



Heriot-Watt University
Research Gateway

Scheduling the repair of aircraft components - a case study

Citation for published version:

Avelino, C, Bourne, D, Ferreira, F, Rasteiro, D & Santos, J 2016, 'Scheduling the repair of aircraft components - a case study', *Mathematics-in-Industry Case Studies*, vol. 7, 5.
<https://doi.org/10.1186/s40929-016-0007-2>

Digital Object Identifier (DOI):

[10.1186/s40929-016-0007-2](https://doi.org/10.1186/s40929-016-0007-2)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Mathematics-in-Industry Case Studies

Publisher Rights Statement:

© 2016 The Author(s). Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

RESEARCH

Open Access



Scheduling the repair of aircraft components - a case study

Catarina Avelino^{1*}, David Bourne², Fátima Ferreira¹, Deolinda Rasteiro³ and Jorge Santos⁴

*Correspondence: cavelino@utad.pt

¹CMAT-UTAD, CEMAT-IST-UL,
Universidade de Trás-os-Montes e
Alto Douro, Vila Real, Portugal
Full list of author information is
available at the end of the article

Abstract

In aircraft components maintenance shops, components are distributed amongst repair groups and their respective technicians based on the type of repair, on the technicians skills and workload, and on the customer required dates. This distribution planning is typically done in an empirical manner based on the group leader's past experience. Such a procedure does not provide any performance guarantees, leading frequently to undesirable delays on the delivery of the aircraft components. Among others, a fundamental challenge faced by the group leaders is to decide how to distribute the components that arrive without customer required dates.

This paper addresses the problems of prioritizing the randomly arriving of aircraft components (with or without pre-assigned customer required dates) and of optimally distributing them amongst the technicians of the repair groups.

We proposed a formula for prioritizing the list of repairs, pointing out the importance of selecting good estimators for the interarrival times between repair requests, the turn-around-times and the man hours for repair.

In addition, a model for the assignment and scheduling problem is designed and a preliminary algorithm along with a numerical illustration is presented.

Keywords: Production planning, Assignment, Scheduling, Optimisation

Introduction

The motivation for this paper is a case study proposed by TAP Maintenance & Engineering, a Portuguese aircraft maintenance, repair and overhaul services provider.

The TAP Maintenance & Engineering aircraft components maintenance shops are responsible for the maintenance and repair of aircraft components of fleets. They perform maintenance on thousands of different aircraft accessories. Annually thousands of components are repaired, tested or modified, such as navigation instruments, displays, valves, heat exchangers, generators, oxygen masks and life vests. Also, for each unit there are distinct types of maintenance, such as overhaul, repair, inspection and testing.

The technicians are divided into groups according to their specialities (hydraulic, mechanic, pneumatic, instruments, radio, testing, cleaning) that receive, at any time, units to be worked on. Within each group, some units can be repaired by any technician, whereas others can only be repaired by few specific skilled technicians. Typically a maintenance shop receives several new components to repair per day, which are divided amongst the groups according to the group specialization. A percentage of the total

units to be maintained arrive with customer required dates (CRDs), i.e., deadlines, pre-established by the planning department. The other units in process do not have a CRD externally provided.

Within each group, the respective team leader has to distribute the units to the technicians based on their skills and workload. The work distribution needs to ensure that the units are processed on time, thus guaranteeing that the maintenance shop meets its objectives. The decision of how to assign the incoming components to the technicians is often based on the experience and intuition of the team leader of each group rather than according to some assignment and scheduling algorithm.

The first challenge we address is to develop a tool to estimate a CRD for the units that arrive without one, based on past data of TAP Maintenance & Engineering. In “Methods” section we present an estimation formula for the CRDs that can be used to prioritize the list of repairs, i.e., determine which units should be tackled first. This formula depends on several parameters that can be obtained from the data, namely the interarrival time between requests (IATs), the number of units in stock and in labour, the turn-around-times (TATs), etc. A review of the literature on setting attainable due dates can be found in [5].

The second challenge we address is to find an efficient way to assign the work to the technicians within each group so that all the components are repaired within the required deadlines.

The units repair times are highly stochastic. Each team leader has to deal with high variability in the (random) arrival of components, the type of components and the required maintenance for them. Also, the number of man hours required to repair a part is typically an order of magnitude smaller than the actual time that part spends in the group. When a component arrives, the required repair type and time is not known. While one unit may be repaired just by adjusting the positioning of an electrical switch, another unit of the same type may need to be completely disassembled. In such cases subparts may need to be sent to other groups. This causes interruptions in the repair process in the original group until all the subparts are returned to be reassembled. The arrival of emergency units, the lack of materials or the need for engineering analysis are other causes of interruption. All these issues make the planning process to guarantee on time delivery of the aircraft components very challenging.

Several authors have analyzed demand forecasting of spare parts, e.g. [4, 6, 11, 13, 14]. In particular, [4] examines techniques to predict spare parts demand for airline fleets. Also, many works have addressed the aircraft engine maintenance scheduling problem (see, e.g., [8]).

In “Methods” section we also explain the issues involved in this assignment and scheduling problem. A mathematical model is presented as well as an algorithm. Preliminary computational results for this problem are given in “Results and discussion” section.

We conclude this paper by stating some conclusions and final comments.

Methods

CRD estimation formula

The first goal of this work is to develop an estimation formula for the CRD, which will prioritize the list of repairs, i.e., determine which units should be tackled first. The CRD assignment affects the timely delivery of the units and the number of repaired units in

stock. Whenever a type of aircraft unit is required, it should be repaired immediately or there must exist a similar unit in stock to replace it. Keeping components in stock allows the aircrafts to be quickly returned to service. On the other hand, the limited space for stock and its associated cost should be taken into account. Therefore it is not desirable to have higher levels of stock or shorter CRDs than necessary. Stock levels should be minimized within a safety level that depends on the costs incurred due to delays on the delivery dates. When the stock of a unit falls below its safety level, an alarm is issued to the stock controller to obtain a new unit as soon as possible. Ideally, if the planning process is carried out more efficiently, i.e., the aircraft components are repaired faster and in a more systematic way, less components may be needed in stock while still providing the same desired low probability of an out-of-stock situation. The problem of determining the best inventory policy in the context of aircraft components was addressed in [1, 3, 7, 12].

Our goal is to estimate a CRD for each aircraft component without a preassigned CRD that arrives for repair in a given work group. This is challenging due to the dynamic nature of the process in which new parts for repair are constantly arriving at highly variable, random times.

To estimate the CRD for each unit of type i , we take into account the date of its entry in the group (CD), the estimated time until the next request of a unit of the same type (α_i) and the number of units of that type that are in stock (S_i) and in labour (L_i) at that time instant. We assume that each unit of type i in labour has probability ρ_i of being successfully repaired. Additionally, we assume that each unit of type i has a safety stock level (SS_i) (whose size depends on the cost arising from the lack of units of this type). If the level of stock falls below the corresponding safety level, then delays in delivery could occur unless the corresponding units are repaired as soon as possible. In this situation, the estimated CRD for a new component that arrives should take into account this urgency. For this reason, we assume that the safety stock level is at least one. If the stock reaches the value one and a new component arrives, the stock becomes zero. This does not cause a delay unless a new request arrives before the stock is replenished.

