Beta
Research School for Operations Management and Logistics

Improving effectiveness of spare part supply by additive manufacturing as dual sourcing option

N. Knofius, M.C. van der Heijden, A. Sleptchenko, W.H.M. Zijm

Beta Working Paper series 530

BETA publicatie	WP 530 (working paper)
ISBN	
ISSN	
NUR	
Eindhoven	May 2017

Improving effectiveness of spare part supply by additive manufacturing as dual sourcing option

N. Knofius*[1], M.C. van der Heijden[1], A. Sleptchenko[2], and W.H.M. Zijm[1]

[1] *Department of Industrial Engineering and Business Information Systems, University of Twente, Enschede, The Netherlands*

[2] *Department of Mechanical and Industrial Engineering,College of Engineering,Qatar University, Doha, Qatar.*

Abstract

The low-volume, high-variety spare parts business is often identified as a potential beneficiary of additive manufacturing (AM) technologies. Short lead times and low setup cost may substantially reduce holding and backorder cost. Unfortunately, high unit cost or low and uncertain reliabilities of AM parts deem the application of AM economical inferior to conventional manufacturing (CM) methods in most cases. In this paper, we investigate the potential to overcome these deficiencies by combining AM and CM methods. For that purpose, we develop an approach that is tailored towards the unique characteristics of dual sourcing with two production methods. In particular and opposed to traditional dual sourcing literature, we consider the different failure behavior of parts produced by AM and CM methods. A case study in the aviation industry and numerical experiments show that dual sourcing is often superior compared to single sourcing, especially for conditions found in the spare part business for advanced capital goods: low demand rates, high backorder cost, and high holding cost. In comparison, single sourcing with AM methods typically leads to higher purchasing cost while single sourcing with CM methods increases backorder or holding cost. Savings of more than 10% compared to the best single sourcing option are likely even if the reliability or unit cost of a part sourced with AM are three times higher than for a CM part. In conclusion, dual sourcing methods may play an important role to exploit the benefits of AM methods while avoiding its drawbacks in the spare part business.

Keywords: Digital manufacturing, 3D printing, Dual Sourcing, Spare parts

*Corresponding author. Tel.: +31534896515. E-mail addresses: n.knofius@utwente.nl (N.Knofius), m.c.vanderheijden@utwente.nl (M.C. van der Heijden), andrei.sleptchenko@qu.edu.qa (A. Sleptchenko), w.h.m.zijm@utwente.nl (W.H.M. Zijm).

1 Introduction

Spare part inventories are essential to keep downtime of advanced capital goods within reasonable limits, cf. Sherbrooke (2004), and Van Houtum and Kranenburg (2015). Investments in spare part inventories can be huge, as the assortment contains many different items, amongst which many expensive parts, slow movers, and long lead time items. This situation holds in particular for parts manufactured using conventional manufacturing (CM) technologies like milling, drilling or injection molding.

Additive manufacturing (AM), also referred to as 3D printing, increasingly matures to an alternative for spare parts production. The potential is exemplified by Airbus, a manufacturer of aircrafts, reporting that its aircraft A350 XWB contains more than 2700 printed parts already today (Airbus, 2016). The unique process of AM in which raw materials are added layer upon layer to build complete parts at once is solely based on digital models. Therefore, part-specific tools or long setup times are untypical for AM methods. Especially in the low-volume, high-variety spare parts business lead times may reduce substantially, and as a consequence decrease or even avoid the need for safety stocks.

The situation we encounter at a service provider in the aerospace industry, whom we revisit later in this paper, may clarify the opportunity of digitization with AM methods. The lifecycle of aircrafts typically spans several decades. Throughout this period, the service provider has to replace deteriorated and failed parts which are often unique for each aircraft type. Currently, the company guarantees supply for more than hundred thousand parts. Customer order response times are deemed to be short, given the high costs of grounded aircrafts and fierce competition in the market. However, short response times typically invoke high stocks around the globe, resulting in high tied-up capital.

Amid the nature of the aerospace business, it is not surprising that, as stated by one of the product managers, AM technologies hold an enormous potential for its service operations. For instance, one can imagine scenarios where spare parts are printed on demand and close to the customer site as studied by Khajavi et al. (2014) and Liu et al. (2014). Another opportunity may arise from AM enabled design changes. For example, Knofius et al. (2017)

study the option to redesign spare parts with fewer but therefore more complex components and find that service costs change significantly.

Despite these opportunities, it is doubtful whether AM technologies will replace CM methods. Instead, it is more likely that AM methods will complement CM methods (Holweg, 2015). As shown by Westerweel et al. (2016), the success in the spare parts business depends on several factors; lead time reductions alone do not compensate for disadvantages of AM technologies. High unit cost or low and uncertain reliabilities of printed parts often rule out the use of AM. Unsolved issues like uncertain liability or non-standardized certification processes cause additional obstacles. While in the latter case significant progress is made, see Scott (2017), high unit cost and reliability concerns are more likely to prevail. Specifically, AM often suffers from additional cost for post-processing and quality checks, high material cost, high equipment cost and high process variability, cf. Book and Sangid (2016), and Frazier (2014).

In this paper, we study the value of sourcing a spare part by combining AM and CM methods. That is, depending on the situation one may decide whether to source a spare part with AM or CM methods. In the traditional dual sourcing literature, costs are higher and lead times are lower for the fast supply option. In our case, this does not necessarily hold true, and more importantly: we take into account that AM parts may show a different failure behavior than CM parts. This extension seems inevitable given that not only the production process but also design and used material are typically different for both sourcing modes, e.g. Wits et al. (2016). As a consequence of the differences in failure behavior, the sourcing decision will impact future demand. This characteristic makes the overall trade-off more complex. To our knowledge, this aspect has not been considered in literature so far.

As we will show, the resulting flexibility of selecting between AM and CM methods, leads to significant cost savings compared to a conventional single sourcing approach (typically more than 10%). This result holds especially under conditions that are often observed in the spare parts business: low demand rates, high holding and downtime cost. Furthermore, it turns out that the benefits of dual sourcing remain high, even if the AM part characteristics are (largely) inferior to the CM version. As we will discuss, this observation may motivate new sourcing concepts where AM parts are used as temporary fix for capital goods operated at remote locations. Also, we will argue that printing parts on demand is typically not suitable for downtime critical spare parts. Instead, under these conditions, our findings demonstrate that a dual sourcing approach where AM is used as a fast emergency source outperforms single sourcing substantially.

The remainder of the paper is organized as follows. In Section 2, we position our work within the literature. Afterwards, in Section 3, we develop a specific dual sourcing model and explain its evaluation and optimization. In Section 4, we conduct numerical experiments to study the value of dual sourcing in the spare parts business. Section 5 revisits the service provider and demonstrates the application of our model in practice. We close with Section 6, in which we suggest directions for future research.

2 Literature review

We discuss two streams of literature. In Section 2.1, we review dual sourcing literature where we focus on selected papers and refer to Minner (2003) and Zhou and Yang (2016) for a more extensive discussion. In Section 2.2, we discuss (the only) two papers considering the combination of AM and CM methods to fulfill demand.

2.1 Dual Sourcing

Dual sourcing models typically distinguish between two supply options: one that is cheap but with a long resupply lead time (regular supply), and one that is expensive but with a short resupply lead time (expedited order). The first contribution to the dual sourcing literature was made by Barankin (1961), who discusses a single-period model with emergency shipments. Whittemore and Saunders (1977) consider the periodic review case with deterministic lead times. They show that as soon as lead times between both supply options differ by more than one period, the optimal policy depends on delivery time and quantity of the in-transit parts. Although dynamic programming methods allow solving such problems, the large state space leads to computational intractability even for medium sized problems. Hence, more recent contributions are devoted to approximations of the optimal policy.

For the periodic-review case, Veeraraghavan and Scheller-Wolf (2008) study a capacitated inventory model with deterministic lead times. They propose a dual-index policy that keeps track of one inventory position for each sourcing option. While the evaluation is partially based on simulation, the dual-index policy provides close to optimal results. Scheller-Wolf et al. (2007) use a single-index policy where they monitor a single inventory position and a target level for each sourcing option. In case the inventory position is below the expedited target level, an order is placed to raise the inventory position to this level. Next, a regular order is used to bring the inventory position up to the regular target level. It turns out that the single-index policy performs comparably well but can be computed 25-60 times faster than the dual-index policy.

For the continuous-review case, Moinzadeh and Schmidt (1991) propose a dual-index policy with deterministic lead times to determine the order quantity for both supply options. Using information about arrival times of outstanding orders, they keep orders from the expensive sourcing option limited. Song and Zipkin (2009) extend the model of Moinzadeh and Schmidt (1991) for multiple supply options and stochastic lead times. Therefore, they construct a queueing network with overflow bypass. Zhou and Yang (2016) introduce a single-index (R,nQ) policy and study a more general setting which regards fixed order cost, batch ordering and compound Poisson demand. Allon and Mieghem (2010) propose a tailored base-surge (TBS) policy with stochastic lead times. Under the TBS policy goods are ordered at a constant rate from the cheap supply source while orders for the expensive supply source are issued according to a base stock policy. Song et al. (2017) find the optimal policy for a continuous-review system for a special case. Therefore they transform the original problem to a simplified queueing system that shares the same optimal policy under specific conditions on the net inventory. For cases in which these conditions are violated, they exploit the results of the queueing system to construct a policy for the original system.

Numerically they show that the resulting heuristic policy is close to optimal and typically outperforms the discussed methods of Song and Zipkin (2009) and Allon and Mieghem (2010). To the best of our knowledge, none of the many papers on dual sourcing addresses the impact of sourcing decisions on future demand which is essential for our analysis (cf. Section 1).

2.2 Combining AM and CM methods

Literature which discusses demand fulfillment with AM and CM production methods is limited. The two papers discussed subsequently are, to the best of our knowledge, the only contributions in this field. Khajavi et al. (2015) establish the possible value of combining AM and CM methods for the production of new products. They conceptualize that, by postponing the setup of CM methods with the usage of AM methods, substantial cost savings appear possible. First, the financial risk in case of market failure or unexpected low demand is kept in check by low setup cost of AM methods. Second, typically short setup times decrease the time to market and thus may secure an early movers advantage. Third, design changes, which are likely for new products, are way less of a burden with AM than with CM methods. Finally, by shifting to CM methods if the product has proven successful, scale effects and shorter throughput times can be exploited.

