Retirement optimization through aircraft transfers and employment

Newcamp, Jeffrey; Verhagen, W. J.C.; Santos, Bruno F.; Curran, Richard

DOI
10.1016/j.jairtraman.2019.101680

Publication date
2019

Document Version
Final published version

Published in
Journal of Air Transport Management

Journal of Air Transport Management 79 (2019) 101680

Contents lists available at ScienceDirect

Journal of Air Transport Management

journal homepage: www.elsevier.com/locate/jairtraman

Retirement optimization through aircraft transfers and employment

Jeffrey Newcamp*, W.J.C. Verhagen, Bruno F. Santos, Richard Curran

Delft University of Technology, Faculty of Aerospace, Kluyverweg 1, 2629 HS, Delft, the Netherlands

ARTICLE INFO

Keywords:
Military aircraft
Aircraft retirement
Mixed-integer linear programming
Aircraft transfers

ABSTRACT

Military aircraft retirements are an afterthought for many lifecycle planners. More active management of end-of-life fleets can yield increased confidence in fleet capability and retirement timelines. This work provides fleet managers with a tool to manage remaining aircraft flight hours to yield a desired fleet retirement pattern. It solves an equivalent flight hour minimization problem using a mixed-integer linear programming model for a military aircraft fleet having a network with basing and mission type constraints. The model minimizes differences in remaining equivalent flight hours for individual aircraft in future years, thereby allowing a fleet manager to alter the timeline for retirement of individual aircraft. A relocation cost is applied to discourage excessive, costly aircraft relocations. The United States Air Force A-10 Thunderbolt II aircraft is used as a case study while disruptions such as deployments are modeled to show the methodology's robustness. This work proves that a fleet of aircraft with dissimilar utilization histories and varying amounts of remaining useful lifetime can be actively managed to change the time at which individual aircraft are ready for retirement. The benefit to fleet managers is the ability to extract additional lifetime out of their aircraft prior to retirement.

1. Introduction

Military aircraft fleets are retired with little regard to remaining flight hours. This leads to unused residual life in multi-million dollar capital assets (Jardine, 2011), (Oakley-Bogdewic and Osman, 2015). An end-of-life fleet's retirement is triggered by political motivation, technological obsolescence or budgetary necessity. These triggers are often outside the control of a fleet manager. Previous work by the authors shows that these triggers can be forecast using aircraft utilization data (Newcamp et al., 2016a). Fleet managers can capture these data and use them to devise methods to extract additional usage from their fleet. One method is to actively manage the transfers of aircraft between bases and the employment of those aircraft at the bases. The Retirement Optimization Through Aircraft Transfers and Employment (ROTATE) tool developed by this research effort gives fleet managers the ability to optimize end-of-life aircraft usage while seeking a desired retirement date profile. The goal of this paper is to provide fleet managers with a tool to manage the remaining flight hours each aircraft in a fleet can fly.

The United States Air Force (USAF) collects large amounts of aircraft utilization data so the motivation for this work is to use those data to provide better fleet lifespan utilization. The USAF manages most of its fleets using equivalent flight hours (EFH). This measure combines flight hours with usage severity information. For example, a particularly strenuous one-hour mission may register as 1.3 EFH while a docile one-hour mission could be 0.8 EFH. Four separate USAF fleets with normalized remaining EFH are shown as cumulative distribution functions (CDF) in Fig. 1. This general CDF shape is similar for other aircraft fleets and is representative of the procurement rates of the aircraft.

The CDFs in Fig. 1 represent a snapshot in time, but as the aircraft are flown, their remaining equivalent flight hours decrease. The general shape of the CDFs shown is called "Ramp." Fig. 2 shows a generalized representation of Ramp. If no intervention occurs, an aircraft fleet would see aircraft reaching zero remaining EFH in a steady stream. In practice it is impractical to frequently retire single aircraft, so like-aged groups are selected for retirement (Jones et al., 1991). This retirement pattern is called "Multi-Step" (Fig. 2). "Cliff" is a profile where all aircraft retire at one forecast time (Fig. 2). It occurs when increased usage is assigned to those assets with less accumulated usage. The Ramp pattern is achieved with little intervention from the status quo while the Multi-Step pattern can be modeled by repeating the Cliff pattern with subsets of the fleet population.

While Ramp retirement patterns require little intervention to the normal aging of a fleet, they are hard for a fleet manager to manage. Continuous dwindling of combat capability and underutilization of fielded support resources are disadvantages of the Ramp philosophy. Multi-Step retirement patterns can effectively retire entire operating units in each step, which makes unit deactivation and replacement

* Corresponding author.
E-mail address: jeffrey.newcamp@us.af.mil (J. Newcamp).

https://doi.org/10.1016/j.jairtraman.2019.101680
Received 19 July 2018; Received in revised form 8 April 2019; Accepted 28 May 2019
Available online 04 June 2019
0969-6997/ Published by Elsevier Ltd.

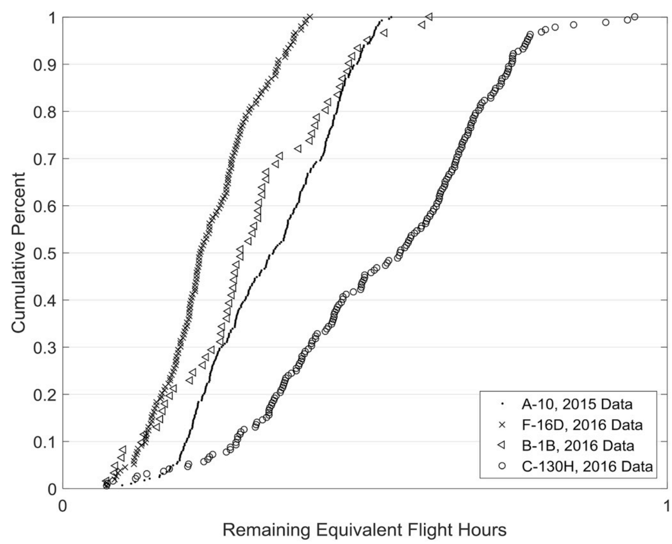

Fig. 1. Remaining EFH curves for four USAF aircraft types.

more efficient than the Ramp. The disadvantage to a Multi-Step retirement is that it requires a great amount of coordination and political buy-in to retire groups of aircraft at the same time and with little remaining useful lifetime. The Cliff pattern is the hardest pattern to achieve because of the disparate nature of aircraft utilization rates across an enterprise. However, it is the most desired because fleet managers and policymakers report that it is the easiest to plan for. An in-depth economic analysis of each pattern has not been conducted.