When a unit of type i arrives to the maintenance shop, if its stock is above its safety level, then the CRD for this unit may be postponed for S_i time intervals of length α_i after the current date, since the stock is sufficient to handle S_i new requests. In addition, if there are already L_i units of the same type in labour, then $\rho_i L_i$ of these units should be successfully repaired before the current unit that arrived ends its repair, so the $\rho_i L_i$ units can be treated as stock. On the other hand, if a unit of type i arrives and its stock is below its safety level, then the repair of the unit is urgent. In this situation we take the estimated CRD to be the current date plus the minimum value between the estimated time until the next request (α_i) and the estimated TAT for this type of unit (τ_i).

These considerations lead to the following simple formula to estimate the CRD:

$$CRD_i = \begin{cases} CD + \alpha_i (S_i + \rho_i L_i) & \text{if } S_i \geq SS_i, \\ CD + \min(\alpha_i, \tau_i) & \text{otherwise,} \end{cases} \tag{1}$$

where:

- CD Current date,
- CRD_i Customer required date for units of type i (estimated at CD),
- α_i Estimated interarrival time between requests of units of type i ,

- τ_i Estimated turn-around-time for units of type i ,
- SS_i Positive safety stock level for units of type i ,
- S_i Number of units of type i in stock at CD ,
- L_i Number of units of type i in labour at CD ,
- ρ_i Repair success probability of units of type i .


Figure 1 illustrates this formula for the case that $S_i = 2 \geq SS_i$.

A crucial question to address is how to estimate the time between requests, α_i , for each component of type i , based on past data. Considering the (random) interarrival times between requests of a given component of type i , this consists of selecting a single number α_i among all the possible IAT data values that represents their “best/most typical” value. A statistical analysis of IAT data samples provides important information about the IAT distribution. Several measures for typical values are pointed out in the statistics literature [9, 15]. They are known as central tendency measures due to their attempt to identify some sort of central or balanced position within the data set. The most common measures of central tendency are the mean (often called the average), the median and the mode, but there are others such as the trimmed means, the trimean, etc. All of them are valid central tendency measures, but we should be aware that, depending on the nature of the data under consideration, some will be more appropriate to use than others. Therefore, while we can select a “best” measure of central tendency for the data we are analyzing, there is no “best” central tendency measure for all data sets.

The estimation of α_i is often determined by using the mean of the observed IATs. However, it is important to note that the sample mean is usually considered as the best measure of central tendency when the data exhibits a symmetric, light-tailed distribution (such as for normally distributed data). In such cases, the values of the mean, median, trimean, trimmed means, or mode (for unimodal distributions) are usually similar. The mean is preferred among the others since it includes all the values in the data set in its computation, minimizing the error in the prediction of any of its values.

On the other hand, when the data has outliers or has a skewed or heavy-tailed distribution, which is the case of the typical data of the aircraft components maintenance shops, the values of the mentioned measures may be quite different and it is not always obvious which value to choose. In such cases, frequently we find that the mean is a poor measure of the typical value. The sample mean is usually quite above or below the wide majority of the data. This is a consequence of the extreme values in the tail that drag the value of the mean in the direction of the skewness. In these situations, the median (less sensitive to extreme values) is generally considered to be a better representative of the central location of the data.

For some data distributions, intermediate (robust) estimators between the sample mean and the sample median may perform better as central location estimators. Examples of



such measures are the trimean and the trimmed means that try to combine the good properties of both measures.

To select convenient estimators for the interarrival time between requests (IATs), we perform an exploratory data analysis on the data provided by the maintenance shop for this case study. The data includes information concerning the repair of 333 distinct types of components that left the same maintenance shop during the period of January 2009 – May 2011 and that arrived without any predefined CRD. Nevertheless, for several types of components, only a few occurrences were observed during this period. Due to this lack of data, we conducted a statistical analysis only for the 66 component types that had at least 10 entry dates in the maintenance shop during the mentioned period.

We plot histograms and boxplots and compute the usual measures of location and spread, which give us a global picture on where and how the data is concentrated and what the shape of its distribution is. Table 1 and Fig. 2 illustrate the results obtained for the IATs of fifteen (typical) components, with $N - 1$ standing for their sample sizes (N corresponds to the total number of repair requests).

The results revealed that the IATs (denoted by T in Figs. 3, 4 and 5) typically exhibit high variability and have positively skewed distributions. As a consequence, the IAT mean tends to be significantly higher than the median. Exceptions are components (i) and (m), for which the IATs are negatively skewed with means lower than their medians, and for components (g) and (n) that exhibit almost symmetric distributions. For several component types (e.g., (b), (e), (h), (j) or (o)) we observe the existence of large extreme values that can be measurement errors or a consequence of a long-tailed distribution. Without any way to confirm this, we point out that it would be desirable to identify outliers and remove them from data before performing the statistical analysis, as they can significantly bias the results.

The high variability of the data means that it is impossible to obtain an accurate single value to represent the IAT distributions. Nevertheless, the skewed data with extreme values implies that using the mean as a typical value of the IATs is not the best option. As can be seen in Table 1, the mean is strongly distorted by the extreme values, tending to overestimate the data. E.g., for components (b), (e) and (o), more than 75 % of their IATs are below the mean ($Q_3 < \text{Mean}$). Therefore, using the mean to obtain an estimate for the IATs and to compute the CRD would frequently imply that new parts arrive to repair before previous parts of the same type are finished, pushing the stock to undesirably low levels.

As previously stated, in such cases the use of robust estimators (less sensitive to the effects of outliers) is undoubtedly a better option. The selection of a single measure among possible options such as the median, trimean or trimmed means is not straightforward. Obviously, the smaller the value we choose to estimate the time until the next request, the smaller the risk we take. Nevertheless, choosing all the values in a conservative fashion leads to shorter CRDs that may be impossible to meet. Therefore, the selection of a central tendency measure should take into account the risk tolerance of each type of unit.

Considering for instance the unit (e), we observe a positive skewed distribution with 7 extreme values (see Fig. 2), which distort the mean value. In this case, the median is the best option since the difference between the IATs and their median value is clustered tightly and almost symmetrically around 0 (see Fig. 3). Similar conclusions can be made, e.g., for units (b), (c), (f) and (h).

