Condition based spare parts supply

Document Version:
Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

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

Condition based spare parts supply

X. Lin, R.J.I. Basten, A.A. Kranenburg, G.J. van Houtum

Beta Working Paper series 371

Condition based spare parts supply

X. Lin[a], R.J.I. Basten[b,*], A.A. Kranenburg[c], G.J. van Houtum[d]

[a]*Dow Benelux, P.O. Box 48, 4530 AA, Terneuzen, The Netherlands*
[b]*University of Twente, P.O. Box 217, 7500 AE, Enschede, The Netherlands*
[c]*Consultants in Quantitative Methods, P.O. Box 414, 5600 AK, Eindhoven, The Netherlands*
[d]*Eindhoven University of Technology, P.O. Box 513, 5600 MB, 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 control, 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., Sherbrooke, 2004).

---

[*]Corresponding author: T: +31 53 489 4007

Email addresses: xiaolin1014@gmail.com (X. Lin), r.basten@utwente.nl (R.J.I. Basten), bram.kranenburg@cqm.nl (A.A. Kranenburg), g.j.v.houtum@tue.nl (G.J. van Houtum)

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. 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. 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 inventory pipeline (stock on hand plus outstanding orders). 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). In an extensive numerical experiment, we find that the optimal policy, which by definition is a state-dependent base stock policy, achieves cost savings of upto more than 30% compared with a state-independent base stock policy. However, using value iteration to obtain the optimal policy is very time consuming, especially if the number of degradation states, the lead time, or the number of machines is high. Therefore, we propose two heuristic policies, and a third heuristic policy that always uses the best of the former two policies. All three heuristics are easy to compute. We show that the third policy is very close to optimal.

Our main contribution is as follows. First, while the use of condition information for maintenance optimization has been studied extensively, there are only a few studies on the effect on spare parts supply. More specifically, we study the effect of condition information on spare parts supply *without* changing the maintenance policy. This is especially relevant if 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). The two papers that are most closely related are by Deshpande et al. (2006) and Louit et al. (2011); for the differences with those and other studies, see Section 2. Second, we show that large savings can be obtained via condition based spare parts supply, and we identify under which circumstances the savings are largest. Third, we derive an efficient and effective heuristic.

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 heuristics that we propose in Section 6. We then perform an extensive numerical experiment in Section 7.

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 Sherbrooke (2004).

Usually, fixed demand rates are assumed, but there has also been some research on varying demand rates and state-dependent (SD) inventory policies. 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 SD inventory policies uses advance demand information (ADI). ADI could result from monitoring the state of the components in the installed base, but this is, to the best of our knowledge, not considered in the literature. Instead, ADI means that customers place orders that will lead to an actual demand only after a certain 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

3

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 lead time, thus making perfect ADI. ADI may also be imperfect, see, for example, Tan et al. (2007), who analyze a model in which there is a certain probability (that is generally less than 1) that an order that is placed will lead to an actual demand.

In our model, we use the degradation states of the critical components as ADI. A related paper is that of Deshpande et al. (2006) in which part-age information is explicitly used in the inventory control policy. The authors assume that a part-age signal can be observed each period, which is then compared with a certain treshold 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 normal 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. For an overview of this stream, including a review of diagnostics and prognostics techniques, see Jardine et al. (2006). Within this stream, also the reducing effect of condition based maintenance on spare parts supply costs has been studied; see Elwany and Gebraeel (2008), Wang et al. (2008, 2009), and Rausch and Liao (2010). 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. To the best of our knowledge, the only paper with the same focus is the one by Louit et al. (2011), but they have different assumptions. They assume a single system for which at most one spare part is kept on stock, and they 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.

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. Time is divided into periods of unit length and we assume an infinite time horizon. We assume that a component can degrade 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_{em} (> 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; 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 $\mathcal{I} = \mathcal{I}' \backslash I$ 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. The latter assumption can be relaxed since the holding costs that are paid for the components in the pipeline are a constant factor, which is independent of the inventory control policy that is used: given that c_{em} is defined as *additional* costs, each demanded spare part will induce those costs.

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 $j \notin \{i, i+1\}$.

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_{em} is defined to be *additional* costs), we may ignore the variable ordering costs per component.

