The multiple-job repair kit problem

Ruud H. Teunter*

August 29, 2003

Abstract

The *repair kit problem* is that of finding the optimal set of parts in the kit of a repairman. An important aspect of this problem, in many real-life situations, is that several job-sites are visited before a kit is restocked. In this paper, we present two heuristics for solving the multiple-job repair kit problem. Both heuristics can be used to determine a solution under the service-objective (minimal holding cost for a required job-fill rate) as well as the cost-objective (minimal expected total cost, including a penalty cost for each 'broken' job). The 'Job Heuristic (JH)' almost always determines the exact optimal solution, as is shown in an extensive numerical experiment. However, it can not (easily) be used in cases where several parts of the same type may be needed on a job, or part failures are dependent, or the number of jobs in a tour varies. The 'Part Heuristic (PH)' is simpler and easy to use in these cases also. In fact, it can be applied in a spreadsheet software package, as we illustrate. The numerical experiments show that it s leads to near-optimal solutions (average 'cost error' of less than 0.1 per cent). Therefore, the PH is an excellent method for solving repair kit problems in practise.

Keywords: Logistics, Inventory, Repair Kit Problem

1   Introduction

This research is motivated by a case study at the service department of a copier manufacturer. The department employs around 40 repair(wo)man, who are clustered into three groups based on the types of copiers (around 25 in total) that they service. Repairmen travel to job-sites in a company car that contains a group-specific set of spare parts, the *repair kit*. A repairman

---

*Erasmus University, PO Box 1738, NL-3000 DR Rotterdam, The Netherlands. E-mail: teunter@few.eur.nl

usually visits 3 or 4 sites per day. If a job cannot be finished because one or more parts are missing, i.e. if a job is 'broken', the repairman will return the next day to finish the job. At the end of each day, a repairman reports which parts were used to carry out repairs and which parts were missing in case there were broken jobs. Those parts are delivered at his home overnight by a courier.

The problem in the above case-study is that of determining how many of parts of each type should be contained in the repair kit of a repairman. Crucial in this determination is the trade-off between holding cost and service. Service can be measured by the *job-fill rate*, i.e. the fraction of jobs that is not broken. In the *service-model* of this *repair kit problem*, the objective is to minimize the holding cost (which is proportional to the value of the repair kit) under a restriction for the job-fill rate. In the *cost-model*, a penalty cost is incurred for each broken job and the objective is to minimize the expected total cost. The penalty cost should include the actual cost for revisiting a job-site (e.g. transportation and labor costs) as well as the artificial cost for loss of goodwill.

An important aspect of the repair kit problem in the above case-study and in most practical situations, is that the repair kit is not restocked after each job. Surprisingly, immediate restocking is assumed in most of the literature (see Teunter and Klein Haneveld [3] for a recent review). In fact, to the best of our knowledge, Heeremans and Gelders [1] are the only authors that allow a multiple (though fixed) number of jobs before a repair kit is restocked. They do assume that part failures are independent and that at most one part of any type is needed on a job. They present an integer (binary) linear programming formulation of the service-model. It is important to remark, however, that their formulation is not fully correct. The formulation imposes a minimum required probability that none of the jobs between two successive re-stockings are broken (*tour-fill rate*), instead of a minimum required job-fill rate.

The binary programming formulation of Heeremans and Gelders [1] is complex and not insightful, and therefore difficult to apply in practise. The authors note this themselves, when discussing a case study. They therefore propose a heuristic solution procedure of the knapsack kind. Again, however, this is based on the tour-fill rate instead of the job-fill rate.

In this paper, we will reformulate the knapsack heuristic proposed by Heeremans and Gelders [1], so that it does consider the job-fill rate. This so-called *Job Heuristic (JH)* starts with an empty kit and adds parts one by one, always selecting the part type with the lowest ratio of increase in holding cost and increase in job-fill rate. This way, a series of kits with increasing job-

fill rates levels is generated. The best solution of this series is picked. An important disadvantage of the JH is that job-fill rates are difficult to calculate. As a result, the JH can not (easily) be used in cases where several parts of the same type may be needed on a job, or part failures are dependent, or the number of jobs in a tour varies.

We therefore propose a second knapsack heuristic which is easy to use in these cases also. This so-called *Part Heuristic (PH)* uses a different part selection rule. It always selects the part type with the lowest ratio of increase in holding cost and increase in part-fill rate (fraction of parts that is not missing). Part-fill rates are much easier to calculate than job-fill rates. In fact, as we illustrate, the PH can be applied in any spreadsheet software package. For each repair kit in the generated series, the PH calculates an estimate of the job-fill rate based on the part-fill rates. That estimate is used to pick the approximately best solution.

Both the JH and the PH can be applied to find a solution for the cost-model as well as the service-model. For the service-model, pick the first kit of the generated series with a sufficient (estimated) job-fill rate. For the cost-model, pick that kit of the series with the lowest (estimated) total cost. Note that for the service-model, the solution of the PH may not satisfy the service requirement if the job-fill rate is overestimated. In the numerical experiments where the JH and the PH are compared, we therefore consider the cost-model only.

