Multiple Period Repair Kit Problem with Capacity Considerations

Michael Gorman

University of Dayton, Dayton, USA
Email: michael.gorman@udayton.edu

Abstract

The repair kit problem seeks an optimal stock of parts for a repair kit for purposes of remote repairs. This problem has often been studied when restocking is possible between each order, or when the number of orders is known between each restocking. This research evaluates a model for the repair kit stocking problem for which there is multiple-period demand with a known time interval between each restocking, but an unknown number of on-site repair visits is required during the restocking period. Most previous work focuses on minimum cost subject to some minimum service level; in the case where the value of technician time and customer service is relatively large, the space constraint of the kit and volume of parts comes into play. This work contributes to the literature by testing the robustness of a heuristic originally proposed by [1] in a field study. We conduct numerical experiments to evaluate the heuristic over a wide range of parameterizations. Our results indicate that the heuristic performs close to optimum, and its performance improves as the problem size grows.

1. Introduction

Many organizations must service products for customers in geographically disperse locations with mobile service parts inventories (kits) which must contain the most useful parts to provide a high service level. Large-scale network industries with remote installations such as utilities, telecommunications companies and railroads face repair kit stocking decisions to support their maintenance and repair operations. Emergency ambulance services in health care face the same stocking decisions. Examples also abound in the corporate and consumer service arenas such as office equipment (e.g.,

copiers, printers and computers) which is the motivating application of much prior research.

The challenge in the kit stocking problem is to decide which parts to stock in repair kits with limited space in order to complete required repairs. Inventory holding costs of parts are substantial, but the parts required to complete a repair are not known until the location is visited. Ideally, the repair is completed by a technician in a single visit to avoid the cost and lost customer satisfaction of a second trip; thus, it is highly valuable to carry the right part to complete the service. The known holding cost of carrying a part is weighted against its possible future use in a repair. If the part is not on-hand for a repair, then a cost is incurred for having to revisit the site and for the added inconvenience to the customer of product down time.

This work is based on a field study described in [1] in which technicians must stock their vehicles for warranty service for household appliances (e.g., washers and dryers, refrigerators and ovens, heaters and air conditioners). [1] solved the problem in a specific case; this research extends that work by testing the approach the more broadly in numerical experiments to ensure its robustness and broad applicability, as well as its performance relative to optimality.

2. Literature Review

This problem is often referred to in the literature as the repair kit problem [2]-[4]. Originally proposed by [5] as a single-visit cost minimization problem, [6] resolves the problem with a minimum service constraint. The difference in these approaches stems from modeling a failure as a cost of revisit and lost customer goodwill (as in [6]), or as a constraint on the minimum service level (as in [6]).

[5] and [6] assume that different parts have independent failure processes and that at most one part of each type is required on a single repair. [7] and [8] relax this assumption, extending the single-period model to multiple periods but deterministic number of visits. [2] demonstrates a generalized approach for the various specifications.

In the single-period model, it is assumed the vehicle is restocked after each job. This assumption does not apply well when repair locations are remote and vehicle restocking costs are high. In this case, a technician may visit multiple repair sites between restocking events. [3] and [4] allow for multiple repair visits (known as a "tour") between vehicle restocking. Both assume that the number of customer visits between restocking opportunities is known and fixed.

[1] and [2] generalized the multiple stop repair kit problem to allow for uncertainty in the number of stops between restocking. Often, there is a known *time interval* between restocking opportunities; the number of visits over the interval varies due to a number of factors. First, the travel time required to the locations of the visits varies. Second, job complexity and repair times vary. Third, the stocking level affects the number of visits per repair; if the part is on hand in the vehicle a single visit is required, but if it is not on hand two visits are required. Thus, the relative frequency of first and second visits which affects the uncertainty of part needs.

[4] and [9] minimize part carrying cost subject to a minimum service level constraint. The heuristic of [4] is based on the ratio of the holding cost of the part and the job or part fill rate, adding parts that have the highest service impact per dollar ratio. More recently, [9] develop a heuristic for meeting a minimum service level requirement at minimum cost for a stochastic number of trips in the tour. Neither approach considers part dimensions or kit capacity in the implemented heuristic. Both formulations and solution approaches naturally minimize vehicle inventory value; thus, a minimum service level constraint is imposed to ensure some inventory is held on the vehicle.