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 will 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 $\mathbf{m} = (m_0, \ldots, m_{I-1})$ 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 $\mathbf{s} = (s_0, \ldots, s_{L-1})$, s_0 denotes the stock on hand, and s_l (for $l \in \{1, \ldots, L-1\}$) 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 will 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 $\mathbf{d} = (d_0, \ldots, d_{I-1})$ 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

$\hat{q}(\mathbf{m}, \mathbf{d})$ 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 $\mathbf{m}$:

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

The subsequent degradation vector, given the current degradation vector $\mathbf{m}$ and the transition vector $\mathbf{d}$ is given by $f(\mathbf{m}, \mathbf{d}) = (f_0(\mathbf{m}, \mathbf{d}), \ldots, f_{I-1}(\mathbf{m}, \mathbf{d}))$, 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 $a \in \mathcal{A}'(\mathbf{m}, \mathbf{s})$ that was chosen (notice that a is chosen before $\mathbf{d}$ is observed) as $g^a(\mathbf{s}, \mathbf{d}) = (g_0^a(\mathbf{s}, \mathbf{d}), \ldots, g_{L-1}^a(\mathbf{s}, \mathbf{d}))$, with:

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

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:

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

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:

$$c_a(\mathbf{m}, \mathbf{s}) = c_{\text{h}} \left(a + \sum_{l \in \{0, \ldots, L-1\}} s_l \right) + c_{\text{em}} \sum_{d \in \mathcal{D}(\mathbf{m})} \hat{q}(\mathbf{m}, \mathbf{d}) (d_{I-1} - s_0)^+.$$

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 $\pi = \{\pi(\mathbf{m}, \mathbf{s}) \mid (\mathbf{m}, \mathbf{s}) \in \mathcal{S}\}$, 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 $(\mathbf{m}, \mathbf{s})$. $V_n^\pi(\mathbf{m}, \mathbf{s})$ is recursively calculated as follows:

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

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

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

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 will first give some properties that an optimal ordering policy satisfies. We can use these properties to reduce the action space. 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. $\qquad\square$

As a result, we know that the action space of actions that may be taken in an optimal policy is finite. We will introduce such an action space formally after Corollary 3.

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:*

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

8

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 expression reduces to the summation over an empty set and the first 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 couple of degradation states. These additional failures are counted in the second term of the expression. $\qquad\square$

Corollary 3. *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 upto 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. But, with a similar argument as we used in Lemma 1, we see that it is better to order such spare parts earliest next period. $\qquad\square$

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

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:

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

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

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

with $e = 10^{-6}$. The stationary policy whose actions minimize the right hand side of Equation 1 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. Therefore, and because the optimal policy has no clear structure, we propose three heuristic policies in the next section.

6. Heuristic policies

In this section, we propose three heuristic policies. First, however, we present a reference policy that we will use to compare our heuristic policies with and that we will use as the basis of our first heuristic.

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 below this value when an ordering decision is taken, then a number of spare parts is ordered such that the inventory position is raised to S_{SID} again. The optimal base stock level, S_{SID}^*, can be found using enumeration, and using the fact that a lower bound on the average costs is given by multiplying the holding costs by the base stock level minus the mean lead time demand

We call our first heuristic the *modified SID heuristic* (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 reason is that we have shown in Corollary 3 that having an inventory position that is higher than $D_{L+1}^{\max}(\mathbf{m})$ can never be optimal. As a result, 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 modified SID heuristic uses a very simple procedure to deviate from the base stock level. As a result, it may sometimes still increase the inventory to a too high level. On the other hand, if the overall condition of the components in the installed base is poor, the level to which the inventory is increased may be too low.

We therefore propose a second heuristic, which we call the *myopic heuristic* (MYO). This heuristic uses the degradation information of the components in the installed base to find an approximated probability distribution of the $(L+1)$-period demand. It next makes an explicit trade-off between holding costs and emergency costs, in a myopic way, as described below.