Table 1 Statistics for the IAT (interarrival time between repair requests) of several components

Aircraft component	<i>N</i>	Mean	Lower quartile (Q_1)	Median (Q_2)	Upper quartile (Q_3)	Trimean $\frac{Q_1+2Q_2+Q_3}{4}$	Trimmed mean (25 %)	Trimmed mean (5 %)	Standard deviation	Coefficient of variation
(a) Valve Ckeck (HTE400115)	12	59.45	25.5	48	96.5	54.5	55.56	—	49.24	0.83
(b) Valve control wing anti-ice (SAS911-002A)	72	12.77	2	5	11	5.75	6.27	10.51	28.64	2.24
(c) Flow control valve (1303A0000-04)	79	11.12	2	7	14.75	7.69	8.57	10.28	13.56	1.22
(d) Valve engine anti-ice (324195-1)	16	54.2	9	29	100.5	41.88	51.85	—	52.74	0.97
(e) Valve starter shutoff (3290064-17)	46	17.84	2	5	15	6.75	10	13.35	36.34	2.04
(f) Valve starter shutoff (3290064-20)	88	9.87	2	7	13.5	7.38	8.07	8.99	10.67	1.08
(g) Valve starter shutoff (3291556-3)	14	28	17	26	33	25.5	27	—	18.23	0.65
(h) Thermostat control temperature (342B040000)	65	11.81	1	4	13.25	5.56	7.58	11.11	16.77	1.42
(i) Auxiliary power unit (3800278-4)	25	34.58	17.5	38	43.75	34.31	33	34.58	24.23	0.7
(j) Actuator (65-20892-13)	82	12	0	0	11	2.75	4.97	8.88	27.12	2.26
(k) Bleed pressure regulating valve (6714D070000)	10	59.11	36	46	104	58	58.86	—	43.98	0.74
(l) Bleed pressure regulating valve (6774E010000)	87	10.42	3	7	16.75	8.44	8.89	9.83	9.93	0.95
(m) Exchanger heat (753A0000-03)	11	89.4	40.25	99	130.25	92.13	91.5	—	53.37	0.6
(n) Element filter (856504-5)	10	29.67	20	29	37	28.75	28.86	—	17.25	0.58
(o) Valve pneu press regulator (898626-3)	11	103.6	7.75	12	88.75	30.13	64.63	—	178.67	1.72

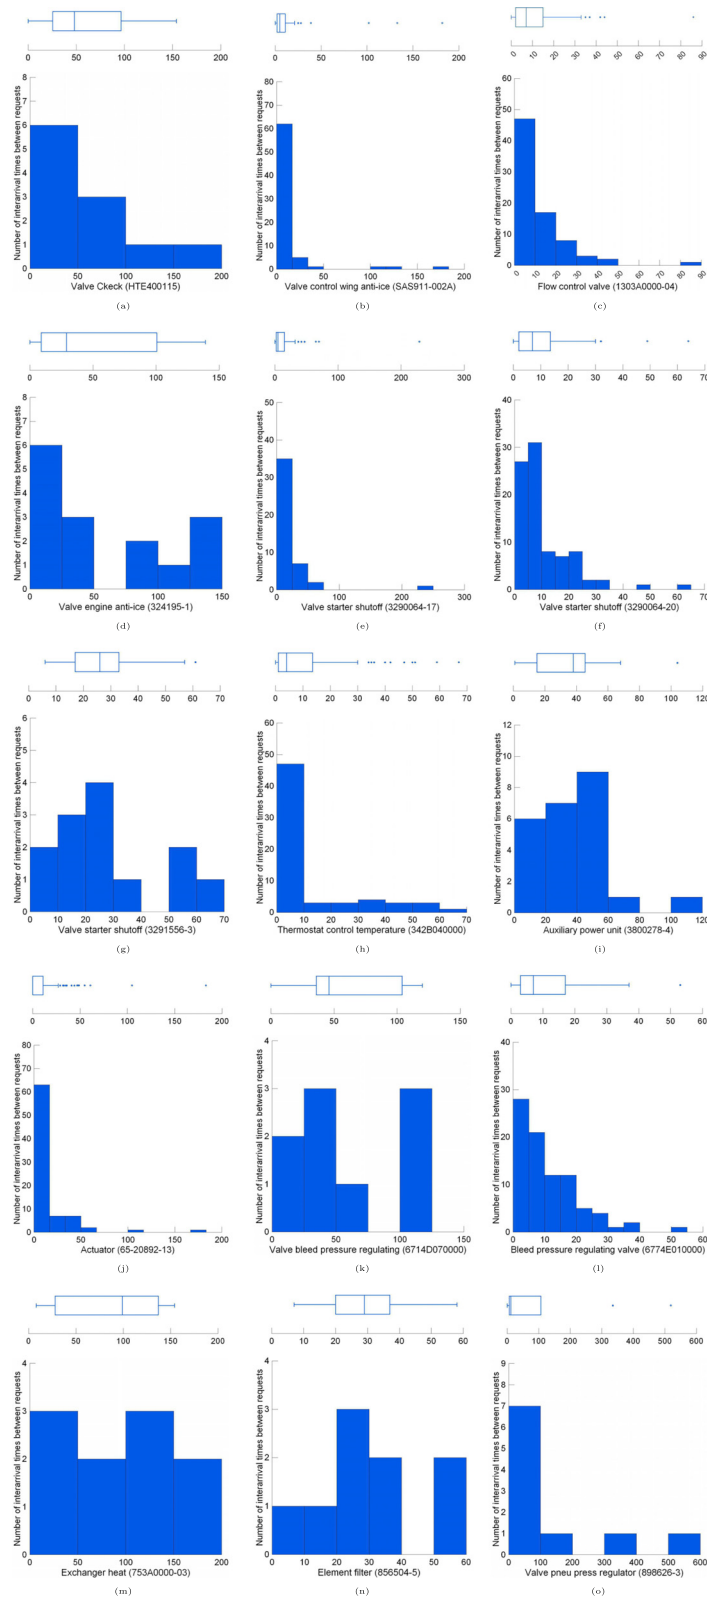
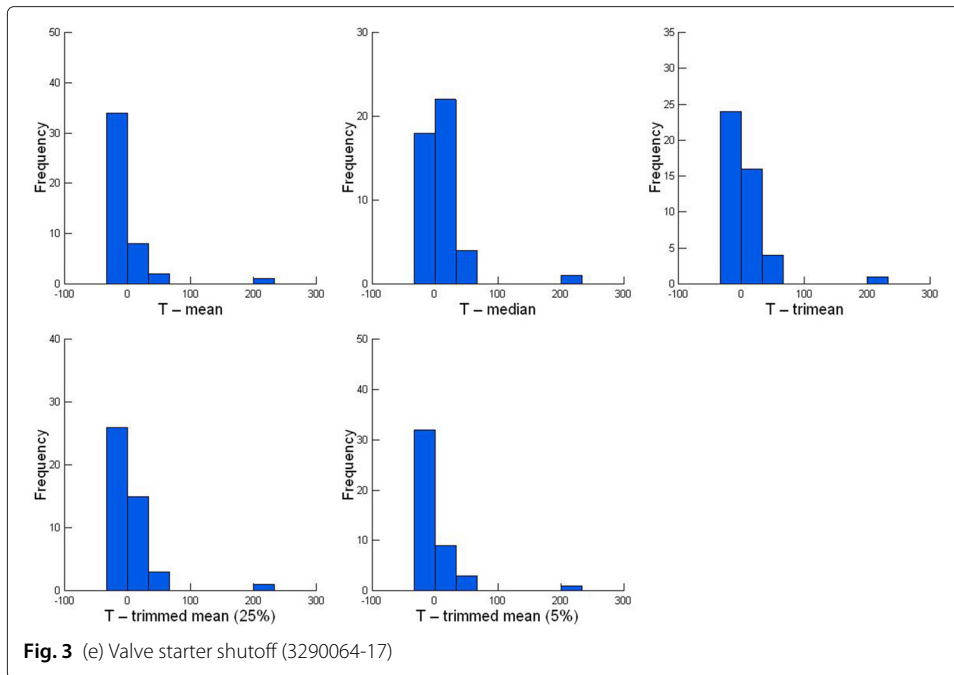
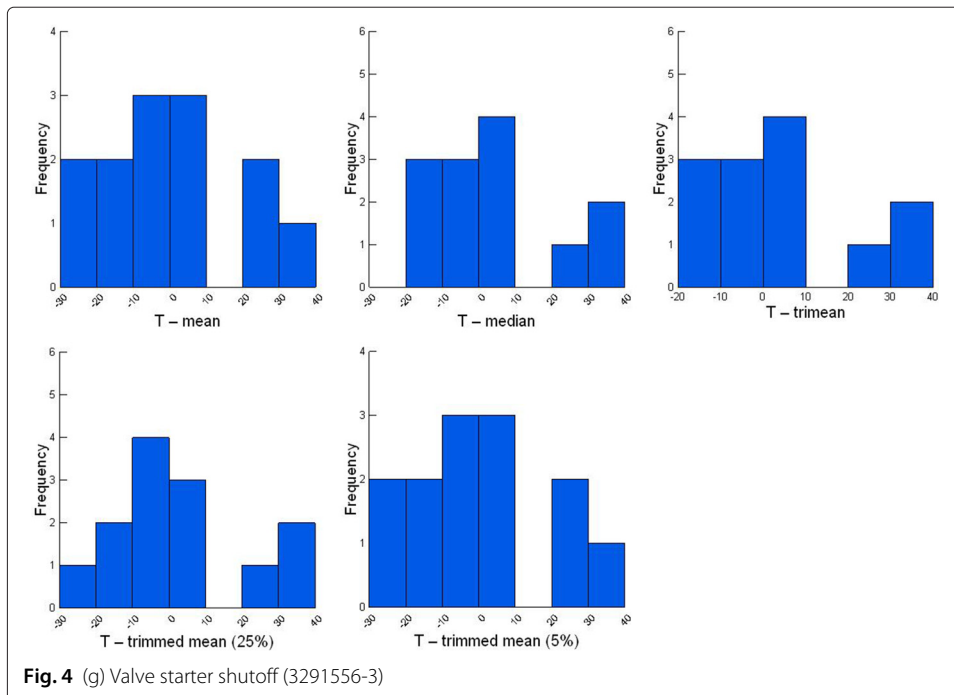