The paper closest related to our work is the one by Song and Zhang (2016), who consider the use of AM methods as an emergency channel that may produce spare parts on-demand. In their model, AM equipment is capacitated (modeled as an M/D/1 queue) but typically allows faster, though more expensive, resupply than the CM source. Also, and this is a fundamental difference compared to our model, they assume that AM parts have the same failure behavior as CM parts. Overall, they find that the production of parts on-demand with AM methods leads to significant cost savings and inventory reductions compared to the application of CM methods only. Especially, for situations with large part variety, these findings were found to hold true. Furthermore, they report that the AM equipment utilization typically remains low and therefore may support the assumption of an uncapacitated printing source.

Our paper contributes to the literature as follows:

1. We develop a new dual sourcing model for a single-item, in which future demand depends upon the sourcing options being used via different failure rates for different spare parts types. We develop an exact algorithm to solve this model.
2. In numerical experiments, we explore the value of using a combination of AM and CM parts compared to single sourcing alternatives and study the structure of the optimal policy in case of dual sourcing.
3. We establish the practical value of dual sourcing using a case study in the aviation industry.

3 Model

3.1 Model description and notation

Consider a single-item inventory system which serves an installed base of k systems, where each system requires one unit of the item to be operational. In principle, the item can be produced either by AM or CM methods. Both item versions exhibit differences in failure behavior, unit cost and lead time. Here we assume that item failures follow a Poisson processes with rates λ_{CM} and λ_{AM} respectively. The unit cost of a new CM or AM part are denoted by c_{CM} and c_{AM} and the resupply rates are distributed exponentially with rates μ_{CM} and μ_{AM}. Even though one may question the validity of exponentially distributed lead times in practice, we apply this assumption for two key reasons. First, it is known that the performance of inventory systems for slow-moving spare parts is not very sensitive to the shape of the lead time distribution, cf. Alfredsson and Verrijdt (1999), and Alvarez et al. (2013). Second, this assumption facilitates the use of a continuous time Markov chain analysis.

The installed base is supported by a single stockpoint carrying S non-repairable spare parts. Upon failure, a (CM or AM) spare part is taken from stock, and a new (CM or

AM) part is ordered immediately. So, we assume one-for-one replenishment, as is common for slow-moving spare parts. This means that the total number of parts in the system (operational, on stock or in resupply) equals $N = k + S$. In case we run out of stock, demand is backordered, and we incur backorder cost b per item per time unit. In fact, these backorder costs can be interpreted as penalty costs for system downtime. Otherwise, if a spare part is available, the replacement of the failed part takes place instantaneously. Holding cost are modeled as a fraction h of the associated unit cost of the items in stock. The state $\mathbf{i}$ of the inventory system is described by the tuple $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ where $nC_{\mathbf{i}}$ ($nA_{\mathbf{i}}$) refers to the number of CM (AM) parts in operation, $rC_{\mathbf{i}}$ ($rA_{\mathbf{i}}$) refers to the number of CM (AM) parts in resupply, and $sC_{\mathbf{i}}$ ($sA_{\mathbf{i}}$) refers to the number of CM (AM) spare parts in stock. The set of feasible states is equal to

$$\begin{aligned} \Omega = \{(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}}) : \\ nC_{\mathbf{i}} + nA_{\mathbf{i}} + rC_{\mathbf{i}} + rA_{\mathbf{i}} + sC_{\mathbf{i}} + sA_{\mathbf{i}} = N \\ sC_{\mathbf{i}} + sA_{\mathbf{i}} = max\{N - k - rC_{\mathbf{i}} - rA_{\mathbf{i}}; 0\} \\ nC_{\mathbf{i}} + nA_{\mathbf{i}} \leq k \\ nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}} \geq 0\} \end{aligned}$$

The definition of the state space excludes degenerated transitions. That is, in case $nC_{\mathbf{i}} + nA_{\mathbf{i}} \leq k$ it is impossible to have items on stock. Consequently, a failed item always has to be replaced immediately if possible. Throughout this paper, we use a consistent order (e.g. lexicographical) of the state probabilities if collected in a vector or matrix.

Upon failure of an item, one has to take two decisions. First, whether to use an AM or a CM item from stock (if possible) to replace the failed item (maintenance decision). Note that the item which failed does not need to be replaced with the same item version. Second, we have to decide whether to reorder an AM or a CM item to replenish the stock (sourcing decision). Optimal decisions are dependent on the state $\mathbf{i}$. For example, it may appear optimal to order the AM version to exploit a typically faster resupply rate of AM if the stock is (nearly) depleted. In other states, though, sufficient stock might be available. Here, despite a long lead time, it may be optimal to order an often more reliable and cheaper CM version of the item instead.

We define a matrix $\mathbf{X}$ to represent the decisions for all states, where each column corresponds to a certain state $\mathbf{i}$, and each row to a decision option $c \in 1, 2, 3, 4$, defined as follows:

- $c = 1$: take AM version from stock (if possible) and order AM version.
- $c = 2$: take CM version from stock (if possible) and order AM version.
- $c = 3$: take AM version from stock (if possible) and order CM version
- $c = 4$: take CM version from stock (if possible) and order CM version

Note that in cases $sC_{\mathbf{i}} + sA_{\mathbf{i}} = 0$ it is indifferent whether to choose $c = 1$ or $c = 2$. The same holds true for the choice between $c = 3$ and $c = 4$. We denote the decisions corresponding to state $\mathbf{i}$ as a column vector $x(\mathbf{i})$ with length 4. Component c of vector $x(\mathbf{i})$ is denoted by the binary variable $x_c(\mathbf{i})$.

The goal of the model is to minimize the long-run average cost by determining the optimal sourcing, maintenance and base stocking policy. In principle we can solve this problem by different methods. The probably most common approach would be the utilization of a continuous-time Markov Decision Process (MDP) that, after transferring it to an equivalent discrete-time MDP, see Heyman and Sobel (1984), can be solved using methods like linear programming, value or policy iteration. Here, we employ a continuous-time Markov Chain analysis in combination with linear programming methods to obtain a more efficient algorithm: instead of four equations per state (for each action one), we can represent the problem with two equations per state (balance equation and policy constraint). In essence our modeling approach follows the lines of Sleptchenko and Johnson (2015). In the next two sections, we discuss how to evaluate and optimize the sourcing and maintenance policy. Afterwards, we focus on the stocking policy.

3.2 Model evaluation

Given $\mathbf{X}$, k and S, the model can be evaluated by means of a continuous-time Markov Chain. In Figure 1 we show the ten possible transitions to and from state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ for the case where an AM part arrives or fails. The ten transitions associated with the arrival or failure of a CM part are omitted but exhibit the same pattern. We use $\mathbf{i}'$ to refer to the associated predecessor states of $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$. Furthermore, if the maintenance decision is indifferent (i.e. no stock available), as we elaborated in the previous section, we use $x_2(\mathbf{i})$ and $x_4(\mathbf{i})$ as default in Figure 1.

To explain the underlying logic of Figure 1, let us focus on the transition in the top-right

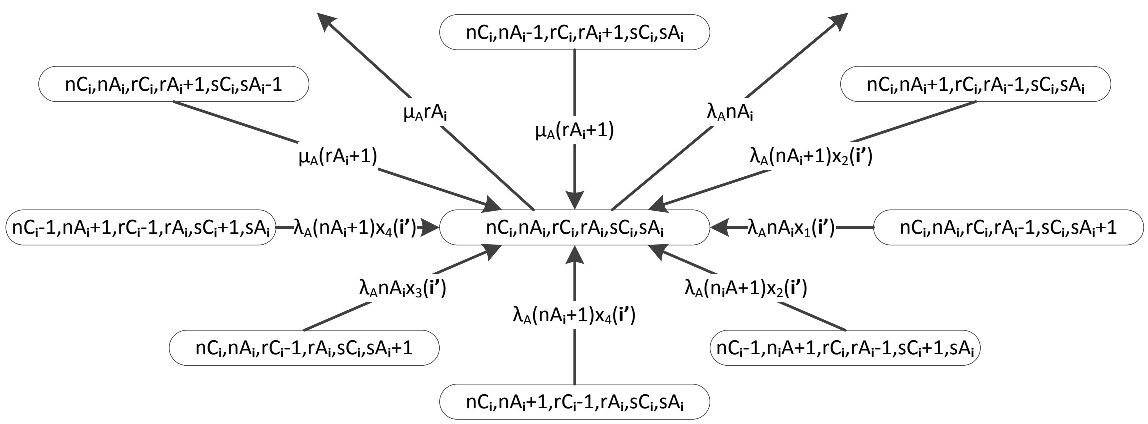

Fig. 1. Transitions of state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ in case an AM item arrives or fails

corner, i.e. with a rate of $\lambda(nA_{\mathbf{i}}+1)x_2(\mathbf{i}')$ we transition from state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}+1, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ to state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$. This transition describes the situation where an AM item fails and we decide to order an AM part while no items are on stock. Hence, $nA_{\mathbf{i}}$ decreases by one unit while $rA_{\mathbf{i}}$ increases by one unit. The balance equations (cf. Appendix 1) directly follow from the transitions illustrated in Figure 1 and, in combination with the normalization equation, allow the computation of the state probabilities $p_{\mathbf{i}}$ with common methods. The result is captured by the column vector $\mathbf{p}$ with elements $p_{\mathbf{i}}$.