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 upto and including $t + L$. We make the first approximation by assuming that each machine will face at most 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 $Q(J)$ to be the probability of exactly J (out of N) components failing in periods t upto and including $t + L$, under the assumption that each machine will face at most one component failure in those periods. For this, we need to consider only those components that are in the last $L + 1$ degradation states at the beginning of period t, assuming for now that $L + 1 \leq I$. From each of these states i, at most m_i components can fail in the periods t upto and including $t + L$. We therefore define the set $\mathcal{J}$ to be the set of all vectors that can lead to exactly J failures in those $L+1$ periods, 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 $t+L+1$), as follows:

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

We can now calculate $Q(J)$ as follows:

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

If $L+1 > I$, the calculations simplify:

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

and

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

Now that we have an approximate demand distribution, we make the second approximation by assuming that all demands are fulfilled from stock in the periods t upto and including $t+L-1$. In other words, we ignore that some demands are satisfied by an emergency supply in those periods, and that then the inventory position is not reduced by those demands. This is quite a reasonable assumption since emergency costs are usually so high that optimal base stock levels will lead to a low probability of lost sales. Using this assumption, we can 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:

$$\Delta C(S(\mathbf{m})) = c_{\mathrm{h}}(L+1) - c_{\mathrm{em}} \left(1 - \sum_{J=0}^{S(\mathbf{m})} Q(J) \right). \tag{2}$$

We have to pay additional holding costs for periods t upto 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 (2) 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:

$$\sum_{J=0}^{S(\mathbf{m})} Q(J) \geq 1 - \frac{c_{\mathrm{h}}(L+1)}{c_{\mathrm{em}}}.$$

We have made two approximations in our calculations. Besides, we consider the demand distribution over the next $L+1$ periods only, and thus ignore the effect of the decision on the next periods. This makes our heuristic myopic.

Evaluation of the reference policy (the SID policy), the modified SID policy, and the myopic policy can be done using value iteration for smaller problem instances, and using simulation for larger problem instances.

The myopic heuristic is expected to lead to lower costs than the modified SID heuristic in most cases. On the other hand, the modified SID heuristic cannot perform worse than the reference policy. Therefore, we propose the *best-of-two heuristic* (BO2), which is the superior of the other two heuristics: we take the results of the former two heuristics, evaluate them as explained above, and choose the one that leads to the lowest costs.

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 1 and 2; 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 practically realistic problem instances. For example, the additional costs for an emergency supply (c_{em}) are much higher than the inventory holding costs (c_{h}) in practice, and thus 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 SD base stock policy instead of the optimal SID base stock policy, and second, we aim to see how much of these cost savings

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 $i \in \mathcal{I}$)	100(1); 100(2); 250	—	—	—
c_{em}	(Emergency costs, in Euros)	10^4; 10^5	—	—	—
c_{h}	(Inventory holding costs, in Euros/unit/week)	1; 200; 10^3	—	—	—

Table 1: Parameters used in each test bed (full factorial); for DPV, see Table 2

Abbreviation	I (Number of degradation states)		
	2	3	5
100(1)	(1/50; 1/50)	(1/50; 1/35; 1/15)	(1/50; 1/20; 1/15; 1/10; 1/5)
100(2)	(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 2: Degradation probability vectors, DPV: $q_{i,i+1}$ for $i \in \mathcal{I}$

are captured by our three heuristics. Test bed 2 is then used to explore the performance of the best heuristic, the best-of-two heuristic, on a wider range of parameter settings. We also use this test bed to investigate in which cases the myopic heuristic 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 chosen such that the expected duration in the perfect state is the same as the total expected duration in the other states. The vector '100(1)' 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 '100(2)' vector and dividing them by 2.5. Notice that when comparing, for example, '100(1)' 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, which is why 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: Y_i for $i \in \{0, \ldots, m-1\}$.

- Divide the run into n batches (we choose $n = 10$) and calculate the mean value for each batch: $\bar{Y}_j = \frac{1}{k} \sum_{i=k \cdot m}^{(k+1) \cdot m - 1} Y_i$ for $j \in \{0, \ldots, n-1\}$, 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 $S^2(n) = \sum_{j=0}^{n-1} \left(\bar{Y}_j - \bar{\bar{Y}} \right)^2 /(n-1)$.

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%.

The cost savings that we show in Tables 3 and 4 are calculated as follows: $\frac{1}{P} \sum_{p=1}^{P} \frac{\text{Costs}_{\text{SID}}(p) - \text{Costs}_{\text{SD}}(p)}{\text{Costs}_{\text{SID}}(p)}$, 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 $1, \ldots, P$, 'SID' referring to the optimal SID base stock policy, and 'SD' referring to either Opt. for the optimal SD base stock policy, MOD, MYO, or BO2.

Table 3 gives the results for test bed 1. Using a degradation state-dependent policy instead of a 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 the best-of-two heuristic (BO2), and it is clear that the performance of that heuristic depends heavily on the performance of the myopic heuristic (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). The reason is that the SID base stock policy improves due to pooling effects.

