University of Groningen

The Repair Kit Problem with positive replenishment lead times and fixed ordering costs

Prak, Derk; Saccani, Nicola; Syntetos, Aris A.; Teunter, Rudolf; Visintin, Filippo

Published in:
European Journal of Operational Research

DOI:
[10.1016/j.ejor.2017.02.019](https://doi.org/10.1016/j.ejor.2017.02.019)

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

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

Publication date:
2017

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

Citation for published version (APA):

Prak, D., Saccani, N., Syntetos, A. A., Teunter, R., & Visintin, F. (2017). The Repair Kit Problem with positive replenishment lead times and fixed ordering costs. *European Journal of Operational Research*, 261(3), 893-902. <https://doi.org/10.1016/j.ejor.2017.02.019>

Copyright

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

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

Take-down policy

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

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

Production, Manufacturing and Logistics

The Repair Kit Problem with positive replenishment lead times and fixed ordering costs

Dennis Prak^{a,*}, Nicola Saccani^b, Aris Syntetos^c, Ruud Teunter^a, Filippo Visintin^d^a Department of Operations, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands^b Rise Lab – Department of Industrial and Mechanical Engineering, University of Brescia, Via Branze 38, 25123 Brescia, Italy^c Panalpina Centre for Manufacturing and Logistics Research, Cardiff Business School, Cardiff University, Cardiff CF10 3EU, United Kingdom^d IBIS Lab – Department of Industrial Engineering, University of Florence, Viale Morgagni 44, 50134 Florence, Italy

ARTICLE INFO

Article history:

Received 21 January 2016

Accepted 14 February 2017

Available online 24 February 2017

Keywords:

OR in service industries

Inventory

Repair kit problem

Lead times

Service parts

ABSTRACT

The Repair Kit Problem (RKP) concerns the determination of a set of items taken by a service engineer to perform on-site product support. Such a set is called a kit. Models developed in the literature have always ignored the lead times associated with delivering items to replenish the kit, thereby limiting the practical relevance of the proposed solutions. Motivated by a real life case, we develop a model with positive lead times to control the replenishment quantities of the items in the kit, and study the performance of (s, S) policies under a service objective. The choice for (s, S) policies is made in order to accommodate fixed ordering costs. We present a method to calculate job fill rates with exact expressions, and discuss a heuristic approach to optimize the reorder level and order-up-to level for each item in the kit. The empirical utility of the model is assessed on real world data from an equipment manufacturer and useful insights are offered to after-sales managers.

© 2017 Elsevier B.V. All rights reserved.

1. Motivation

After sales service is an area of strategic importance in today's competitive manufacturing environment. Manufacturers (or service providers) of durable goods that service many clients heavily invest in field service workforce, equipment and inventories of service parts to sustain their businesses (Gebauer et al., 2005; Lay, 2014). The management of such resources is challenging (Cohen et al., 2006). There are considerable opportunities to reduce costs for companies in the after-sales industry through improved service parts management (Syntetos et al., 2009; Guajardo & Rönnqvist, 2015). The tremendous investments in service parts signify that small improvements in this area may translate into dramatic financial benefits. In this paper, we are concerned with the Repair Kit Problem (RKP).

In field service, engineers typically perform daily tours along various job sites, and carry in their vehicles a kit that consists of one or more units of several items that can be required to complete the service jobs. The RKP aims at finding the optimal set of items to be stocked in the vehicle, and the order strategies for

the different items. In its most extended form, the engineer needs to complete a stochastic number of jobs, where each job requires stochastic numbers of units of different items. There are two model formulations for the RKP. *Cost models* consider the minimization of the total cost consisting of a penalty cost for each incomplete job (Smith, 1980), and holding costs for each item in the repair kit. *Service models* consider the minimization of the holding costs subject to a service level constraint to be satisfied. The main difference is that cost models attach a monetary cost to incomplete jobs, while service models limit the number of incomplete jobs through a threshold value, i.e. a target service level has to be achieved. Service level models bear greater relevance to real-world applications. Especially the job fill rate measure of service, defined as the fraction of jobs that can be completed at the first visit by using items from the repair kit, closely resembles the service perception from a customer point of view.

To the best of our knowledge all models in the literature are based on the assumption that the replenishment lead times for all items in the kit are zero. This assumption is only justified in practice if all items that are used from a kit on a tour can always be delivered to the kit before the start of the next tour. For daily tours, it implies overnight replenishments of all required items, including expensive slow-moving items. However, to limit inventory costs, large service organizations typically stock the most expensive and least demanded items centrally (e.g. at the European or U.S. central

* Corresponding author:

E-mail addresses: d.r.j.prak@rug.nl, dennisprak@ziggo.nl (D. Prak), nicola.saccani@unibs.it (N. Saccani), syntetosa@cardiff.ac.uk (A. Syntetos), r.h.teunter@rug.nl (R. Teunter), filippo.visintin@unifi.it (F. Visintin).

Table 1
Literature overview.

	Single unit	Multi-unit
Single job	Smith et al. (1980), Graves (1982)	Mamer and Smith (1982), March and Scudder (1984), Mamer and Shogan (1987), Brumelle and Granot (1993)
Multi-job	Heeremans and Gelders (1995)	Teunter (2006), Bijvank et al. (2010), Saccani et al. (2017)

warehouse) rather than in each of the operating countries. Transporting those items from the central location to a local warehouse and from there to the repair kit typically cannot be done overnight. This limits the applicability of the models that have been proposed in the literature.

Another important assumption in those models is that ordering costs are ignored. In a system where items have to be ordered, picked and shipped from central warehouses, this assumption is not valid. Instead, there needs to be a fixed ordering cost for each item that is ordered, associated with the replenishment activities (e.g., filling out an order form and picking an item from its stocking location in the warehouse). Note that, for the following two reasons, it is less relevant to include a fixed cost component per order that is shared across items (which is mostly studied in joint replenishment literature). First, items may be shipped from different warehouses. Second, there are typically many different items in a repair kit, which requires that at least one item needs to be replenished in each period. Consequently any fixed cost per order is incurred in every period and so becomes a constant than can be disregarded without loss of generality.