After $\mathbf{p}$ was determined, the long-run average costs C are computable by $\mathbf{g}^{\mathbf{T}}\mathbf{p}$ with $\mathbf{g}$ representing an $|\Omega|$-dimensional column vector of cost $g_{\mathbf{i}}$ in state $\mathbf{i}$. Given that we consider unit cost, holding cost and backorder cost we have

$$g_{\mathbf{i}} = \mu_{TM} c_{TM} r C_{\mathbf{i}} + \mu_{AM} c_{AM} r A_{\mathbf{i}} + h(c_{TM} s C_{\mathbf{i}} + c_{AM} s A_{\mathbf{i}}) + max\{k - nC_{\mathbf{i}} - nA_{\mathbf{i}}; 0\} b$$

3.3 Optimization of the sourcing and maintenance policy

Formally, the optimization problem can be expressed as Problem 1:

$$\begin{aligned} \underset{\mathbf{X},\mathbf{p}}{\text{minimize}} \quad & \mathbf{g^T p} \\ \text{subject to} \quad & \mathbf{Q(X)p} = \mathbf{0} \\ & \mathbf{1}^{\mathbf{T}}_{|\mathbf{\Omega}|}\mathbf{p} = 1 \\ & \mathbf{p} \geq \mathbf{0} \\ & \mathbf{1}^{\mathbf{T}}_{|\mathbf{c}|}\mathbf{X} = \mathbf{1}^{\mathbf{T}}_{|\mathbf{\Omega}|} \\ & x_c(\mathbf{i}) \in \{0,1\} \quad \forall c, \mathbf{i} \end{aligned} \tag{1}$$

where $\mathbf{Q(X)p} = \mathbf{0}$ represents the balance equations with matrix $\mathbf{Q(X)}$ describing the generator of the Markov process and $\mathbf{0}$ denoting an $|\Omega|$-dimensional vector of zeros. Furthermore, $\mathbf{1}^{\mathbf{T}}_{|\mathbf{\Omega}|}\mathbf{p} = 1$ defines the normalization equation with $\mathbf{1_m}$ being an m-dimensional column vector of ones.

Given the product $x_c(\mathbf{i})p_{\mathbf{i}}$ in the balance equations (cf. Appendix 1), problem formulation (1) is nonlinear and computationally intractable. Therefore, we transform (1) into an equivalent linear formulation. Here we follow the approach of Sleptchenko and Johnson (2015) who encounter a comparable problem structure. Key step of the transformation is the rearrangement of the balance equations and the substitution of the product $x_c(\mathbf{i})p_{\mathbf{i}}$ with the variable $y_c(\mathbf{i})$. One may interpret $y_c(\mathbf{i})$ as the long run fraction of the time that the system is in state $\mathbf{i}$ and action c is chosen. This operation allows us to redefine Problem (1) as follows:

$$
\begin{aligned}
\underset{\mathbf{Y},\mathbf{p}}{\text{minimize}} \quad & \mathbf{g^T p} \\
\text{subject to} \quad & \sum_c \mathbf{L_c y_c} + \mathbf{Mp} = \mathbf{0} \\
& \mathbf{1^T_{|\Omega|} p} = 1 \\
& \mathbf{p} \geq \mathbf{0} \\
& \mathbf{1^T_{|c|} Y} = \mathbf{p^T} \\
& y_c(\mathbf{i}) \geq 0 \quad \forall c, \mathbf{i}
\end{aligned}
\tag{2}
$$

where matrix $\mathbf{L_c}$ contains all transition rates dependent on policy c, and matrix $\mathbf{M}$ contains all transition rates independent of policy c. For example, as shown in Appendix 1, $\mathbf{M}$ contains the last four terms in the balance equations. Furthermore, we use $\mathbf{y_c}$ to describe an $|\Omega|$-dimensional column vector with elements $y_c(\mathbf{i})$ and $\mathbf{Y}$ to describe a matrix with row vectors $\mathbf{y_c}$. As a consequence, (2) allows us to find the optimal values of $\mathbf{Y}$ and $\mathbf{p}$ with linear programming methods. Afterwards, the optimal decisions $\mathbf{X}^*$ are recovered with the relation $x_c(\mathbf{i}) = y_c(\mathbf{i})/p_{\mathbf{i}}$. Note that in case $0 < y_c(\mathbf{i}) < p_{\mathbf{i}}$ we would obtain a randomized policy which violates the constraint $x_c(\mathbf{i}) \in \{0,1\}$ $\forall c, \mathbf{i}$ of Problem (1). Sleptchenko and Johnson (2015) however, show that, given a linear cost function, $x_c(\mathbf{i}) \in \{0,1\}$ $\forall c, \mathbf{i}$ holds, i.e. we always obtain a deterministic policy with problem (2).

3.4 Stocking policy

To determine the stocking policy, we apply a greedy approach. That is, we determine the optimal long-run average cost $C^*(0)$ given that $S = 0$. Afterwards, we set $S = 1$ and determine $C^*(1)$. In case $C^*(0) < C^*(1)$, $S = 0$ is the optimal base stock level. Otherwise, we continue to increment S until $C^*(S) < C^*(S+1)$. This procedure leads to the optimal base stock level because of the linear structure of the cost function, see e.g. Van Houtum and Kranenburg (2015).

It is also possible to jointly optimize the stocking policy with the sourcing and maintenance policy (cf. Appendix 2). In the joint optimization approach, we set an upper bound on the number of spare parts required, and include the option not to order any part upon failure in a certain state. Despite leading to the same results, numerical experiments show that the join optimization is computationally inferior to the greedy approach. The key reason for this characteristic is that it is difficult to find a tight upper bound on the number of spare parts, in combination with the fact that the computation times grow rapidly with the size of the state space. On the other hand, an advantage of the joint optimization approach is that this model is extendable to allow a dynamic stocking policy (cf. Appendix 3). Given the different failure rates of the AM and CM version, it is likely that the optimal inventory level depends on the mixture of AM parts and CM parts in the installed base. So, the inventory level will be state-dependent. The impact of such a dynamic spare parts inventory policy is a topic for further research.

4 Numerical experiments

To gain insights on the value of dual sourcing with AM and CM methods, we analyze the impact of key input parameters to our model. We keep the parameters of the CM item constant with $\lambda_{CM} = 0.1$, $\mu_{CM} = 1$ and $c_{CM} = 10$. All other parameters are varied as shown in Table 1, which results in 26.460 problem instances. The computation time of an instance amounts to a few seconds, see Appendix 4 for further details on the computation time.

Parameter	**Range**	**Step size**
k	$[2, 10]$	2
λ_{AM}	$[0.05, 0.3]$	0.05
μ_{AM}	$[1, 25]$	4
c_{AM}	$[5, 30]$	5
b	$[20, 500]$	80
h	$[0.15, 0.25]$	0.05

Tab. 1. Parameter settings numerical experiments

In Section 4.1, we study the effect of the different parameter settings on the cost saving potential with dual sourcing compared to the single sourcing alternatives. In Section 4.2, we examine the structure of the optimal policy in case dual sourcing turns out to be the superior sourcing mode. Section 4.3 summarizes the key findings.

4.1 Cost savings with dual sourcing

In general, the cost savings are highest if the resupply lead time is short with AM (and hence the resupply rate high). In extreme cases, we find instances with cost savings of more than 30% compared to the best single sourcing option. The result for these cases is remarkable because the failure rate or unit cost of the AM part is often two to three times higher than for the CM part. Below, we study these relations in more in detail.

Figure 2 shows the effect of the resupply rate on the cost saving potential with dual sourcing. For that purpose we illustrate the cost savings compared to single sourcing with AM, single sourcing with CM and the best single sourcing approach as average over all analyzed instances.

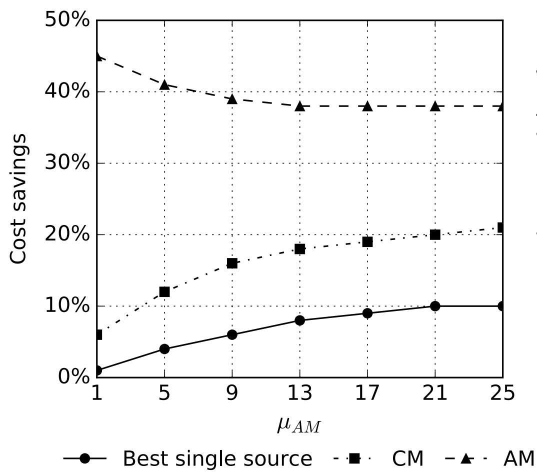

Fig. 2. Dual sourcing compared to single sourcing dependent on μ_{AM} where $\mu_{CM} = 1$

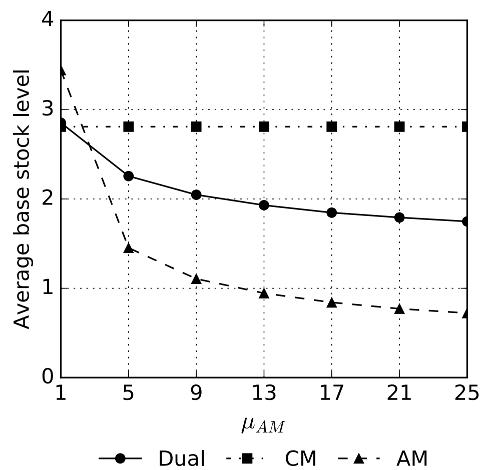

Fig. 3. Average base stock level for different values of μ_{AM} where $\mu_{CM} = 1$

As we observe, with increasing μ_{AM} the cost saving potential of dual sourcing growths to 10% on average compared to the best single sourcing approach. This finding appears reasonable, because a high μ_{AM} causes a reduction of holding and backorder cost as we explain in further detail later. Less intuitive appears the result that the cost savings with dual sourcing compared to sourcing with AM only remain high (>35%) even if $\mu_{AM} = 25$. Instead, one may expect that a very fast resupply rate would diminish the cost savings with dual sourcing. Figure 3 sheds light on this ambiguity and shows the average base stock level for the different supply options. As is illustrated, on average, even in case $\mu_{AM} = 25$ we typically do not print spare parts on demand but keep stock with AM only. Consequently, holding cost prevail while high unit cost and failure rate of AM parts lead to additional purchasing cost compared to the other sourcing options. This finding establishes, opposed to common belief, that printing spare parts on demand may turn out unrealistic for downtime-critical parts despite very short resupply lead times.

On average, the opposite approach, to rely on CM methods only, appears inferior to dual sourcing as well. As shown in Figure 2, in case $\mu_{AM} = 25$ dual sourcing leads to cost savings of more than 20% compared to CM only. This result follows from the necessity of keeping a higher base stock level compared to the dual sourcing option (cf. Figure 3). Also, encountering backorder cost is more likely given that, in the event where stock is nearly depleted, an emergency source by means of AM supply is not available.

Figure 4 illustrates the cost savings with dual sourcing as a function of λ_{AM} averaged over all analyzed instances. Again, we compare to single sourcing with AM, single sourcing with CM and the best single sourcing approach.