We perform two extensive numerical experiments. In the first, we consider 1000 (randomly generated) examples for the cost-model with up to 4 jobs and up to 8 parts. For each example, we compare the solutions of the JH and the PH to the exact solution (full search). It turns out that the JH determines the exact solution for all 1000 examples. The PH also performs excellent. It determines the exact solution for 985 examples. For the other 15 examples, the total cost for the solution of the PH is less than 1 per cent higher than the total cost for the exact solution.

In the second numerical experiment, we consider 1000 (randomly generated) examples for the cost-model with up to 10 jobs and up to 100 parts. For these large instances, it is not possible to determine the exact solution by full search. But we compare the solution of the PH to that of the JH. It turns out that the total cost for the solution of the PH is at most 1.4 per cent higher and only 0.03 per cent higher on average.

We conclude from the numerical results that the JH generally determines the optimal solution and that the solution of the PH is always close to optimal. Since the PH is much easier to use and more broadly applicable, we recommend it for use in practise.

2 A model of the repair kit problem

The notations that will be used in the remainder are listed in Table 1.

TABLE 1

A repairman visits job sites. The number of different part types that may be needed on a job is denoted by N. Multiple parts of the same type may be needed on a job. Let L_i^{job} denote the maximum number of parts of type i, $i = 1, 2, \ldots, N$, that may be needed on a job. Let $p_i^{\mathrm{job}}(j)$, $i = 1, 2, \ldots, N$, $j = 1, \ldots, L_i^{\mathrm{job}}$, denote the probability that j parts of type i are needed on a job. The numbers of parts of different types that are needed on a job may be dependent.

A *tour* is defined as a sequence of one or more jobs before a repairman is restocked. The number of jobs in a tour may be stochastic and is denoted by M. Let $L_{\mathrm{job}}^{\mathrm{tour}}$ denote the maximum number of jobs in a tour. Let $p_{\mathrm{job}}^{\mathrm{tour}}(j)$, $j = 1, 2, \ldots, L_{\mathrm{job}}^{\mathrm{tour}}$ denote the probability that the number of jobs in a tour is j, i.e. $\Pr(M = j) = p_{\mathrm{job}}^{\mathrm{tour}}(j)$. In cases where M is deterministic, its fixed value is denoted by $\overline{M}$.

Let L_i^{tour} denote the maximum number of parts of type i, $i = 1, 2, \ldots, N$, that may be needed on a tour. Let $p_i^{\mathrm{tour}}(j)$, $i = 1, 2, \ldots, N$, $j = 1, \ldots, L_i^{\mathrm{tour}}$, denote the probability that j parts of type i are needed on a tour. We remark that the (maximum) number of jobs in a tour may depend on the job types, i.e. more time is needed is more parts have to be replaced. In cases where M is not deterministic, it is therefore not necessarily true that $L_i^{\mathrm{tour}} = L_{\mathrm{job}}^{\mathrm{tour}} L_i^{\mathrm{job}}$, and the relation between probabilities $p_i^{\mathrm{tour}}(j)$ and $p_i^{\mathrm{job}}(j)$ may be complex.

A repair kit is defined as the set of parts after restocking, i.e. at the beginning of a tour. A repair kit with n_i parts of type i, $i = 1, \ldots, N$, is denoted by $[n_1, \ldots, n_N]$. For each part of type i in the repair kit, a holding cost H_i is incurred (per tour). The holding cost for repair kit $S = [n_1, \ldots, n_N]$ is denoted by $C_h(S)$, i.e. $C_h(S) = \sum_{i=1}^{N} n_i H_i$. The job-fill rate, defined as the probability that a random job is broken, of a repair kit S is denoted by $\gamma^{\mathrm{job}}(S)$.

The repair kit problem is that of determining the optimal repair kit. There are two possible objectives. The service-objective is to find the repair kit S with the smallest holding cost $C_h(S)$ for which the job-fill rate $\gamma^{\mathrm{job}}(S)$ is equal to or larger than some pre-described level $\gamma_{\mathrm{min}}^{\mathrm{job}}$. The cost-objective is to include a penalty cost P for each broken job and find the repair kit with the smallest expected total cost $C_t(S) := C_h(S) + P\gamma^{\mathrm{job}}(S)$. The optimal repair kit (under both objectives) is denoted by S^*.

The above model of the repair kit problem is very general. It encompasses all important elements of real-life repair kit problems: a multiple and stochastic number of jobs per tour, multiple parts of the same type needed on a job, and dependent demands for parts of different types. To the best of our knowledge, we are the first to study such a general model of the repair kit problem. One of the main results of this paper is the development of a heuristic that can be applied to this general model and finds near-optimal solutions. It is called the Part Heuristic (PH). The repair kit determined by the PH is denoted by S^{PH}.

A second heuristic that will be developed is the Job Heuristic (JH). The JH is more sophisticated than the PH, but it can only (easily) be applied if the following assumptions are satisfied.

As1 The number of jobs in a tour is fixed.

As2 At most one part of each type is needed on a job.