[1] is unique in the literature because the size of the part is considered explicitly in the decision to stock it. (This extension was called for by [4] in the conclusion of that work.) The objective function is unique in [1]; they maximize the net benefit of service savings of parts in the kit less their holding cost. This research extends [1] by numerically evaluating the heuristic's performance in a broader setting with various model parameterizations, testing its robustness in a variety of settings.

[1] proposed a heuristic method based on a novel measure—Net Benefit per unit Volume (NBV) which balances the service value of inventory, its cost, and the space it consumes in the kit. They estimate the cost of failure and maximize the net benefit of part carrying (failure avoidance less part carrying cost) per cubic volume consumed by the part, which captures the larger opportunity cost of large parts. The heuristic estimates the costs and benefits of carrying a part and is the basis for improved service vehicle stocking recommendations. Because volume is explicitly included in the heuristic, the formulation requires the inclusion of a vehicle space constraint and individual part volumes to ensure the capacity of the service vehicle is not exceeded ([3] included a space constraint in their formulation, but ignore the restriction in their solution technique because their cost minimizing approach naturally minimizes inventories).

The heuristic employed here is the same as in [1]. Where [1] describes the value of implementing the heuristic in a particular field setting, this research numerically evaluates the quality of the solutions attained relative to optimum with various settings for demand, different counts of parts (SKUs), different part dimensions and various costs for part holding, service and revisit.

3. Model

Notation and Model Assumptions:

Repair kits are restocked on a periodic basis. We define the period between restocking as the lead time, T. The lead time is equal to the order placement frequency plus the delivery lead time, where the reorder placement frequency is a function of the volume of part usage and the distance and cost of the delivery. The minimum leadtime for a part is the transit time; the expected order frequency delay is half the order frequency. For purposes of this modeling endeavor, we assume that the expected lead time is transit time plus one-half the order frequency, and it is treated as a constant.

Replacement part needs are a function of the random arrival of part breakdowns. We model part needs as a Poisson arrival process with mean demand during lead time of d_i

because the Poisson naturally describes arrivals of part orders. The probability is a function of its mean demand and the Poisson distribution. Any demand process with non-increasing probability P_{ij} of use of part i in demand volume of j would work as well. As is common in the literature, we assume that part failure rates probabilities are stochastically independent. Mean demand d_i results in some probability of any level of demand, j, for the part from $j = 0 \cdots M$, where M is the maximum level of demand, or the level for which there is a near-zero probability of M orders for part i over the restocking horizon ($P(D_i \geq M) \leq \varepsilon$, where $\varepsilon = 1\text{E}{-}6$). We assume over the horizon of the repair kit setting decision, demand for each part is stationary. Any seasonality or change in demand can easily be incorporated by resetting the planned repair kit on a seasonal or regular basis.

We note that there are significant costs of secondary visits due to stock outs such as lost productivity of technicians resulting from multiple visits, setup costs for rescheduling and planning each visit, the incremental costs of small-volume, expedited shipments, and the customer inconvenience for the revisit, all of which increase the incentives to carry additional vehicle stock. Additionally, we could include the very real additional cost of customer dissatisfaction from multiple technician visits. Customer dissatisfaction may cause downstream loss of product and service revenue because the service component of the product reflects negatively on the product and causes lost future sales. As [3] points out, this cost is difficult to quantify. Although this cost is not easily quantified, it is quite real. In any case, it applies equally to all parts (a failed visit is a failed visit), so the size of the cost does not affect which parts are stocked, but rather its inclusion increases the incentive to stock parts at all. [1] found that technician time savings alone was enough to justify increased inventories in their application; customer dissatisfaction costs only add to that incentive. Exclusion of revisit costs (or treatment of service requirements as a constraint) naturally leads to an inventory cost minimization approach as described in [3] and [4].

