

Optimal policies for aircraft fleet management in the presence of unscheduled maintenance

J.R. Looker^a, **V. Mak-Hau**^b and **D.O. Marlow**^a

^a *Joint and Operations Analysis Division, Defence Science and Technology, Victoria*

^b *School of Information Technology, Deakin University, Victoria*

Email: jason.looker@dst.defence.gov.au

Abstract: This paper presents a fleet management model for military aircraft that generates daily flying and maintenance allocations. This model has two novel features that distinguish it from previous models of fleet management in the literature. Most notably, random unscheduled maintenance is integrated with an optimisation approach. This is important, as unscheduled maintenance is a regular occurrence in day-to-day military aircraft fleet management, thus limiting the ability of deterministic models to provide useful guidance. Furthermore, the model is designed to discover policies that dynamically adapt to random events to deliver optimal fleet serviceability. Other features of the model include flying-hour based phased maintenance with variable induction (day when maintenance begins), day-based regular inspections, and separate phased and flight-line maintenance facilities with both manpower and line capacity constraints.

In the model, serviceable aircraft fly up to a specified number of hours per day. Once a certain number of flying hours has been achieved, those aircraft must begin phased maintenance within an allowable induction range. Regular inspections are an additional type of scheduled maintenance based on elapsed time. Phased maintenance typically takes a few hundred maintenance man hours, whereas regular inspections take tens of man hours. Two types of maintenance facilities are included. Phased maintenance can only be undertaken on the phased maintenance line, and regular inspections are typically undertaken on the flight line. If a regular inspection is due at the same time as a phased service, these can be performed concurrently. Unscheduled maintenance is undertaken in flight line maintenance if incurred on serviceable aircraft, or in either facility if discovered during the course of scheduled maintenance.

We use Approximate Dynamic Programming (Powell, 2011) to formulate the decision making model. The objective is to find the policy that maximizes the expected value of the contribution function. Our contribution function seeks to maximize the flying hours achieved and maintenance throughput of the fleet over a time period, with a weighting factor balancing these priorities. Decision variables include the allocation of flying hours and maintenance man hours per aircraft per day, and when aircraft are inducted into phased maintenance. The model consists of 75 sets of constraints describing: the allocation of flying hours to serviceable aircraft and maintenance man hours to aircraft in maintenance; the allocation of different maintenance types to aircraft in different maintenance facilities; and state transition tests to determine when aircraft are eligible or required to begin or end types of maintenance. The unscheduled maintenance constraints include an exogenous stochastic variable that represents the daily amount of new unscheduled maintenance generated randomly from a lognormal probability distribution. Aircraft are allowed to carry forward a limited amount of unscheduled maintenance while remaining serviceable. However, once a specified level of unscheduled maintenance is exceeded, the aircraft must undergo maintenance at an appropriate maintenance facility until the outstanding unscheduled maintenance is reduced to a required level.

In Approximate Dynamic Programming, a policy is a function that returns a decision, given the current state of the fleet (in the present context). We consider a deterministic look-ahead policy that depends on a discount factor and a look-ahead time horizon. We undertake computational experiments with a squadron of 12 aircraft over a 180 day period, with time horizons ranging from 1 to 14 days and discount factors from 0.8 to 1. Indicative simulation results are presented, where the time horizon substantially dominates the discount factor with respect to influence on the mean total flying hours achieved. While longer time horizons produce larger optimal objective values for a given parameter set and random data, simulation results may not significantly differ when compared with shorter time horizons due to random variation. This is apparent in our results, demonstrating the potential for substantial savings in computational time.