As3 Demands for parts of different types are independent.

The repair kit determined by the JH is denoted by S^{JH}.

2.1 Estimating model parameters

The model input parameters L_i^{job}, $p_i^{\mathrm{job}}(j)$ and L_i^{tour}, $p_i^{\mathrm{tour}}(j)$ are the most difficult to estimate. In general, these estimations will be based on historic data for the number of parts needed per job (job demand) and per tour (tour demand), respectively. The tour demand data are relatively easy to gather, since a repairman places a replenishment order at the end of each tour. In order to obtain data on job demands, repairman have to record and communicate part demands for each job separately. As will appear in later sections, the JH requires job demand date, whereas tour demand data are sufficient to apply the PH.

3 Heuristics

Clearly, the number of parts for type i, $i = 1, \ldots, N$, in the repair kit should be at least 0 and at most L_i^{tour}. So the number of possible solutions for the repair kit problem is $\prod_{i=1}^{N}(L_i^{\mathrm{tour}} + 1)$. In practise, there are usually hundreds and often thousands of different part types. Obviously, a full search for the optimal solution is out of the question then.

In this section, we propose two heuristics for solving the repair kit problem. Both are of the knapsack kind. They start with an empty kit, and add parts one by one until the kit is complete, i.e. until the number of parts for each type i is L_i^{tour}. The generated series of kits with increasing job-fill rates can be used to solve the repair kit problem from a service-objective as well as a cost-objective.

3.1 Job Heuristic

The first heuristic that we propose always adds that part to the repair kit which results in the largest increase in job-fill rate per added unit of holding cost. We call it the 'Job Heuristic (JH)'. A mathematical description is given in the frame below (see Table 1 for notations).

Job Heuristic

Step 1: $k := 0$ (first kit). For all $i = 1, 2, \ldots, N$: $n_i := 0$.

$S_k := [n_1, \ldots, n_N] = [0, \ldots, 0]$ (first kit is empty).

$\gamma^{\text{job}}(S_k) := \prod_{i=1}^{N} p_i^{\text{job}}(0)$. $C_h(S_k) := 0$. $C_t(S_k) := P\left(1 - \gamma^{\text{job}}(S_k)\right) E[M]$.

Step 2: For all $i = 1, 2, \ldots, N$ for which $n_i < L_i^{\text{tour}}$:

$\Delta_i^{\text{job}} := \gamma^{\text{job}}([n_1, \ldots, n_{i-1}, n_i + 1, n_{i+1}, \ldots, n_N]) - \gamma^{\text{job}}([n_1, \ldots, n_N])$

(increase in the *job-fill rate* if one part of type i is added to the kit).

Step 3: $i^* := \arg \max\limits_{i : n_i < L_i^{\text{tour}}} \frac{\Delta_i^{\text{job}}}{H_i}$ (part type with the largest ratio $\frac{\Delta_i^{\text{job}}}{H_i}$).

$k := k + 1$ (next kit). $n_{i^*} := n_{i^*} + 1$ (add one part of type i^*).

$S_k := [n_1, \ldots, n_N]$. $\gamma^{\text{job}}(S_k) := \gamma^{\text{job}}(S_{k-1}) + \Delta_{i^*}^{\text{job}}$.

$C_h(S_k) := C_h(S_{k-1}) + H_{i^*}$. $C_t(S_k) := C_t(S_{k-1}) + H_{i^*} - P\Delta_{i^*}^{\text{job}} E[M]$.

Step 4: If $k = \sum_{i=1}^{N} L_i^{\text{tour}}$ ($n_i = L_i^{\text{tour}}$ for all $i = 1, 2, \ldots, N$), then go to Step 5. Otherwise, go to Step 2.

Step 5: Service-objective: $S^{\text{JH}} := S_{k^{\text{JH}}}$, where $k^{\text{JH}} := \min\{k | \gamma^{\text{job}}(S_k) > \gamma_{\min}^{\text{job}}\}$.

Cost-objective: $S^{\text{JH}} := S_{k^{\text{JH}}}$, where $k^{\text{JH}} := \arg\min_k C_t(S_k)$.

A disadvantage of the JH is, that the calculations of the job-fill rates in Step 2 are complex. Those calculations cannot (easily) be done for the general version of the repair kit problem that we presented in Section 2. However, if assumptions As1, As2, and As3 (see Section 2) are satisfied, then the job-fill rates can be calculated using a closed-form expression that we derive next.