Overall, in case the failure rate is high, the incentive to rely on AM methods only is

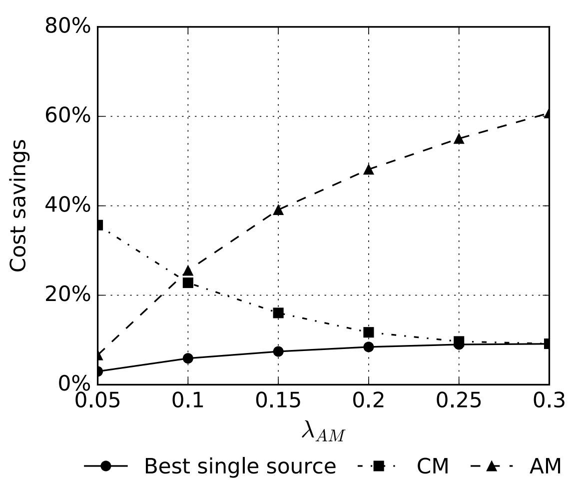

Fig. 4. Dual sourcing compared to single sourcing dependent on λ_{AM} where $\lambda_{CM} = 0.1$

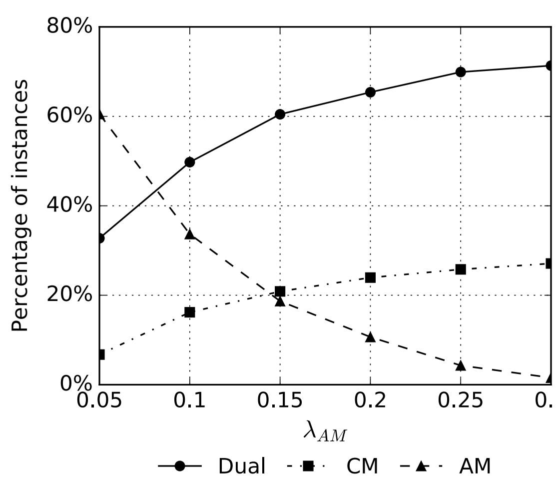

Fig. 5. Percentage of instances in which one sourcing approach dominates for different values of λ_{AM} where $\lambda_{CM} = 0.1$

low. For instance, in case $\lambda_{AM} = 3\lambda_{CM}$, the cost savings with dual sourcing exceed 60% on average. More surprising is the fact that even if λ_{AM} is high, dual sourcing remains valuable compared to single sourcing with CM. As such if $\lambda_{AM} = 0.2$ or 0.3 dual sourcing leads to cost saving of 8-10% on average. Figure 5 further details this finding where we plot the best sourcing approach dependent on λ_{AM}. As we observe, only in case $\lambda_{AM} < \lambda_{CM}$, single sourcing with AM becomes the dominating approach. Otherwise, dual sourcing turns out to be the optimal sourcing strategy. An explanation for this outcome can be derived from Figure 6 where we illustrate the distribution of the cost factors for different λ_{AM}. For comparability purposes, we also show the cost distribution with CM only and set it equal to 100%.

On average, for each value of λ_{AM} backorder and holding costs remain lower than for the

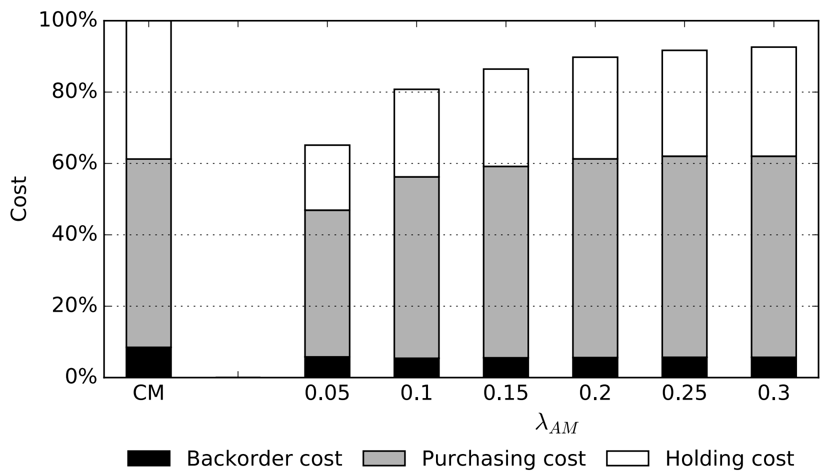

Fig. 6. Cost distribution with dual sourcing dependent on λ_{AM} compared to CM only (100%) where $\lambda_{CM} = 0.1$

CM only sourcing approach. The benefit to use dual sourcing even in case of high λ_{AM} is motivated by the reduction of backorder and holding cost as we observe in Figure 6. If we execute the same analysis for the unit cost c_{AM}, we find a comparable effect (cf. Appendix 5). We conclude that dual sourcing has the potential to substantially extend the operating range of AM methods in the spare parts business.

This result may even give rise to new sourcing concepts. For example, AM parts may function as a so-called temporary fix where locally producible AM spare parts are used to service capital goods at remote locations. Even though it appears likely that locally manufactured spare parts are less reliable, our results indicate that holding and downtime costs may reduce with such a concept. Today, first considerations for these types of applications can be found in defense organizations which experiment with mobile AM production facilities (McLearen, 2015). A small sub-experiment may substantiate the potential further. Therefore we consider the situation of very fast resupply lead times with AM in combination with low reliabilities of AM parts. Accordingly, we set $\mu_{AM} = 25$ and $\lambda_{AM} = 10\lambda_{CM}$. Other parameter values remain unchanged. Even under these conditions dual sourcing leads to cost savings of 6% on average. In the most extreme case, where backorder and holding cost are high ($b = 500$, $h = 0.25$), we even find instances with 28% cost savings compared to the best single sourcing approach. Based on these outcomes, it comes without surprise that we argue that a temporary fix with AM methods demands additional attention in the literature.

Table 2 shows the results for the remaining parameters. Also, for the purpose of concise presentation, we depict the AM usage which we discuss in the next section.

Parameter	**Value**	**CM**	**AM**	**Best**	**AM usage**
	0.15	14%	41%	4%	9%
h	0.2	16%	39%	5%	10%
	0.25	18%	38%	6%	11%
	2	20%	35%	7%	12%
	4	17%	38%	6%	10%
k	6	15%	40%	5%	10%
	8	14%	41%	4%	10%
	10	13%	42%	4%	9%
	20	11%	38%	1%	25%
	100	15%	39%	4%	12%
	180	17%	39%	6%	10%
b	260	17%	40%	6%	9%
	340	17%	39%	7%	9%
	420	17%	39%	7%	9%
	500	17%	39%	7%	9%

Tab. 2. Dual sourcing compared to single sourcing and AM usage where *Best* indicates the best single sourcing approach

In line with the findings so far, the holding cost fraction h has a positive effect on the cost saving potential with dual sourcing. As elaborated before, due to a typically high resupply rate for AM parts, holding cost with dual sourcing are often lower compared to sourcing with CM only. Thus, the higher h, the higher the cost saving potential with dual sourcing. The installed base size k provides insights about the impact of the total demand rate on the cost saving potential with dual sourcing. That is, for each installed part we observe an additional demand stream. As we observe in Table 2, the demand rate has a negative impact on the cost saving potential with dual sourcing. Figure 7 allows an explanation where we show, scaled to 100%, the cost distribution for different k.

In the case of increasing k purchasing cost progressively becomes the dominating cost

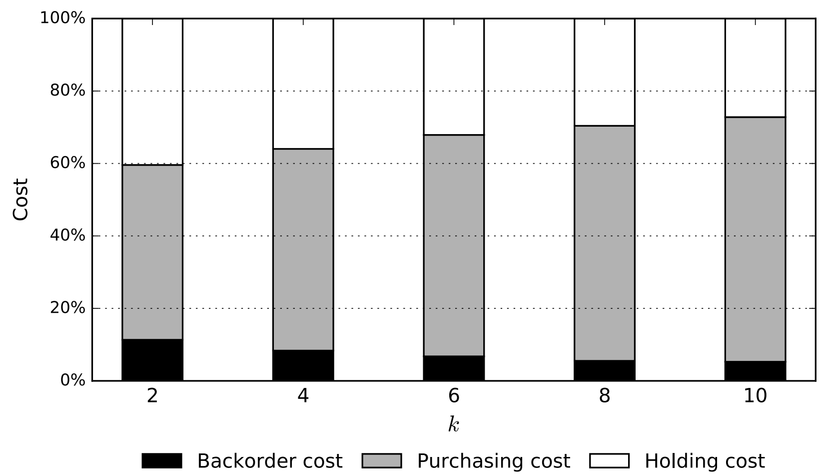

Fig. 7. Cost distribution dependent on the installed base size, scaled to 100%

factor (>65%). Given that an AM part typically has higher unit cost than a CM part, this condition reduces the value of dual sourcing. We conclude that dual sourcing – very much as AM technology in general - is most valuable for low demand rates that are often observed in the spare part business.

Finally, the cost saving potential proves rather independent of the backorder cost b for high values. Only in the event that backorder cost become low, does the benefit of dual sourcing diminish (Table 2). The explanation for this finding follows a comparable reasoning as for a high demand rate. In the case that the backorder cost are low, the incentive to hold stock reduces and thus leads to a reduction of the holding cost. Also, the backorder cost decrease and thus render the purchasing cost the dominating cost factor again. We conclude that dual sourcing is not suitable if the backorder cost are low. This situation, however, typically does not apply to downtime critical spare parts.

4.2 Structure of optimal policy

In this section, we analyze the structure of the optimal policy. Therefore, we restrict ourselves to instances where dual sourcing is the dominant sourcing approach. Figure 8 shows the usage of the AM source and the number of instances where dual sourcing is the dominating approach for different values of μ_{AM}.

As shown, in case $\mu_{AM} = \mu_{CM}$, the inventory level does not influence the sourcing decision. Consequently, the median base stock level and critical level are the same. In case $\mu_{AM} > \mu_{CM}$ however, the AM source functions as an emergency source only. That is, for most instances we order the first time from the AM source if the inventory level is 0. This finding corresponds to the dual sourcing literature where the fast and expensive supply mode is typically used as an emergency source.