In this paper we extend previous RKP formulations with positive replenishment lead times and fixed ordering costs per item. We study the behaviour of multi-item (s , S) policies and derive exact expressions to calculate the job fill rate for a given policy. Finding the best values for reorder level s_i and order-up-to level S_i for each item i with exact expressions for the RKP requires an exhaustive search procedure which is computationally challenging. Therefore, we propose a heuristic solution procedure. We benchmark its solutions to the best found solution for small instances of the problem, and also compare its performance to that of a heuristic inspired by the current literature that assumes a zero lead time. We then analyse real data from an Italian-based office equipment manufacturer, which has 13 engineers who carry out 1 to 15 jobs per day, and use hundreds of different items to complete their jobs. These items are shipped from two different warehouses and have prices ranging from €0.01 to almost €5,000. Application of the proposed solution procedure to this data set generates interesting insights from a practitioner perspective.

The remainder of the paper is organised as follows: in Section 2 the background literature is presented along with its limitations that are addressed in our work. Section 3 presents our model and the solution procedure. The description of the case study is presented in Section 4, followed by the results of our benchmarks, case study investigation and relevant insights in Section 5. We conclude in Section 6 with the implications of our work and an agenda for further relevant research.

2. Background

Since the early 1980's, the RKP has been studied in various formulations. The original formulation aimed at the minimization of the costs associated with inventory holding and penalties for item shortages (Smith et al., 1980). An item shortage implies a second visit by the field engineer to the customer location to complete the job, which entails extra travel and labour costs as well as loss of customer goodwill. The solution proposed by Smith et al. (1980) ranks items according to the increasing ratio of their inventory holding costs divided by usage frequency. That is, items that are more expensive or demanded less frequently are less likely to

be included in the kit. Smith et al. (1980) only consider the decision of which items to carry to a single job. The job can only require a single unit of each item. Other studies have extended this work. These are listed in Table 1, where they are classified based on the main modelling assumptions: replenishments occur after every job (single job) or after a number of jobs (multi-job), and only one unit of each item can be used during a job (single unit) or multiple units can be used (multi-unit).

The problem, in its most realistic representation, comprises a (stochastic) number of jobs in a tour after which the kit can be replenished. Each job can require a (stochastic) number of different items, where each item can potentially be needed a (stochastic) number of times, i.e. more than one unit may be needed. Table 1 shows that only Teunter (2006), Bijvank et al. (2010) and Saccani et al. (2017) have considered a realistic multi-unit scenario in a multi-job horizon. According to Bijvank et al. (2010), the study by Teunter (2006) may be regarded as a single unit (though multi-item) problem, because a single-unit assumption is made when a closed-form expression for the job fill rate is derived. In fact, although modelling developments rely upon every item being potentially needed in excess of one unit, when calculating a closed-form expression of the job fill rates a single unit assumption is being made. Saccani et al. (2017), instead, have modelled the problem as a deterministic one: they propose an integer linear programming model to determine the composition of the repair kit and assess via simulation the cost incurred if the model solution was applied in a stochastic environment.

There also exist different approaches on how to model the length of a tour in the multi-job problem. Heeremans and Gelders (1995) consider a fixed tour length. Teunter (2006) assumes the tour length to be variable in the problem formulation, but this is restricted to fixed lengths in the expressions to approximate the job fill rate. Bijvank et al. (2010) consider a variable tour length.

To assess the service level, three measures are mentioned in the literature: i) part fill rate: fraction of items needed to perform a service repair that are also available in the repair kit (see e.g., Smith et al., 1980; Teunter, 2006); ii) job fill rate: fraction of jobs completed by using items from the repair kit (completed jobs, as opposed to uncompleted jobs which require that the engineer returns to the customer since some items are missing, see e.g., Smith et al., 1980; Graves, 1982; Mamer & Smith, 1982; March & Scudder, 1984; Mamer & Shogan, 1987; Brumelle & Granot, 1993; Teunter, 2006; Bijvank et al., 2010; Saccani et al., 2017); iii) tour fill rate: fraction of tours composed of only completed jobs. The tour fill rate has been analyzed by Heeremans and Gelders (1995) when a fixed tour size is assumed, but is less relevant in practice compared to the job fill rate, since the latter better corresponds to the service level perceived by customers (Bijvank et al., 2010).

All papers discussed thus far assume that the replenishment lead times are zero. That means that a repair (wo)man always starts a new tour with a completely full repair kit. Equivalently, a kit is always fully stocked after a tour. This implicitly assumes that the local warehouse that replenishes the repair kit has sufficient spare parts inventory available to replenish the kits. This assumption is unrealistic in many business settings. Heeremans and Gelders (1995) only mention in their case study the need to consider (positive) lead times, but they did not include this in their model formulation.

In single-tour problem formulations, some studies consider dependent demand processes for different items (Mamer & Smith, 1982; March & Scudder, 1984; Brummelle & Granot, 1993). In multi-job problems, this is typically not the case. Bijvank et al. (2010) note that dependence is of little relevance in several practical settings. A related stream of literature is that on order fulfillment in multi-item settings. The analogy is that an order for multiple items can only be fulfilled if all types are available in sufficient quantities, whereas a repair can only be completed if all items are available in sufficient quantities. Song (1998) therefore argues that if orders can consist of different items, then the job fill rate (which she calls order fill rate) is a better measure for customer satisfaction than the individual item fill rates. Like in our research, she incorporates positive, constant lead times. However, her model is a simplification of the model we present in two ways. First, Song assumes that orders can be placed after every job (i.e., a single-job setting). Second, she restricts the analysis to base-stock policies, whereas we study more general (s, S) policies.

Our proposed model presents the following features: it is multi-job (with a stochastic number of jobs for each tour), multi-item and multi-unit, with positive lead times, with the assumption of independent demand for different items, and with a job fill rate service measure. We provide a method to exactly compute the job fill rate that is achieved by a given setting of the reorder levels and order-up-to levels, and subsequently discuss a solution procedure to set the reorder level s_i and order-up-to level S_i for every item i . The proposed model and the relevant solution procedure are presented analytically in the next section.