The key observation in deriving the job-fill rate of a repair kit $[n_1, \ldots, n_N]$ (under assumptions As1, As2, and As3) is the following. The m-th, $m = 1, 2, \ldots, \overline{M}$, job will not be broken if and only if one of the two following conditions is satisfied for each part type i, $n = 1, 2, \ldots, N$: (1) a part of type i is not needed on the m-th job, (2) a part of type i is needed on the m-th job and on at most $n_i - 1$ out of the first $m - 1$ jobs. So we get (for ease of notation, an empty summation $\sum_0^{-1}$ equals 0)

$$
\begin{aligned}
&\gamma^{\text{job}}([n_1, \ldots, n_N]) \\
&= \frac{1}{\overline{M}} \sum_{m=1}^{\overline{M}} \left(\prod_{i=1}^{N} \left(1 - p_i^{\text{job}}(1) + p_i^{\text{job}}(1) \sum_{l=0}^{\max\{n_i-1, m-1\}} \left(\binom{m-1}{l} (p_i^{\text{job}}(1))^l (p_i^{\text{job}}(0))^{m-1-l} \right) \right) \right).
\end{aligned}
$$

In Section 4 we will the use the above expression to apply the JH and compare its solution to the optimal solution (determined by full search) for a large number of randomly generated examples. It will turn out that the solution of the JH is equal to the optimal solution for each of those examples. To show that this is not true in general, we end this section with a (carefully constructed) example for which the solution determined by the JH is suboptimal, and provide an intuitive explanation.

Example 3.1 (for which the solution determined by the JH is suboptimal).

Consider the following situation. A tour always consists 2 jobs, i.e. $\overline{M} = 2$. There are $N = 3$ types of parts, and at most one of each is needed on a job. So, $L_i^{\text{job}} = 1$ and $L_i^{\text{tour}} = 2$ for $i = 1, 2, 3$. The probabilities that part 1, 2, and 3 are needed on a job are $p_1 = 0.1$, $p_2 = 0.9$, and $p_3 = 0.9$, respectively, and part failures are independent. The holding costs for parts 1, 2, and 3 are $H_1 = 0.001$, $H_2 = 1$, and $H_3 = 1.01$, respectively. There is a cost-objective, and the penalty cost for each broken job is $P = 2000$.

Table 2 illustrates the generation of a series of repair kits by the JH.

<div align="center">TABLE 2</div>

So $S^{\mathrm{JH}} = [2, 1, 1]$ with total cost 3939.8. But the optimal solution (full search) is $S^* = [1, 1, 1]$ with total cost 3939.5. The sub-optimality S^{JH} is caused by the myopic behavior of the JH. In each step, the JH adds a single item to the repair kit, without considering the effect of that addition on future steps. Since both items 2 and 3 are needed to complete most jobs (both are needed with probability 0.90), adding just one of them to a repair kit results in a relatively small increase of the job-fill rate. But adding the other in the next step does cause a large increase in the job-fill rate. The JH does not look ahead, and adds an item of type 1 in the first two steps. As a result, the optimal solution with one item of each type is not in the generated series. □

3.2 Part Heuristic

The second heuristic that we propose is the 'Part Heuristic (PH)'. The PH is a knapsack heuristic similar to the JH. It uses the same 5-step approach. However, it avoids complicated calculations by focussing on the part-fill rate instead of the job-fill rate. The part-fill rate is defined as the fraction of parts that is not missing, i.e. that is available directly from the kit. E.g., consider a tour that consists of three jobs and a kit with 7 parts of a certain type. If 1, 5, and 4 parts of that type are needed on the first, second, and third job, respectively, then the part-fill rate is $7/(1 + 5 + 4) = 0.7$.

The PH is inspired by the study of Smith, Chambers, and Shlifer [2]. Their results imply that the PH generates the same series of kits as the JB under the following assumptions: restocking after each job (more restrictive than As1), at most one part of each type is needed on a job (As2), and demands for parts of different types are independent (As3). It seems reasonable to assume that the PH and JH generate 'comparable' series of repair kits if these assumptions are not satisfied. This is validated by the results that are reported in the next section.

It is easy to see that the part-fill rate is equal to the ratio of the expected number of parts used on a tour and the expected number of parts needed on a tour, i.e., the part-fill rate a repair kit $[n_1, \ldots, n_N]$ can be expressed as

$$\gamma^{\mathrm{part}}([n_1, \ldots, n_N]) = \frac{\sum_{i=1}^{N} \sum_{j=1}^{n_i} \Pr(D_i^{\mathrm{tour}} \geq j)}{\sum_{i=1}^{N} \sum_{j=1}^{\infty} \Pr(D_i^{\mathrm{tour}} \geq j)} = \frac{\sum_{i=1}^{N} \sum_{j=1}^{n_i} \Pr(D_i^{\mathrm{tour}} \geq j)}{\sum_{i=1}^{N} E[D_i^{\mathrm{tour}}]}.$$

It follows that the increase in part-fill rate from adding the $n_i + 1$-th part of type i to a repair kit $[n_1, \ldots, n_N]$ is

$$\gamma^{\mathrm{part}}([n_1, \ldots, n_{i-1}, n_i + 1, n_{i+1}, \ldots, n_N]) - \gamma^{\mathrm{part}}([n_1, \ldots, n_N]) = \frac{\Pr(D_i^{\mathrm{tour}} > n_i)}{\sum_{i=1}^{N} E[D_i^{\mathrm{tour}}]}.$$

8

So the selection process for the PH is much easier than that for the JH. Since $\sum_{i=1}^{N} E[D_i^{\text{tour}}]$ is independent of $n_i, \ i = 1, \ldots, N$, the PH simply adds the part type i with the largest ratio $\Pr(D_i^{\text{tour}} > n_i)/H_i$ (compare to Steps 2 and 3 of the JH).