To alter a fleet's CDF shape to more closely mimic a desired retirement pattern, a fleet manager may employ two approaches. Aircraft may be transferred from one base to another and aircraft may be assigned to a different mix of mission types. Previous work showed that aircraft experience different EFH demands at each base in a fleet's network and that mission types flown also impact EFH accumulation (Newcamp et al., 2016b). A fleet manager may therefore choose to transfer aircraft between bases and alter the expected mission type assignments to change the aircrafts' expected utilization.

This work proposes a single-period mixed integer linear programming model to alter the remaining EFH CDF of a fleet. The model is used to transform a fleet from an existing Ramp pattern to a Cliff (and by association Multi-Step) pattern. Cliff was chosen because the customer for this work, the USAF, views Cliff as the historically most common method of retirement. The scope of this optimization problem is:

1. Only one fleet considered during the simulation.
2. Aircraft transfers only considered once per simulation period.
3. The number of aircraft, bases and required number of missions only changes once per simulation period.

In this problem, demand is modeled as the set of mission requirements at an air base. Supply is modeled as the set of capital assets and their corresponding remaining EFH. Because the network demands change with time, the single-period model is employed in a multi-period simulation. Inputs to the model are free to change for each simulation period. The problem is stated as follows: given an existing fleet of aircraft and an existing network of basing locations, minimize the distribution of EFH subject to realistic operational constraints. A relocation cost (in EFH) is included in the objective function to realistically model the trade-off fleet managers encounter when deciding to relocate aircraft. This methodology uses mixed integer linear programming to influence the remaining useful life of a fleet of aircraft. The idea of relocating aircraft to impact utilization represents a new way to view the lifecycle of aging aircraft.

The remainder of this article is split into four sections. The Literature Review describes similar work in this field. Then the Methodology section presents the mathematical formulation and describes the inputs to the model. The Results and Discussion section shows actual A-10 Thunderbolt II case study results and also highlights the model's robustness given unplanned disruptions to the model. Lastly, the Conclusions section synthesizes the findings and highlights areas for further research using this approach.

2. Literature review

No existing literature discusses shaping aircraft retirements nor is there a substantial amount of literature on relocating and changing mission types for aircraft as a way to prolong their lifetime. However, there are plenty of related works addressing commercial aircraft fleet management as well as capital equipment replacement.

This work focuses on relocating aircraft as a way to impact their utilization. This work is analogous to Başdere and Bilge's work on the aircraft maintenance routing problem. They consider it inefficient to conduct maintenance activities on aircraft that possess remaining useful time before requiring those maintenance activities (Başdere and Bilge, 2014). Their model tracks remaining time on aircraft for the purpose of maximizing remaining useful time utilization. Their integer linear programming model accounts for operational considerations for commercial aircraft fleets. One such consideration is the cost of asset relocations. Sriram and Haghani include aircraft relocation costs into their model by penalizing unnecessary or duplicate assignments (Sriram and Haghani, 2003). Similarly, in this work, relocation costs are considered as a way to discourage unnecessary relocations. Relocating assets to other bases impacts what is flown as well as the costs associated with operations (Clark, 2007), (Robbert, 2013).

Litvinchev et al. use a Lagrangian heuristic for solving the many-to-many assignment problem (Litvinchev et al., 2010). This work is important because it allows for agent and task capacity limits necessary for a military fleet assignment problem.

Fleet size impacts a fleet's capacity while the specificity of aircraft roles impacts the ability of a fleet to meet demand. Beaujon and Turnquist explore this interaction between fleet size and utilization decisions, observing that while demand can exhibit regular changes over time, future demand forecasting is difficult and requires a management-based solution (Beaujon and Turnquist, 1991).

Since little work has been published for military fleet base and mission pairing optimization given a realistic network architecture, literature from the airline industry and for other capital assets is a vital link. In the commercial aircraft field, a large amount of work is conducted on assigning aircraft to origin-destination pairing and maintenance routing in: (Safaei and Jardine, 2017), (Abara, 1989), (Clarke et al., 1997), (Verhoeff et al., 2015), (Salazar-González, 2014), (Jansen and Perez, 2016). Sherali et al.'s review of fleet assignment work is a sufficient introduction to the field (Sherali et al., 2006).

Military aircraft are capital assets, making machine replacement studies a valuable contribution to understanding the context of this work. Sethi and Chand develop algorithms for generalized machine replacement given technological improvements through time (Sethi and Chand, 1979). Their models emphasize cost minimization but the real impact of their work is the recognition that an optimal first-period decision does not require accurate all-period forecasting. In aircraft fleet management, first-period knowledge is high but full lifecycle knowledge is low. Similarly, Narisetty et al. develop a model to optimize empty railroad freight car assignment across the Union Pacific network given first-period demand information (Narisetty et al., 2008). Hopp and Nair emphasize that using minimal forecast data for capital equipment replacement decisions could reduce future uncertainty (Hopp and Nair, 1991). Jin and Kite-Powell conclude that parallel replacement problems must be informed by first optimizing utilization levels (Jin and Kite-Powell, 2000). Only then can effective lifecycle

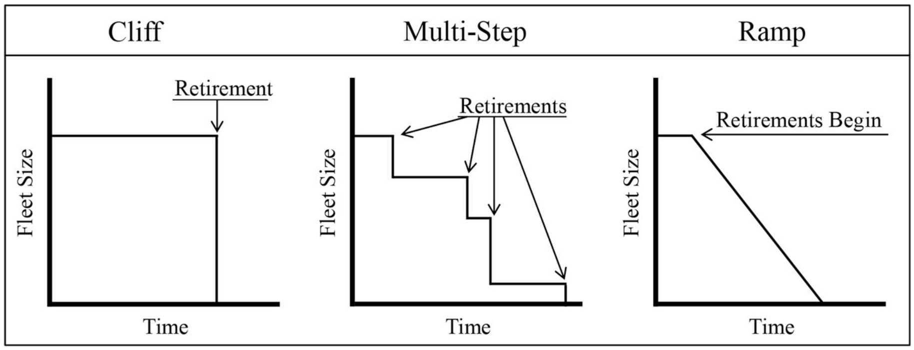

Fig. 2. Retirement philosophy patterns.

planning take place.