3. Model and analysis

3.1. Model description

During a tour, a stochastic number of jobs can occur according to some probability distribution, where the maximum number of jobs is J . The probability that j jobs occur in a tour is denoted by $q^T(j)$. There is a set of N items i of which a stochastic number of units can be needed at a single job. The probability that k units of item i are needed at a job is denoted by $p_i^J(k)$. The item requirements of the various jobs are i.i.d. At the end of each tour, a replenishment order can be placed. An order placed for item i at the end of tour t arrives at the end of tour $t + L_i$. The holding costs are incurred after the end of the tour and, if applicable, after order arrival.

The objective is to minimize the sum of the average holding and ordering costs per tour under the constraint that a minimum job fill rate is achieved. The unit holding cost per tour for item i is denoted by h_i and the fixed ordering cost for item i is f_i . Corresponding to our cost structure, we consider policies specified by a reorder level s_i and an order-up-to level S_i . This means that at the end of each tour an order is placed to raise the inventory position of item i to level S_i , if the inventory position reached to or below s_i .

The probability $p_i^{T|j}(k)$ that k units of item i are needed on the first j jobs of a tour is

$$p_i^{T|j}(k) = \sum_{k_1, \dots, k_j \mid k_1 + \dots + k_j = k} \left(\prod_{l=1}^j p_i^J(k_l) \right), \\ 1 \leq i \leq N, \quad 1 \leq j \leq J, \quad k \geq 0.$$

The probability $p_i^T(k)$ that k items of type i are needed on a tour is

$$p_i^T(k) = \sum_{j=1}^J \left(q^T(j) p_i^{T|j}(k) \right), \quad 1 \leq i \leq N, \quad k \geq 0.$$

The probability $p_i^L(k)$ that k units of item i are needed during the L_i tours of the (item-specific) lead time is

$$p_i^L(k) = \sum_{k_1, \dots, k_{L_i} \mid k_1 + \dots + k_{L_i} = k} \left(\prod_{l=1}^{L_i} p_i^T(k_l) \right), \quad 1 \leq i \leq N, \quad k \geq 0.$$

The derivation of the inventory level distribution and the job fill rate follows the logic of that on page 113 of Axsäter (2006) for single-item (s, S) policies. We extend it to allow that the inventory levels of different items together influence the probability that a job is completed successfully. We furthermore incorporate the fact that a tour can consist of multiple jobs. Let us refer to the time interval between two replenishment orders for item i as an order cycle. We first derive the steady-state distribution of the inventory position. Let $\pi_i(k)$ denote the probability that the inventory position reaches level k , $k = s_i + 1, s_i + 2, \dots, S_i$, (at the start of one or more tours) during an order cycle. Obviously, the inventory position at the start of the first tour of an order cycle (after the order is placed) is S_i and so we have $\pi_i(S_i) = 1$. The other probabilities can be determined recursively (in reverse order from $k = S_i - 1$ to $k = s_i + 1$) using:

$$\pi_i(k) = \sum_{l=k+1}^{S_i} \pi_i(l) \frac{p_i^T(l-k)}{1 - p_i^T(0)}, \quad 1 \leq i \leq N, \quad k = s_i + 1 \leq k \leq S_i - 1.$$

where the division by $1 - p_i^T(0)$ is needed because a new level is only reached after a tour with positive demand.

To derive the steady-state distribution of the inventory position, observe the following. The number of jobs per tour and the number of units per item that is demanded on a job both have a fixed distribution. Therefore, the average number of tours between two demands for a certain item is fixed, and we can proceed similar to Axsäter (2006) to obtain the probability $\pi_i^{IP}(k)$ that the inventory position of item i equals k units. We find

$$\pi_i^{IP}(k) = \pi_i(k) \left/ \sum_{l=s_i+1}^{S_i} \pi_i(l) \right., \quad 1 \leq i \leq N, \quad k = s_i + 1 \leq k \leq S_i.$$

We finally derive the inventory level distribution from the inventory position distribution. Recall from the system description that an order placed at the end of tour t arrives in tour $t + L_i$. Hence, the inventory level (after the order arrival, as specified in the order of events) in tour $t + L_i$ equals the inventory position in tour t minus the lead time demand over the L_i tours in interval $[t + 1, t + L_i]$. Let $\pi_i^{IL}(k)$ denote the probability that the inventory level of item i equals k units at the end (or equivalently, at the beginning) of a tour. This gives

$$\pi_i^{IL}(k) = \sum_{l=\max\{s_i+1, k\}}^{S_i} \pi_i^{IP}(l) p_i^L(l-k), \quad 1 \leq i \leq N, \quad k \leq S_i.$$

This calculation holds under the assumption that each requested item is taken from the kit, also if the job cannot be completed immediately. This assumption is also made in Teunter (2006). The expected holding cost per tour is obtained as

$$EHC = \sum_{i=1}^N \left[h_i \left(\sum_{k=1}^{S_i} k \pi_i^{IL}(k) \right) \right].$$

To derive the expected ordering cost per tour, we first derive the probability $\pi_i^O(k)$ that k units of item i are ordered at the end of a tour. An order for item i can only be placed if the inventory position in the preceding tour is at level $l > s_i$ and at least $l - s_i$ units of item i are needed in the current tour. Moreover, if $l - s_i + k$ items are needed in the current tour, then the order

quantity is $S_i - l + (l - S_i + k) = k$. Hence,

$$\pi_i^Q(k) = \sum_{l=S_i+1}^{S_i} \pi_i(l) p_i^T(l - S_i + k), \quad 1 \leq i \leq N, \quad k \geq S_i - s_i.$$

The expected number of orders per tour is equal to the ratio of the expected demand per tour and the expected order quantity, and so we get an expected ordering cost of

$$EOC = \sum_{i=1}^N \left[f_i \frac{E[T]E[D_i^l]}{E[Q]} \right],$$

where $E[T] = \sum_{j=1}^J j q^T(j)$ is the expected tour size, $E[D_i^l] = \sum_{k=1}^{\infty} k p_i^l(k)$ is the expected demand for item i per job, and $E[Q] = \sum_{k \geq S_i - s_i} k \pi_i^Q(k)$ is the expected order size for item i .