- The cost savings reduce with an increasing lead time (L). 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.

- 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_{em}) 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 and that they increase if the ratio increases. However, at a certain point, they start decreasing again.

The results for test bed 2 can be found in Table 4. They basically confirm our findings on test bed 1 on a broader range of parameter values (this is why we have used simulation for this test bed). In addition, from this test bed we learn that there are some problem instances on which the myopic (MYO) heuristic performs badly (not shown in the table). For example, for $N = 1$, $L = 5$, $I = 5$, $DPV = 100(2)$, $c_{\mathrm{em}} = 10^4$, and $c_{\mathrm{h}} = 10^3$, we see that the costs of the MYO heuristic are almost five times the costs of the SID policy. This is caused by one of the approximations underlying MYO: it ignores lost sales during $L+1$ periods. When the ratio between emergency costs and holding costs is low, a significant number of demands will lead to lost sales. Another interesting example is that of $N = 1$, $L = 5$, $I = 2$, $DPV = 100(1)$, $c_{\mathrm{em}} = 10^5$, and $c_{\mathrm{h}} = 1$. Here, the costs resulting from using MYO are almost 2.5 times those of the SID policy. This is caused by another approximation in MYO: it assumes that a machine may fail at most once during $L + 1$ periods. However, in this case, a machine may fail upto 3 times during $L + 1$ periods ($(L + 1)/I = 3$). Such an unanticipated

16

Parameter	Value	# Problem instances	Costs SID	Cost savings			
				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	100(1)	48	327.9	21.6%	5.1%	20.0%	20.0%
	100(2)	48	327.9	19.5%	5.1%	18.4%	18.5%
	250	48	201.0	17.5%	3.6%	17.0%	17.3%
$(c_{\mathrm{em}}, c_{\mathrm{h}})$	$(10^4, 10^3)$	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 3: Results for test bed 1

failure leads to huge costs if the emergency costs are high. In total, out of the 216 problem instances in test bed 2, 10 lead to a cost increase of more than 10% when using the myopic heuristic instead of the SID policy. That is why our best-of-two heuristic also considers using the results of the modified SID policy.

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

Parameter	Value	# Problem instances	Costs SID	Cost savings BO2
N	1	72	193.8	24.8%
	5	72	428.4	17.9%
	10	72	650.3	13.2%
L	2	108	393.8	21.6%
	5	108	454.5	15.6%
I	2	108	424.4	5.4%
	5	108	423.9	31.9%
DPV	100(1)	72	493.8	20.7%
	100(2)	72	493.8	17.7%
	250	72	284.9	17.5%
(c_{em}, c_h)	$(10^4, 10^3)$	36	426.7	0.5%
	$(10^4, 200)$	36	232.5	11.2%
	$(10^4, 1)$	36	2.7	23.5%
	$(10^5, 10^3)$	36	1485.0	24.3%
	$(10^5, 200)$	36	395.0	27.8%
	$(10^5, 1)$	36	3.1	24.6%
Average		216	424.2	18.6%

Table 4: Results for test bed 2

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.

Acknowledgements

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

cation, training and research worldwide for the benefit of all.

References

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.

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.

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.

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.

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.

Song, J.-S. and Zipkin, P. H. (1993). Inventory control in a fluctuating demand environment. *Operations Research*, 41(2):351–370.

Tan, T., Güllü, R., and Erkip, N. K. (2007). Modelling imperfect advance demand information and analysis of optimal inventory policies. *European Journal of Operational Research*, 177:897–923.

Tijms, H. C. (1986). *Stochastic Modeling and Analysis*. John Wiley & Sons, New York (NY).

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.