To account for realistic demand in a parallel replacement study, Hartman's integer programming model accepts a population of assets that have varying ages and histories (Hartman, 2000). His model contains a decision point after each period, asking whether or not each asset should be retired, based on the available lifetime. Hartman's work permits storage of unneeded assets, which is only economical for an aircraft application wherein the forecast period is multiple years. Parallel replacement decisions are generally economic decisions so utilization rates become a factor for predicting useful lifetime (Hartman, 1999).

Karabakal et al.'s work with vehicle fleet replacement illustrates the differences between serial replacement (Ramp) and parallel replacement (Cliff, Multi-Step) and shows the challenges of executing parallel replacement strategies (Karabakal et al., 1994). Karabakal's later work deals with realistically sized problems whose budget considerations force a portfolio-level perspective (Karabakal et al., 2000).

This methodology emphasizes the necessity of a fleet manager who should actively manage the military aircraft fleet. Zak concludes that despite ample mathematical tools, there is still no surrogate for a fleet's decision maker (Zhao et al., 2015). Also, life cycle cost estimation is necessary for managers to make informed decisions (Karabakal et al., 1994). These ideas are important to consider as the methodology for ROTATE is described in the following section.

3. Methodology

This optimization model assumes an available pool of capital assets at an initial state. All assets possess dissimilar utilization histories. This methodological approach assigns aircraft to bases to satisfy demand as represented in Fig. 3. Here, only two bases are shown, the first having a maximum number of seven aircraft and the second having a maximum number of four aircraft. The minimum number of aircraft are six and two, respectively. Both bases require a minimum, known number of flights of five different mission types to be flown. Actual mission types and amount flown may be greater than or equal to the required, but not less.

Simulation periods can represent any timeframe, but this paper treats each simulation period as one calendar year. In simulation period one, all aircraft are assigned to bases and to a number of missions of varying types. In each subsequent simulation period, aircraft are permitted to relocate to a different base to perform a different amount and mix of mission types. All relocations are assessed using a relocation cost, dependent on the origin-destination pairing. Actual flown EFH in a simulation period are deducted from each aircraft's remaining EFH (Sarac et al., 2006). Any aircraft that reach zero EFH are removed from the fleet. More preference is given to fly aircraft possessing higher remaining EFH than aircraft possessing lower remaining EFH because this reduces the standard deviation in EFH among the aircraft. This aligns a fleet more closely to a Cliff retirement philosophy. The Multi-Step and

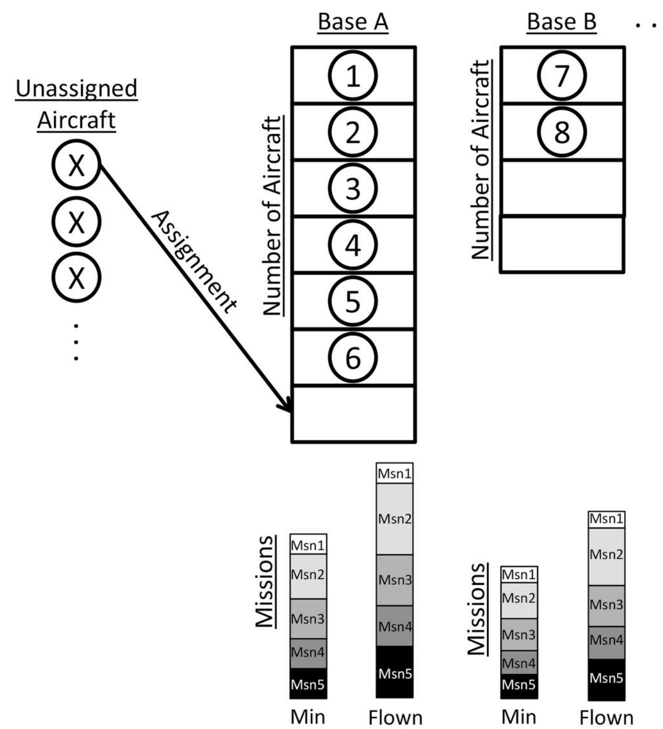

Fig. 3. ROTATE fleet assignment logic.

Ramp philosophies can also be implemented using this methodology by merely changing the denominator of the objective function.

It is important to build the methodology in a way to allow fleet managers to input their fleet's peculiarities. For example, not all aircraft in a fleet can be located at all the bases in a network nor can all aircraft fly all mission types. Realistic concerns like bases that are forecast to close in the future must also be modeled. These complex relational dependencies are formatted as matrices for the solver.

There are four core assumptions made in the formulation of this methodology:

1. Each asset is able to perform its assigned tasks during a simulation period.
2. Mission type quantity and base capacity are known.
3. Deployment usage mimics home station usage.
4. The decisions made for the fleet being studied do not impact the remainder of a larger fleet or enterprise.

These assumptions are reasonable for this application. Assumption one may fail when an aircraft is broken, which is an unpredictable event. Assumption two breaks down when an air force's defense architecture changes, but that would also dramatically impact the validity of this methodology. Knowing the number and type of missions flown is well defined through a flying hour program that budgets each unit for flying. Small variations year-to-year are very unpredictable. Assumption three is in opposition to reality. Deployment usage can actually be more docile than home station usage. Lastly, assumption four suggests that the fleet under study is in a vacuum, which is not true. However, a model that encompasses all operationally relevant aircraft missions would be untenable.

This methodology is limited by several simplifications made during its development. It is a single-period model which cannot provide an optimum solution for the entire planning horizon. This limits the applicability and usefulness of the results because the model does not account for approaching retirement dates. This methodology should be viewed as a baseline for future improvements, which can further the work and eliminate the single-period limitation. Another model limitation is that maintenance is excluded since the maintenance dataset was incomplete. Lastly, this methodology does not allow a user-defined retirement date, which would be useful to a fleet manager who is given a retirement date mandate.