Keywords: *Aircraft fleet management, computational stochastic optimization, unscheduled maintenance*

1. INTRODUCTION

1.1. Military aircraft fleet management

Managing a fleet of military aircraft is an inherently complex task. The life of a military aircraft fleet is typically 20-30 years. Over the fleet life, there are numerous requirements to be met, both by the whole fleet and for the individual squadrons within the fleet. Every year, each squadron will need to meet a number of flying requirements, for instance, to train aircrew and meet operational requirements. Furthermore, if there are emergencies such as natural disasters, governments may call on the military to assist, at which time it is expected that the military will be able to respond as requested. Consequently there is an implicit requirement for squadrons to have a sufficient number of aircraft available to be deployed at short notice in such situations.

These flying requirements need to be achieved while also meeting the fleet's maintenance requirements. The scheduled maintenance regime consists of several types, based on both elapsed time and achieved flying hours. Regular inspections are both time-based (e.g., daily, weekly, etc.) and flying hours-based and usually take 1-2 days to complete; only time-based regular inspections are considered here. Phased maintenance is flying hours-based and may take a few hundred maintenance man hours (MMH) to complete, equating to a period of days or weeks. Finally, deep or depot maintenance usually occurs every few years or several hundred flying hours, and may take weeks or months to complete; this is not considered in the present model. In addition, unscheduled maintenance arises regularly to disrupt the plans made by fleet managers. Unscheduled maintenance requirements may be discovered before an aircraft is scheduled to fly, thus rendering the aircraft unserviceable, or it may be discovered during the course of scheduled maintenance, thus delaying the return of the aircraft to a serviceable state. Different maintenance facilities handle the different maintenance types, and some types of services may be combined. Given all of these factors, fleet planners and managers make daily decisions about which aircraft to fly and how much, and which to maintain and how much in order to meet ongoing requirements.

1.2. Literature review

The literature in this area is relatively new compared with civilian applications. Civilian and military aircraft fleet managers have different objectives and constraints. Civilian aircraft fleet managers are typically concerned with maximizing profits, with aircraft flying point-to-point between various cities. The objectives in military aircraft fleet management are more concerned with maximizing aircraft availability or the number of flying hours achieved, with decisions regarding which aircraft to fly each day and how much.

Most of the literature on military aircraft fleet management concentrates on purely deterministic applications, using mathematical programming techniques to develop a flying and maintenance plan over a given time horizon. A Masters thesis from the Naval Postgraduate School (Pippin, 1998) developed a mixed integer-linear program (MIP) in order to minimize the deviation from the 'goal line' for a US Army helicopter fleet. The 'goal line' is a visual planning aid that tracks and sorts individual aircraft and their flying hours remaining until their next phased service, to enable a steady flow of aircraft into maintenance. This has been extended (e.g., Kozanidis *et al.*, 2012) by seeking to generate flying and maintenance plans over multiple time horizons (here, a monthly plan for six months) while maximizing aircraft availability (i.e., the total fleet hours remaining until phased maintenance). These authors (Gavranis and Kozanidis, 2015) then exploited the limited number of aircraft combinations that can begin and end phased maintenance to develop an exact solution to this problem. A US Coast Guard study (Hahn and Newman, 2008) used a MIP to generate an optimal weekly plan for 12 weeks for a helicopter fleet. Phased maintenance is included and helicopters may be deployed to two separate locations. The objective is to minimize helicopter movements (e.g., between deployed and non-deployed) with penalties for not meeting flying requirements. Other papers apply MIPs to plans that consider daily requirements: one paper (Cho, 2011) develops a twice-daily plan with two fixed-sortie types for 15 aircraft over a one-year horizon while minimizing the maximum number of aircraft in maintenance; while another paper (Marlow and Dell, 2017) develops a daily plan for various time periods (20 and 30 days) and aircraft numbers (12 and 24), incorporating both phased maintenance and regular inspections with the aim to balance multiple requirements, including daily and total flying hours requirements as well as a type of 'goal line' maintenance stagger at specified days.