Our fill rate is based on part fill rate rather than a job completion rate as in [4], who considers both job and part fill rates. The job completion rate specifies that all required parts for a job must be in stock on the vehicle in order for the revisit cost to be avoided. A part fill rate specification more than suffices for a number of reasons. Often, a single part is needed in a repair, and thus the complication of job completion rate calculation is unnecessary. When multiple parts are required for a repair, we note (as does [4]) that the parts are often codependent (*i.e.*, they are often installed in tandem) and can thus be combined in the analysis as a single "part". Also, supplemental parts for many repairs are often more general-purpose parts (clamps, bolts, etc.) with relatively high demand and low cost, making them more obvious candidates for stocking. As noted by [4] the part fill rate is a reasonable approximation of the job fill rate when the number of parts on each job is low. Finally, we find that given these stylized facts, that the job completion rate is only slightly lower than the part availability rate. Following the suggestion of [4] that the part fill rate be used, we focus on part fill rates. By focusing on the part fill rate, we were able to take a simplified approach to describing part order ar-

rivals instead of repair job arrivals.

Each part i has some cost, C_i, which directly contributes to its annual holding cost via the company's annual holding cost rate, HC_i. We assume no holding cost is incurred for parts that are not held in stock; essentially all special ordered parts are installed immediately. This is a reasonable assumption as special ordered parts are for a waiting incomplete job and are likely to be served immediately.

We assume that the part volume, v_i, includes the space consumed, including packaging. Further, while we do not explicitly model how parts are fit into the kit, one could easily penalize odd shaped parts or keep a buffer space allowance of the total space V reserved for use for such parts. The cubic area for the kit V includes only the effective storage space of the storage area, adjusted for shelving, part access, and the like. Such considerations are an implementation issue and are not relevant to our efforts here to evaluate the heuristic performance in an experimental setting.

Optimization Model Formulation:

We define the optimization model formulation in Equations (1)-(4).

$$\text{Maximize} \sum_{i=1}^{N} \sum_{j=1}^{M} P_{ij} \cdot RC \cdot \left(1 - x_{ij}\right) - HC_i \cdot x_{ij} . \tag{1}$$

Subject to:

$$\sum_{i=1}^{N} \sum_{j=1}^{M} v_i x_{ij} \leq V \tag{2}$$

$$\sum_{i=1}^{N} \sum_{j=1}^{M} P_{ij} \cdot x_{ij} >= \lambda \cdot \sum_{i=1}^{N} \sum_{j=1}^{M} P_{ij} \tag{3}$$

$$x_{ij} \varepsilon \{0, 1\} \quad \text{for all} \quad x_{ij}. \tag{4}$$

Indices:

N = number of candidate parts,

M_i = maximum units in vehicle inventory for each part,

i = index for parts, $i = 1 \cdots N$,

j = index for unit j of a part, for $i = 1 \cdots M_i$.

Decision Variables:

x_{ij} = Binary decision variable, part i, unit j.

Model Parameters:

V = Total allowable vehicle stocking volume,

RC = Revisit cost,

HC_i = holding cost for part i during lead time, L,

v_i = Unit volume of part I,

P_{ij} = probability of demand level j for part i during lead time, L,

λ = minimum required part fill rate (0% to 100%).

The binary decision variable x (constraint 4) is whether to carry the unit j ($j = 0 \cdots M$ units) of part i (out of $i = 1 \cdots N$ candidate parts); there are $N * M$ total binary variables in the problem. A binary variable is used to allow different coefficients on each subsequent unit of each part corresponding to falling probability of use. Note that although

x_{ij} is a binary variable depicting the decision to carry unit j of part i, no additional constraints on the decision variables j in the same part i are necessary. With non-decreasing probability of use of each incremental unit of part i and a constant cost of each unit, there is a non-increasing benefit of each successive unit of inventory. Thus, no solution will opt for unit $j + 1$ before opting for unit j, which has a higher probability of use and the same holding cost. No additional constraints are necessary to ensure that if inventory unit $j + 1$ is carried, then part j will also be carried.