The expected total cost per tour is

$$ETC = EHC + EOC.$$

Let p_j^c denote the probability that the j -th job of a tour is completed. Note that the first job of a tour is completed if and only if, for each item i , the quantity needed is at most equal to the quantity in the kit at the start of the tour (i.e. the inventory level at the end of the previous tour). Let $p_i^l(l) = \sum_{j=0}^l p_i^l(j)$ denote the probability that at most l items of type i are needed on a single job. We obtain the probability of completing the first job:

$$p_1^c = \prod_{i=1}^N \left[\sum_{l=0}^{S_i} \pi_i^{ll}(l) p_i^l(l) \right].$$

The j -th job of a tour is completed if the number of units of all items that are needed for the j -th job does not exceed the number of units in the kit at the start of the tour minus the number of units that were needed in the previous $j - 1$ jobs. Let $p_i^{Tlj}(k) = \sum_{j=0}^k p_i^{Tlj}(j)$ denote the probability that at most k units of item i are needed for the first j jobs. Consequently,

$$p_j^c = \prod_{i=1}^N \left[p_i^l(0) + \sum_{l=1}^{S_i} \sum_{n=0}^{S_i-l} \sum_{m=l+n}^{S_i} p_i^l(l) p_i^{Tlj-1}(n) \pi_i^{ll}(m) \right], \quad 2 \leq j \leq J.$$

The overall job fill rate now results as the ratio of the expected number of completed jobs on a tour and the expected number of jobs on a tour. We use the relationship $E[X] = \sum_{j=0}^{\infty} P(X \geq j)$, and the fact that the probability that at least j jobs are completed is given by the product of the probability that at least j jobs occur and the probability that the on hand stock is enough to satisfy the demand for j jobs. Let $Q^T(k) = \sum_{j=0}^k q^T(j)$. This results in

$$JFR = \frac{\sum_{j=1}^J [(1 - Q^T(j-1)) p_j^c]}{\sum_{j=1}^J (1 - Q^T(j-1))}.$$

We end this section by listing the notations that have been introduced, for ease of reference in what follows.

Policy parameters:

s_i : Reorder level of item i
 S_i : Order-up-to level of item i

System parameters:

L_i : Lead time of item i in tours
 h_i : Holding cost of item i per tour
 f_i : Fixed ordering cost of item i
 J : Maximum number of jobs per tour
 $q^T(j)$: Probability of j jobs in a tour
 $p_i^l(k)$: Probability that k units of item i are needed to complete a single job

Other notations (useful for the analysis):

$p_i^{Tlj}(k)$: Probability that k units of item i are needed in total for j jobs
 $p_i^T(k)$: Probability that k units of item i are needed in a tour
 $p_i^l(k)$: Probability that k units of item i are needed in the L_i tours of the lead time
 p_j^c : Probability that the j -th job of a tour is directly completed
 $\pi_i(k)$: Probability that the inventory position of item i reaches level k between two orders
 $\pi_i^{lp}(k)$: Probability that the inventory position of item i equals k units before the tour starts
 $\pi_i^{ll}(k)$: Probability that the inventory level of item i equals k units after the tour ends
 $\pi_i^Q(k)$: Probability that replenishment order size for item i equals k units
 EHC : Expected holding cost per tour
 EOC : Expected ordering cost per tour
 ETC : Expected total cost per tour

3.2. Heuristic solution procedure

With the model discussed in Section 3.1, it is possible to calculate the expected cost per tour and job fill rate for a particular repair kit, i.e. with a particular setting of order levels s_i and order-up-to levels S_i , for items $i = 1, \dots, N$. However, finding the optimal solution involves solving a constrained integer maximization problem with $2N$ decision variables, with both the objective cost function and the job fill rate constraint non-linear. Performing an exhaustive search procedure will quickly become impossible when the number of items increases. Recall from Section 2 that (in a different setting) Song (1998) studies a simplified version of our model, and discusses that even for that model exact evaluation of job fill rates is hard, let alone optimizing reorder levels and order-up-to levels.

In this section we therefore present a heuristic solution procedure based on what others have proposed in the literature for simpler systems. As is common in the repair kit literature (see e.g. Graves, 1982; Teunter, 2006; Bijvank et al., 2010), we propose a greedy marginal analysis procedure and generalize this idea to the dynamic setting where both the order level and order-up-to level of all items have to be defined.

It is common practice to calculate s_i and S_i jointly (see, e.g., Silver et al., 1998 and Axsäter, 2006). We will follow the same approach and first decide on the difference $S_i - s_i$, which we denote by Q_i for ease of writing. We set this lower bound on the order size equal to the well-known and often used economic order quantity, but ensure that every order has at least the size of the minimum possible positive demand for an item per tour, denoted by $K_i^{\min}$. So, we set

$$Q_i \equiv S_i - s_i = \max \left\{ \left\lceil \sqrt{\frac{2f_i \sum_{k=0}^{\infty} k p_i^l(k)}{h_i}} \right\rceil, K_i^{\min} \right\},$$

where $\lceil \cdot \rceil$ is the rounding operator.

The algorithm is described as follows. Define for reorder levels $s = s_1, \dots, s_N$ and order-up-to levels $S = S_1, \dots, S_N = s_1 + Q_1, \dots, s_N + Q_N$, the achieved JFR by $JFR(s)$ and the expected

holding cost by $EHC(s)$. Define the maximum demanded quantity of item i by K_i^{max} , the desired JFR by JFR^* , and the i -th unit vector by e_i .

```
# Initialization:
Set  $s_i = -Q_i$ ,  $S_i = 0$  for all items  $i$ 
Calculate  $JFR(s)$  and  $EHC(s)$ 

# Repetition:
while  $JFR(s) < JFR^*$ 
  for all items  $i$ 
    for all quantities  $j=1, \dots, K_i^{max}$ 
      set  $s_i^{new} = s + je_i$ ,  $S_i^{new} = S + je_i$ 
      calculate  $incr(i, j) = (JFR(s_i^{new}) - JFR(s)) / (EHC(s_i^{new}) - EHC(s))$ 
    find  $i^*, j^* = \text{argmax}(incr)$ 
    set  $s = s + je_{i^*}$ ,  $S = S + je_{j^*}$ 
```

Note that once the job fill rate and expected costs for the initial solution are computed, it is not necessary to completely redo all computations at every iteration of the algorithm. First, we observe that $p_i^{Tij}(k)$, $p_i^T(k)$ and $p_i^L(k)$ do not depend on the value of s_i and S_i , so that they have to be computed only once and can then be stored for later reference. Also, when changing s_i and S_i for one item, only the probability that enough units of that particular item are in stock changes. The probabilities corresponding to the other items are not affected. Likewise, the cost associated with the new solution only changes due to the extra holding cost related to the item of which the order level is changed.