As discussed above, the PH avoids the complex calculations of job-fill rates. However, it does calculate approximations of the job-fill rates, since those are needed to decide which repair kit from the generated series is the best one (Step 5). The approximation of the job-fill rate $\gamma^{\text{job}}([n_1, \ldots, n_N])$ is denoted by $\widehat{\gamma^{\text{job}}}([n_1, \ldots, n_N])$ and is defined as follows.

$$\widehat{\gamma^{\text{job}}}([n_1, \ldots, n_N]) := \left(\prod_{i=1}^{N} \Pr(D_i^{\text{tour}} \leq n_i) \right)^{1/E[M]}, \tag{1}$$

where $E[M]$ denotes the expected number of jobs per tour. The idea underlying this approximation is that the probability $\prod_{i=1}^{N} \Pr(D_i^{\text{tour}} \leq n_i)$ that no parts are missing on a tour is equal to the probability that none of the, on average $E[M]$, jobs of a tour are broken, which is approximately $(\gamma^{\text{job}})^{E[M]}$. From (1) we get

$$\widehat{\gamma^{\text{job}}}([n_1, \ldots, n_{i-1}, n_i + 1, n_{i+1}, \ldots, n_N]) = \widehat{\gamma^{\text{job}}}([n_1, \ldots, n_N]) \left(\frac{\Pr(D_i^{\text{tour}} \leq (n_i + 1))}{\Pr(D_i^{\text{tour}} \leq n_i)} \right)^{1/E[M]}. \tag{2}$$

This recursive equation can be used to calculate job-fill rates for the series of repair kits that the PH generates. A mathematical description of the PH is given in the frame below (see Table 1 for notations).

Part Heuristic

Step 1: $k := 0$ (first kit). For all $i = 1, 2, \ldots, N$: $n_i := 0$.

$S_k := [n_1, \ldots, n_N] = [0, \ldots, 0]$ (first kit is empty).

$\widehat{\gamma^{\text{job}}}(S_k) := \left(\prod_{i=1}^{N} p_i^{\text{tour}}(0) \right)^{1/E[M]}$. $C_h(S_k) := 0$. $\widehat{C_t}(S_k) := P \left(1 - \widehat{\gamma^{\text{job}}}(S_k) \right) E[M]$.

Step 2: For all $i = 1, 2, \ldots, N$ for which $n_i < L_i^{\text{tour}}$: $\Delta_i^{\text{part}} := \Pr(D_i^{\text{tour}} > n_i)$

Step 3: $i^* := \arg \max_{i:n_i < L_i^{\text{tour}}} \frac{\Delta_i^{\text{part}}}{H_i}$ (part type with the largest ratio $\frac{\Delta_i^{\text{part}}}{H_i}$).

$k := k + 1$ (next kit). $n_{i^*} := n_{i^*} + 1$ (add one part of type i^*).

$S_k := [n_1, \ldots, n_N]$. $\widehat{\gamma^{\text{job}}}(S_k) := \widehat{\gamma^{\text{job}}}(S_{k-1}) \left(\frac{\Pr(D_i^{\text{tour}} \leq n_i^*)}{\Pr(D_i^{\text{tour}} \leq n_i^*-1)} \right)^{1/E[M]}$.

$C_h(S_k) := C_h(S_{k-1}) + H_{i^*}$. $\widehat{C_t}(S_k) := C_h(S_k) + P \left(1 - \widehat{\gamma^{\text{job}}}(S_k) \right) E[M]$.

Step 4: If $k = \sum_{i=1}^{N} L_i^{\text{tour}}$ ($n_i = L_i^{\text{tour}}$ for all $i = 1, 2, \ldots, N$), then go to Step 5.
Otherwise, go to Step 2.

Step 5: Service-objective: $S^{\text{PH}} := S_{k^{\text{PH}}}$, where $k^{\text{PH}} := \min\{k | \widehat{\gamma^{\text{job}}}(S_k) > \gamma_{\min}^{\text{job}}\}$.
Cost-objective: $S^{\text{PH}} := S_{k^{\text{PH}}}$, where $k^{\text{PH}} := \arg \min_k \widehat{C_t}(S_k)$.

Steps 2 and 3 and the above discussion show that the PH always adds the part type i with the largest ratio $\Pr(D_i^{\text{tour}} > n_i)/H_i$ to the repair kit, where n_i is the number of parts of that type which are already in the kit. Note that this selection ratio does not depend on the numbers of parts of other types which are already in the kit. As a consequence, the sequential approach for determining the order in which parts enter the kit can be replaced by the following direct approach: (i) calculate the selection ratio for the 1st, 2nd, ..., L_i^{tour}th part of each type and order those $\sum_{i=1}^{N} L_i^{\text{tour}}$ ratios from the highest to the lowest.

The direct approach for determining the order in which parts enter the kit, combined with the recursive equations for calculating the holding cost and the approximate job-fill rate in Step 3, can be used to apply the PH in any spreadsheet software package. We will illustrate this for a simple example.