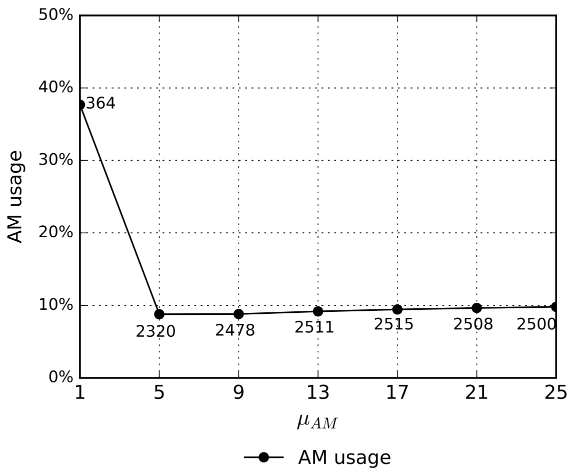

Fig. 8. Usage of the AM source dependent on μ_{AM} where $\mu_{CM} = 1$ and number of instances plotted next to the data points

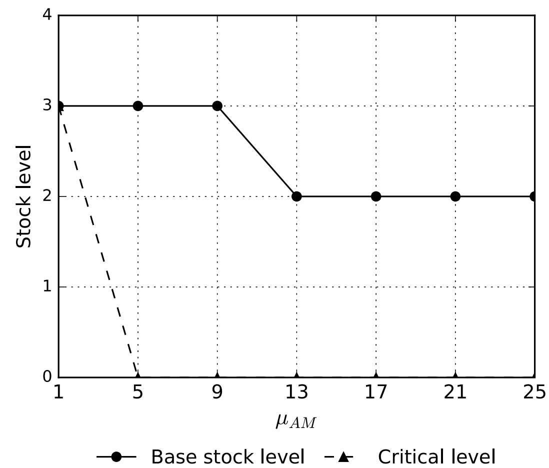

Fig. 9. Median of base stock and critical level dependent on μ_{AM} where $\mu_{CM} = 1$

Next, we analyze the effect of c_{AM} on the AM source usage. The results are shown for instances where dual sourcing is the dominating sourcing mode in Figure 10.

As one may expect in case of $c_{AM} < c_{CM}$ the AM source is dominating. Otherwise the

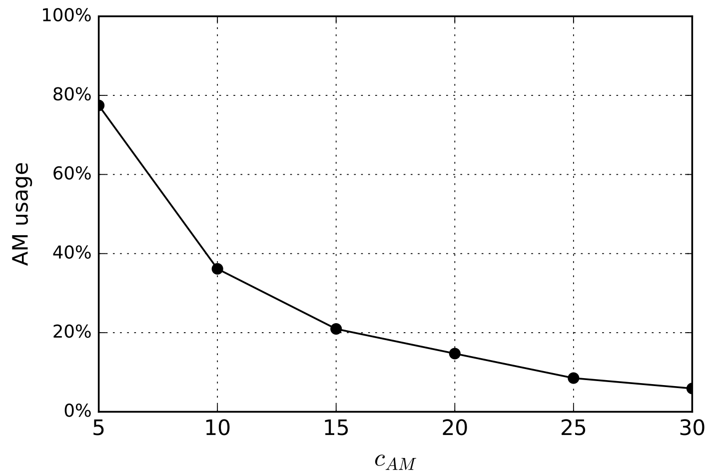

Fig. 10. Usage of the AM source dependent on c_{AM} where $c_{CM} = 10$

CM source is used primarily. It is remarkable though, that even in case $c_{AM} = 3c_{CM}$ the AM source is still used in about 6% of cases. We find comparable results for λ_{AM} where even if $\lambda_{AM} = 3\lambda_{CM}$ the AM usage exceeds 5% (cf. Appendix 6). These findings confirm the observation that the value of the AM source primarily stems from the ability to provide emergency supply in case the inventory is (nearly) depleted.

The last column of Table 2 (Section 4.1) shows the results for the remaining parameters. Overall, the holding cost rate appears to have a marginal influence on the AM usage. The same holds for the installed base size k (i.e. demand rate). For both parameters though, it turns out that the AM source is typically used as emergency source and typically used in less than 10% of cases. We find slightly different result for the backorder cost b. If the backorder cost are low ($b < 100$) we use the AM source in about 25% cases on average. Otherwise, the AM usage approximately drops to 10% on average. An explanation for this difference follows the same line of reasoning as for μ_{AM}: in the few instances where dual sourcing is the dominating sourcing mode with $b < 100$, the motivation to use AM stems from the potential to reduce purchasing cost, i.e. either $c_{AM} < c_{CM}$ or $\lambda_{AM} < \lambda_{CM}$. Otherwise, if $b \geq 100$, the AM source is used primarily as an emergency source and thus for the reduction of holding and/or backorder cost. We conclude, in case dual sourcing is the superior sourcing mode, the AM source is typically used as an emergency source.

4.3 Summary of key findings

The numerical experiments motivate the following conclusions:

1. The concept to solely print spare parts on demand does not appear suitable for downtime critical spare parts. Instead, keeping stock remains necessary to reduce the risk of expensive downtime. As a result, dual sourcing largely outperforms single sourcing with AM methods. On average we find cost savings of more than 35% even if $\mu_{AM} = 25$ (Figure 2 and Figure 3).

2. In the case of an undesirable failure behavior ($\lambda_{AM} > 3\lambda_{CM}$) or high unit cost ($c_{AM} > 3c_{CM}$) of the AM part, dual sourcing allows one to maintain the benefits of AM while mitigating drawbacks. On average, even under these unfavorable conditions we find cost savings of about 10% compared to conventional single sourcing (Figure 4-6, Figure 10 and Figure A4-A7).

3. In line with dual sourcing literature, the AM source typically functions as an emergency source if $\mu_{AM} > \mu_{CM}$ and $b \geq 100$ (Figure 8 and 9, Table 2).

4. Motivated by the reduction of stock with dual sourcing, we find that the higher the holding cost fraction h, the higher the cost saving potential with dual sourcing (Table 2).

5. In the case of a high demand rate dual sourcing appears less valuable. Under these conditions purchasing cost become the dominating cost factor which diminishes the value of a fast emergency source. That is backorder and holding cost reductions do not justify the additional purchasing cost (Table 2).

6. If the backorder cost are low, the value of dual sourcing is limited. This finding relates to the same observation as in the previous point. If the backorder cost are low, purchasing cost become the dominating cost factor and thus reduce the value of dual sourcing (Table 2).

5 Case study

To gain further insights into the practical implications of combining AM and CM sourcing modes, we conduct a case study at a service provider in the aerospace industry. More explicitly, we consider a hinge bracket that is used for connecting the rudder and the aircraft. According to the Federal Aviation Administration, the hinge is categorized as a Class 2 product. That is, a failure may jeopardize the safety of an aircraft and thus is considered critical. To avoid devastating consequences, aircraft manufacturers typically work with redundancies. Nevertheless, a failure of a hinge has to be corrected upon discovery. The most common failure modes are fractures of the hinge that are caused by fatigue, tensile stress or corrosion.

Figure 11 illustrates the CM and AM design of the hinge. While the CM hinge is an aluminum machined part, the AM hinge is built from titanium powder (Ti-6AL-4V) with Selective Laser Melting (SLM). Next to account for the different material properties, the AM design is topology optimized to reduce the weight of the hinge. Overall, despite titanium being heavier than aluminum, the topology optimization leads to a weight reduction of about 25%. Cost reductions resulting from fuel savings though were not sufficient to motivate the production of the hinge with AM methods only. In particular, this is caused by the higher unit cost of the AM hinge compared to the CM hinge. Missing standardization of the certification process for flight critical AM parts further hinders the production of the hinge with AM methods. Overall, the company argued that the AM hinge unit costs have to decrease significantly before the production with AM methods becomes economically feasible.

We make following simplifications for our analysis: first, we do not take fuel savings into account, but only focus on the service cost. Second, we assume that each hinge is exposed to the same load profile and fails according to a Poisson process. In the field, this may not hold true, and thus we may observe a higher demand variability. Table 3 shows the input parameter for the hinge case that, if not mentioned otherwise below, were obtained from company records.

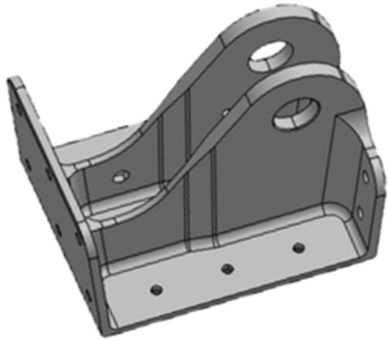

(a)

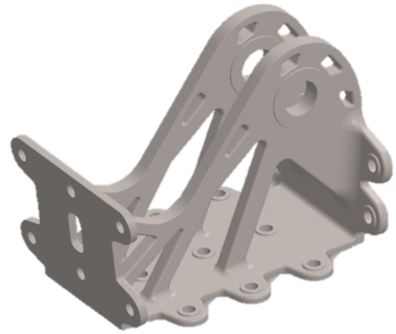

(b)

Fig. 11. CM (a) and AM hinge (b) design

Parameter	**Value**	**Unit**
c_{AM}	1197	Euro
c_{CM}	480	Euro
μ_{AM}	1	Month
μ_{CM}	0.032	Month
λ_{AM}	0.00225	Failure/month
λ_{CM}	0.003	Failure/month
k	382	Units
b	32500	Euro
h	0.017	Euro/Euro/month

Tab. 3. Model input data hinge case

The failure rate λ_{AM} is based on following insights: mechanical tests revealed that the AM hinge exhibits superior static strength compared to the CM hinge. Also, given that the AM hinge is produced with titanium powder rather than with aluminum, failures caused by corrosion can be ruled out. On the other hand, the surface roughness and porosity expected from the SLM production process may concentrate more tensions in the hinge. Hence, it is likely that the AM hinge fails more often under cyclic load caused by fatigue. Based on these observations, AM experts estimate $\lambda_{AM} = 0.75\lambda_{CM}$ in the best case. We use this estimate as base case, but later also consider the case where $\lambda_{AM} = \lambda_{CM}$.