4. Case study description

The problem addressed in this study is motivated by a real case of a multinational company that manufactures and sells printing equipment for personal, office and professional use. Spare parts and services (including pay-per-use contracts, where customers are charged based on usage) account for about 50% of the company turnover. The field service organization in Europe consists of around 3,000 engineers. For the empirical application we focus on the Italian subsidiary of the company, and in particular on the 13 engineers operating in the Milan-area. This region is characterized by a large number of products in the field and by the presence of several key customers.

The geographical area is divided into zones, each one assigned to a specific engineer, although some overlap occurs occasionally: when products requiring specific skills have to be served, or there is a high concentration of demand in some zones, engineers may move from their regular zone to a different one. Engineers carry a stock of service parts (repair kit) to perform on-site product support. The items in the repair kits, their reorder levels and reorder quantities are currently based on the usage in the past six months and extra technical or market information (e.g. failure rates and sales of new products). These data differ across engineers, since they serve different zones with a different product mix. Thus repair kits are specific to each engineer.

If an engineer cannot fix a product failure during her/his first visit to a customer due to a service parts shortage, the engineer places an order for the missing item(s) and performs a return visit after restocking. This generates additional labour and travel costs, and may even entail a penalty (which is currently the case for very few of the company's customers).

Replenishment orders are fulfilled from a national warehouse if possible, and from a central European warehouse located in the Netherlands, otherwise. Items that are rarely used and very expensive are stored only in the central European warehouse. This allocation decision is made by the European headquarters based on an 80–20 rule, which means that around 80% of the volumes are replenished locally and the remaining 20% from the central European warehouse. This does not imply that 80% of the items are kept in

the local warehouse, since many items have sporadic demand patterns, whereas a small number of items are demanded more frequently. Lead times are 2 days from the national warehouse and 3 days from the central one.

The dataset includes the job requests handled by the 13 engineers, the day/time the job occurred and the items required during the years 2010 and 2011. The jobs have been carried out over a set of 3,079 machines, belonging to 1,375 different customers. The 13 engineers carried out between 1 and 15 jobs per day in the studied period, with an average of 1.82 jobs per day (tour) and a standard deviation of 0.35. On average, each engineer used 511 different items, with an average usage rate of 3.52 items per job and a standard deviation of 0.82. Overall, the job requests in the dataset concern 3,584 items: 507 are stored at the national warehouse and the remaining at the central European warehouse. The average item price is €55.4, with a standard deviation of €153.9: 40% of items have a unit cost below €5. Based on discussions with the company, for the holding cost h_i we assume a value equal to 5% of the item price on a yearly basis and so 5/365% per day (tour), and a fixed ordering cost $f_i = €1$ for every item i .

The company does not implement lateral transshipments, i.e. an engineer cannot be replenished from the repair kit of another engineer. Therefore, we consider engineers as separate entities.

5. Results

In this section, we first study the performance of the solution procedure developed in Section 3.2 in several set-ups and then demonstrate it using the case study data described in Section 4. The performance is analysed in Section 5.1. We compare, in (fictive) test instances, the cost resulting from the solution procedure to the cost corresponding to the best solution that is found by an extensive search over all possible combinations of s_i and S_i for all items i . These test instances can only contain a small number of different items, as the search space increases exponentially in this number. We furthermore compare the results of our solution procedure (denoted “H1”) to those of a simpler procedure that, in each iteration, weighs increase in the job fill rate increase against the item price (without computing the holding cost increase by evaluating the effect on the inventory level distribution, denoted “H2”). This simpler heuristic is inspired by earlier literature that assumes a zero lead time and therefore considers a simpler type of inventory model (e.g. Smith et al. 1980). In Section 5.2, we apply solution procedure H1 to the case described in Section 4, again comparing its performance to that of the simpler heuristic H2, and elaborating on the results to gain further insights into the solution.

5.1. Performance of the proposed solution procedure

We first compare the performance of the solution procedures H1 and H2 to the best found solution in 192 different test instances. The test instances comprise of all possible combinations of the following parameter values: ordering cost 0 and 1, lead time 1 and 2 tours, maximum number of jobs per tour 2 and 5, number of different items 2 and 3, maximum item demand per job 1, 2, 3 and 4 units, and target job fill rate 80%, 95%, and 99%. The number of jobs per tour and the number of units per item demanded per job are uniformly distributed. The results are summarized in Table 2. Every entry in Table 2 comprises 12 scenarios, with maximum 1, 2, 3, and 4 items demanded per item, and with target job fill rates 80%, 95%, and 99%. In every category, “H1” corresponds to the solution procedure from Section 3, whereas “H2” corresponds to the simpler heuristic. The first entry is the average percentage cost difference between H1 and H2, respectively, and the best found solution. The second entry, placed between brackets, is the number of instances (out of the 12 instances corresponding to that entry)

Table 2

Cost difference of both heuristics with the best solution found in small instances.

				Lead time (tours)							
				1				2			
				Nr. of different items				Nr. of different items			
				2		3		2		3	
				H1	H2	H1	H2	H1	H2	H1	H2
Ordering cost (€)	0	Max. jobs/tour	2	3.6% (3)	3.6% (3)	4.6% (2)	4.6% (2)	1.8% (5)	3.6% (4)	3.3% (3)	3.6% (2)
			5	2.5% (3)	3.7% (3)	2.5% (1)	2.5% (1)	2.0% (1)	6.2% (1)	1.4% (1)	6.7% (1)
	1	Max. jobs/tour	2	2.2% (0)	2.8% (0)	1.3% (2)	1.4% (2)	2.3% (1)	4.0% (1)	1.4% (2)	3.0% (1)
			5	1.4% (0)	2.7% (0)	0.9% (0)	1.9% (0)	1.5% (1)	6.8% (0)	1.1% (0)	7.2% (0)

in which H1 and H2, respectively, found the solution that was also found by the exhaustive search procedure.