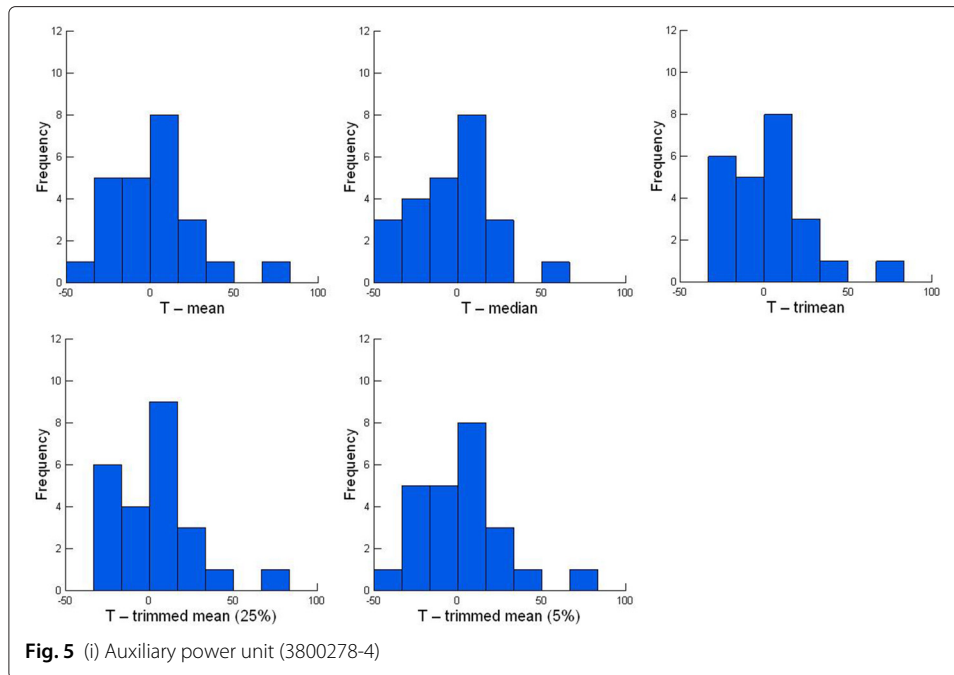
Fig. 2 Histograms for the time between requests (in days)



For units (g) and (n), with near symmetric distributions, the values of the mean, median, trimean and trimmed means are very similar, as expected, and so any of them can be used to estimate the IATs (see Fig. 4).

A similar but different situation arises for unit (i), in which the IATs distribution is negatively skewed. In this case, the mean is only slightly lower than the median due to





the existence of an upper extreme value and the mean, trimean and trimmed means have very similar values. Any of these values can be used to estimate the IATs.

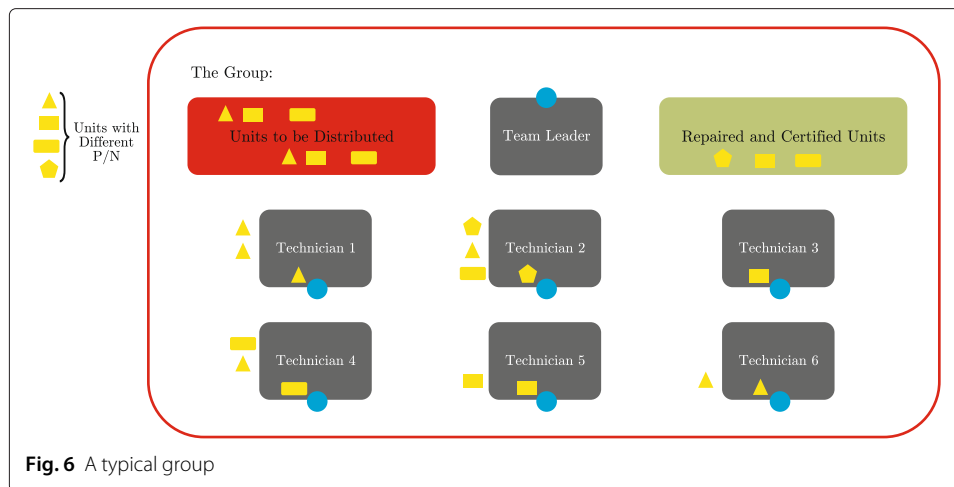
Another important issue is that components can arrive for repair in batches (see (j) in Fig. 2). This situation leads to several null IATs and usually to a null median which, even when the IATs are skewed distributed, should not be used as the typical value to estimate the CRDs. A better option to estimate α_i could be to consider the IAT between batches instead of units. Also, in such cases the safety stock level should take into account the mean batch size.

