Please select category below:
Normal Paper ☒
Student Paper ☐
Young Engineer Paper ☐

Optimal plans and policies for the management of military aircraft fleets

David O. Marlow, Jason R. Looker and Joyanto Mukerjee

Joint and Operations Analysis Division, Defence Science and Technology Group, 506 Lorimer Street, Fishermans Bend, Victoria, 3207, Australia

Abstract

Managing military aircraft fleets is a challenging enterprise. Fleet planners have to balance day-to-day tasking requirements with the long-term management of the fleet through to life-of-type. They also have to deal with unscheduled maintenance events, and short-turnaround high-priority tasking, which can render their short-term plans redundant. A sequence of poor short-term decisions can have cascading effects in the medium- and long-term: for example, leading to the premature retirement of individual aircraft from a fleet. In this paper we summarise the application of various optimisation techniques to help address these problems in military aircraft fleet management. The techniques are applied to both day-to-day planning and medium-term forecasting that includes unscheduled maintenance effects, and long-term fleet planning out to life-of-type including fatigue effects. Aspects of this work are being developed into a software tool for use by Royal Australian Air Force fleet planners.

Keywords: military aircraft fleet planning, unscheduled maintenance, life-of-type management, mathematical programming, sequential stochastic optimisation, data sciences.

Introduction and problem description

Military aircraft fleet management is a complex, multi-dimensional problem with often-competing priorities. Prosecuting its many aspects successfully is a challenging undertaking for any fleet planner or manager. Fleet planning decisions that seem appropriate in the short-term may have adverse longer-term consequences. While numerous decisions need to be made each day, the aim of fleet planners is to manage the fleet such that it can meet any requirement on any day of a fleet life, which is typically 20-30 years but can be longer.

All of these decisions are constrained by maintenance requirements. Each fleet follows a documented maintenance schedule such as a Technical Maintenance Plan (TMP). The TMP specifies maintenance intervals and durations for particular services, and identifies the maintenance organisations that are eligible to undertake those services. Some services may have intervals based on flying hours, others based on elapsed time, and others based on achieving whichever of those comes first. Additional services may be triggered by other factors such as engine cycles, or are condition-based. Furthermore, unscheduled maintenance is a frequent disruptor to fleet planning and can result in significant changes to any plan.

Table 1 illustrates some of the issues faced by fleet planners at various intervals.

Table 1: Fleet planning issues at various intervals

Interval	Fleet planning issues
Daily	Flying and maintenance allocations; short-interval maintenance; unscheduled maintenance; immediate high-priority tasking
Monthly	Training exercises; deployments/operations; longer-interval maintenance
Multi-month	Deployment durations; aircraft rotations between squadrons
Annually	Flying and other targets (e.g. training)
Multi-year	Modification programs; contractor maintenance; usage rates
Life-of-type	Phased withdrawal from service

There have been several notable instances when fleet management practices have led to poor outcomes such as deficient readiness levels. For a maritime example, none of the Royal Australian Navy's fleet of three amphibious warfare ships were available to provide support to the Cyclone Yasi relief operation in 2011 [1]. A battalion of US Army UH-60 Blackhawk helicopters could only deploy 9 of the 24 aircraft to support peace operations in Bosnia in 1995 despite reporting 89% of aircraft as fully mission capable [2]. An optimisation-based strategy could have saved 1.88 aircraft lifetimes (> US$83M in acquisition costs) on a 41 aircraft sub-fleet of the United States Air Force's A-10 Thunderbolt II fleet over the last several years of its life [3]. These examples both typify the lack of decision-support tools available to military aircraft fleet planners and highlight the need for same.

The military aircraft fleet management literature is still relatively young but growing. Most of the literature has focused on deterministic optimisation applications over particular timeframes, such as a daily plan lasting for a month [4], a weekly plan for 12 weeks [5], or a monthly plan for 6 months [6]. Algorithms have been developed that exploit the objective and structure for particular problems that provide exact solutions rapidly [7]. These papers only consider scheduled maintenance. Other simulation-based approaches have been developed that include multiple types of scheduled and unscheduled maintenance ([8], [9]). When applied using simulation experimental design approaches, these can provide insights into the key influences on fleet metrics. Other papers explore the application of machine learning techniques to these problems, both for simulation models [10] and optimisation models [11].