The more involved heuristic H1, as expected, performs as least as well as the simpler heuristic H2. The difference is largest in instances where cost calculations become more complex, e.g. where the lead time is larger, where there are positive ordering costs, and where the number of jobs per tour can be large. Also, especially in the scenarios with a longer lead time, H1 finds the optimal solution more often than H2. There is no scenario in which H2 finds the optimal solution, whereas H1 does not. The computation times of both heuristics are similar and less than a minute, typically less than 10 seconds (executed in R, using a single core of a 2.5 GHz Intel® Xeon® E5 2600v3 processor). For these small instances, there is no significant difference measurable between the computation times of both heuristics.

5.2. Case study results

From the case study data, for each engineer, we set the system parameters as indicated in Section 3 and subsequently derive the probability distributions $p_{i,j}^{DP}(k)$, $p_i^{DT}(k)$, and $p_i^{DL}(k)$. Then we perform the solution procedure from Section 3.2, denoted “H1”. For comparison, we also include the simpler heuristic described in Section 5.1, denoted “H2”.

Table 3 shows the achieved job fill rate, expected holding and ordering costs, and the average on-hand stock investment, for all engineers. We first observe that the “fill rate overshoot” of H1 is at most 0.5%, and typically 0.2% or less. Especially for higher target job fill rates this overshoot is small. The average on-hand stock investment (and thus the average daily holding cost) increases with the job fill rate, and the return on investment when targeting high fill rates is decreasing. The average on-hand stock investment required to achieve a job fill rate of 80% ranges from around €9,000 to €22,000, with an average of €16,000, whereas for achieving a job fill rate of 99%, the stock investments grow up to €55,000, with an average of €37,000, on average more than doubling the stock required compared to the 80% job fill rate target. Since we keep $S_i - s_i$ fixed for every item, the average daily ordering cost does not differ between the solutions for various job fill rates.

The simpler heuristic H2 leads to larger average on-hand stock investments (and thus larger holding costs) for every engineer and for every job fill rate setting. The relative cost increase compared to H1 is largest for the cases with a target job fill rate of 80%. Indeed, on average over all engineers, the on-hand stock investment more than doubles for those cases if H2 is applied instead of H1. For a target job fill rate of 99%, the average cost gap is 28%. The simpler heuristic H2 does not take into account that if an item is more likely to be used on a tour, it adds less to the expected holding cost. As a result, it is too “reluctant” to add expensive items to the kit, even if those items are very likely to be needed. For larger required job fill rates, this is less of an issue as relatively many items need to be added to the kit, even expensive items that are

not very likely to be needed. For such additions (in the last few iterations), the added holding cost approaches that of always having an item in the kit.

H2 does have a smaller job “fill rate overshoot” on average, which can also be explained by the fact that it favours adding inexpensive items (repeatedly). The ordering costs resulting from both heuristics are equal, as both use the same difference $S_i - s_i$, namely the EOQ. Computation times vary from 22 minutes (H1) and 24 minutes (H2) for engineer 13, to 145 minutes (H1) and 156 minutes (H2) for engineer 2. Whereas the simpler heuristic does fewer computations in every iteration (since the expected holding costs do not need to be computed), it typically takes more iterations before the desired job fill rate is achieved. Therefore, interestingly, for most engineers the simpler benchmark heuristic H2 has a longer computation time than the proposed heuristic H1. Observe that there is no one-to-one relationship between the number of different items and the computation time. The major determinant of the computation time is the number of iterations that is needed to achieve the target job fill rate. The number of iterations depends not only on the number of different items, but also on the demands for the different items by the jobs. Since the more sophisticated heuristic H1 clearly outperforms the simpler heuristic H2, we use H1 in the remainder of this section to gain more insights from the case study data.

In Fig. 1 we graphically depict the relationship between the average on-hand stock investment and the achieved job fill rate for two specific engineers, exemplary of the two types of relationships found in the sample, and for the sample average.

Fig. 1 shows an S-type curve for engineer 1, illustrating that the first few iterations of the solution procedure lead to a relatively small job fill rate increase, then the return on investment (in terms of job fill rate increase for a stock increase) grows, and it decreases again as the average investment value further increases. This can be explained by the fact that many jobs need multiple different items, so that, for low stock levels, adding only one or a few items does not significantly increase the job fill rate. When the repair kit becomes larger on average, increasing the inventory on hand leads to a greater improvement in the job fill rate, as other items that may be needed for completing a job are already available. Finally, to achieve very high job fill rates, the stock levels have to be large enough to satisfy job numbers and job requirements in the tail of the distribution, so that a relatively large investment is needed. For the typical job fill rate settings of 50% and larger, we see the familiar relationship with decreasing returns on investment.

For engineer 2, Fig. 1 instead shows no S-type curve, indicating that the first additions to the repair kit lead to the largest increase in job fill rate, and the return on investment decreases as the investment value increases. This reflects the fact that engineer 2 typically handles jobs that require few items: therefore, even for very low stock levels a stock increase generates an important job fill rate improvement. This relationship reflects the one illustrated in the case example by Bijvank et al. (2010). The dynamics of the

Table 3
Case study analysis results.