The methodology is implemented using MATLAB version 2015b with all optimization tasks computed by IBM's CPLEX Optimization Studio version 12.6.3. Fig. 4 shows the flow chart for ROTATE. As shown at the bottom left of Fig. 4, retirement philosophy is an input to the methodology. While this article discusses the Cliff philosophy, the model's objective function can be changed to accommodate any desired retirement philosophy. Other inputs include information about the base network, historical aircraft utilization levels and parameters like minimum and maximum numbers of aircraft at each base. Then the mathematical model outputs where each individual aircraft should be assigned and what missions that aircraft should fly in each simulation period. The decision diamond asks if the fleet possesses enough remaining EFH to execute the required mission assignments. The loop permits simulation periods to continue until EFH runs out. When the fleet no longer has enough remaining flight hours, the simulation ends.

3.1. Mathematical formulation

This single-period mixed-integer linear programming model is formulated with the objective of changing the retirement timeline for a fleet of aircraft. As a generalized assignment problem that links assets (aircraft) to tasks (bases, mission types), the approach is formulated using two sets of decision variables. Flight hours are continuous variables but base assignments are binary, thereby making the problem harder to solve.

This work focuses on the Cliff retirement philosophy, which is achieved through the denominator of the objective function. By utilizing aircraft with higher remaining EFH, the objective function can impact the lifetime estimate of aircraft. Table 1 shows the mathematical notation used by ROTATE for each simulation period.

The mathematical formulation is outlined in Equation (1) through Equation (8). Equation (1) shows the objective function necessary for achieving the Cliff retirement philosophy. Placing the difference between CSL and initial aircraft EFH in the denominator encourages higher utilization for aircraft possessing larger remaining EFH (Sarac et al., 2006). While not addressed in this paper, the Ramp philosophy is roughly achieved either with no intervention or more precisely achieved by minimizing the delta between the current CDF slope and the desired slope. The Multi-Step is achieved by selecting similar remaining EFH groups, segregating those aircraft and implementing the Cliff philosophy for the groups.

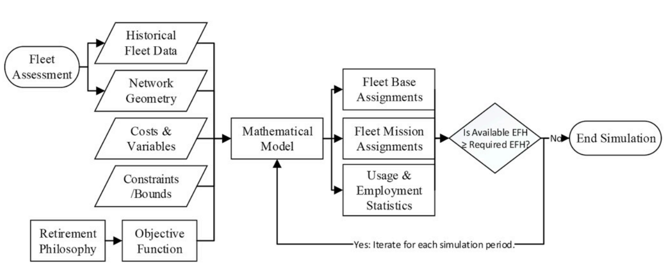

Fig. 4. ROTATE methodology flow chart.

Table 1
Mathematical notation.

Indices:	
a	index for aircraft
b	index for base
m	index for mission type
Basic Sets:	
A	fleet
B	bases
M	mission types
Decision Variables:	
L_{ba}	if aircraft a is assigned to base b
X_{mb}^{a}	number of flight hours flown of type m at base b by aircraft a
Parameters:	
AC_a^b	administrative cost for moving aircraft a to base b (translated into) EFH
$\overline{CSL}_a$	maximum certified service life in flight hours for aircraft a
EFH_a	equivalent flight hours for aircraft a
$iEFH_a$	initial equivalent flight hours for aircraft a
$\underline{FH}$	minimum flight hours per aircraft
$\underline{FH}_{mb}$	minimum flight hours for mission m at base b
$\overline{FH}$	maximum flight hours per aircraft
FHR_b^a	flight hours required for relocation of aircraft a to base b
SF_{mb}	severity factor for mission m at base b
$\underline{W}_b$	minimum number of aircraft at base b
$\bar{W}_b$	maximum number of aircraft at base b

The initial EFH is the total accumulated EFH for an aircraft, a, at the beginning of a simulation period. For a subsequent simulation period, the initial EFH is reduced by the EFH flown in the previous period and reduced by the relocation flight hours, if applicable. The administrative costs and the flight hours for relocation of each aircraft are precomputed for each simulation period, based on the location of the aircraft at the end of the previous simulation period. The assignment variable, L, is represented in the objective function to apply the relocation cost, in EFH. If the aircraft remains at the same base, both the relocation flight hours and the administrative costs are assumed to be equal to zero.

Equation (2) mandates assigned EFH to be less than the remaining EFH for a particular aircraft. The remaining EFH is the maximum certified service life less previously allocated EFH in the simulation. EFH is calculated within Equation (2) by multiplying the number of flight hours flown by the severity factor for each. Equation (3) ensures each aircraft flies within the bounds of allowed flight hours in a simulation period. Equation (4) ensures that the flight hour requirement (demand) is met for each base/mission type combination. Equation (5) links the decision variables to ensure an aircraft can only fly missions at a base if it is assigned to that base. Equation (6) bounds the number of aircraft assigned to a base and Equation (7) states that an aircraft can only be assigned to one base in each simulation period. Lastly, Equation (8) stipulates that negative flight hour assignments are not permitted.

The decision variables are:

$X_{mb}^{a}\ (continuous)$

$$L_{ba} = \begin{cases} 1, & assigned, \\ 0, & not\ assigned. \end{cases}$$

The objective function is shown as Eq. (1):

$$\min:\ Z = \sum_{a \in A} \frac{\sum_{m \in M} \sum_{b \in B} (X_{mb}^{a} \times SF_{mb}) + \sum_{b \in B} (FHR_b^a + AC_b^a) \times L_{ba}}{\overline{CSL}_a - iEFH_a} \quad (1)$$

Subject to:

$$EFH_a < \left(\overline{CSL}_a - \sum_{m\in M}\sum_{b\in B}(X^a_{mb}\times SF_{mb})\right),\ \forall\ a\in A \tag{2}$$

$$\overline{FH} \ge \sum_{m\in M}\sum_{b\in B} X^a_{mb} \ge \underline{FH},\ \forall\ a\in A \tag{3}$$

$$\sum_{a\in A} X^a_{mb} \ge \underline{FH_{mb}},\ \forall\ b\in B, \forall\ m\in M \tag{4}$$

$$X^a_{mb} \le \underline{FH_{mb}} \times L_{ba},\ \forall\ a\in A,\ \forall\ b\in B, \forall\ m\in M \tag{5}$$

$$\bar{W}_b \ge \textstyle\sum_{a\in A} L_{ba} \ge \underline{W}_b, \forall\ \ b\in\ \ B \tag{6}$$

$$\sum_{b\in B} L_{ba} = 1,\ \forall\ a\in A \tag{7}$$

$$X^a_{mb} \ge 0,\ \forall\ a\in A,\ \forall\ b\in B,\ \forall\ m\in M \tag{8}$$

4. Results and discussion

ROTATE's decision variable output shows which aircraft are assigned to each base during each simulation period and the number of flight hours assigned to each aircraft for each mission type at each base. These data are cataloged for each simulation period in the simulation. Further, the MATLAB interface calculates the number of aircraft relocations per simulation period and the standard deviation of EFH in the fleet.