Working Papers Beta 2009 - 2012

nr.	Year	Title	Author(s)
371	2012	Condition based spare parts supply	X.Lin, R.J.I. Basten, A.A. Kranenburg, G.J. van Houtum
370	2012	Using Simulation to Assess the Opportunities of Dynamic Waste Collection	Martijn Mes
369	2012	Aggregate overhaul and supply chain planning for rotables	J. Arts, S.D. Flapper, K. Vernooij
368	2012	Operating Room Rescheduling	J.T. van Essen, J.L. Hurink, W. Hartholt, B.J. van den Akker
367	2011	Switching Transport Modes to Meet Voluntary Carbon Emission Targets	Kristel M.R. Hoen, Tarkan Tan, Jan C. Fransoo, Geert-Jan van Houtum
366	2011	On two-echelon inventory systems with Poisson demand and lost sales	Elisa Alvarez, Matthieu van der Heijden
365	2011	Minimizing the Waiting Time for Emergency Surgery	J.T. van Essen, E.W. Hans, J.L. Hurink, A. Oversberg
364	2011	Vehicle Routing Problem with Stochastic Travel Times Including Soft Time Windows and Service Costs	Duygu Tas, Nico Dellaert, Tom van Woensel, Ton de Kok
363	2011	A New Approximate Evaluation Method for Two-Echelon Inventory Systems with Emergency Shipments	Erhun Özkan, Geert-Jan van Houtum, Yasemin Serin
362	2011	Approximating Multi-Objective Time-Dependent Optimization Problems	Said Dabia, El-Ghazali Talbi, Tom Van Woensel, Ton de Kok
361	2011	Branch and Cut and Price for the Time Dependent Vehicle Routing Problem with Time Window	Said Dabia, Stefan Röpke, Tom Van Woensel, Ton de Kok
360	2011	Analysis of an Assemble-to-Order System with Different Review Periods	A.G. Karaarslan, G.P. Kiesmüller, A.G. de Kok
359	2011	Interval Availability Analysis of a Two-Echelon, Multi-Item System	Ahmad Al Hanbali, Matthieu van der Heijden

extenstion of the Erlang loss function

343 2011 Vehicle refueling with limited resources — Murat Firat, C.A.J. Hurkens, Gerhard J. Woeginger

342 2011 Optimal Inventory Policies with Non-stationary Supply Disruptions and Advance Supply Information — Bilge Atasoy, Refik Güllü, TarkanTan

341 2011 Redundancy Optimization for Critical Components in High-Availability Capital Goods — Kurtulus Baris Öner, Alan Scheller-Wolf Geert-Jan van Houtum

339 2010 Analysis of a two-echelon inventory system with two supply modes — Joachim Arts, Gudrun Kiesmüller

338 2010 Analysis of the dial-a-ride problem of Hunsaker and Savelsbergh — Murat Firat, Gerhard J. Woeginger

335 2010 Attaining stability in multi-skill workforce scheduling — Murat Firat, Cor Hurkens

334 2010 Flexible Heuristics Miner (FHM) — A.J.M.M. Weijters, J.T.S. Ribeiro

333 2010 An exact approach for relating recovering surgical patient workload to the master surgical schedule — P.T. Vanberkel, R.J. Boucherie, E.W. Hans, J.L. Hurink, W.A.M. van Lent, W.H. van Harten

332 2010 Efficiency evaluation for pooling resources in health care — Peter T. Vanberkel, Richard J. Boucherie, Erwin W. Hans, Johann L. Hurink, Nelly Litvak

331 2010 The Effect of Workload Constraints in Mathematical Programming Models for Production Planning — M.M. Jansen, A.G. de Kok, I.J.B.F. Adan

330 2010 Using pipeline information in a multi-echelon spare parts inventory system — Christian Howard, Ingrid Reijnen, Johan Marklund, Tarkan Tan

329 2010 Reducing costs of repairable spare parts supply systems via dynamic scheduling — H.G.H. Tiemessen, G.J. van Houtum

328 2010 Identification of Employment Concentration and Specialization Areas: Theory and Application — F.P. van den Heuvel, P.W. de Langen, K.H. van Donselaar, J.C. Fransoo

A combinatorial approach to multi-skill workforce scheduling — Murat Firat, Cor Hurkens

327 2010

326 2010 | Stability in multi-skill workforce scheduling | Murat Firat, Cor Hurkens, Alexandre Laugier

325 2010 | Maintenance spare parts planning and control: A framework for control and agenda for future research | M.A. Driessen, J.J. Arts, G.J. v. Houtum, W.D. Rustenburg, B. Huisman

324 2010 | Near-optimal heuristics to set base stock levels in a two-echelon distribution network | R.J.I. Basten, G.J. van Houtum

323 2010 | Inventory reduction in spare part networks by selective throughput time reduction | M.C. van der Heijden, E.M. Alvarez, J.M.J. Schutten

322 2010 | The selective use of emergency shipments for service-contract differentiation | E.M. Alvarez, M.C. van der Heijden, W.H. Zijm

321 2010 | Heuristics for Multi-Item Two-Echelon Spare Parts Inventory Control Problem with Batch Ordering in the Central Warehouse | B. Walrave, K. v. Oorschot, A.G.L. Romme

320 2010 | Preventing or escaping the suppression mechanism: intervention conditions | Nico Dellaert, Jully Jeunet.