In this paper we summarise a multi-layered approach to addressing questions in military aircraft fleet planning, from day-to-day management through to life-of-type. Our exemplar fleet is the Royal Australian Air Force (RAAF)'s C-130J fleet of 12 aircraft. This fleet was chosen to demonstrate the capability and utility of our approach because it has been operational for many years and thus provides a relatively rich and mature source of data. In the following sections, we provide descriptions of our approach, targeted at different levels of the fleet planning problem; how they might be integrated; and a short outline of the current status of a software tool being developed from this work. We conclude with a brief summary and possible further applications and extensions.

Solution approaches

Due to the multi-dimensional nature of military aircraft fleet planning, we tailor our approach accordingly. Our models have been developed and refined following extensive consultation with C-130J stakeholders from Defence and industry. Both employ optimisation-based approaches. The first approach explicitly includes the impact of unscheduled maintenance in forecasting future fleet performance within a stochastic optimisation framework. The second approach considers issues of multi-year fleet planning out to life-of-type within a deterministic optimisation framework.

Managing uncertainty

We describe the incorporation of unscheduled maintenance effects into the military aircraft fleet planning space as "managing uncertainty". Aside from an earlier instantiation [12] upon which we build in this work, we are unaware of comparable work in this domain that seeks to include random effects within the framework of an optimisation approach.

The managing uncertainty model includes all types of maintenance that impact the fleet at the chosen resolution (generally, daily). This incorporates intervals, durations, induction windows, and the eligible maintenance organisations. More frequent services are typically undertaken by uniformed personnel at a squadron, while less frequent but longer duration services are typically performed by contractors. The model also includes deployments and exercises, as well as deployed maintenance and pre- and post-deployment maintenance. The treatment of deployments in the model integrates the maintenance system with the fleet's operational capability. The model is targeted to assist both squadron maintenance coordinators with short-term decisions, and wing and squadron fleet planners with longer-term trends and forecasts.

To incorporate unscheduled maintenance effects, the model uses probability distributions for intervals and durations. These distributions are derived from C-130J fleet data using data analytics techniques such as data mining and Maximum Likelihood Estimation to fit the distributions. Consequently, distributions have been fitted for each individual aircraft, and are extensible to provide distributions by mission type.

The model is formulated as a sequential stochastic optimisation problem. Figure 1 illustrates the sequential stochastic optimisation paradigm, whereby a stochastic optimisation problem is solved each day as new information arrives: in this case the realisation of unscheduled maintenance drawn from those probability distributions.

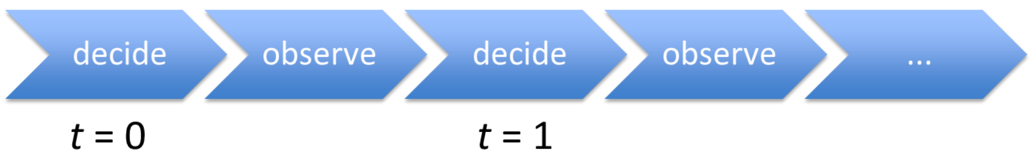

Figure 1: Sequential stochastic optimization paradigm

There are two phases to the sequential stochastic optimisation paradigm. The first is the "learning" phase, where we discover the policy (as opposed to a fixed plan) that provides the best overall outcomes against the objective over a number of simulated possible futures. Having determined this policy, the second "doing" phase implements this policy through daily decisions as in Figure 1.

There are five fundamental elements of these problems [13]: states (S_t), decisions (x_t), exogenous information (in our case, random unscheduled maintenance events), the transition function (which moves the system from state S_t to S_{t+1}), and the objective function. A typical objective function for this type of sequential stochastic optimisation problem takes the following form:

$$\max_{\pi \in \Pi} E\{\sum_{t \epsilon T} C(S_t, X_t^{\pi}(S_t)) | S_0\} \tag{1}$$

In equation (1), Π is the set of allowable policies, $C(S_t, X_t^{\pi}(S_t))$ is the contribution function (which captures the reward for making a decision x_t in state S_t) and $X_t^{\pi}(S_t)$ is a policy that returns feasible decisions x_t. Equation (1) therefore seeks to find the policy that maximises the expected value of the contribution function over the desired time horizon. The expectation is required due to the explicit inclusion of unscheduled maintenance effects in the model.

Consider an example where the objective is to determine the optimal maintenance policy that maximises long-term fleet serviceability (where a 'serviceable' aircraft is one that can fly). The "policy" options to achieve this may be the preference to "induct aircraft into maintenance early in their maintenance window" or "induct aircraft into maintenance later in their maintenance window". Each policy would then be tested over a statistically-significant number of simulated future scenarios, each providing different instances of unscheduled maintenance. In each scenario, aircraft inductions dates would follow those maintenance policies. The policy that provides the best overall outcomes against the objective (equation (1)) would be chosen. Once chosen, that maintenance policy would be implemented in real-life (i.e. a single future) fleet planning.