This methodology is scalable to large network sizes and large fleet sizes. The scaling configurations and associated run times for sample fleet configurations are shown in Table 2. Run times are computed for a Microsoft Windows 7 machine operating dual 2.93 GHz processors with 16 GB of memory. The input data for this table are from historical USAF data (1992–2015). Because utilization forecasts do not extend far into the future, historical trends are used to project future needs. This includes aggregate numbers of missions each year and standard fluctuations from lost aircraft. The last entry represents the F-35 Joint Strike Fighter acquisition (Gertler, 2012). The USAF plans to order 1763 F-35s, making it the foreseeable natural limit for this class of problems. Assuming 15 base locations and 6 mission types for the F-35 fleet yields 185,115 decision variables.

Big O computational complexity is $O(n^2)$ due to nested iterations in the methodology. The number of decision variables is calculated by Equation (9) while constraints are calculated by Equation (10).

$$DV \propto (Aircraft \times Bases \times Missions) + (Aircraft \times Bases) \tag{9}$$

$$C \propto (Bases \times Aircraft \times Missions) + Aircraft \times Missions \tag{10}$$

4.1. Case study

The USAF's A-10 Thunderbolt II is chosen for study because it is nearing end-of-life (Pendleton and GAO, 2016). The fleet's EFH CDF in 2015 was roughly aligned with the Ramp retirement philosophy. The goal of this case study is to show that ROTATE can optimize the A-10 fleet's usage over time to produce an EFH CDF that mimics the Cliff retirement philosophy. Table 3 shows the settings for this case study. The settings are derived from A-10 fleet metrics from 2015 data provided by the USAF. The relocation cost consists of two parts: the flight hour expenditure for a relocation and an administrative cost. The flight hour expenditure is zeroed out for an aircraft remaining at its origin base. The administrative cost for this case study includes the origin base's ground inspection of the aircraft, a destination base's ground inspection and a two-hour induction sortie.

Table 2
Network scaling computation run times.

Bases	Aircraft	Mission Types	Variables	Run Time (s)
1	1	1	2	0.4397
2	8	2	48	0.5151
3	10	3	120	0.7006
4	10	4	200	0.8327
6	30	6	1260	1.392
8	50	8	3600	2.806
12	100	12	15600	17.39
20	400	20	168000	942.6
15	1763*	6	185115	3529.7

*Projected data used. Represents future F-35 fleet.

Table 3
ROTATE settings.

Parameters	Setting
Number of bases	9
Number of aircraft	283
Number of mission types	6
Max flight hours per aircraft per sim. period	504
Min flight hours per aircraft per sim. period	50
Min/Max aircraft per base, bounds	[14,84]
Administrative cost in EFH	8
Permitted moves per aircraft per sim. period	1
Severity factors for mission types	Per Data*

*Large matrices are used but are not reproduced here.

With the actual A-10 fleet architecture and future-years utilization forecast input (obtained from the USAF's Logistics, Installations, and Mission Support-Enterprise View database), ROTATE optimizes the base and mission assignment for each aircraft. For example, the output data show that aircraft X is assigned to base Y in simulation period one where it will fly Z_1 EFH of mission type Q, Z_2 EFH of mission type R and Z_3 EFH of mission type S. For this study, each simulation period represents one calendar year for the A-10 fleet.

Fig. 5 shows the remaining EFH for the A-10 case study for each aircraft in the fleet for each simulation period. The various slopes of the lines from left to right show that the methodology acts to utilize the low EFH outliers more in the first simulation periods of the simulation, within the maximum flight hours constraint. Once all aircraft possess roughly the same number of EFH (occurring between simulation periods 15 and 20), the methodology then rotates aircraft between bases and missions to continue utilizing the fleet's aircraft at similar levels. The right-side axis shows the standard deviation for the EFH of the fleet. Because the objective function seeks to minimize the variation in the CDF to result in a Cliff, the standard deviation for the fleet declines

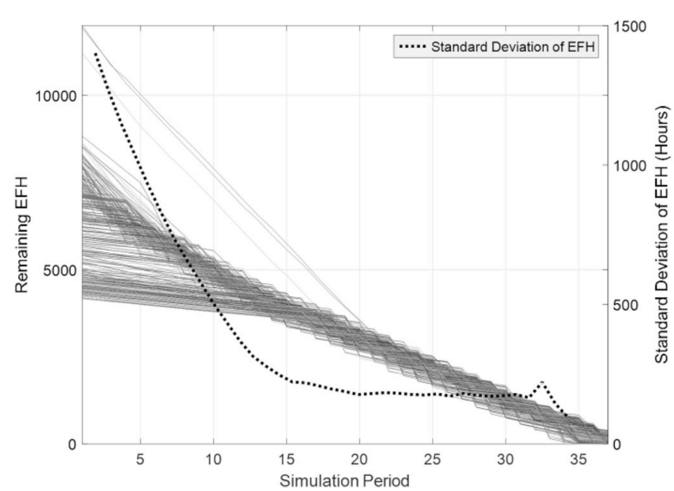

Fig. 5. Remaining EFH and standard deviation of EFH change for each simulation period.

