost of adding a spare part, which is why more information does not help a lot. The potential cost savings increase if the ratio increases, but at a certain point, they start decreasing again. The reason may be that the costs resulting from the optimal SID policy are already very low, since preventing an emergency is worthwhile almost against any costs. Again, more information then does not help a lot.

Parameter	Value	# Problem instances	Costs	Cost savings		
			SID	MOD	MYO	BO2
N	1	72	193.8	10.1%	11.9%	24.8%
	5	72	428.4	4.7%	16.6%	17.9%
	10	72	650.3	2.0%	12.6%	13.2%
L	2	108	393.8	11.2%	19.2%	21.6%
	5	108	454.5	0.0%	8.1%	15.6%
I	2	108	424.4	0.0%	2.5%	5.4%
	5	108	423.9	11.2%	24.8%	31.9%
DPV	100v1	72	493.8	7.0%	8.9%	20.7%
	100v2	72	493.8	5.7%	15.3%	17.7%
	250	72	284.9	4.1%	16.8%	17.5%
(c_e, c_h)	$(10^4, 10^3)$	36	426.7	0.0%	-19.3%	0.5%
	$(10^4, 200)$	36	232.5	0.6%	10.3%	11.2%
	$(10^4, 1)$	36	2.7	10.3%	23.0%	23.5%
	$(10^5, 10^3)$	36	1485.0	4.9%	24.3%	24.3%
	$(10^5, 200)$	36	395.0	7.5%	27.8%	27.8%
	$(10^5, 1)$	36	3.1	10.3%	16.5%	24.6%
Average		216	424.2	5.6%	13.7%	18.6%

Table 5: Results for test bed 2

The results for test bed 2 can be found in Table 5. They confirm our findings on test bed 1 for MOD and BO2 on a broader range of parameter values (some of which lead to problem instances that cannot be solved to optimality, which is why we have used simulation for this test bed). However, the gap between MYO and BO2 is much larger now, and we even see a cost *increase* using MYO for the set of problem instances with $c_e = 10,000$ and $c_h = 1,000$.

Table 6 shows some more details on the performance of MYO. It is easily seen that two of the approximations underlying this policy cause problems in some cases:

- The policy ignores the fact that a machine may fail more than once during the lead time. However, if $L = 5$ and $I = 2$, a machine may fail up to three times. As a result, MYO may propose to stock too few spare parts. If then the emergency costs are extremely high compared to the holding costs (10^5 and 1, respectively), we get a percentagewise huge increase in the total costs. Remember from Table 5 that the costs resulting from the SID policy in this case are quite low, so in absolute terms, the increase in costs is minor.
- MYO further ignores lost sales during the lead time. However, if the emergency costs are

		(c_e, c_h)					
		$(10^5, 1)$	$(10^4, 1)$	$(10^5, 200)$	$(10^5, 10^3)$	$(10^4, 200)$	$(10^4, 10^3)$
(L, I)	$(2, 5)$	52.6%	51.9%	49.0%	41.8%	22.3%	-25.6%
	$(5, 5)$	36.0%	36.1%	39.6%	33.2%	14.0%	-53.3%
	$(2, 2)$	6.6%	2.1%	14.7%	11.7%	3.3%	0.0%
	$(5, 2)$	-29.1%	1.9%	7.7%	10.5%	1.5%	-0.3%

Table 6: Detailed cost savings for the myopic heuristic on test bed 2

	Additional costs	
	$I = 3$	$I = 2$
Average	52.2%	137.2%
# instances with add. costs in range 0-25%	17	6
# instances with add. costs in range 0-100%	10	12
# instances with add. costs in range 100-567%	9	18

Table 7: Results for test bed 3 (average costs for $I = 5$ are 142.1)

relatively low compared to the holding costs (10^4 and 10^3 , respectively), we may face quite a number of lost sales. As a result, MYO can perform much worse than the SID policy. Fortunately, we know from our results on test bed 1 (see Table 4), that in a situation with relatively low emergency costs, there is not much to gain from using condition information anyway.

In total, out of the 216 problem instances in test bed 2, 10 lead to a cost increase of more than 10% when using MYO instead of SID. That is why BO2 also considers using the results of MOD. In that way, a robust policy has been obtained that performs good in all instances.