Example 3.2 (illustrates the application of the PH in a spreadsheet package).

Consider the following situation. A tour always consists 3 jobs, i.e. $\overline{M} = 3$. There are $N = 2$ types of parts, and at most one of each is needed on a job. So, $L_i^{\text{job}} = 1$ and $L_i^{\text{tour}} = 3$ for $i = 1, 2$. The probabilities that part 1 and 2 are needed on a job are $p_1 = 0.1$ and $p_2 = 0.3$, respectively, and part failures are independent. It easily follows that: $p_1^{\text{tour}}(0) = 0.729$, $p_1^{\text{tour}}(1) = 0.243$, $p_1^{\text{tour}}(2) = 0.027$, $p_1^{\text{tour}}(3) = 0.001$, $p_2^{\text{tour}}(0) = 0.343$, $p_2^{\text{tour}}(1) = 0.441$, $p_2^{\text{tour}}(2) = 0.189$, $p_2^{\text{tour}}(3) = 0.027$. We remark that in practice, it is easier to estimate the tour demand probabilities directly using historic data on tour demand (see Section 2.1). The holding costs for parts 1 and 2 are $H_1 = 1$ and $H_2 = 5$, respectively. There is a service objective, and the minimum required job-fill rate is 95%.

Using the above data, a worksheet is constructed as follows. The first row (Row 1) is reserved for column titles. The second row (Row 2) corresponds to the empty repair kit. Initially, this row is empty. Rows 3, 4, and 5 correspond to the first, second, and third ($L_1^{\text{tour}} = 3$) part of type 1 that is added to the repair kit. Rows 6, 7, and 8 correspond to the first, second, and third ($L_2^{\text{tour}} = 3$) part of type 2 that is added to the repair kit. For all rows 3-8, that correspond to a part being added to the repair kit, the following is given: the part type i (column A), the number n_i of parts of the same type already in the kit (column B), the holding cost H_i (column C), the probability $\Pr(D_i^{\text{tour}} > n_i)$ (column D), the probability $\Pr(D_i^{\text{tour}} > n_i + 1)$ (column E), and the selection ratio $\Pr(D_i^{\text{tour}} > n_i)/H_i$ (column F) of the PH are given. This results in the worksheet that is given below. For rows 3-8, the spreadsheet formula for calculating the selection ratio is given between parenthesis.

	A	**B**	**C**	**D**	**E**	**F**
1	i	n_i	H_i	$\Pr(D_i^{\text{tour}} > n_i)$	$\Pr(D_i^{\text{tour}} > n_i + 1)$	$\Pr(D_i^{\text{tour}} > n_i)/H_i$
2						
3	1	0	1	0.2710	0.0280	0.2710 (E3/D3)
4	1	1	1	0.0280	0.0010	0.0280 (E4/D4)
5	1	2	1	0.0010	0.0000	0.0010 (E5/D5)
6	2	0	5	0.6570	0.2160	0.1314 (E6/D6)
7	2	1	5	0.2160	0.0270	0.0432 (E7/D7)
8	2	2	5	0.0270	0.0000	0.0054 (E8/D8)

Next, the selection ratio's in (column F) are set from formula to number, and those numbers are

used to sort Rows 3-8 in descending order. This leads to the sequence of repair kits that the PH generates, starting with the empty kit in Row 2 and adding one part (of type i in column A) for each row 3-8. To be able to select the optimal kit, three columns are added: column G with the number k of parts in a kit, column H with the approximate job-fill rate (formulae based on the recursive equation (2)), and column I with the holding cost.

	A	**B**	**C**	**D**	**E**	**F**	**G**	**H**	**I**
1	i	n_i	H_i	...	...	...	k	$\widehat{\gamma^{\text{job}}}(S_k)$	$C_h(S_k)$
2							0	0.6300 (0.9*0.7)	0
3	1	0	1	0.2710	0.0280	0.2710	1	0.6934 (H2*((1-F3)/(1-E3))^(1/3))	1 (I2+D3)
4	2	0	5	0.6570	0.2160	0.1314	2	0.9134 (H3*((1-F4)/(1-E4))^(1/3))	6 (I3+D4)
5	2	1	5	0.2160	0.0270	0.0432	3	0.9816 (H4*((1-F5)/(1-E5))^(1/3))	11 (I4+D5)
6	1	1	1	0.0280	0.0010	0.0280	4	0.9906 (H5*((1-F6)/(1-E6))^(1/3))	12 (I5+D6)
7	2	2	5	0.0270	0.0000	0.0054	5	0.9997 (H6*((1-F7)/(1-E7))^(1/3))	17 (I6+D7)
8	1	2	1	0.0010	0.0000	0.0010	6	1.0000 (H7*((1-F8)/(1-E8))^(1/3))	18 (I7+D8)

Since the minimum required job-fill rate is 95%, the PH picks the repair kit in Row 5 which contains 1 part of type 1 (Column A, added in Row 3) and 2 parts of type 2 (Column A, added in Rows 4 and 5). So $S^{\text{PH}} = [1, 2]$ for this example, and the estimated job-fill rate is 0.9816.