Engineer ID	# of items (N)	Target JFR	Achieved JFR		Average daily holding cost (€)		Average daily ordering cost (€)		Average on-hand stock investment (€)	
			H1	H2	H1	H2	H1	H2	H1	H2
1	N = 526	80%	80.3%	80.1%	2.87	6.79	2.14	2.14	20,949	49,562
		95%	95.1%	95.0%	5.03	7.57	2.14	2.14	36,753	55,295
		99%	99.1%	99.0%	7.14	8.95	2.14	2.14	52,121	65,367
2	N = 486	80%	80.2%	80.0%	1.62	3.95	1.65	1.65	11,848	28,812
		95%	95.2%	95.1%	2.80	4.42	1.65	1.65	20,474	32,231
		99%	99.2%	99.0%	3.77	4.95	1.65	1.65	27,497	36,120
3	N = 409	80%	80.5%	80.1%	1.86	4.04	1.73	1.73	13,614	29,528
		95%	95.1%	95.0%	2.93	4.35	1.73	1.73	21,395	31,780
		99%	99.0%	99.0%	3.99	4.83	1.73	1.73	29,093	35,281
4	N = 415	80%	80.2%	80.1%	1.56	3.34	1.47	1.47	11,393	24,356
		95%	95.1%	95.0%	2.62	3.66	1.47	1.47	19,137	26,747
		99%	99.2%	99.0%	3.40	4.15	1.47	1.47	24,825	30,277
5	N = 578	80%	80.1%	80.1%	1.85	3.64	1.89	1.89	13,509	26,536
		95%	95.0%	95.0%	2.80	4.18	1.89	1.89	20,406	30,509
		99%	99.0%	99.0%	3.65	4.84	1.89	1.89	26,638	35,329
6	N = 829	80%	80.0%	80.0%	2.92	8.13	2.54	2.54	21,336	59,331
		95%	95.0%	95.0%	5.40	8.86	2.54	2.54	39,415	64,692
		99%	99.0%	99.0%	7.60	9.92	2.54	2.54	55,447	72,431
7	N = 326	80%	80.0%	80.0%	1.28	3.05	1.58	1.58	9,340	22,265
		95%	95.0%	95.0%	2.05	3.51	1.58	1.58	14,985	25,602
		99%	99.1%	99.1%	2.85	4.32	1.58	1.58	20,790	31,569
8	N = 359	80%	80.4%	80.0%	2.14	4.37	1.66	1.66	15,648	31,867
		95%	95.1%	95.1%	3.44	5.21	1.66	1.66	25,104	38,007
		99%	99.1%	99.0%	4.89	6.20	1.66	1.66	35,670	45,228
9	N = 468	80%	80.4%	80.3%	3.02	5.95	2.08	2.08	22,064	43,442
		95%	95.2%	95.1%	5.16	7.09	2.08	2.08	37,633	51,723
		99%	99.0%	99.0%	6.61	8.16	2.08	2.08	48,236	59,556
10	N = 497	80%	80.2%	80.0%	2.39	4.87	1.87	1.87	17,463	35,524
		95%	95.1%	95.0%	3.86	5.55	1.87	1.87	28,174	40,498
		99%	99.1%	99.0%	4.98	7.16	1.87	1.87	36,375	52,279
11	N = 768	80%	80.1%	80.0%	2.88	5.51	2.35	2.35	21,018	40,220
		95%	95.0%	95.1%	4.57	6.03	2.35	2.35	33,383	43,996
		99%	99.1%	99.0%	5.80	6.96	2.35	2.35	42,325	50,829
12	N = 477	80%	80.1%	80.0%	2.61	5.83	2.05	2.05	19,020	42,550
		95%	95.0%	95.0%	4.65	6.51	2.05	2.05	33,947	47,521
		99%	99.0%	99.0%	6.13	7.31	2.05	2.05	44,760	53,352
13	N = 449	80%	80.1%	80.6%	2.03	4.89	1.79	1.79	14,804	35,690
		95%	95.0%	95.0%	3.66	5.29	1.79	1.79	26,717	38,584
		99%	99.1%	99.0%	5.04	6.50	1.79	1.79	36,791	47,464
Average		80%	80.2%	80.1%	2.23	4.95	1.91	1.91	16,308	36,130
		95%	95.1%	95.0%	3.77	5.55	1.91	1.91	27,502	40,553
		99%	99.1%	99.0%	5.07	6.48	1.91	1.91	36,967	47,314

curve of the engineer sample average are similar to that of the curve of engineer 2, showing that on average the marginal benefit of adding an extra item to the repair kit is decreasing. As the jobs typically require few items, it is not surprising that the sample average does not show an S-type curve, but rather a concave increasing curve.

The order in which items are added to the repair kit is not trivial. Fig. 2 shows that the majority of added items are relatively inexpensive in the first iterations, and the emphasis shifts towards more expensive items when the job fill rate is already high in the final iterations. However, Fig. 2 also shows that there are exceptions to this rule. Therefore, a policy that would focus on item prices only would be suboptimal. This finding is in line with the underperformance of the simpler heuristic H2 in the benchmark instances. Contrarily, the order in which items should be added (the order in which items contribute maximally to the job fill rate at minimum extra cost) depends in a non-trivial manner on the item price and the item's demand distribution, as elaborated in Section 3. Furthermore, this relationship is dynamic, as items that have already been added have to be reconsidered for adding again

in later stages, but their current order levels influence the job fill rate increase that can be established by adding them again.

An important feature of the heuristic algorithm discussed in Section 3 is that it allows for increasing the order levels by values larger than one in a single iteration, thereby acknowledging the possibility that items may typically be needed in certain multiples. This contrasts with previous literature on the single-unit variant of the RKP. For engineer 1, when aiming for a job fill rate of 99%, 534 of in total 3968 iterations (13.5%) increase s_i and S_i by more than one unit. To illustrate the relevance of allowing for adding multiple units of an item to the repair kit at once, consider Fig. 3.

Fig. 3 shows that for the company studied in this research, a heuristic that increases s_i and S_i by only one unit at a time does not yield desirable results. Firstly, the investment value needed to achieve a given job fill rate is substantially higher. In several iterations, adding multiple units of a certain item is the preferred option, whereas adding only one unit of that item yields a lower increase in the ratio of job fill rate and holding cost than adding a single unit of another item. Then, a heuristic that only considers adding a single unit in every iteration selects the other item

Fig. 1. Fill rate versus investment value.

Fig. 2. Item prices and order of addition to the repair kit (Engineer 1).

to be added to the repair kit, which eventually leads to a worse solution. Even more importantly, items of which several units are always needed simultaneously are never added to the kit and fill rates higher than some threshold (of 71% in this case) will never be achieved as a result.