The model finds the level of inventory that optimizes the vehicle stock and service level trade-off from a maximization of the net benefit of inventory perspective, including both inventory holding and (avoided) revisit costs. The objective function (1) maximizes the net benefit of kit inventory: the probability P_{ij} of needing part I, unit j times the revisit cost RC, less the holding cost of each part stocked HC_i. If a part is carried ($x_{ij} = 1$), then the potential revisit cost (RC) is avoided based on the probability (P_{ij}) of needing unit j of part i during the lead time interval between restocking events, and a holding cost (HC_i) is incurred with certainty.

Constraint (2) is the vehicle space constraint, in which both part size and vehicle capacity are specified in cubic volume. Constraint (3) reflects a minimum service level constraint, where λ is the part fill rate measured as a percentage. As is discussed in [4], we note that with a high frequency of single-part repairs, high part fill rates, and combining parts that are demanded in common, that the part fill rate only slightly overstates the job fill rate.

It is highly unlikely that both constraints 2 and 3 will be binding constraints in an optimal solution. This would imply that a minimally acceptable service level can be achieved through a completely fully stocked service vehicle. Either inventory holding costs are relatively high, driving service levels down to a minimally acceptable threshold, or revisit costs are high, driving inventory levels to vehicle capacity (or, a non-binding solution somewhere in between). In the model developed by [3] and [4], it is implicitly assumed (through omission of the revisit cost component of Equation (1)) that inventory holding costs dominate revisit costs resulting in an inventory minimizing objective. The approach of [4] is more appropriate in the case where inventory costs are large relative to revisit costs and the economic solution leads to an unacceptable service level.

We develop a heuristic that focuses on the vehicle capacity constraint (3).

In our numerical experiments, we focus on the case where space is scarce, or where the value of the parts pales in comparison to repairman revisit cost and customer satisfaction. In this case, a minimally acceptable service level could leave parts with positive economic value out of the stock, or force parts with negative economic value into inventory. To the extent that revisit (and service) costs dominate inventory costs, the tendency will be to fill the kit, not to achieve a minimally acceptable service level at a minimal cost. If sufficient kit capacity exists, the resulting service level would be based on the relative costs and benefits of holding inventory (*i.e.*, the "economic service level").

4. Heuristic Solution Methodology

The above formulation is a classic knapsack problem which can grow quickly for realistic sized problems. Similar to previous research, we explore the accuracy of heuristic methods as an alternative to exact solutions. In general, the heuristic considers the known inventory carrying cost (HC_i) of each part, relative to its expected benefit, which is a function of the cost of a revisit (RC) and the probability of the jth unit of part i(P_{ij}) being demanded over the lead time. This net expected benefit is divided by the total volume (v_i) of the part to capture the opportunity cost of the space taken in the kit. Parts are selected in descending order of their net expected benefit per unit volume, $NBV_{ij} = (RC*P_{ij} - HC_i)/v_i$. Parts enter the proposed inventory until the kit's capacity is consumed (space constraint met) or NBV_{ij} becomes negative (economic return does not justify stocking).

The greedy Net Benefit per unit Volume (NBV) heuristic implemented follows.

1) Find maximum units, M_i, for each part, i

For $i = 1 \cdots N$

$\quad M_i = 0$

$\quad j = 1$

While $j <> M_i$

$EB_{ij} = RC*P_{ij}$

$\quad$ If $EB_{ij} < HC_i$

$\quad\quad M_i = j$

$\quad\quad i = i + 1$

$\quad\quad j = 1$

$\quad$ else $j = j + 1$

Next

Next i

2) Calculate net benefit per unit volume (NBV)

For $i = 1 \cdots N$

For $j = 1$ to M_i

$\quad NBV_{ij} = (EB_{ij} - HC_i)/v_i$

3) Establish stocking levels

$\quad CPV = 0$

$\quad$ Sort NBV_{ij} in decreasing order

$\quad$ While $NBV_{ij} \geq 0$ and $CPV \leq V$

$\quad\quad$ If $CPV + v_i < V$ then $CPV = CPV + v_i$

$\quad\quad x_{ij} = 1$

$\quad$ Next.