The backorder cost b follow from the criticality of the hinge. According to company representatives, it is likely that a stock-out leads to additional downtime of the aircraft. Here, we assume that a stock-out leads to downtime of one additional day at maximum. Hence, depending on the aircraft type, backorder cost b vary between 15000 to 50000 euros. For our analysis, we use $b = 32500$ as base case but note that our results are not very sensitive to changes of b in that range: the long-run average service cost per month differ by less than 1% if we compare $b = 15000$ and $b = 50000$.

Currently, the case company services an installed base k of 382 hinges each with an expected mean time between failure of 27.78 years. To obtain computationally tractable data, we need to consolidate demand streams. For that purpose, we consider a smaller k instead and multiply the failure rates λ_{AM} and λ_{CM} with the fraction $382/k$ in order to keep the total demand rate unchanged. Given that, as soon as k is sufficiently large, the demand variability becomes nearly independent of k, this transformation leads to proportionally the same results. We illustrate this effect by plotting the long-run average costs C per month with dual sourcing for different values of k in Figure 12. After the results are computed, we simply multiply the resulting cost factors with $382/k$ to obtain the cost for the case $k = 382$. In the reminder, we use $k = 50$ for our analysis.

The service costs of the hinge are dominated by the purchasing cost and thus reduces

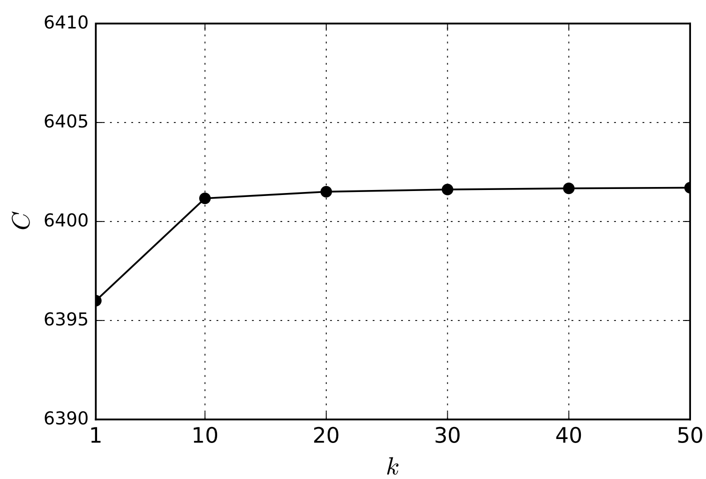

Fig. 12. Effect of consolidating demand streams on the long-run average costs C per month with dual sourcing where $\lambda_{AM} = 382/k\lambda_{AM}$ and $\lambda_{CM} = 382/k\lambda_{CM}$

the value of dual sourcing (cf. Section 4.1). Nevertheless, we find in both cases that dual sourcing is the cheapest sourcing approach. The cost savings with dual sourcing are a consequence of the option to decrease the base stock level by one unit while the backorder cost nearly remain unchanged. Furthermore, as expected, high unit cost of AM ($\sim 2.5c_{CM}$) diminish the value of souring with AM methods only and thus confirm the results of the case company. Overall, the hinge case gives further evidence that dual sourcing is more valuable in case of high holding or backorder cost. Nevertheless, the results confirm that dual sourcing offers the potential to exploit the short resupply lead time of AM methods for a higher part variety.

Figure 13 shows the long-run average service cost per month of the three sourcing options for the cases $\lambda_{AM} = 0.75\lambda_{CM}$ and $\lambda_{AM} = \lambda_{CM}$.

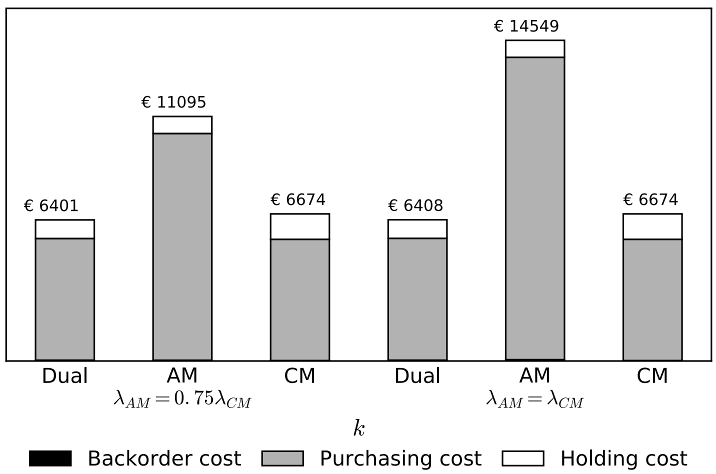

Fig. 13. Service cost for the three sourcing options with $\lambda_{AM} = 0.75\lambda_{CM}$ and $\lambda_{AM} = \lambda_{CM}$

Finally, and this represents a key learning for us from this case study, we emphasize that today's attention for AM technologies is mainly motivated by the prospect of design improvements. While this certainly has its value, it dilutes the value of AM technologies in the low-volume, downtime-critical spare part business. In our opinion this has two reasons: first, design improvements often scale with quantity. For example, if topology optimization leads to weight reductions of a few grams only, the effort is typically justified by a large installed base size. Accordingly, AM demonstrators frequently exhibit "high" demand rates (say, >20/year) or low backorder cost. In both cases, as discussed in Section 4.1, savings of service cost become a less dominant factor. Second, the prospect of design improvements and technological complexity of AM methods typically puts engineers in the lead to identify parts worth for the production with AM methods. Unfortunately, as a result, logistic opportunities may be undervalued/overlooked. Instead, logisticians should become more involved in this task. In conclusion, we believe that it is important to carry out more case-based research in the low-volume, downtime-critical spare part business to exemplify the value of AM methods in this field to practitioners.

6 Conclusion

Demand fulfillment with a mix of AM and CM production methods has not been sufficiently discussed in the literature yet. In this paper, we have addressed this gap by studying a dual sourcing concept where AM and CM methods are used in parallel to fulfill spare part demand. A key aspect of our model is that we account for the different failure behavior of parts sourced from AM and CM methods. In this paper, we assume a base stock policy. State-dependent demand rates, however, may expose a dynamic inventory policy optimal. Discussed extensions of our model may facilitate the analysis of the value of dynamic inventory policy in future research.

Overall, dual sourcing turns out to reinforce the value of AM methods in the spare part business. In particular, this holds true if backorder cost is high, demand rates are low or holding cost are high. Moreover, our work clarifies that the concept to print spare parts on demand is not necessarily suitable for downtime-critical parts. Instead, stock remains necessary even if the resupply lead time is very short. In our opinion, the most remarkable finding, however, is that dual sourcing offers an approach to profit from the fast resupply lead time of AM technologies even if the AM part unit cost or failure rate are high compared to the CM part. Consequently, dual sourcing may extend the operating range of AM methods in the spare parts business significantly. In the light of this finding, new sourcing concepts are likely. For instance, our results indicate that supplementing CM supply with less reliable, but locally producible AM parts may reduce operating cost of capital goods at remote locations considerably. It appears valuable to investigate the value of this approach more in detail.

Using a case study in the aerospace industry, we were able to obtain further evidence that dual sourcing primarily benefits sourcing of low-volume, downtime-critical spare parts. The case study also revealed questions that may stimulate further research. First, how does the demand variability influence the trade-off between dual and single sourcing concepts with AM? Second, if we consider operation and service cost jointly, does the AM source remain the emergency source and when do we refrain from a dual sourcing concept? Finally, observations made during the case study exposed that more case-based research is required to demonstrate the value of AM methods in the low-volume, downtime-critical spare part business to practitioners.

Acknowledgments

This research is part of the project “Sustainability Impact of New Technology on After sales Service supply chains (SINTAS)” and has been sponsored by the Netherlands Organization for Scientific Research under project number 438-13-207.

References

Airbus (2016). *Innovative 3D printing solutions are “taking shape” within Airbus.* URL: http://www.airbus.com/newsevents/news-events-single/detail/innovative-3d-printing-solutions-are-taking-shape-within-airbus/ (visited on 03/02/2017).

Alfredsson, P. and J. Verrijdt (1999). “Modeling Emergency Supply Flexibility in a Two-Echelon Inventory System”. In: *Management Science* 45(10), pp. 1416–1431. DOI: 10.1287/mnsc.45.10.1416.

Allon, G. and J. A. V. Mieghem (2010). “Global Dual Sourcing: Tailored Base-Surge Allocation to Near- and Offshore Production”. In: *Management Science* 56(1), pp. 110–124. DOI: 10.1287/mnsc.1090.1099.

Alvarez, E., M. van der Heijden, and W. Zijm (2013). “The selective use of emergency shipments for service-contract differentiation”. In: *International Journal of Production Economics* 143(2), pp. 518–526. DOI: 10.1016/j.ijpe.2012.02.01.

Barankin, E. W. (1961). “A delivery-lag inventory model with an emergency provision (the single-period case)”. In: *Naval Research Logistics Quarterly* 8(3), pp. 285–311. DOI: 10.1002/nav.3800080310.

Book, T. A. and M. D. Sangid (2016). "Evaluation of Select Surface Processing Techniques for In Situ Application During the Additive Manufacturing Build Process". In: *JOM* 68(7), pp. 1780–1792. DOI: 10.1007/s11837-016-1897-y.

Frazier, W. E. (2014). "Metal Additive Manufacturing: A Review". In: *Journal of Materials Engineering and Performance* 23(6), pp. 1917–1928. DOI: 10.1007/s11665-014-0958-z.

Heyman, D. P. and M. J. Sobel (1984). *Stochastic Models in Operations Research, Vol. I: Stochastic Processes and Operating Characteristics*. McGraw-Hill Inc., New York.

Holweg, M. (2015). *The limits of 3-D printing*. URL: https://hbr.org/2015/06/the-limits-of-3d-printing (visited on 03/02/2017).

Khajavi, S. H., J. Partanen, and J. Holmström (2014). "Additive manufacturing in the spare parts supply chain". In: *Computers in Industry* 65(1), pp. 50–63. DOI: 10.1016/j.compind.2013.07.008.

Khajavi, S. H. et al. (2015). "Risk reduction in new product launch: A hybrid approach combining direct digital and tool-based manufacturing". In: *Computers in Industry* 74, pp. 29–42. DOI: https://doi.org/10.1016/j.compind.2015.08.008.