We end this example by noting that $[1, 2]$ is indeed the optimal solution (checked by full search), and that the exact job-fill rate is 0.9815. □

4   Numerical results

We perform two extensive numerical experiments. In Section 4.1, we consider 1000 (randomly generated) examples for the cost-model with up to 4 jobs and up to 8 parts. For each example, we compare the solutions of the JH and the PH to the exact solution (full search). In Section 4.2, we consider 1000 (randomly generated) examples for the cost-model with up to 10 jobs and up to 100 parts. For these large instances, it is not possible to determine the exact solution by full search. But we compare the solutions of the JH and the PH.

Recall from Section 3.1 that applying the JH requires complex calculations of job-fill rates. In fact, we were only able to derive a closed-form expression for the job-fill rate under the assumptions that the number of jobs in a tour is fixed (As1), at most one part of each type

is needed on a job (As2), and demands for parts of different types are independent (As3). We therefore only consider examples where these assumptions are satisfied. So for each example, we have $L_i^{\text{job}} = 1$ and $L_i^{\text{tour}} = L$ for all $i = 1, \ldots, N$. Moreover, we have

$$p_i^{\text{tour}}(j) = \binom{\overline{M}}{j} \left(p_i^{\text{job}}(1)\right)^j \left(p_i^{\text{job}}(0)\right)^{\overline{M}-j}, \quad j = 0, 1, \ldots, \overline{M}.$$

4.1   Small instances

We consider 1000 examples for the cost-model. The number of part types N, number of jobs per tour L, failure probabilities $p_i^{\text{job}}(1)$, $i = 1, \ldots, N$, holding costs H_i, $i = 1, \ldots, N$, and broken job cost P are all drawn from (Discrete_)Uniform distributions. The lower and upper bounds of these distributions are given in Table 3.

TABLE 3

By generating examples randomly, large variation is ensured. The job-fill rate for the exact optimal solution varies between 29% and 100% and is 87% on average, which indicates that the cost settings in Table 3 are reasonable.

It turns out that the JH determines the exact solution for all 1000 examples. The PH also performs excellent. It determines the exact solution for 985 examples. For the other 15 examples, the total cost for the solution of the PH at most 0.6% higher than the total cost for the exact solution.

4.2   Large instances

Again, 1000 examples for the cost-model are randomly generated. See Table 4.

TABLE 4

For these examples, the job-fill rate for the solution of the JH varies between 0% and 100%, and is 84% on average.

For each example, the solution of the PH is close to that of the JH. The total cost for the solution of the PH is at most 1.4 per cent higher and only 0.03 per cent higher on average.

13

Moreover, the performance of the PH is not deteriorating for increasing numbers of part types N or increasing numbers of jobs per tour, as is illustrated by Table 5.

TABLE 5

This indicates that the PH will perform well for even larger instances also.

5   Conclusion

We presented two knapsack heuristics for solving the multiple-job repair kit problem. The 'Job Heuristic (JH)' is the most exact. In fact, it determined the optimal solution for each of the examples that were considered in an extensive numerical experiment (1000 small instances with up to 8 part types and up to 4 jobs per tour). A disadvantage of the JH is that it can only be (easily) applied if at most one part of each part type is needed on a job, part failures are independent, and the number of jobs in a tour is fixed.

The 'Part Heuristic (PH)' does not require any of the above assumptions, and can easily be applied in a spreadsheet software package. Furthermore, the PH determines near-optimal solutions. In the before mentioned numerical experiment, it turned out that the total cost for the solution of the PH is at most 0.6% higher than the total cost for the optimal solution, and only 0.004% higher on average. Another experiment with larger instances (up to 100 part types and up to 10 jobs per tour) showed that the total cost for the solution of the PH is at most 1.4% higher than the total cost for the solution of the JH, and only 0.03% higher on average. So, the PH is easily and broadly applicable, and it leads to near-optimal solutions. Therefore, the PH is an excellent method for solving repair kit problems in practise.

A direction for further research is to consider the optimization of repair kits as a subproblem of optimizing stock levels in a multi-echelon inventory system. Repair kits can be interpreted as the end stock-points of such systems. Combined with order level, order-up-to policies for other stock-points, (modified versions of) the heuristics that we presented could be tested for such systems. Another direction for further research is to consider volumes of different part types and include space restrictions.

References

[1] D. Heeremans and L.F. Gelders. Multiple period repair kit problem with a job completion criterion: a case study. *European Journal of Operational Research*, 81(2):239–248, 1995.

[2] S.A. Smith, J.C. Chambers, and E. Shlifer. Optimal inventories based on job completion rate for repairs requiring multiple items. *Management Science*, 26(8):849–852, 1980.

[3] R.H. Teunter and W.K. Klein Haneveld. Inventory contol of service parts in the final phase: a central depot and repair kits. *European Journal of Operational Research*, 138:76–86, 2002.