We define three possible outcome conditions for the kit problem: binding service level constraint, binding kit size constraint, or no binding constraints. In the first case, a relatively high inventory cost and low revisit costs leads to a low inventory, low service solution for which a minimum service constraint is binding. In the second case, inventory is relatively less costly to hold than incurring low service and revisit costs

leading to the full kit situation, for which the service value of additional inventory is positive, but the kit size constraint is binding. In the third case, the economic service level is derived endogenously by including all parts that have a positive net profit impact (avoided expected revisits less inventory holding costs) with no binding constraints necessary to arrive at this solution. In this case, the optimal stock naturally leads to a level of inventory greater than some deemed minimum, but the optimal total inventory is less than the total volume of the kit.

As discussed, in the binding service constraint case, the heuristic method suggested by [4] is appropriate because it focuses on achieving the minimum service level at least cost, where the heuristic presented here would put too much emphasis on the volume of the parts in this case. In the second case, it is critical to assess the accuracy of the heuristic in selecting the best parts from both service level and economic return perspectives; this assessment is carried out numerically below. In the no binding constraint case, this heuristic is guaranteed to arrive at the optimum solution because all parts with positive economic value are included in the final kit via heuristic or full search; conversely, the heuristic suggested by [4] would not arrive at the same solution as would be indicated by the objective function above. We note that [3] implemented the inventory cost objective subject to a minimum service constraint, but interestingly in their case study the client chose to improve customer service rather than reduce costs.

5. Small-Scale Example Problem

The illustration depicted in Table 1 shows an example of this heuristic with a comparison to that proposed by [4] based on 5 parts. This example is based on the case with no binding service or space constraints, in which the NBV heuristic achieves the optimum economic service level. The results are sorted in the order that the part i and unit j are admitted into the kit. Note that because the space constraint is not exceeded (91 cubic feet used out of 100 cubic feet available), all parts with positive NBV are accepted. (One part with a negative NBV is shown for illustrative purposes.)

Focusing on part 5 reveals the primary difference between [4] and the NBV heuristic. Although part 5 is expensive to hold, it is small and in demand; thus, the NBV heuristic admits it as the 4th and 5th parts. The heuristic of [4] admits part 5 much later (12th and 13th) due to its cost. In this example, if a 50% service level was set as the minimum, the heuristic of [4] would not have included part 5 in the kit, having stopped after admitting the 11th part.

6. Numerical Analysis

We test the heuristic in numerical experiments with randomly generated values of key model parameters that lead to exhausting kit capacity to see how well it performs relative to the optimum in allocating scarce storage capacity. We hold the V, the volume of the kit, L, the leadtime, and RC, the revisit cost constant, and let d_i, v_i and HC_i vary (these values are only relevant on a relative basis). In effect, d_i is scalable with L, c_i relative to RC, and v_i relative to V.

Table 1. Illustrative example of NBV heuristic results.

Part Num (i)	Mean Demand (d_i)	Part Size (v_i)	Cost/ Period (HC_i)	Optimal Solution:	
1	2	3	$10.00	Objective Function	$241.93
2	2	4	$15.00	Space Used	91 Cubic ft
3	3	8	$20.00	Service Level	75%
4	4	10	$25.00		
5	5	2	$40.00		

Part number (i)	Demand level (j)	Prob (D >= j)	NB/V	Space	Teunter ($\Delta j/HC_i$)	Rank (NB/V)	Rank (Teunter)	Svlv
1	1	86%	$11.08	3	0.09	1	1	5%
2	1	86%	$7.06	7	0.06	2	3	15%
1	2	59%	$6.57	10	0.06	3	2	9%
5	1	99%	$4.83	12	0.02	4	12	58%
5	2	96%	$3.99	14	0.02	5	13	64%
2	2	59%	$3.67	18	0.04	6	6	29%
3	1	95%	$3.44	26	0.05	7	4	20%
3	2	80%	$2.51	34	0.04	8	5	25%
4	1	98%	$2.41	44	0.04	9	7	35%
1	3	32%	$2.06	47	0.03	10	9	43%
4	2	91%	$2.04	57	0.04	11	8	41%
5	3	88%	$1.88	59	0.02	12	15	73%
4	3	76%	$1.31	69	0.03	13	10	48%
3	3	58%	$1.11	77	0.03	14	11	51%
4	4	57%	$0.33	87	0.02	15	14	67%
2	3	32%	$0.29	91	0.02	16	16	75%
~~3~~	~~4~~	~~35%~~	~~$(0.30)~~	~~96~~	~~0.02~~	~~17~~	~~17~~	~~81%~~