Finally, it is important to point out that the larger the variability of the data, the larger the sample size should be in order to get accurate estimators. As a consequence, in atypical cases like for part (o) (see Fig. 2), where there are a small number of observations and they are highly variable, it is impossible to infer any good measure of central tendency. In these cases larger samples should be collected.

Assignment and scheduling problem

The second challenge was to find an efficient way to assign the work within each group so that all the components are repaired on time and as quickly as possible. Figure 6 illustrates a typical group.

Not every technician within a group is qualified to repair every type of component. Furthermore, the time required to repair a part varies greatly with the serial number, and the number of man hours required to repair a part is typically an order of magnitude smaller than the actual time that the part spends in the group. For example, an actuator (part number 65-20892-13) takes on average 9 man hours to repair but has an average turn-around-time of 68 days, i.e., there are on average 68 days between the time when the actuator enters and leaves the group. See Tables 2 and 3. This discrepancy is due to interruptions in the repair process; rarely does a technician start working on a



part and work on it until completion without interruptions. Interruptions occur, e.g., if the part or subparts of it need to be sent to other groups (specialising in other types of repair/maintenance), due to lack of materials, lack of resources (e.g., if a machine is unavailable because another technician is using it), or because of the arrival of more urgent components. All this leads to a challenging assignment problem.

Model

We formulate this problem as a binary constrained nonlinear optimisation problem [10, 16]. To do this we make several simplifying assumptions.

We assume that each component arriving in the group is assigned to exactly one technician, i.e., that once a part is assigned to a technician then he is the only person within the group that will ever work on it. In practice this is not the case and technicians within a group may rotate between work benches so that several different technicians end up working on the same part. Note that we still allow for the possibility that a part, or subparts of it, be sent to other groups for other technicians to work on.

We also assume that a technician can work on only one part at once. In practice a technician may work on several parts at once if they are small and of the same type (same part number).

For a given technician, we divide the parts assigned to him into three categories:

- (1) The (unique) part that he is working on right now,
- (2) The parts waiting in queue to be worked on,
- (3) The interrupted parts.

See Fig. 7.

The interrupted parts are those that the technician was forced to stop working on, e.g., due to the lack of materials or the dispatch of the parts to other groups. We assume that when an interruption is over and a part is ready to be worked on again, the part joins the queue of waiting parts. Thus in our model the queue of waiting parts consists of those that have yet to be started as well as those that were interrupted but are now ready to be worked on again. We assume that every part entering the group has a CRD, either predefined or estimated using formula (1).

Table 2 Statistics for the MdO (man hours required for repair) of several components

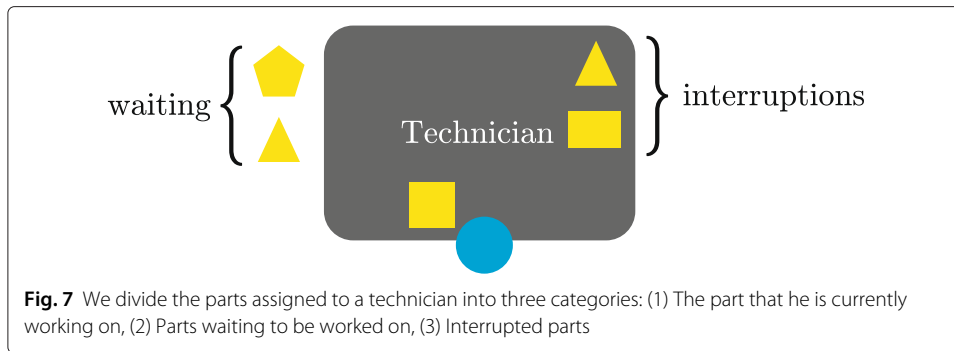
Aircraft component	<i>N</i>	Mean	Lower quartile (Q_1)	Median (Q_2)	Upper quartile (Q_3)	Trimean $\frac{Q_1+2Q_2+Q_3}{4}$	Trimmed mean (25 %)	Trimmed mean (5 %)	Standard deviation	Coefficient of variation
(a) Valve Ckeck (HTE400115)	12	5	3.51	5.03	6.41	5	4.98	_____	2.28	0.46
(b) Valve control wing anti-ice (SAS911-002A)	72	11	5.19	7.47	14.36	8.62	9.2	10.96	10.25	0.9
(c) Flow control valve (1303A0000-04)	79	14	7.33	11.7	17.74	12.12	12.42	13.36	10.54	0.76
(d) Valve engine anti-ice (324195-1)	16	25	6.46	25.85	34.89	23.26	22.07	_____	21.85	0.86
(e) Valve starter shutoff (3290064-17)	46	13	5.64	12.19	18.73	12.18	11.95	12.72	9.1	0.7
(f) Valve starter shutoff (3290064-20)	88	14	6.65	12.26	17.38	12.13	12.18	13.14	10.76	0.77
(g) Valve starter shutoff (3291556-3)	14	10	7.64	9.19	11.05	9.26	9.79	_____	5.1	0.49
(h) Thermostat control temperature (342B040000)	65	6	0.75	1.5	11.75	3.88	4.67	6.02	7.64	1.21
(i) Auxiliary power unit (3800278-4)	25	15	6.08	10.93	17.57	11.38	11.44	15.49	19.27	1.24
(j) Actuator (65-20892-13)	82	9	1.65	8.7	13.89	8.24	8.28	8.67	7.42	0.82
(k) Bleed pressure regulating valve (6714D070000)	10	22	5.55	12.62	41.66	18.11	20.91	_____	19.46	0.89
(l) Bleed pressure regulating valve (6774E010000)	87	4	0.75	3.14	5.73	3.19	3.36	3.97	4.16	0.99
(m) Exchanger heat (753A0000-03)	11	24	5.62	9.94	42.69	17.05	21.28	_____	22.33	0.93
(n) Element filter (856504-5)	10	3	2.12	2.83	3.58	2.84	3.05	_____	2.11	0.64
(o) Valve pneu press regulator (898626-3)	11	38	14.95	22.97	49.65	27.63	30.67	_____	37.77	0.98

N is the sample size

Table 3 Statistics for the TAT (turn-around-time in days) of several components

Aircraft component	<i>N</i>	Mean	Lower quartile (Q_1)	Median (Q_2)	Upper quartile (Q_3)	Trimean $\frac{Q_1+2Q_2+Q_3}{4}$	Trimmed mean (25 %)	Trimmed mean (5 %)	Standard deviation	Coefficient of variation
(a) Valve Ckeck (HTE400115)	12	43	6	22	57.75	26.94	37.6	—	48.26	1.13
(b) Valve control wing anti-ice (SAS911-002A)	72	23	3	9.5	20.25	10.56	12.81	20.3	37.1	1.64
(c) Flow control valve (1303A0000-04)	79	16	3	6	20	8.75	11.43	15.23	20.7	1.28
(d) Valve engine anti-ice (324195-1)	16	56	7.5	24.5	77.5	33.5	40.17	—	68.53	1.22
(e) Valve starter shutoff (3290064-17)	46	11	4.25	7	17	8.8	10.03	10.8	9.72	0.86
(f) Valve starter shutoff (3290064-20)	88	11	4	7	14.25	8.06	8.83	10.39	11.26	1
(g) Valve starter shutoff (3291556-3)	14	8	4.25	7	8.75	6.75	6.75	—	5.83	0.76
(h) Thermostat control temperature (342B040000)	65	32	21	24	35	26	27.78	31.7	22.87	0.71
(i) Auxiliary power unit (3800278-4)	25	48	6	17	51	22.75	28.16	48.08	77.80	1.62
(j) Actuator (65-20892-13)	82	68	22.5	49	75	48.88	51.47	62.24	71.09	1.05
(k) Bleed pressure regulating valve (6714D070000)	10	32	7	24	61.25	29.06	31.13	—	29.86	0.92
(l) Bleed pressure regulating valve (6774E010000)	87	32	22.5	35	42	33.63	31.93	31.77	18.19	0.56
(m) Exchanger heat (753A0000-03)	11	39	28.5	35	52	37.63	38.33	—	21.89	0.56
(n) Element filter (856504-5)	10	57	45.75	56.5	70.75	57.38	57.5	—	14.97	0.26
(o) Valve pneu press regulator (898626-3)	11	147	56	132	150	117.5	111.56	—	155.24	1.05