The managing uncertainty model can operate at multiple levels. It can provide decision support over short timescales (using the policy in the previous paragraph): which aircraft to deploy and when to deploy them; when to induct an aircraft into maintenance; when to move aircraft between contractor and squadron maintenance; and how to allocate maintenance manpower to unserviceable aircraft. It can also provide long-term forecasts of fleet performance, such as the probability of achieving a certain number of serviceable aircraft over a chosen time horizon.

For further information on the approach, the reader is referred to [13]. A more comprehensive description of the methodology will be provided in a future publication.

Managing to life-of-type

The managing to life-of-type model considers issues pertaining to multi-year fleet planning out to the fleet Projected Withdrawal Date (PWD). It is a deterministic optimisation model and thus excludes any random effects. The model operates at coarser time steps (no shorter than one week) for multi-year timeframes and is targeted to assist fleet planners at an aircraft wing. Inputs include deployment and exercise requirements (including pre- and post-deployment maintenance), stand-down periods for the workforce, TMP maintenance (intervals, durations and induction windows) at the appropriate resolution, modification programs, and flying requirements for the fleet and each squadron, as well as model parameters. Related long-term fleet planning work is more focussed on mission assignment (e.g. [3], [11]).

There may be multiple objectives for the life-of-type model, reflecting the reality of competing priorities for fleet planners. Therefore the objective function for the managing to life-of-type model takes the following form:

$$\text{Minimise } Z = \sum_{p \in P} w_p E_p + \left(\sum_{v \in V, l \in L} w_{vl}^{-} D_{vl}^{-} + \sum_{v \in V, l \in L} w_{vl}^{+} D_{vl}^{+} \right) \quad (2)$$

Equation (2) seeks to minimise variation from fleet planning targets. If all targets are met, the value of the objective function is zero. In equation (2), the E terms represent the difference from the target, the D terms represent the penalties (explained below) and the w terms represent the weights on those differences and penalties. The weights may be varied according to fleet planning priorities. The model is constructed in this way to both allow flexibility and avoid infeasibilities. Because of the multiple competing priorities, it may be simply impossible to exactly meet all targets, which would result in no solution.

There are two groups of terms in the objective function equation (2). The first group covers the set of all variables P that includes terms that reflect targets where shortfalls may arise. These may include meeting annual flying targets for a fleet and each squadron; meeting targets for the number of deployed aircraft; and flying aircraft into maintenance (i.e. such that their flying hours interval and elapsed time interval coincide). In these instances, the target may be exceeded without penalty (as in the case of flying hour targets) or be a hard limit (for maintenance inductions). The second group in parentheses covers the set of all variables V that includes terms that reflect targets that should be met precisely: i.e. not have a shortfall nor be exceeded (hence the negative and positive superscripts). These include meeting availability targets (where an aircraft is considered available if it is not in contractor maintenance). It should be noted that these requirements can vary throughout the time horizon: e.g. there may be different deployment requirements at different times of year; availability targets may be lowered during a block upgrade program; etc.

Any variation from the target can be penalised in two ways. For flying hour target shortfalls, the amount of the shortfall is used directly (the E terms in equation (2) grouped by set P). For variations around availability targets (grouped by set V), these are penalised piecewise linearly: the greater the difference, the greater the penalty [2]. This allows some variation when targets simply cannot be met exactly, but increasingly discourages greater variation from that target. The typical form of these constraints is shown in Equations (3) and (4):

$$E_v \leq \sum_l D_{vl} \qquad \forall v \in V \quad (3)$$

$$D_{vl} \leq \delta_l \qquad \forall v \in V, l \in L \quad (4)$$

Equation (3) splits the difference from the target (the E_v terms) into component terms D_{vl} based on levels l. The values of D_{vl} are constrained by δ_l in equation (4). For example, if there were a flying hour shortfall E of 100 hours, equations (3) and (4) may give D_l = 20 + 30 +50 hours for $l = 1 \dots 3$, and the respective weights w_l on each in the objective function (2) may be 0, 5 and 15. In this example, we completely tolerate a shortfall of 20 hours, somewhat tolerate a shortfall of up to 50 hours, but discourage a shortfall of greater than 50 hours.

The multi-objective nature of equation (2) can make the solution time prohibitive. To address this, a user may wish to omit some terms that are not deemed suitably important. Alternatively, equation (2) can be addressed heuristically by first solving for the highest-priority objective(s), and then using those results as inputs to solve again for other objectives. The latter approach is generally adopted as it directly reflects fleet planner priorities and is faster to execute.