RC and V are set arbitrarily at 100. C_i range is set to drive an optimal service level above 0.9 for each random generation of demand. Mean demand is uniformly distributed between zero and 5. Demand is assumed to follow a Poisson process. v_i and N are set so that the $v * N$ coupled with demand distribution results in space demand for 150% of V. When the kit utilization is 100%, the feasible service levels in this simulation range between 60% and 95% with a mean of 76% - 78%, depending on the randomly generated demand values. For demand with a maximum mean demand of 5, M is set to 12 for all parts, leaving at most 0.3% chance of a 12th unit of demand in any period, which equates to a negative or very low NBV for any part; $M = 12$ is a reasonable maximum for this experiment. (In no cases was the 12th unit of any part in the optimal stock.)

In the numerical exploration here, in all cases the kit will face a tight space constraint

and be essentially full in the optimum case. In the case with no binding space or service level constraints, the optimum is achieved in every case with this heuristic, and is not of interest to analyze. In the case where the optimal service level is less than the required service level, the heuristic proposed by Teunter [4] should be used (This case is not explored in this research).

In summary, the parameters for the numerical experiment are as follows.

$RC = 100$,

$V = 100$,

$D_i = \text{Poisson}(U(0,5))$,

$N = (5, 10, 30, 30, 100, 1000)$,

$HC_i = U(0, 0.1RC)$,

$v_i = U(0, 1.5 * V/N)$.

For each of six different part type counts, 1000 random replications of mean demand are generated, the model solved via the heuristic and optimally (For $N = 1000$, 250 replications are run for time reasons). Results of the global search and the heuristic are compared based on objective function value, service level, and space utilization.

Results Discussion:

The results of the numerical experiments are presented in Table 2. The variable of interest is in the number of parts under consideration for stocking, N, which drives the number of decision variables and problem complexity. Other random variables were to create problem variation. Note that the optimal service level ranges between 75% and 78% and the space constraint is tight; smaller parts and 100% service level is uninteresting because it is guaranteed optimal. (All parts with some value greater than their cost are included.) Lower optimal service levels would require an approach more like that proposed by [4], which focuses on minimizing holding cost subject to a minimum service level constraint.

Table 2. Numerical results from simulation as number of parts grows.

Parts	5 Parts		10 Parts		20 Parts		30 Parts		100 Parts		1000 Parts	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Opt-Obj	$435.76	$155.67	$891.77	$229.16	$2,762.36	$402.29	$2,761.49	$403.02	$9,252.44	$691.78	$89,266.11	$2,212.53
Heur-Obj	$425.65	$153.34	$878.40	$226.89	$2,745.96	$399.87	$2,745.10	$400.59	$9,233.70	$690.40	$89,194.60	$2,219.74
Pct Difference	2.30%	3.30%	1.50%	1.70%	0.90%	0.90%	0.60%	0.50%	0.19%	0.16%	0.10%	0.20%
Opt Svlv	75%	15%	77%	11%	78%	6%	78%	6%	78%	4%	78%	1%
Slv-heur	74%	16%	76%	12%	77%	7%	77%	7%	78%	4%	78%	1%
Pct Difference	1.6%	2.1%	1.1%	1.2%	0.5%	0.4%	0.5%	0.4%	0.2%	0.1%	0.0%	0.1%
Space-Opt	96.38	5.12	99.19	1.31	99.91	0.08	99.91	0.08	100.00	0.0001	100.00	0.01
Space-Heur	90.21	7.01	94.83	3.62	98.31	1.17	98.31	1.17	99.58	0.40	99.48	0.36

n = 1000 for all part levels but 1000 parts, where n = 250

We observe that the NBV heuristic averages within 2.3% of optimum on average in all cases. Importantly, the deviation between the heuristic and optimum *shrinks* as the problem size grows; conveniently, when exact methods become too computationally burdensome to undertake, the heuristic is most accurate. For values of N ≥ 100 parts, the NBV heuristic and optimum solutions are essentially the same. The reason for this is intuitive; the more parts that fit in the kit, the less the impact of the parts on the margin that an exact optimization finds but a heuristic does not. The vast majority of the parts to stock are "easy" decisions; as the number of parts grows, marginal parts are a small percentage relative to the number of parts under consideration. The same observations can be made about space utilization and service levels; as the number of parts rises (and their size falls in this experiment to keep service levels constant across scenarios), the heuristic consistently finds better combination of parts that fills the kit and achieves the highest service level.