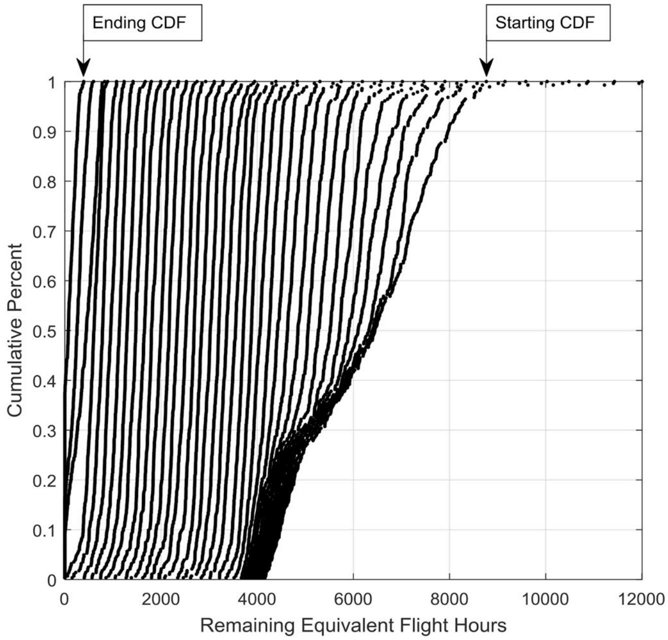

Fig. 6. CDF for each simulation period in simulation.

after each simulation period. The standard deviation begins high but then decreases as the outlier aircraft expend or conserve EFH to align more closely with the median usage rate. This phenomenon can alternatively be observed in the decreasing 'bandwidth' of the remaining EFH set of curves. Because the case study fleet uses real inputs instead of a uniform demand, there is no perfect convergence of the standard deviation to the ideal value of zero despite there being a cost for relocations. Zero standard deviation would mean that all aircraft in the fleet have the same remaining EFH, which would be a perfect match of the Cliff philosophy. Basing restrictions, aircraft model types, software versions and other network peculiarities prevent the achievement of the ideal Cliff and at times can cause spikes.

To visually check the fleet's adherence to the desired retirement philosophy shown in Fig. 2, a CDF representing each aircraft's remaining EFH can be generated. Fig. 6 shows a CDF for each simulation period produced by ROTATE. The A-10 fleet's initial EFH CDF, labeled "Starting CDF" is shown on the right. Each successive simulation period's CDF flows to the left. With the Cliff philosophy as the goal, the bulk of change occurs in this simulation in the first ten simulation periods. The bunching effect seen at the bottom of the CDFs is a visual depiction of the low remaining EFH aircraft flying the minimum number of flight hours allowed per simulation period. A vertical line would map perfectly to the desired shape from Fig. 2 but that is not achieved for aforementioned reasons.

ROTATE's ability to match the Cliff retirement shape is evaluated using the mean percent deviation between desired and achieved. Shown in Fig. 7, three sets of data are represented. "Forecast" shows how the case study fleet would look at the end of the simulation with no intervention. This assumes future utilization mimics past utilization and no network changes over time. The "Forecast" fleet size decreases because aircraft reach their maximum EFH and must be retired. "Desired" shows the ideal, benchmark retirement shape, which is Cliff in this simulation. Lastly, "Achieved" shows ROTATE's results. The "Forecast" results mimic historical patterns, are reasonable and give a mean percent deviation of 18.86%. ROTATE's "Achieved" solution reduces the deviation to 1.65%. ROTATE cannot match a desired shape perfectly for a real fleet because of the constraints inherent to the problem. In this case study, some bases required very high utilization rates of very damaging mission types. This caused a residual delta in any simulation period after rough EFH convergence is accomplished, thereby resulting

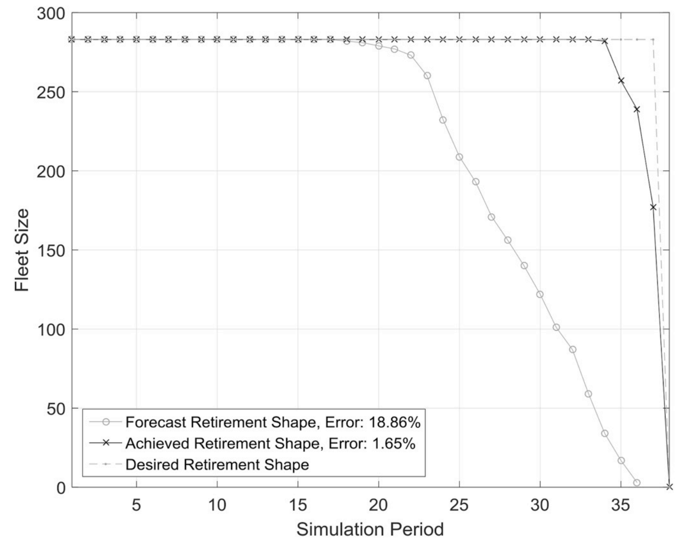

Fig. 7. Fleet size per simulation period; forecast, achieved and desired.

in non-perfect matching of the desired retirement shape.

The number of aircraft relocations represents a benefit or cost for the retirement shape improvement. While the baseline historical transfer rate is 0.1110 relocations per aircraft per year, the ROTATE solution requires only a transfer rate of 0.1061 relocations per aircraft per year. With a fleet size of 283 aircraft and a lifetime of 35 simulation periods until the fleet can no longer meet demand, ROTATE's solution for this case study reduces the number from 1099 transfers to 1050 transfers.

The managers in the USAF surveyed for this study represent fighter, attack and cargo aircraft fleets. Each believes their fleet's intricacies must be addressed in a rotation model. Aircraft models, software versions and special maintenance procedures, however, can all be modeled using this methodology's approach. While results will vary, there exists no fleet too complicated to be represented by quantitative input data.

4.2. Disruption management

Testing the methodology using sample data results in a broad study of the simulator's sensitivity. ROTATE successfully optimizes fleets within the range of reasonable inputs. ROTATE's robustness is also tested using real-world scenario inputs such as deployments, base realignments and closures, aircraft mishaps and fleet groundings. In each scenario, ROTATE is able to continue optimizing the retirement shape. This section showcases one example, represented in Table 4. Four disruption periods are chosen within which 32 randomly selected aircraft are assigned a one-year deployment that increases usage by 200 EFH. All other variables are set to the values shown in Table 3.

Fig. 8 shows the remaining EFH burndown for the entire fleet. The disruptions are clearly visible at simulation periods {6, 11, 13, 23}, represented by decreases in the remaining EFH traces as well as increases in the EFH standard deviation trace. Similarly sized disruptions have a larger impact on the fleet's EFH standard deviation when they occur in later simulation periods. This is relevant to fleet managers and

Table 4
Disruption timing, size and impact.