PWD management is a key aspect of the life-of-type model. Analogous to methods used in previous papers (e.g. [2], [4]), it pre-calculates desired positions (based on targets) for the fleet at the end of the chosen time horizon, and allocates each aircraft uniquely to one position. The targets reflect some measure of overall aircraft usage such as total flying hours. The user can choose the positions based on achieving some desired distribution and/or range. The intent is to achieve a phased withdrawal of a fleet from service at PWD to accommodate the phased introduction of a new fleet. This will avoid individual aircraft retiring prior to PWD through over-utilisation, or conversely having aircraft still with hours remaining at PWD (as in [3]).

In accordance with meeting the objectives, the model generates decisions regarding: which aircraft to fly in each time period, how much they fly, and the squadron to which they are allocated; which aircraft to send on deployments or exercises, when to induct aircraft into maintenance given their maintenance window; and the schedule for a modification program or programs. It thus provides an optimal flying, deployment, maintenance and modification plan against the chosen objectives each time period for the chosen time horizon (assuming zero unscheduled maintenance).

Integration of the two models

The models both utilise two types of inputs. One type is general input data such as TMP information, flying limits, deployment and exercise requirements, etc. The second type is the initial status of the fleet: the current location of each aircraft, the elapsed time and flying hours since its last service of each maintenance type, etc. The models are designed to include this information so that it can produce outputs based on real time periods. Thus the models can include actual dates for stand-downs and exercises as inputs, and produce actual maintenance induction or deployment dates, or financial year-reporting for flying hours achieved.

Table 2 compares the modelling approaches. Clearly there is overlap between what the models include and what they can produce. While they are targeted at different questions in the broader fleet planning space, there is commonality in such areas as the choices of aircraft to deploy and maintenance inductions.

Table 2: Comparison between modelling approaches

	Uncertainty model	Life-of-type model
Resolution	Daily	Weekly-monthly
Time horizon	Up to a few years	Up to life-of-type (20-30 yr)
Scheduled maintenance	Intervals from daily to longest	From weekly to longest
Unscheduled maintenance	Yes	No
Maintenance organisations	All	Only contractor
Deployments/exercises	Yes, including maintenance	Yes
Modifications	Can be fixed to test impacts	Schedule output from model
Usage rates	Yes (not optimisation output)	Yes (optimisation output)

There are various ways that the models can be utilised together to provide a "complete" picture for fleet planners. One option is the following:

- Run the life-of-type model to generate a "plan" that incorporates all aspects of multi-year planning such as deployments, maintenance, modifications and PWD management;
- Using the outputs from the life-of-type model as a guide, run the uncertainty model to provide a "sanity check". This may reveal, for example, that an aircraft identified for deployment by the life-of-type model may not be the best candidate due to expected unscheduled maintenance;
- Re-run the life-of-type model, using the most-likely results from the "sanity check" uncertainty model run as fixed inputs.

If such an approach is taken, greater weighting should be placed on likely outcomes in the immediate future when utilising them in the life-of-type model. The further into the future, the less reliable the forecast, especially if there are very high levels of uncertainty.

Current status

The managing uncertainty model is currently being converted into a software tool for usage by the C-130J fleet planning community. The tool, tentatively named Delphi, is being developed by the Augmented Aviation Asset Intelligence (A3I) program within the Air Domain Centre of Capability Acquisition and Sustainment Group (CASG).

The development of the Delphi system is being conducted in partnership with industry and Chief Information Officer Group (CIOG) to enable rapid deployment of Delphi onto Defence networks. Additional stakeholders include other Defence groups and Defence industry.

Conclusion and future work

This paper has described how optimisation models have been used to address the range of problems faced by fleet planners, from day-to-day through to life-of-type. We have summarised how the two models operate: both independently, and how they may be used in conjunction.

Future work will extend the models to larger fleets and different fleet types (than transport aircraft) to test the robustness of the methods to these different circumstances. Other extensions may include fleets in different domains (land and maritime), as well as to further sub-problems within the fleet planning space, such as maintainer-to-task allocation and aircrew-to-aircraft assignment.

As we have described, the managing uncertainty model explicitly incorporates unscheduled maintenance within its modelling framework, using observed unscheduled maintenance data from the fleet under consideration. The methodology therefore examines how best to respond to unscheduled maintenance events. We do not attempt to use this work to predict when individual unscheduled maintenance events may occur (e.g. as for US Navy platforms in [14]). However, the impact of alternative maintenance philosophies can be tested within the modelling framework: e.g. when changing the elapsed time or flying hour interval between services. It could also be used to make predictions on the best fleet maintenance policies to use during different periods of the bathtub curve [15]. If the intent is to determine the optimal value of the “probability of failure” at which aircraft in a fleet should be inducted into unscheduled maintenance in order to maximise overall serviceability, an alternative formulation would be required.