Knofius, N., M. Van der Heijden, and W. Zijm (2017). *Consolidating spare parts for asset maintenance with additive manufacturing*. Working Paper. URL: http://onderzoeksschool-beta.nl/wp-content/uploads/wp_527.pdf (visited on 05/02/2017).

Liu, P. et al. (2014). "The impact of additive manufacturing in the aircraft spare parts supply chain: supply chain operation reference (scor) model based analysis". In: *Production Planning & Control* 25(13), pp. 1169–1181. DOI: 10.1080/09537287.2013.808835.

McLearen, L. (2015). Calhoun, the NPS Institutional Archive. URL: https://calhoun.nps.edu/handle/10945/45903.

Minner, S. (2003). "Multiple-supplier inventory models in supply chain management: A review". In: *International Journal of Production Economics* 81, pp. 265–279. DOI: http://dx.doi.org/10.1016/S0925-5273(02)00288-8.

Moinzadeh, K. and C. P. Schmidt (1991). "An (S - 1, S) Inventory System with Emergency Orders". In: *Operations Research* 39(2), pp. 308–321. DOI: 10.1287/opre.39.2.308.

Scheller-Wolf, A., S. Veeraraghavan, and G. J. Van Houtum (2007). *Effective Dual Sourcing with a Single Index Policy*. Working Paper, Carnegie-Mellon University, Pittsburgh.

Scott, C. (2017). *Norsk Titanium Produces First FAA-Approved 3D Printed Structural Titanium Commercial Airplane Components for Boeing Dreamliner*. URL: https://3dprint.com/170652/norsk-titanium-boeing-aircraft/ (visited on 05/02/2017).

Sherbrooke, C. C. (2004). *Optimal Inventory Modeling of Systems*. Springer US. DOI: 10.1007/b109856.

Sleptchenko, A. and M. E. Johnson (2015). "Maintaining Secure and Reliable Distributed Control Systems". In: *INFORMS Journal on Computing* 27(1), pp. 103–117. DOI: 10.1287/ijoc.2014.0613.

Song, J.-S. J. and Y. Zhang (2016). *Stock or Print? Impact of 3D Printing on Spare Parts Logistics*. Working Paper. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2884459 (visited on 05/02/2017).

Song, J.-S. and P. Zipkin (2009). "Inventories with Multiple Supply Sources and Networks of Queues with Overflow Bypasses". In: *Management Science* 55(3), pp. 362–372. DOI: 10.1287/mnsc.1080.0941.

Song, J.-S. et al. (2017). "Optimal Policies for a Dual-Sourcing Inventory Problem with Endogenous Stochastic Lead Times". In: *Operations Research* 65(2), pp. 379–395. DOI: 10.1287/opre.2016.1557.

Van Houtum, G.-J. and B. Kranenburg (2015). *Spare parts inventory control under system availability constraints*. Springer US. DOI: 10.1007/978-1-4899-7609-3.

Veeraraghavan, S. and A. Scheller-Wolf (2008). "Now or Later: A Simple Policy for Effective Dual Sourcing in Capacitated Systems". In: *Operations Research* 56(4), pp. 850–864. DOI: 10.1287/opre.1080.0552.

Westerweel, B., R. Basten, and G. J. Van Houtum (2016). *Traditional or Additive Manufacturing? Assessing component design options through lifecycle cost analysis*. Working Paper. URL: http://onderzoeksschool-beta.nl/wp-content/uploads/wp_519.pdf (visited on 05/02/2017).

Whittemore, A. S. and S. C. Saunders (1977). "Optimal Inventory Under Stochastic Demand with Two Supply Options". In: *SIAM Journal on Applied Mathematics* 32(2), pp. 293–305. DOI: 10.1137/0132023.

Wits, W. W., J. R. R. García, and J. M. J. Becker (2016). "How Additive Manufacturing Enables more Sustainable End-user Maintenance, Repair and Overhaul (MRO) Strategies". In: *Procedia CIRP* 40, pp. 693–698. DOI: 10.1016/j.procir.2016.01.156.

Zhou, S. X. and C. Yang (2016). "Continuous-Review (R, nQ) Policies for Inventory Systems with Dual Delivery Modes". In: *Operations Research* 64(6), pp. 1302–1319. DOI: 10.1287/opre.2016.1538.

Appendix

A1 Balance equations basic model

Subsequently, we present the balance equations, where $p_\bullet$ refers to the state probability under consideration. For example, in case $x_2(nC_\mathbf{i}, nA_\mathbf{i}+1, rC_\mathbf{i}, rA_\mathbf{i}-1, sC_\mathbf{i}, sA_\mathbf{i})$ then $p_\bullet = p_{(nC_\mathbf{i}, nA_\mathbf{i}+1, rC_\mathbf{i}, rA_\mathbf{i}-1, sC_\mathbf{i}, sA_\mathbf{i})}$. Furthermore, if the maintenance decision is indifferent (i.e. no stock available) we use $x_2(\mathbf{i})$ and $x_4(\mathbf{i})$ as default.

$$
\begin{aligned}
(\lambda_{CM} nC_{\mathbf{i}} + \lambda_{AM} nA_{\mathbf{i}} + \mu_{CM} rC_{\mathbf{i}} + \mu_{AM} rA_{\mathbf{i}}) p_{\mathbf{i}} = \\
\lambda_{AM}(nA_{\mathbf{i}}+1) x_2(nC_{\mathbf{i}}, nA_{\mathbf{i}}+1, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{AM} nA_{\mathbf{i}} x_1(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}, sA_{\mathbf{i}}+1) p_{\bullet} + \\
\lambda_{AM}(nA_{\mathbf{i}}+1) x_2(nC_{\mathbf{i}}-1, nA_{\mathbf{i}}+1, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}+1, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{AM}(nA_{\mathbf{i}}+1) x_4(nC_{\mathbf{i}}, nA_{\mathbf{i}}+1, rC_{\mathbf{i}}-1, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{AM} nA_{\mathbf{i}} x_3(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}-1, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}}+1) p_{\bullet} + \\
\lambda_{AM}(nA_{\mathbf{i}}+1) x_4(nC_{\mathbf{i}}-1, nA_{\mathbf{i}}+1, rC_{\mathbf{i}}-1, rA_{\mathbf{i}}, sC_{\mathbf{i}}+1, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{CM}(nC_{\mathbf{i}}+1) x_2(nC_{\mathbf{i}}+1, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{CM}(nC_{\mathbf{i}}+1) x_1(nC_{\mathbf{i}}+1, nA_{\mathbf{i}}-1, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}, sA_{\mathbf{i}}+1) p_{\bullet} + \\
\lambda_{CM} nC_{\mathbf{i}} x_2(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}-1, sC_{\mathbf{i}}+1, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{CM}(nC_{\mathbf{i}}+1) x_4(nC_{\mathbf{i}}+1, nA_{\mathbf{i}}, rC_{\mathbf{i}}-1, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}}) p_{\bullet} + \\
\lambda_{CM}(nC_{\mathbf{i}}+1) x_3(nC_{\mathbf{i}}+1, nA_{\mathbf{i}}-1, rC_{\mathbf{i}}-1, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}}+1) p_{\bullet} + \\
\lambda_{CM} nC_{\mathbf{i}} x_4(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}-1, rA_{\mathbf{i}}, sC_{\mathbf{i}}+1, sA_{\mathbf{i}}) p_{\bullet} + \\
\mu_{AM}(rA_{\mathbf{i}}+1) p_{(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}+1, sC_{\mathbf{i}}, sA_{\mathbf{i}}-1)} + \\
\mu_{AM}(rA_{\mathbf{i}}+1) p_{(nC_{\mathbf{i}}-1, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}+1, sC_{\mathbf{i}}, sA_{\mathbf{i}})} + \\
\mu_{CM}(rC_{\mathbf{i}}+1) p_{(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}+1, rA_{\mathbf{i}}, sC_{\mathbf{i}}-1, sA_{\mathbf{i}})} + \\
\mu_{CM}(rC_{\mathbf{i}}+1) p_{(nC_{\mathbf{i}}-1, nA_{\mathbf{i}}, rC_{\mathbf{i}}+1, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})}
\end{aligned}
$$

A2 Joint optimization

To jointly optimize the base stock level, maintenance policy and sourcing policy, we increase the decision space of the sourcing decision and add the possibility to order nothing upon failure of an item. This extension enables the transition to states where N decreases.

Accordingly, we obtain the following updated decision space:

$$
\begin{aligned}
\Omega = \{(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}}) : \\
LB \geq N \leq UB \\
nC_{\mathbf{i}} + nA_{\mathbf{i}} + rC_{\mathbf{i}} + rA_{\mathbf{i}} + sC_{\mathbf{i}} + sA_{\mathbf{i}} = N \\
sC_{\mathbf{i}} + sA_{\mathbf{i}} = max\{N - k - rC_{\mathbf{i}} - rA_{\mathbf{i}}; 0\} \\
nC_{\mathbf{i}} + nA_{\mathbf{i}} \leq k \\
nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}} \geq 0\}
\end{aligned}
$$

where LB and UB describe the lower and upper bound on N respectively. We set $LB = k$ which permits the extreme case to manufacture parts on demand only. In case of the UB, we compute the optimal base stock level $\hat{S}$ of a single sourcing model with $\lambda = max\{\lambda_{CM}; \lambda_{AM}\}$, $\mu = min\{\mu_{CM}; \mu_{AM}\}$ and $c = min\{c_{CM}; c_{AM}\}$. Next, we set $UB = k+\hat{S}$. As this is a worst case scenario, we always find an upper bound, even though it may not be very tight. An alternative is to use as a heuristic upper bound which is equal to the base stock level of the better performing single sourcing option. Unfortunately, our numerical experiments reveal that this bound is not always sufficient. In this case we set $S = UB - k$ and follow the greed heuristic as explained in Section 3.4.

To include the decision option to order nothing we increase the action space of the sourcing decision . Accordingly, we add two components to vector $\mathbf{x}(\mathbf{i})$:

- $c = 5$: take AM version from stock (if possible) and order nothing.
- $c = 6$: take CM version from stock (if possible) and order nothing.

The increase of the action space leads to six additional transitions. These are illustrated for state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ in Figure A1. Despite these changes, the optimization procedure remains the same.