Simulation Period	Number of Aircraft	Impact
6	32	−200 EFH
11	32	−200 EFH
13	32	−200 EFH
23	32	−200 EFH

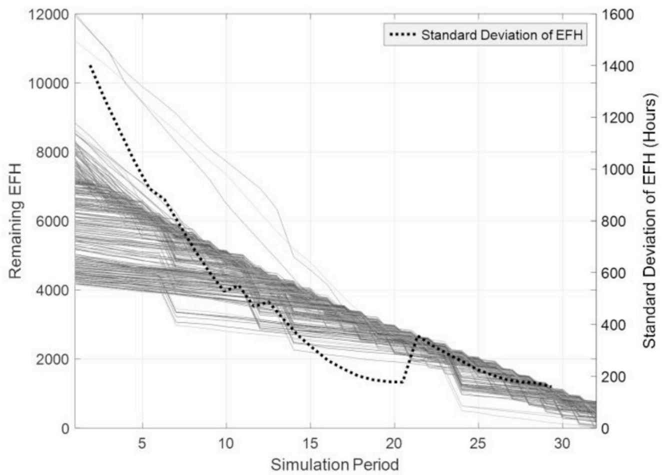

Fig. 8. Remaining EFH and EFH standard deviation changes for each simulation period with four simulated deployment disruptions.

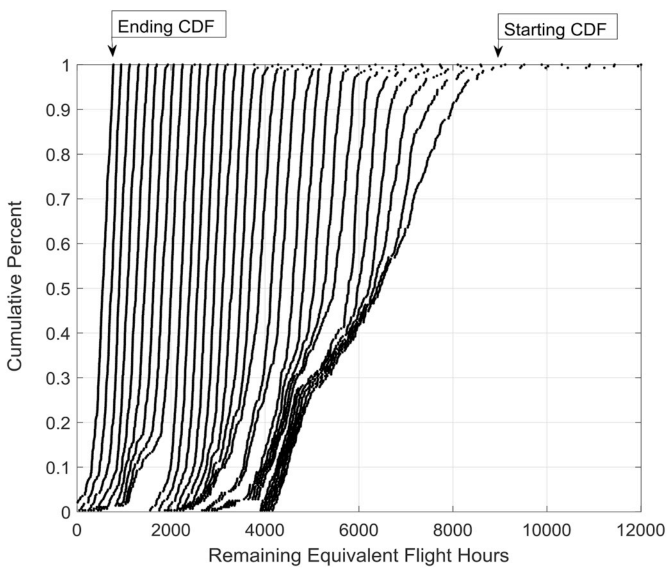

Fig. 9. CDF for each simulation period with four simulated deployment disruptions.

analysts – reducing disruption uncertainty in far-afield simulation periods can improve retirement shape convergence.

The cumulative distribution function shows a disruption as a shift in a portion of a trace. Fig. 9 shows the four disruptions for this simulation. Disruptions that involve more assets have a larger impact on the fleet and require more simulation periods for the objective function to correct the usage discontinuity.

Disruptions are not limited to deployment scenarios that increase EFH usage. Historically, some deployments have actually reduced yearly usage rates so ROTATE can also model a slower EFH accumulation rate. Base realignments and closures at future periods have the effect of shifting the demand profile for the simulation. Closures require a redistribution of the fleet, which increases transfer costs that are accounted for in this simulation. Aircraft mishaps are simulated by removing assets during the simulation. Since the demand profile is driven by the base and mission requirement inputs, any asset loss decreases the supply margin for the objective function. Lastly, fleet groundings (for impoundments, mishap investigation or time-compliant technical orders) effectively decrease usage levels in one simulation period. A secondary effect of increased usage levels in a secondary simulation period may be seen, but can also be modeled.

5. Conclusions

This work describes a methodology and tool fleet managers can use to manage aircraft flight hours within a fleet. This active management can yield specific retirement patterns for a fleet. This work focused on the “Cliff” philosophy, where the goal is to retire an entire fleet at one time. A realistic base network and forecast mission demand are used as model inputs. This methodology handles actual fleet-sized problems and effectively alters fleet CDFs to more closely mimic the Cliff retirement philosophy. At each simulation period, aircraft are permitted to relocate. The additional cost of these aircraft relocations is considered within the single-period optimization. A sensitivity analysis shows that the calculated network average relocation cost impacts the relocation frequency. A disruption management study shows this methodology's robustness despite planned or unplanned changes to fleet utilization. The A-10 fleet case study shows that ROTATE could achieve a retirement shape approaching a perfect CLIFF while decreasing the aircraft relocation frequency by a small amount from the baseline.

It is shown that this concept is feasible. Also, ROTATE is a powerful tool with which to model future usage plans. Lastly, this work shows that one can achieve a desired retirement shape within reasonable accuracy. Herein, the Cliff philosophy is proven feasible and by proxy, the Multi-Step.

This work may lead to savings for fleets either from the perspective of aligning a fleet to a retirement plan or by ensuring less useful life remains in a fleet at retirement. Better lifespan forecast information can aid decision makers in their procurement and divestment planning.

This was a proof-of-concept, so future work is needed. The next step includes applying ROTATE to a multi-period optimization problem. This would allow a fleet manager to optimize the usage and relocations for each aircraft for the remaining useful life of the fleet and could increase utilization (Haouari et al., 2012). Future work will also focus on investigating the transfer of this methodology to other fields. These ideas extend beyond fighter aircraft to fleets where similar ideas have been proposed and some are in use. Additionally, researchers interested in this topic can test the validity of this model through time with a candidate fleet of capital assets. More work can be done using the Ramp methodology, potentially implementing the Gini Coefficient from the field of economics as a quantitative measure for EFH equality.

Disclaimer

The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government. This research is funded by the United States Air Force. The authors declare no conflict of interest.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jairtraman.2019.101680.

References

Abara, J., 1989. Applying integer linear programming to the fleet assignment problem. Interfaces 19 (4), 20–28.

Başdere, M., Bilge, Ü., 2014. Operational aircraft maintenance routing problem with remaining time consideration. Eur. J. Oper. Res. 235 (1), 315–328.

Beaujon, G.J., Turnquist, M.A., 1991. A model for fleet sizing and vehicle allocation. Transport. Sci. 25 (1), 19–45.

