ARTICLE IN PRESS
JID: RESS [m5GeSdc;June 8, 2017;11:32]
Reliability Engineering and System Safety 000 (2017) 1–9
Contents lists available at ScienceDirect
Reliability Engineering and System Safety
journal homepage: www.elsevier.com/locate/ress

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, Netherlands
[b] Eindhoven University of Technology, P.O. Box 513, 5600 MB, Eindhoven, Netherlands
[c] Consultants in Quantitative Methods, P.O. Box 414, 5600 AK, Eindhoven, Netherlands

ARTICLE INFO

Keywords:
Inventory
Spare parts
Condition monitoring

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.

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., 1].

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., [2–4], for an 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., 5]. 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

* Corresponding author.
E-mail 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).

http://dx.doi.org/10.1016/j.ress.2017.05.035
Available online xxx
0951-8320/© 2017 Elsevier Ltd. All rights reserved.

Please cite this article as: X. Lin et al., Reliability Engineering and System Safety (2017), http://dx.doi.org/10.1016/j.ress.2017.05.035

(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 [6]. 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 [7] and Muckstadt [8] or the review by Basten and Van Houtum [1].

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 [9]. We discuss three papers in more detail, all considering a discrete-time inventory model with lost sales: Bijvank and Johansen [10] and Zipkin [11], [12]. Bijvank and Johansen [10] 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" [10, 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 [11] discusses various heuristic policies, one of which is the myopic policy [based on 13]. 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 [12] 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 [14] 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 [15]. 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. [16] 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. [16] 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. [17]. 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 [18]. 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 whorthwhile 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. [19]. Two more recent overviews of the literature on CBM are those by Alaswad and Xiang [20] and Olde Keizer et al. [21]. 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. [22], Elwany and Gebraeel [23], Van Horenbeek and Pintelon [24], Rausch and Liao [25], Wang et al. [26], Wang et al. [27], Wang [28], Wang et al. [29] or Wang et al. [30]. Van Horenbeek et al. [31] 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 Louit et al. [32], 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 [33]. 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 [33] 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 $\mathcal{I}' = \{0, \ldots, I\}$, 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 2, 3, 4]. 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