Test bed 3 consists of 36 problem instances with three different numbers of degradation states (I): 2, 3 or 5. Table 7 shows the results, with the additional costs being computed as follows: $\frac{1}{36} \sum_{p=1}^{36} \frac{\text{Costs}_{\text{OPT}, I \in \{2, 3\}}(p) - \text{Costs}_{\text{OPT}, I = 5}(p)}{\text{Costs}_{\text{OPT}, I = 5}(p)}$, with the problem instances being numbered $1, \dots, 36$, ‘OPT, $I \in \{2, 3\}$ ’ referring to the optimal SD policy with $I \in \{2, 3\}$ and ‘OPT, $I = 5$ ’ referring to the optimal SD policy with $I = 5$.

We see that huge additional costs are incurred if not all degradation information is used. For example, if we distinguish 3 degradation states only, whereas there are 5 present, this leads to additional costs of more than 25% in more than half of the problem instances. Distinguishing only 2 degradation states leads to maximum additional costs of 567%. This means that if more degradation information is available, it is very costly not to use it. The other way around, this

means that it may be worthwhile to invest in condition monitor equipment if this leads to the ability to distinguish more failure states.

8. Conclusions

We have proposed to incorporate degradation state information in the spare parts inventory control policy. Using value iteration, we are able to find the optimal policy. Since this is computationally intensive, we have proposed three heuristic policies.

Using an extensive numerical experiment, we have found various interesting insights. First, using a degradation state-dependent policy instead of a state-independent policy may lead to costs savings of about 20% on average and over 70% at maximum. Second, some parameters greatly influence the cost savings. It is especially interesting to see that if the number of machines (N) or the lead time (L) increases, then the cost savings decrease. The main reason in the former case is that the demand for spare parts gets more stable, since chances increase that if some machines have a high degradation level, others have a low degradation level; there are fewer ‘extreme’ states for which it really helps to adapt the base stock policy. The reason in the latter case is that if the lead time is 1, the SD policies can stock spare parts as soon as any machine reaches the last degradation state. If the lead time is higher, the SD policies cannot wait so long and will look more like the SID policy. Third, it is very important to use the degradation information that is present. This means that it may be worthwhile to invest in condition monitor equipment if this leads to the ability to distinguish more failure states. It also means that if more degradation information is available, it is very costly not to use it. For example, if we distinguish 3 degradation states only, whereas there are 5 present, this leads to additional costs of more than 25% in more than half of the problem instances in our numerical experiment. Fourth, the performance of one of our heuristics is close to optimal.

References

- Alaswad, S. and Xiang, Y. (2017). A review on condition-based maintenance optimization models for stochastically deteriorating system. *Reliability Engineering & System Safety*, 157:54 – 63.
- Basten, R. J. I. and Van Houtum, G. J. (2014). System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science*, 19(1):34–55.

- Bijvank, M. and Johansen, S. G. (2012). Periodic review lost-sales inventory models with compound poisson demand and constant lead times of any length. *European Journal of Operational Research*, 220:106–114.
- Bijvank, M. and Vis, I. F. A. (2011). Lost-sales inventory theory: A review. *European Journal of Operational Research*, 215:1–13.
- Bjarnason, E. T., Taghipour, S., and Banjevic, D. (2014). Joint optimal inspection and inventory for a k-out-of-n system. *Reliability Engineering & System Safety*, 131:203–215.
- Christer, A. H. (1982). Modelling inspection policies for building maintenance. *The Journal of the Operational Research Society*, 33(8):723–732.
- Deshpande, V., Iyer, A. V., and Cho, R. (2006). Efficient supply chain management at the U.S. Coast Guard using part-age dependent supply replenishment policies. *Operations Research*, 54(6):1028–1040.
- Downs, B., Metters, R., and Semple, J. (2001). Managing inventory with multiple products, lags in delivery, resource constraints, and lost sales: A mathematical programming approach. *Management Science*, 47(3):464–479.
- Elwany, A. H. and Gebraeel, N. Z. (2008). Sensor-driven prognostic models for equipment replacement and spare parts inventory. *IIE Transactions*, 40(7):629–639.
- Feeney, G. J. and Sherbrooke, C. C. (1966). The $(s - 1, s)$ inventory policy under compound poisson demand. *Management Science*, 12(5):391–411.
- Giorgio, M., Guida, M., and Pulcini, G. (2011). An age- and state-dependent markov model for degradation processes. *IIE Transactions*, 43:621–632.
- Hariharan, R. and Zipkin, P. H. (1995). Customer-order information, leadtimes, and inventories. *Management Science*, 41(10):1599–1607.
- Jardine, A. K. S., Lin, D., and Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20:1483–1510.
- Law, A. M. (2007). *Simulation Modeling & Analysis*. McGraw-Hill, New York (NY), international edition.