Clark, P., 2007. Buying the Big Jets: Fleet Planning for Airlines. Ashgate Publishing, Ltd.

Clarke, L., et al., 1997. The aircraft rotation problem. Ann. Oper. Res. 69, 33–46.

Gertler, J., 2012. F-35 Joint Strike Fighter (JSF) Program. Library of Congress, Congressional Research Service, Washington, D.C, pp. 46.

Haouari, M., Shao, S., Sherali, H.D., 2012. A lifted compact formulation for the daily aircraft maintenance routing problem. Transport. Sci. 47 (4), 508–525.

Hartman, J.C., 1999. A general procedure for incorporating asset utilization decisions into replacement analysis. Eng. Econ. 44 (3), 217–238.

Hartman, J.C., 2000. The parallel replacement problem with demand and capital budgeting constraints. Nav. Res. Logist. 47 (1), 40–56.

Hopp, W.J., Nair, S.K., 1991. Timing replacement decisions under discontinuous technological change. Nav. Res. Logist. 38 (2), 203–220.

Jansen, P.W., Perez, R.E., 2016. Coupled optimization of aircraft families and fleet allocation for multiple markets. J. Aircr. 53 (5), 1485–1504.

Jardine, A.K., 2011. Optimizing Maintenance and Replacement Decisions. Asset Management Excellence, pp. 259.

Jin, D., Kite-Powell, H.L., 2000. Optimal fleet utilization and replacement. Transport. Res. E Logist. Transport. Rev. 36 (1), 3–20.

Jones, P.C., Zydiak, J.L., Hopp, W.J., 1991. Parallel machine replacement. Nav. Res. Logist. 38 (3), 351–365.

Karabakal, N., Lohmann, J.R., Bean, J.C., 1994. Parallel replacement under capital rationing constraints. Manag. Sci. 40 (3), 305–319.

Karabakal, N., Bean, J.C., Lohmann, J.R., 2000. Solving large replacement problems with budget constraints. Eng. Econ. 45 (4), 290–308.

Litvinchev, I., et al., 2010. Lagrangian heuristic for a class of the generalized assignment problems. Comput. Math. Appl. 60 (4), 1115–1123.

Narisetty, A.K., et al., 2008. An optimization model for empty freight car assignment at Union Pacific railroad. Interfaces 38 (2), 89–102.

Newcamp, J., Verhagen, W., Curran, R., 2016a. Time to retire: indicators for aircraft fleets. Int. J. Aviat. Manag. 3 (4), 221–233.

Newcamp, J., Verhagen, W., Curran, R., 2016b. Correlation of mission type to cyclic loading as a basis for agile military aircraft asset management. Aero. Sci. Technol. 55, 111–119.

Oakley-Bogdewic, L.K., Osman, E.M., 2015. An Analytic Model for DoD Divestments. DTIC Document.

Pendleton, J.H., 2016. In: GAO (Ed.), Force Structure - Better Information Needed to Support Air Force A-10 and Other Future Divestment Decisions, pp. 69 Washington D.C.

Robbert, A.A., 2013. Costs of Flying Units in Air Force Active and Reserve Components. DTIC Document.

Safaei, N., Jardine, A.K., 2017. Aircraft Routing with Generalized Maintenance Constraints. Omega.

Salazar-González, J.-J., 2014. Approaches to solve the fleet-assignment, aircraft-routing, crew-pairing and crew-rostering problems of a regional carrier. Omega 43, 71–82.

Sarac, A., Batta, R., Rump, C.M., 2006. A branch-and-price approach for operational aircraft maintenance routing. Eur. J. Oper. Res. 175 (3), 1850–1869.

Sethi, S., Chand, S., 1979. Planning horizon procedures for machine replacement models. Manag. Sci. 25 (2), 140–151.

Sherali, H.D., Bish, E.K., Zhu, X., 2006. Airline fleet assignment concepts, models, and algorithms. Eur. J. Oper. Res. 172 (1), 1–30.

Sriram, C., Haghani, A., 2003. An optimization model for aircraft maintenance scheduling and re-assignment. Transport. Res. Pol. Pract. 37 (1), 29–48.

Verhoeff, M., Verhagen, W., Curran, R., 2015. Maximizing operational readiness in military aviation by optimizing flight and maintenance planning. Transport. Res. Procedia 10, 941–950.

Zhao, X., Verhagen, W.J., Curran, R., 2015. In: Aircraft bi-level life cycle cost estimation. ISPE CE.

Dr. Jeffrey Newcamp is a Lieutenant Colonel in the United States Air Force and the Chief Test Engineer at the 413th Flight Test Squadron at Hurlburt Field, Florida, United States. He holds a PhD from the Delft University of Technology, in Delft, Netherlands and is a 2008 graduate of the United States Air Force Test Pilot School. His research foci include aging aircraft, aircraft operations and aircraft safety. Prior to his current work, Major Newcamp was an Assistant Professor of Aeronautics and Systems Engineering at the United States Air Force Academy and has been a flight test engineer for the F-35 Joint Strike Fighter. Major Newcamp holds three masters degrees, is a FAA certified private pilot and has flight time in 28 military and civilian aircraft types.

Dr. W.J.C. Verhagen is the Assistant Professor of Maintenance Operations in the Air Transport and Operations section at the Delft University of Technology. His research focuses on aircraft maintenance operations, with specific attention to development of knowledge-based maintenance systems and decision support tools to improve the efficiency of aircraft maintenance planning and execution.

Dr. Bruno F. Santos is an Assistant Professor at the Air Transport and Operations section at Delft University of Technology, in the Netherlands. His research focuses on airline operations planning problems, including fleet and network planning, operations scheduling, dynamic disruptions recovery, and aircraft maintenance scheduling. He is also part of the International Editorial Board of the Transport Policy Journal (Elsevier) and member of several research societies, such as INFORMS and WCTRS.

Professor Ricky Curran is Full Professor at Delft University of Technology and head of the Air Transport and Operations section. He is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA). He is also a member of the Economics Technical Committee, the Value Driven Design Programme Committee and the *Progress in Aerospace Sciences* Editorial Board. Among various editorial positions he is also the Editor in Chief of the *Journal of Aerospace Operations* and General Chair and founder of the Air Transport and Operations Symposium (ATOS)