A3 Dynamic inventory policy

Given that the part is available with two versions of different reliability, the expected number of failures is state dependent. Same holds for the expected number of arrivals given the

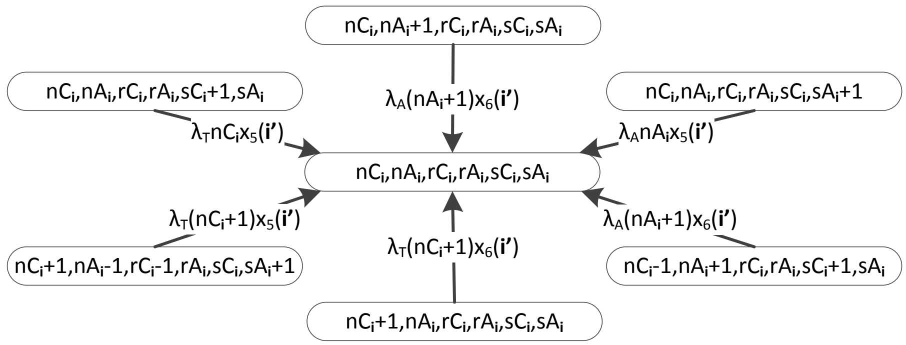

Fig. A1. Additional transitions for state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ for joint optimization

difference in resupply rate. As a consequence, the stocking policy might depend on the state. For example, consider the situation where the installed base is mainly equipped with AM items. In this case, a higher failure frequency is likely, and thus we may want to increase the base stock level to avoid backorder cost. Then, if the number of CM components increases, we may find that it is cost efficient to reduce the base stock level to reduce holding costs. Given that we already added the possibility to decrease N (cf. Appendix 2), we only need to include transitions that allow increasing N to facilitate a dynamic base stock level. Again, we realize this extension by increasing the action space to allow ordering more than one part. Given that typically $\lambda_{CM} < \lambda_{AM}$, the possibility to order more than one part is most valuable if a CM parts fails. Also, it is not reasonable to order more than 2 parts because this would imply that it would have been useful to already order 2 parts at a previous failure of a CM part. Accordingly, we add six components to vector $\mathbf{x}(\mathbf{i})$:

- $c = 7$: take TM part from stock (if possible) and order AM and TM part.
- $c = 8$: take TM part from stock (if possible) and order 2 AM parts.
- $c = 9$: take AM part from stock (if possible) and order AM and TM part.
- $c = 10$:take AM part from stock (if possible) and order 2 AM parts.
- $c = 11$: take TM part from stock (if possible) and order 2 TM parts.
- $c = 12$: take AM part from stock (if possible) and order 2 TM parts.

We can incorporate these actions by regarding nine additional transitions. These are illustrated for state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ in Figure A2. Note that in case the "take AM or TM part from stock" decision is indifferent (i.e. no stock available) we use $x_7(\mathbf{i})$, $x_8(\mathbf{i})$ and $x_{11}(\mathbf{i})$ as default. Despite these changes, the solution procedure remains the same.

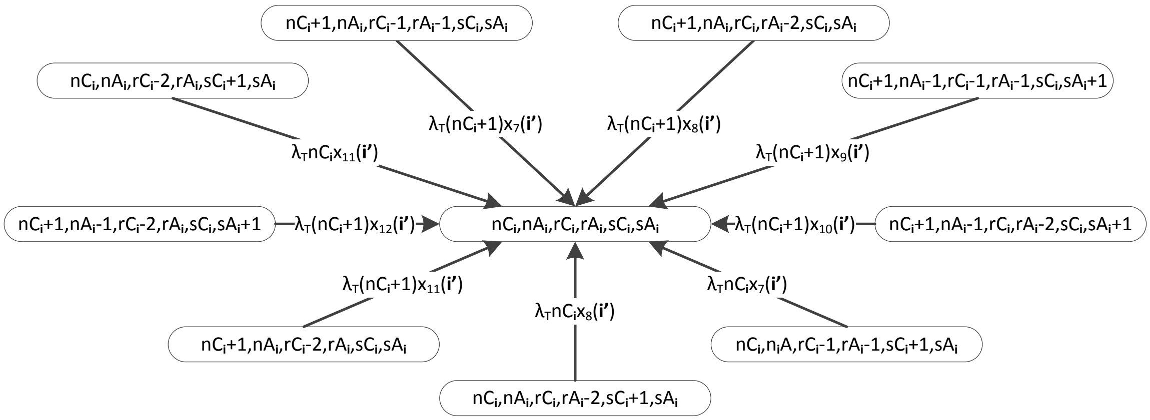

Fig. A2. Additional transitions for state $(nC_{\mathbf{i}}, nA_{\mathbf{i}}, rC_{\mathbf{i}}, rA_{\mathbf{i}}, sC_{\mathbf{i}}, sA_{\mathbf{i}})$ for dynamic inventory policy

A4 Average run time with greedy optimization

In Table A.1 we show the test bed we used to determine the runtime.

Parameter	Range/Values	Step size
k	$[2, 30]$	2
c_{AM}	$5, 30$	-
μ_{AM}	$1, 25$	-
λ_{AM}	$0.05, 0.3$	-
b	$20, 500$	-
h	$0.15, 0.20$	-

Tab. A.1. Test bed for the runtime analysis

Figure A3 shows the runtime of the greedy optimization approach as a function of the installed base size k.

A5 Cost savings as function of c_{AM}

Figure A4-A6 follow similar explanations as for Figure 4-6, except that this time they show the data as function of c_{AM}.

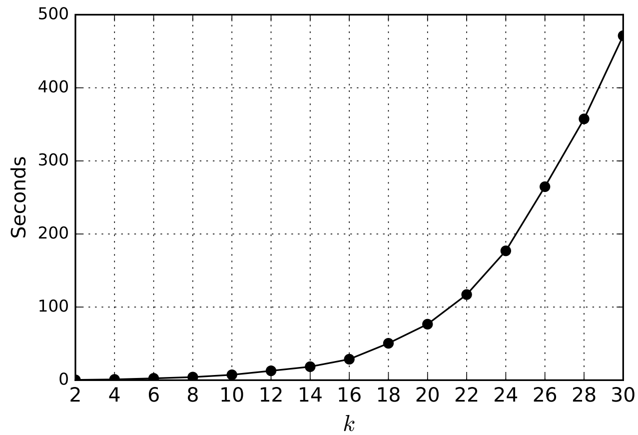

Fig. A3. Run time with greedy optimization for different k

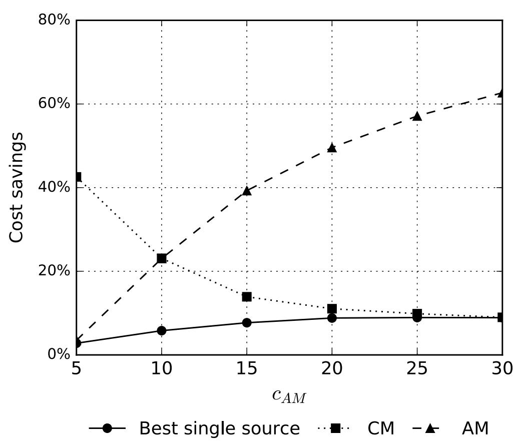

Fig. A4. Dual sourcing compared to single sourcing dependent on c_{AM} where $c_{CM} = 10$

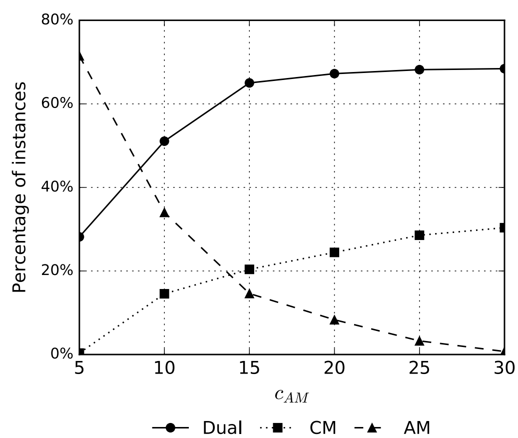

Fig. A5. Percentage of instances in which one sourcing approach dominates for different values of c_{AM} where $c_{CM} = 10$

A6 Usage of the AM source as function of λ_{AM}

Figure A7 follows similar explanations as Figure 10, except that this time they show the data as function of λ_{AM}.

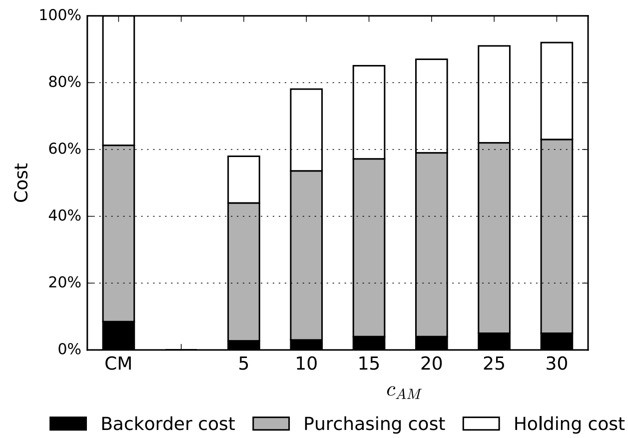

Fig. A6. Cost distribution dependent on c_{AM} compared to CM only (100%) with $c_{CM} = 10$

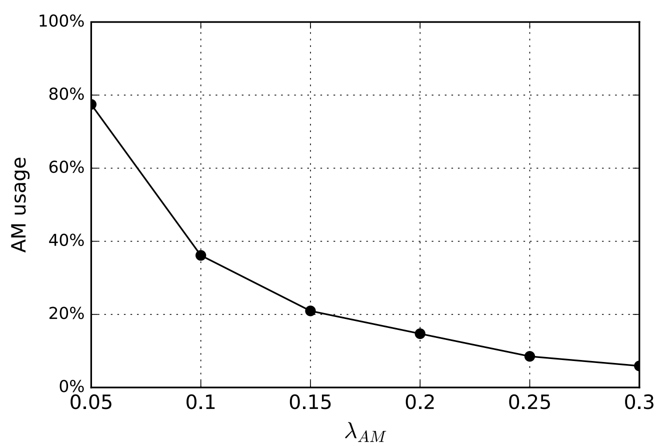

Fig. A7. Usage of the AM source dependent on λ_{AM} where $\lambda_{CM} = 0.1$