N is the sample size



Algorithm

In this subsection we present an algorithm for determining the distribution of new components amongst the technicians within a single group, and a general methodology for treating interruptions and the interaction between all the groups.

Events. For each group, we identify four types of event at which a decision about assignment or scheduling must be made:

1. Arrival of new components.
2. An interruption occurs for some technician.
3. A part that was interrupted is returned (to the same technician).
4. A technician finishes repairing a part.

Event 1. In the case of Event 1 we solve a binary constrained nonlinear optimisation problem. For a given group, let m be the number of technicians and let n be the number of arriving components. Let $\mathbb{B}^{n \times m}$ be the set of n -by- m binary matrices, i.e., matrices with entries 0 or 1. Let $A \in \mathbb{B}^{n \times m}$ represent an assignment of the arriving n parts to the m technicians, i.e., A_{jk} is nonzero if and only if part j is assigned to technician k . The assignment matrix A is subject to the following constraints:

- (i) Each part must be assigned to exactly one worker, i.e., $\sum_{k=1}^m A_{jk} = 1$ for all $j \in \{1, \dots, n\}$.
- (ii) Part j can only be assigned to technician k if technician k is qualified to repair it.
- (iii) Let CRD_j be the customer required date for part $j, j \in \{1, \dots, n\}$. Let f_j be the estimated date by which part j will be finished (repaired and ready to leave the group). This depends not only on the assignment matrix A but also on the current workload of the technicians and in fact even on the unknown future events. Thus f_j is difficult to compute. Different models for the function $f_j(A)$ are suggested below. The third set of constraints is that part j is finished on time, i.e., that $f_j(A) \leq CRD_j$ for all $j \in \{1, \dots, n\}$.

Let $\mathcal{K} \subset \mathbb{B}^{n \times m}$ be the set of binary matrices satisfying these constraints. Below we describe how to choose an optimal assignment matrix $A \in \mathcal{K}$.

Once the new parts have been assigned to the technicians, we schedule the work as follows. For each technician, his new parts are added to his waiting queue and then the waiting queue is ordered by customer required date (CRD), with the most urgent part being at the top of the queue. We stipulate that each technician finishes working on his

current part (or works on it until it is interrupted) before moving onto the first part in his waiting queue, even if the first part in the waiting queue is more urgent than his current part (this simplifying assumption is to avoid situations where, due to an urgent arrival, a technician stops working on his current part when it is almost finished).

Now we return to the problem of assigning the components. We seek a feasible assignment matrix $A \in \mathcal{K}$ that minimises some given cost function c :

$$\min_{A \in \mathcal{K}} c(A). \tag{2}$$

The cost function c should be chosen so that minimising it corresponds to finishing the repairs as fast as possible. Here are two possible choices for c : Let P be a subset of all the parts in the group (the exact choice of P depends on the choice of f_j , see below). We could take the cost function c to be

$$c(A) = \sum_{j \in P} (f_j - CRD_j) \tag{3}$$

so that we minimise the total (signed) delay (this corresponds to maximising the total slack time). An alternative choice for c is

$$c(A) = \max_{j \in P} (f_j - CRD_j). \tag{4}$$

The constraint (iii) above could be dropped if another choice of c is made that heavily penalizes delays.

In general solving binary constrained nonlinear optimisation problems is costly (in fact NP hard [2]), but in this case the number of parts to assign and the number of technicians is typically a small number. Thus the optimisation problem (2) can be solved quickly. See the example in “Results and discussion” section.

Now we discuss different ways to compute f_j , the estimated date by which part j will be finished. The difficulty here, as already described, is the large variance in repair times for each component due to interruptions and new arrivals.

Thus, given a queue of parts waiting to be repaired, it is very difficult to estimate when work will start and finish for each part, and thus whether each part will be finished on time.

First we introduce some notation. We divide all the parts in the group into the following sets:

- C The parts that are currently being worked on,
- N The new parts that have arrived in the group,
- I The interrupted parts,
- \widehat{S} The parts in the waiting queues that have not been started yet,
- S The parts in the waiting queues that have already been started.

Note that the size of the set C equals the number of technicians. Here we assume that the new parts have already been assigned so that $N \subset \widehat{S}$.

For part j , let τ_j be the estimated turn-around-time (time between when the part enters and leaves the group). Depending on the type of part, this value may be chosen to be, e.g., the mean, median, trimean or any other typical value of the turn-around-time (as for the problem of estimating the IAT, see the discussion in “CRD estimation formula” subsection). In the example in “Results and discussion” section we choose the median value.

For $j \in \widehat{S}$, let μ_j be the estimated MdO (man hours required for repair) for part j . Note that before work starts on a part, we may not know what type of repair is necessary. Thus μ_j must be computed by taking the estimated MdO over all repair types. For $j \in C \cup S \cup I$, the set of parts that have already been started, let μ_j be the estimated *remaining* number of man hours required for repair. This allows for input from the technicians based on their progress. If these values are too difficult to estimate, then the original estimated MdO values can simply be used instead.

An approximate upper bound for f_j can be computed as follows. If work has already started on part j , i.e., $j \in C \cup S \cup I$, then f_j is taken to be the starting date of repair s_j plus the estimated turn-around-time τ_j :

$$f_j = s_j + \tau_j, \quad j \in C \cup S \cup I. \tag{5}$$

Note that if work has already started on part j , then the assignment matrix A has no effect on the value of f_j computed by Eq. (5). Thus we do not include these parts in the set P appearing in the definition of c (see Eqs. (3) and (4)); we define $P = \widehat{S}$. In practice of course the assignment of new parts to a technician can affect the finish times of parts that he has already started. This effect is hidden in the turn-around-time τ_j in (5), which includes the time when the technician is working on other parts.