Terminology

tour	sequence of jobs before a repair kit is restocked
job-fill rate	fraction of jobs not broken
part-fill rate	fraction of parts not missing

Input parameters

N	number of different part types
$H_i,\ i = 1, \ldots, N$	holding cost per part of type i per tour
P	penalty cost per broken job
$\gamma_{\min}^{\mathrm{job}}$	minimum required job-fill rate under the service-objective
$L_{\mathrm{job}}^{\mathrm{tour}}$	maximum number of jobs in a tour
$p_{\mathrm{job}}^{\mathrm{tour}}(j),\ j = 1, 2, \ldots, L_{\mathrm{job}}^{\mathrm{tour}}$	probability that the number of jobs in a tour is j
$L_i^{\mathrm{job}},\ i = 1, \ldots, N$	maximum number of parts of type i needed on a job
$p_i^{\mathrm{job}}(j),\ i = 1, \ldots, N,$	probability that j parts of type i are needed on a job
$\qquad j = 1, \ldots, L_i^{\mathrm{job}}$	
$L_i^{\mathrm{tour}},\ i = 1, \ldots, N$	maximum number of parts of type i needed on a tour
$p_i^{\mathrm{tour}}(j),\ i = 1, \ldots, N,$	probability that j parts of type i are needed on a tour
$\qquad j = 1, \ldots, L_i^{\mathrm{tour}}$	

Other notation

M	number of jobs in a tour $(\Pr(M = j) = p_{\mathrm{job}}^{\mathrm{tour}}(j))$
$\overline{M}$	fixed value of M in cases where it is deterministic (if $\Pr(M = \overline{M}) = 1$)
$D_i^{\mathrm{tour}},\ i = 1, \ldots, N$	number of parts of type i needed on a tour $(\Pr(D_i^{\mathrm{tour}} = j) = p_i^{\mathrm{tour}}(j))$
$[n_1, \ldots, n_N]$	repair kit with n_i parts of type i, $i = 1, \ldots, N$
$\gamma^{\mathrm{job}}(S)$	job-fill rate for repair kit S
$\gamma^{\mathrm{part}}(S)$	part-fill rate for repair kit S
$C_h(S)$	holding cost for repair kit S
$C_t(S)$	total expected (holding and penalty) cost for repair kit S
S^{JH}	repair kit determined by the Job Heuristic (JH)
S^{PH}	repair kit determined by the Part Heuristic (PH)
S^*	optimal repair kit
$E[.]$	expected value

Table 1: Notations.

k	S_k	$\gamma^{\mathrm{job}}(S_k)$	$C_h(S_k)$	$C_t(S_k)$	i	Δ_i^{job}	$\Delta_i^{\mathrm{job}}/H_i$	i^*
0	[0,0,0]	0.00900	0.000	3964.0	1	0.0019	1.900	1
					2	0.0891	0.089	
					3	0.0891	0.088	
1	[1,0,0]	0.00995	0.001	3961.2	1	0.0001	0.100	1
					2	0.0989	0.099	
					3	0.0989	0.098	
2	[2,0,0]	0.01000	0.002	3962.0	1	0.0000	0.000	
					2	0.0990	0.099	2
					3	0.0990	0.098	
3	[2,1,0]	0.05950	1.002	4764.0	1	0.0000	0.000	
					2	0.0810	0.081	
					3	0.9171	0.908	3
4	[2,1,1]	0.51805	2.012	3939.8	1	0.0000	0.000	
					2	0.1539	0.154	2
					3	0.1539	0.152	
5	[2,2,1]	0.59500	3.012	4632.0	1	0.0000	0.000	
					2	0.0000	0.000	
					3	0.8100	0.802	3
6	[2,2,2]	1.00000	4.022	4022.0				

Table 2: Results of applying the JH for a specific example.

Model parameter	Value
N	Discrete_Uniform[1,8]
$\overline{M}$	Discrete_Uniform[1,4]
$p_i^{\mathrm{job}}(1),\ i = 1,\ldots,N$	Uniform[0,0.20]
$H_i,\ i = 1,\ldots,N$	Uniform[0,1]
P	Uniform[0,10]

Table 3: Distributions from which the model parameters are drawn in the numerical experiments (small instances).

Model parameter	Value
N	Discrete_Uniform[1,100]
$\overline{M}$	Discrete_Uniform[1,10]
$p_i^{\text{job}}(1),\ i=1,\ldots,N$	Uniform[0,0.20]
$H_i,\ i=1,\ldots,N$	Uniform[0,1]
P	Uniform[0,100]

Table 4: Distributions from which the model parameters are drawn in the numerical experiments (large instances).

N	1...10	11...20	21...30	31...40	41...50	51...60	61...70	71...80	81...90	91...100
Error	0.06%	0.05%	0.02%	0.02%	0.02%	0.02%	0.02%	0.04%	0.02%	0.03%
L	1	2	3	4	5	6	7	8	9	10
Error	0.00%	0.08%	0.01%	0.02%	0.01%	0.01%	0.04%	0.03%	0.04%	0.07%

Table 5: Average increase in total cost for the solution of the PH compared to that of the JH ('Error') for different ranges of the number of part types N and for different values of the number of jobs per tour L (large instances).