Following a minimum cost heuristic like that of [4] or [9] with a tight capacity constraint leads to lower service levels and higher revisit costs (though lower holding costs) because no consideration is given to the tight space constraint and oversized parts consume space that multiple smaller parts could have used.

7. Conclusions

This research considers the kit stocking problem for multi-period stocking and uncertain number of jobs. A Net Benefit per unit Volume (NBV) heuristic is proposed as a fast and effective greedy heuristic for finding a near-optimal kit inventory based on the relative costs of inventory holding and repair visit failures, as well as the space that the part consumes in the kit of limited capacity. We suggest that when repairman time and customer service are valued heavily relative to inventory holding costs, the proposed objective function that maximizes net benefit is more appropriate than one that minimizes cost subject to a minimum service constraint (This approach is not valid in the case of valuable inventory and low valued time).

Numerical simulations demonstrate the effectiveness of the approach. We show in the case of a non-binding space constraint (inventory levels are determined strictly by the relative costs and benefits of each unit of inventory), the heuristic arrives at the optimal solution all the time. In the case of a binding space constraint, the heuristic regularly achieves an objective function and service level within 2% of optimum. Importantly, the heuristic performance improves as the problem size grows as the integrality constraint on the part stocking decision plays a lesser roll. We suggest that the NBV heuristic is a highly appropriate substitute for exact methods on large-scale problems with high inventory storage demands on scarce space.

References

[1] Gorman, M.F. and Ahire, S. (2006) A Major Appliance Manufacturer Rethinks Its Inventory Policies for Service Vehicles. *Interfaces*, **36**, 407-419.

[2] Brumelle, S. and Granot, D. (1993) The Repair Kit Problem Revisited. *Operations Research*,

41, 994-1006. http://dx.doi.org/10.1287/opre.41.5.994

[3] Heeremans, D. and Gelders, L. (1995) Multiple Period Repair Kit Problem with a Job Completion Criteria: A Case Study. *European Journal of Operational Research*, **81**, 239-248. http://dx.doi.org/10.1016/0377-2217(94)00177-E

[4] Teunter, R.H. (2006) The Multiple-Job Repair Kit Problem. *European Journal of Operational Research*, **175**, 1103-1116. http://dx.doi.org/10.1016/j.ejor.2005.05.023

[5] Smith, S.A., Chambers, J.C. and Shlifer, E. (1980) Optimal Inventories Based on Job Completion Rate for Repairs Requiring Multiple Items. *Management Science*, **26**, 849-852. http://dx.doi.org/10.1287/mnsc.26.8.849

[6] Graves, S.C. (1982) A Multiple-Item Inventory Model with a Job Completion Criterion (Note). *Management Science*, **28**, 1334-1337. http://dx.doi.org/10.1287/mnsc.28.11.1334

[7] Mamer, J.W. and Smith, S.A. (1985) Optimizing Field Repair Kits Based on a Job Completion Rate. *Management Science*, **28**, 1328-1333. http://dx.doi.org/10.1287/mnsc.28.11.1328

[8] March, S.T. and Scudder, G.D. (1984) On "Optimizing Field Repair Kits Based on a Job Completion Rate". *Management Science*, **30**, 1025-1028. http://dx.doi.org/10.1287/mnsc.30.8.1025

[9] Bijvank, M., Koole, G. and Vis, I. (2010) "Optimising a General Repair Kit Problem with a Service Constraint. *European Journal of Operational Research*, **204**, 76-85. http://dx.doi.org/10.1016/j.ejor.2009.10.004