For parts $j \in P = \widehat{S}$, which have not been started, we compute f_j as follows. Suppose that the technician who is assigned part j is currently working on a part with an estimated remaining MdO of μ_c hours. Suppose that the technician started working on this part on date s_c . Suppose also that in the queue before part j , there are l parts with estimated (original or remaining) MdOs $\mu_1, \mu_2, \dots, \mu_l$. Then we estimate s_j , the starting date for part j , by taking it to be the latest date possible, i.e., we assume that all the parts in the queue before part j are finished without interruptions before part j is started:

$$s_j = s_c + (\mu_c + \mu_1 + \mu_2 + \dots + \mu_l)/8, \quad j \in \widehat{S}. \tag{6}$$

Here we divide by 8 to convert hours into working days. Then we can estimate the finish date by

$$f_j = s_j + \tau_j = s_c + (\mu_c + \mu_1 + \mu_2 + \dots + \mu_l)/8 + \tau_j, \quad j \in \widehat{S}. \tag{7}$$

Equation 7 will overestimate the finishing date for part j in general. As time progresses, however, as the number of items in the queue before part j decreases, this upper bound will improve.

An alternative formula for f_j , which underestimates the true finishing dates in general, can be obtained by replacing the estimated turn-around-time τ_j in (7) with the estimated number of man hours required for repair μ_j , i.e., by ignoring interruptions. In this case, the estimated finishing date f_j of parts $j \in S$, parts that have already been started, should be computed using Eq. (7) (with τ_j replaced by μ_j) rather than using (5). Thus the set $P = \widehat{S} \cup S$.

Since this alternative formula for f_j ignores interruptions it would grossly underestimate the finishing dates in general and so we do not discuss it further. Obviously there are many other possibilities for f_j other than the approximate upper bound and the lower bound presented here, e.g., a combination of the two could be used.

Events 2–4. In the case of Events 2–4, there is no assignment problem to be solved. For Events 2 and 4 the technician just starts working on the next most urgent part, as determined by the CRD. For Event 3, the returned part is added to the waiting queue in the appropriate position. As with the arrival of new parts, we specify that the technician finishes working on his current part, even if the returned part has a more urgent due date.

In the case of Event 2, an interruption, the interrupted part (or a subpart of it) may be sent to another group. In this case, this part (or subpart) could be treated as a new part within the receiving group, i.e., triggers an Event 1 within the receiving group (the CRD for this part, however, should be modified within the receiving group so that the part is returned to the original group in time for the rest of the repairs to be completed). In this way we could develop an algorithm for the assignment of parts within all the groups, taking into account the interactions between them.

Results and discussion

In this section we give a partial implementation of the algorithm, implementing it in the case of just one group and one event, the arrival of new parts. A full scale implementation is beyond the scope of this paper and would be a major programming task. Moreover, before doing so our algorithm should be generalized to address the limitations discussed in “Model” subsection and to better meet the needs of an aircraft components maintenance shop.

Due to the nonsmooth nonlinear nature of the problem addressed, a preliminary implementation was undertaken using the evolutionary solver of Excel 2010 that combines genetic algorithms and local search optimization methods. This solver cannot determine whether a given solution is optimal, nevertheless good solutions can be attained under heuristic rules or other predefined stopping criteria.

Figure 8 shows the results of our algorithm for a small example, for a group of 4 technicians, the arrival of 4 new parts, and 20 parts in total within the group. The rows containing the new parts are highlighted. The columns are sorted by the CRDs of the parts, with the most urgent part at the top. All the repairs are completed in time and the value of the cost function *c* is given in the bottom-right, highlighted box. The minimisation problem (2) was solved in about 10 sec by the Excel evolutionary solver.

Part Name	Arrival Date	CRD	Status of Part					Waiting Parts	Estimated TAT (days)	Estimated MSD (hours)	Estimated Remaining MSD (C & S)	Technician Skills				Assignments to Technicians				Constraint	Remaining MSD of Current Part	Total MSDs of Parts Ahead in Queue	Estimated Finish Date (N & S)	Delay (days)			
			C	N	S	S	I					1	2	3	4	1	2	3	4								
Valve Starter Shutoff	10-07-2011	31-07-2011	1					0	7,00	9,19	6,00	1	0	1	1	1	0	0	0	1	1	4	3	0	0	0	0
Valve Starter Shutoff	10-07-2011	31-07-2011					1	0	7,00	12,26	12,26	1	0	1	1	0	0	0	1	1	4	3	0	0	0	0	0
Valve Starter Shutoff	10-07-2011	31-07-2011					1	0	7,00	12,26	12,26	1	0	1	1	0	0	1	0	1	3	5	0	0	0	0	0
Valve Starter Shutoff	10-07-2011	31-07-2011	1					0	7,00	12,19	3,00	1	0	1	1	0	0	1	0	1	3	5	0	0	0	0	0
Valve-Regulating, H.P.	08-07-2011	10-08-2011	1					0	22,00	28,29	35,00	1	1	1	0	1	0	0	1	2	15	0	0	0	0	0	0
Safety Valve	13-07-2011	13-08-2011					1	0	6,00	4,24	4,24	0	1	0	1	0	1	0	0	1	2	15	0	0	0	0	0
Safety Valve	13-07-2011	13-08-2011	1					0	6,00	4,24	3,00	0	1	0	1	0	0	1	1	4	3	0	0	0	0	0	0
Safety Valve	13-07-2011	13-08-2011		1				1	6,00	4,24	4,24	0	1	0	1	0	1	0	1	2	15	0	0	0	0	0	0
Safety Valve	13-07-2011	13-08-2011					1	0	6,00	4,24	4,24	0	1	0	1	0	1	0	1	2	15	0	0	0	0	0	0
Accessory Power Unit	14-07-2011	14-08-2011	1					1	13,00	10,93	10,93	1	1	0	0	0	1	0	1	3	4	0	0	0	0	0	0
Valve Control Wing Anticice	11-07-2011	15-08-2011				1		1	9,90	7,47	2,00	1	1	1	0	0	1	0	1	2	5	10,93	0	0	0	0	0
Valve Control Wing Anticice	11-07-2011	15-08-2011				1		1	9,90	7,47	7,47	1	1	1	0	0	0	1	1	4	3	0	0	0	0	0	0
Flow Control Valve	14-07-2011	25-08-2011	1					1	6,00	11,70	11,70	1	1	1	0	0	1	1	4	3	0	0	0	0	0	0	0
Flow Control Valve	14-07-2011	25-08-2011	1					1	6,00	11,70	11,70	1	1	1	0	1	0	1	2	15	0	0	0	0	0	0	0
Flow Control Valve	14-07-2011	25-08-2011	1					1	6,00	11,70	11,70	1	1	1	0	0	1	0	1	3	5	12,93	0	0	0	0	0
Thermostat Control Temperature	12-07-2011	09-09-2011				1		1	24,00	1,50	1,00	0	1	1	1	0	0	1	0	1	3	5	24,63	0	0	0	0
Thermostat Control Temperature	12-07-2011	09-09-2011				1		1	24,00	1,50	1,00	0	1	1	1	0	0	1	1	4	3	11,7	0	0	0	0	0
Thermostat Control Temperature	12-07-2011	09-09-2011				1		1	24,00	1,50	1,50	0	1	1	1	0	0	1	1	4	3	12,3	0	0	0	0	0
Valve-Pressure Regulating	09-07-2011	09-10-2011				1		1	11,00	12,70	12,70	1	0	0	1	1	0	0	1	1	6	0	0	0	0	0	0
Valve-Pressure Regulating	09-07-2011	09-10-2011				1		1	11,00	12,70	12,70	1	0	0	1	1	0	0	0	1	1	6	12,7	0	0	0	0
Valve-Pressure Regulating	09-07-2011	09-10-2011				1		1	11,00	12,70	12,70	1	0	0	1	1	0	0	0	1	1	6	25,4	0	0	0	0

Date of new arrivals: 14-07-2011

Cost function: -390,80375

Fig. 8 Results of the algorithm for a small example

The column *Status of Part* indicates whether the part is a current part for some technician (C), is a new arrival (N), is in the waiting queue of some technician and has not been started (\widehat{S}), is in the waiting queue and has been started (S), or is interrupted (I).