319 2010 | Hospital admission planning to optimize major resources utilization under uncertainty | R. Seguel, R. Eshuis, P. Grefen.

318 2010 | Minimal Protocol Adaptors for Interacting Services | Tom Van Woensel, Marshall L. Fisher, Jan C. Fransoo.

317 2010 | Teaching Retail Operations in Business and Engineering Schools | Lydie P.M. Smets, Geert-Jan van Houtum, Fred Langerak.

316 2010 | Design for Availability: Creating Value for Manufacturers and Customers | Pieter van Gorp, Rik Eshuis.

315 2010 | Transforming Process Models: executable rewrite rules versus a formalized Java program | Bob Walrave, Kim E. van Oorschot, A. Georges L. Romme

314 2010 | Getting trapped in the suppression of exploration: A simulation model | S. Dabia, T. van Woensel, A.G. de Kok

A Dynamic Programming Approach to Multi-Objective Time-Dependent Capacitated Single Vehicle Routing Problems with Time Windows

294	2009	Fujaba hits the Wall(-e)	Pieter van Gorp, Ruben Jubeh, Bernhard Grusie, Anne Keller
293	2009	Implementation of a Healthcare Process in Four Different Workflow Systems	R.S. Mans, W.M.P. van der Aalst, N.C. Russell, P.J.M. Bakker
292	2009	Business Process Model Repositories - Framework and Survey	Zhiqiang Yan, Remco Dijkman, Paul Grefen
291	2009	Efficient Optimization of the Dual-Index Policy Using Markov Chains	Joachim Arts, Marcel van Vuuren, Gudrun Kiesmuller
290	2009	Hierarchical Knowledge-Gradient for Sequential Sampling	Martijn R.K. Mes; Warren B. Powell; Peter I. Frazier
289	2009	Analyzing combined vehicle routing and break scheduling from a distributed decision making perspective	C.M. Meyer; A.L. Kok; H. Kopfer; J.M.J. Schutten
288	2009	Anticipation of lead time performance in Supply Chain Operations Planning	Michiel Jansen; Ton G. de Kok; Jan C. Fransoo
287	2009	Inventory Models with Lateral Transshipments: A Review	Colin Paterson; Gudrun Kiesmuller; Ruud Teunter; Kevin Glazebrook
286	2009	Efficiency evaluation for pooling resources in health care	P.T. Vanberkel; R.J. Boucherie; E.W. Hans; J.L. Hurink; N. Litvak
285	2009	A Survey of Health Care Models that Encompass Multiple Departments	P.T. Vanberkel; R.J. Boucherie; E.W. Hans; J.L. Hurink; N. Litvak
284	2009	Supporting Process Control in Business Collaborations	S. Angelov; K. Vidyasankar; J. Vonk; P. Grefen
283	2009	Inventory Control with Partial Batch Ordering	O. Alp; W.T. Huh; T. Tan
282	2009	Translating Safe Petri Nets to Statecharts in a Structure-Preserving Way	R. Eshuis
281	2009	The link between product data model and process model	J.J.C.L. Vogelaar; H.A. Reijers
280	2009	Inventory planning for spare parts networks with delivery time requirements	I.C. Reijnen; T. Tan; G.J. van Houtum
279	2009	Co-Evolution of Demand and Supply under Competition	B. Vermeulen; A.G. de Kok
278	2010	Toward Meso-level Product-Market Network Indices for Strategic Product Selection and (Re)Design Guidelines over the Product Life-Cycle	B. Vermeulen, A.G. de Kok
277	2009	An Efficient Method to Construct Minimal Protocol Adaptors	R. Seguel, R. Eshuis, P. Grefen
276	2009	Coordinating Supply Chains: a Bilevel Programming Approach	Ton G. de Kok, Gabriella Muratore
275	2009	Inventory redistribution for fashion products under demand parameter update	G.P. Kiesmuller, S. Minner
274	2009	Comparing Markov chains: Combining aggregation and precedence relations applied to sets of states	A. Busic, I.M.H. Vliegen, A. Scheller-Wolf

Working Papers published before 2009 see: http://beta.ieis.tue.nl