We also treat an individual aircraft as our ‘quantum’, meaning that we do not consider specific failures of various sub-components such as parts of the airframe or the engine. Rather, any failures are aggregated to the aircraft level. This does not preclude future work from pursuing the viability of using these concepts at the sub-component level. Before proceeding however, the trade-off should be explored between any benefits obtained by operating at that level against the costs of increased model complexity.

References

1. Oakes, D., “Neglect of amphibious fleet triggers Defence inquiry,” *The Sydney Morning Herald*, February 26, 2011.
2. Pippin, B., “Allocating flight hours to Army helicopters”, M. Sc. Thesis, Naval Postgraduate School, 1998.
3. Newcamp, J.M., Verhagen, W.J.C. and Curran, R., “Validation of the SmartBasing aircraft rotation and retirement strategy”, *CEAS Aeronautical Journal*, Vol. 10, 2019, pp. 875–883.
4. Marlow, D.O. and Dell, R.F., “Optimal Short-Term Military Aircraft Fleet Planning”, *Journal of Applied Operational Research*, Vol. 9, No. 1, 2017, pp. 38-53.
5. Hahn, R. and Newman, A., “Scheduling United States Coast Guard helicopter deployment and maintenance at Clearwater Air Station, Florida”, *Computers & Operations Research*, Vol. 35, 1829-1843, 2008, pp. 1829-1843.

6. Kozanidis, G., Liberopoulos, G. and Pitsilkas, C., “Flight and maintenance planning of military aircraft for maximum fleet availability”, *Military Operations Research*, Vol. 15, No. 1, 2010, pp. 53-73.
7. Gavranis, A. and Kozanidis, G., “An exact solution algorithm for maximising the fleet availability of a unit of aircraft subject to flight and maintenance requirements”, *European Journal of Operational Research*, Vol. 242, 2015, pp. 631-643.
8. Mattila, V., Virtanen, K. and Raivio, T., “Improving maintenance decision-making in the Finnish Air Force through Simulation”, *Interfaces*, Vol. 38, No. 3, 2008, pp. 187-201.
9. Marlow, D.O., Sanchez, S.M. and Sanchez, P.J., “Testing policies and key influences on long-term aircraft fleet management using designed simulation experiments”, *Military Operations Research*, Vol. 24, No. 3, 2019, pp. 5-25.
10. Mattila, V. and Virtanen, K., “Scheduling fighter aircraft maintenance with reinforcement learning”, *Proceedings of the 2011 Winter Simulation Conference (WSC)*, Phoenix, 2011, Jain, S., Creasey, R. R., Himmelspach, J., White, K. P. and Fu, M., Eds, pp. 2535-2546.
11. Pescheira, F., Dell, R., Royset, J., Haït, A., Dupin, N. and Battaïa, O., “A novel solution approach with ML-based pseudo-cuts for the Flight and Maintenance Planning problem”, *OR Spectrum*, Vol. 43, No. 2, 2021, pp. 635-664.
12. Looker, J.R., Mak-Hau, V. and Marlow, D.O., “Optimal policies for aircraft fleet management in the presence of unscheduled maintenance”, *MODSIM2017, 22nd International Congress on Modelling and Simulation*, Hobart, 2017, Syme, G., Hatton MacDonald, D., Fulton, B. and Piantadosi, J., Eds, pp. 1392-1398.
13. Powell, W.B., *Approximate dynamic programming: solving the curses of dimensionality*, Wiley Series in Probability and Statistics (2nd ed.), Hoboken, New Jersey, 2011.
14. Byington, C.S., Roemer, M.J., Kacprzymki, G.J. and Galie, T., “Prognostic enhancements to diagnostic systems for improved condition-based maintenance”, *Proceedings of the IEEE Aerospace Conference*, 2002, pp.6:2815-6:2824.
15. Marlow, D.O., Sanchez, S.M. and Sanchez, P.J., “Assessing the impact of the bathtub curve failure rate on fleet performance using designed simulation experiments”, *MODSIM2017, 22nd International Congress on Modelling and Simulation,* Hobart, 2017, Syme, G., Hatton MacDonald, D., Fulton, B. and Piantadosi, J., Eds, pp. 716-722.