The column *Estimated Remaining MdO (C & S)* gives the estimated man hours of work left for those parts that have already been started (the C and S parts) and the original estimated MdO (the median MdO in this example) otherwise. The values of the estimated remaining man hours would be entered by the technicians. If they are too difficult to estimate, then the original estimated MdOs could simply be used instead.

The column *Technician* indicates which technician is working on each part. The column *Constraint* indicates that constraint (i) (see “Algorithm” subsection) is satisfied (this is indicated by a 1). The constraints (ii) and (iii) are also enforced by the solver but are not shown in the spreadsheet.

The column *Remaining MdO of Current Part* evaluates μ_c from Eq. (7). The column *Total MdOs of Parts Ahead in Queue* evaluates $\mu_1 + \dots + \mu_l$ from Eq. (7). The column *Estimated Finish Date (N & \widehat{S})* computes f_j using Eq. (7). The last column computes $f_j - CRD_j$.

Conclusions

In “CRD estimation formula” subsection we derived a formula for estimating the CRD of the components that arrive in the groups without one. As suggested by TAP Maintenance & Engineering, this formula was kept simple and easy to understand. This formula requires an estimated interarrival time (IAT) for each type of component. We performed a statistical analysis of data provided by an aircraft components maintenance shop to determine a good estimator for the IATs. We found that, in general, the data is positively skewed so that a robust estimator such as the median performs better than the mean (with a few exceptions, discussed in “CRD estimation formula” subsection, such as for components that tend to arrive in large batches). We also point out that the estimated CRD must be updated over time as more information becomes available, rather than being a fixed value, to take into account, e.g., the frequency at which components actually arrive (the actual IATs) rather than the predicted values. Formula (1) for the CRD should be tested by comparing it with real data.

In “Assignment and scheduling problem” subsection we designed an algorithm to assign the components to the technicians within a single group, and we implemented it in Excel for a small example. Several limitations need to be addressed before the algorithm could be used by an aircraft components maintenance shop, such as our assumptions that a technician can only work on one part at once and that technicians do not rotate between work benches. We indicated how the algorithm could be extended to take into account the interactions between all the groups. This needs further development, however, and its implementation would be a major programming task. The assignment algorithm should be tested by comparing the assignments produced by our algorithm with those produced manually by team leaders.

We solved the binary constrained nonlinear optimisation problem in Excel. Since the assignment problem is NP hard, finding an optimal solution is difficult and time consuming in general and for larger examples than the one given in “Assignment and scheduling problem” subsection, the Excel solver may not be able to find an optimal solution. Other

possible optimisation packages include AIMMS and GAMS. Alternatively, instead of solving the optimisation problem exactly, a heuristic could be developed.

Many specialized programs are available for assignment and scheduling, such as SIMUL8 and the TORSCHÉ Scheduling Toolbox for MATLAB. It would be worth investigating whether their capabilities could be applied to the optimisation problem presented in “Assignment and scheduling problem” subsection.

Authors' contributions

All authors made contributions to the paper conception, namely on the statistical data manipulation, analysis and interpretation, and on the mathematical model formulation and algorithmic implementation. CA, DB and FF were additionally involved in the writing of this paper. All authors read and approved the final manuscript.

Acknowledgements

This research was partially supported by Fundação para a Ciência e a Tecnologia (FCT) through projects UID/MAT/00013/2013 and UID/Multi/04621/2013. The authors would like to thank TAP Maintenance & Engineering for providing the problem and data set treated in this case study and for all the valuable discussions.

Author details

¹CMAT-UTAD, CEMAT-IST-UL, Universidade de Trás-os-Montes e Alto Douro, Vila Real, Portugal. ²Department of Mathematical Sciences, Durham University, Durham, UK. ³Departamento de Física e Matemática, Instituto Superior de Engenharia de Coimbra, Coimbra, Portugal. ⁴Departamento de Matemática, Universidade de Évora, Évora, Portugal.

Received: 19 May 2014 Accepted: 26 May 2016

Published online: 22 July 2016

References

1. Cornelissen F (2010) Shop floor control in repair shops. Which shop floor control method can be used for which repair shop?, Master Thesis, University of Twente
2. Garey M, Johnson D (1979) Computers and intractability: a guide to the theory of NP-completeness. W. H. Freeman and Company, New York
3. Ghobbar AA, Friend CH (2002) Sources of intermittent demand for aircraft spare parts within airline operations. *J Air Transp Manag* 8:221–231
4. Ghobbar AA, Friend CH (2003) Evaluation of forecasting methods for intermittent parts demand in the field of aviation: a predictive model. *Comput Oper Res* 30:2097–2114
5. Gordon V, Proth J, Chu C (2002) A survey of the state-of-the-art of common due date assignment and scheduling research. *Eur J Oper Res* 139:1–25
6. Gutierrez RS, Solis AO, Mukhopadhyay S (2008) Lumpy demand forecasting using neural networks. *Int J Prod Econ* 111:409–420
7. Jaarsveld W, Dollevoet T (2011) Spare parts inventory control for an aircraft component repair shop. Report EI2011-24, Econometric Institute, Erasmus University Rotterdam
8. Kleeman MP, Lamont GB (2005) Solving the aircraft engine maintenance scheduling problem using a multi-objective evolutionary algorithm. *Lect Notes Comput Sci* 3410:782–796
9. Loether HJ, Mctavish DG (1976) Descriptive and Inferential Statistics: an Introduction. Boston: Allyn and Bacon
10. Pardalos PM, Pitsoulis LS (eds) (2000) Nonlinear assignment problems: algorithms and applications. Kluwer Academic Publishers, Dordrecht
11. Romeijnders W, Teunter RH, Jaarsveld W (2012) A two-step method for forecasting spare parts demand using information on component repairs. *Eur J Oper Res* 220:386–393
12. Simao H, Powell W (2009) Approximate dynamic programming for management of high-value spare parts. *J Manuf Technol Manag* 20(2):147–160
13. Syntetos AA, Boylan JE (2001) On the bias of intermittent demand estimates. *Int J Prod Econ* 71(1–3):457–466
14. Teunter RH, Duncan L (2009) Forecasting intermittent demand: a comparative study. *J Oper Res Soc* 60(3):321–329
15. Tukey JW (1977) Exploratory Data Analysis. Addison-Wesley Publishers, Massachusetts
16. Wolsey LA (1998) Integer Programming. John Wiley & Sons Publishers, New York