Finally, we consider the sensitivity of the solution with respect to the choice of the difference $S_i - s_i$. Recall that we set $S_i - s_i$ equal to the economic order quantity (rounded to the nearest multiple of the minimum positive item requirement in a job). Fig. 4 depicts the relationship between the job fill rate and average daily cost for various order quantity settings. The average daily cost consists of holding costs (5/365% of the item price) and ordering costs (€1 per order). Clearly, the costs needed to achieve a given fill rate

when $S_i - s_i$ is set lower than the EOQ, are higher than the cost achieved when it is set equal to the EOQ. Contrarily, if $S_i - s_i$ is set higher than the EOQ, then low job fill rates (lower than 62%) are achieved at lower costs, but high job fill rates (higher than 62%) are achieved at higher costs than if the EOQ was chosen. Arguably, job fill rates higher than 62% are most relevant from a practical perspective. We conclude that setting $S_i - s_i$ equal to the EOQ is (nearly) optimal for high job fill rates.

6. Conclusion

When performing on-site product support, completing repair interventions at the first visit leads to satisfied customers and sav-

Fig. 3. Fill rate versus investment value if only one unit may be added per iteration (Engineer 1).

Fig. 4. Fill rate versus average daily cost for various settings of $S_l - s_l$ (Engineer 1).

ings for the service provider. Therefore, the (optimal) composition of the repair kit used by field engineers is of utmost importance. In practical terms, such optimization should always take into account the replenishment lead times of the different items, as they are rarely negligible. This paper develops a new service model for the repair kit problem with positive replenishment lead times and fixed ordering cost and provides an exact expression for calculating the job fill rates for (s, S) policies. We also develop a greedy marginal analysis procedure.

When analysing the relationship between the level of stock on-hand and the job fill rate, two different kinds of behaviour

emerge. For most engineers, a typical curve is found with decreasing marginal returns from stock level increases. For some engineers that service complex systems, however, an S-shaped curve is obtained, showing that a “critical mass” of items need to be added to the repair kit to see greater benefits. A further investigation of this finding, in order to assess the conditions under which one or the other type of relationship occurs, is an important direction for future research.

Our analysis has also shown that the order in which items are added to the repair kit is very important. That order should depend in a dynamic manner on the item price and the item's demand dis-

tribution. Further, the necessity of allowing multiple units of some item to be added at once to the kit was established and the sensitivity of the results with respect to the chosen order quantity was explored, showing that the EOQ provides a good choice for the difference between reorder and order-up-to level.

The application of the proposed model on other empirical datasets would help further assess its utility for real world applications. An important direction for further research is to augment the problem formulation to consider also space constraints; i.e. assessing whether the set of chosen items actually fits in a kit. Finally, a further research stream could investigate the role of judgmental adjustments (or even purely judgmental decisions) by after-sales managers concerning repair kits definition, their interaction with what stems from (algorithm-based) decision support systems and their (comparative) performance.

Acknowledgements

We acknowledge the financial assistance provided by the Università degli Studi di Brescia that facilitated an invited research visit of Aris Syntetos and Ruud Teunter to Italy in June 2014. The research project described in this paper was conceived during that visit. Also, the research conducted by Dennis Prak was supported by The Netherlands Organisation for Scientific Research (NWO) [grant number 406-15-014]. Finally, we thank the three anonymous referees for their valuable comments that have improved the paper significantly.

References

- Axsäter, S. (2006). *Inventory control*. New York: Springer (2nd edition).
- Bijvank, M., Koole, G., & Vis, I. F. A. (2010). Optimising a general repair kit problem with a service constraint. *European Journal of Operational Research*, 204, 76–85.
- Brumelle, S., & Granot, D. (1993). The repair kit problem revisited. *Operations Research*, 41, 994–1006.
- Cohen, M. A., Agrawal, N., & Agrawal, V. (2006). Winning in the aftermarket. *Harvard Business Review*, 84, 129–138.
- Gebauer, H., Fleisch, E., & Friedli, T. (2005). Overcoming the service paradox in manufacturing companies. *European Management Journal*, 23, 14–26.
- Graves, S. C. (1982). A multiple-item inventory model with a job completion criterion. *Management Science*, 28, 1334–1337.
- Guajardo, M., & Rönnqvist, M. (2015). Cost allocation in inventory pools of spare parts with service-differentiated demand classes. *International Journal of Production Research*, 53, 220–237.
- Heeremans, D., & Gelders, L. F. (1995). Multiple period repair kit problem with a job completion criterion: a case study. *European Journal of Operational Research*, 81, 239–248.
- Lay, G. (2014). *Servitization in industry*. Heidelberg: Springer.
- Mamer, J. W., & Smith, S. A. (1982). Optimising field repair kits based on job completion rate. *Management Science*, 28, 1328–1333.
- Mamer, J. W., & Shogan, A. W. (1987). A constrained capital budgeting problem with applications to repair kit selection. *Management Science*, 33, 800–806.
- March, S. T., & Scudder, G. D. (1984). On optimizing field repair kits based on job completion rate. *Management Science*, 30, 1025–1028.
- Saccani, N., Visintin, F., Mansini, R., & Colombi, M. (2017). Improving spare parts management for field services: a model and a case study for the repair kit problem. *IMA Journal of Management Mathematics*, 28, 185–204.
- Silver, E. A., Pyke, D. F., & Peterson, R. (1998). *Inventory management and production planning and scheduling*. New York: John Wiley & Sons (3rd edition).
- Smith, S. A., Chambers, J. C., & Shlifer, E. (1980). Optimal inventories based on job completion rate for repairs requiring multiple items. *Management Science*, 26, 849–854.
- Song, J. (1998). On the order fill rate in a multi-item, base stock inventory system. *Operations Research*, 46, 831–845. doi:10.1287/opre.46.6.831.
- Syntetos, A. A., Keyes, M., & Babai, M. Z. (2009). Demand categorisation in a European spare parts logistics network. *International Journal of Operations & Production Management*, 29, 292–316.
- Teunter, R. H. (2006). The multiple-job repair kit problem. *European Journal of Operational Research*, 175, 1103–1116.