- Li, R. and Ryan, J. K. (2011). A Bayesian inventory model using real-time condition monitoring information. *Production and Operations Management*, 20(5):754–771.
- Louit, D., Pascual, R., Banjevic, D., and Jardine, A. K. S. (2011). Condition-based spares ordering for critical components. *Mechanical Systems and Signal Processing*, 25:1837–1848.
- Morton, T. E. (1971). The near-myopic nature of the lagged-proportional-cost inventory problem with lost sales. *Operations Research*, 19(7):1708–1716.
- Muckstadt, J. A. (2005). *Analysis and Algorithms for Service Parts Supply Chains*. Springer, New York (NY).
- Neves, M. L., Santiago, L. P., and Maia, C. A. (2011). A condition-based maintenance policy and input parameters estimation for deteriorating systems under periodic inspection. *Computers & Industrial Engineering*, 61:503–511.
- Olde Keizer, M. C., Flapper, S. D. P., and Teunter, R. H. (2017). Condition-based maintenance policies for systems with multiple dependent components: A review. *European Journal of Operational Research*. Accepted for publication.
- Rausch, M. and Liao, H. (2010). Joint production and spare part inventory control strategy driven by condition based maintenance. *IEEE Transactions on Reliability*, 59(3):507–516.
- Sherbrooke, C. C. (2004). *Optimal inventory modelling of systems. Multi-echelon techniques*. Kluwer, Dordrecht (The Netherlands), second edition.
- Si, X., Wang, W., Hu, C., and Zhou, D. (2011). Remaining useful life estimation – a review on the statistical data driven approaches. *European Journal of Operational Research*, 213(1):1–14.
- Song, J.-S. and Zipkin, P. H. (1993). Inventory control in a fluctuating demand environment. *Operations Research*, 41(2):351–370.
- Tijms, H. C. (1986). *Stochastic Modeling and Analysis*. John Wiley & Sons, New York (NY).
- Topan, E., Tan, T., Van Houtum, G. J., and Dekker, R. (2016). Using imperfect advance demand information in lost-sales inventory systems. *BETA Working paper*, 512.
- Van Horenbeek, A., Buré, J., Cattrysse, D., Pintelon, L., and Vansteenwegen, P. (2013). Joint maintenance and inventory optimization systems: A review. *International Journal of Production Economics*, 143:499–508.

- Van Horenbeek, A. and Pintelon, L. (2015). A joint predictive maintenance and inventory policy. In Tse, W. P., Mathew, J., Wong, K., Lam, R., and Ko, C., editors, *Engineering Asset Management - Systems, Professional Practices and Certification: Proceedings of the 8th World Congress on Engineering Asset Management (WCEAM 2013) & the 3rd International Conference on Utility Management & Safety (ICUMAS)*, pages 387–399. Springer International Publishing, Cham.
- Wang, K., Chu, J., and 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.
- Wang, L., Chu, J., and Mao, W. (2008). A condition-based order-replacement policy for a single-unit system. *Applied Mathematical Modeling*, 32(11):2274–289.
- Wang, W. (2011). A joint spare part and maintenance inspection optimisation model using the delay-time concept. *Reliability Engineering & System Safety*, 96(11):1535–1541.
- Wang, Z., Hu, C., Wang, W., Kong, X., and Zhang, W. (2015a). A prognostics-based spare part ordering and system replacement policy for a deteriorating system subjected to a random lead time. *International Journal of Production Research*, 53(15):4511–4527.
- Wang, Z.-Q., Wang, W., Hu, C.-H., Si, X.-S., and Zhang, W. (2015b). A prognostic-information-based order-replacement policy for a non-repairable critical system in service. *IEEE Transactions on Reliability*, 64(2):721–735.
- Zipkin, P. H. (2000). *Foundations of Inventory Management*. McGraw-Hill, Boston (MA).
- Zipkin, P. H. (2008a). Old and new methods for lost-sales inventory systems. *Operations Research*, 56(5):1256–1263.
- Zipkin, P. H. (2008b). On the structure of lost-sales inventory models. *Operations Research*, 56(4):937–944.