Other work in this field has included unscheduled maintenance. The initial focus (Mattila *et al.*, 2008) has been developing discrete-event simulation models incorporating various types of scheduled maintenance, with the focus on average availability. These authors next apply a reinforcement learning approach (Mattila and Virtanen, 2011) in order to either maximize the availability over an infinite time horizon, or exceed a

target availability level over a finite time horizon, in conflict conditions. They test various optimization policies based on determining the best time to induct an aircraft into maintenance (which may include delaying entry). They also seek to develop a decision support tool (Mattila and Virtanen, 2014) combining their discrete-event simulation model with a multi-objective simulated annealing algorithm, with the aim to both maximize average aircraft availability and minimize deviation from the desired maintenance schedule.

1.3. Aim of this work

The literature review describes deterministic MIP models used to generate flying and maintenance plans. When unscheduled maintenance arises, these plans will have to be adjusted and the shorter the planning horizon, the greater the impact on these plans. The simulation models surveyed incorporate unscheduled maintenance, but do not focus on generating flying and maintenance plans.

Hence the primary aim of this work is to present a model that seeks to optimally manage a military aircraft fleet while explicitly including unscheduled maintenance. We therefore require an approach that combines simulation (to model random unscheduled maintenance) and mathematical programming in order to determine optimal policies for an aircraft fleet, similar to the testing of maintenance policies described previously (Mattila and Virtanen, 2011). Such a model can provide fleet planners with the ability to make decisions that improve outcomes while explicitly incorporating uncertainty.

All features of the present model cannot be described in this paper due to space limitations; instead, the key elements are introduced.

2. MODEL DESCRIPTION

We consider a single squadron containing a fixed number of aircraft operating in peacetime. Aircraft may fly up to a maximum number of hours per day when they are serviceable. Unserviceable aircraft (i.e., those in scheduled or unscheduled maintenance) are unable to fly.

We include two types of scheduled maintenance: regular inspections (RI) (based on elapsed time) and phased maintenance (PM) (based on flying hours). These maintenance events may be combined where possible. Aircraft become unserviceable when a RI is due, whereas aircraft may be inducted into PM with some flexibility around the specified flying hours since the previous PM. Depot maintenance and services that can be completed within a day, and not affect the flying program, are not included in the model.

The different types of maintenance are performed at separate maintenance facilities. PM occurs at the phased line maintenance (PLM) facility, while regular inspections occur at the flight line maintenance (FLM) facility. When PM occurs at the PLM facility, RI (if due) may also be combined with the phased service. Scheduled maintenance must be completed in blocks; i.e., once work on scheduled maintenance begins, it cannot stop until the maintenance is complete. Aircraft queue for maintenance if there is no spare capacity at a maintenance facility.

Unscheduled maintenance (UM) has two components: time between failures and time to repair (duration). Both are modelled as lognormal probability distributions, based on previous analysis of aircraft data (Marlow and Novak, 2013). UM may be undertaken at the FLM facility (either when an aircraft fails, or is combined with a regular inspection) or at the PLM facility (either when an aircraft that is eligible for PM fails, or when UM is discovered on that aircraft during a phased service).

New UM MMH requirements for each aircraft are revealed at the start of each day according to the input probability distributions. Aircraft are permitted to accumulate a specified amount of carried forward unserviceability (CFU) in MMH while remaining serviceable. However, once the CFU MMH threshold is exceeded, these accumulated MMH become the UM requirement and the aircraft must undergo maintenance in order to return to a serviceable state. All UM requirements are considered equally important in the present model, and hence urgent UM requirements cannot be identified.

3. MATHEMATICAL METHOD AND MODEL FORMULATION

3.1. Approximate dynamic programming

We use Approximate Dynamic Programming (ADP) (Powell, 2011) to formulate the decision making model. ADP is a family of methods for solving sequential stochastic decision making problems by finding the best policy (decision function) from a given class of policies, with respect to an objective function. The present ADP formulation incorporates aspects of mathematical programming, including an objective function,

constraints and decision variables. In addition, it contains state variables (that hold the information needed to make a decision), a transition function (that governs how, in this case, the system of constraints transition the state variables from their current values to their next at each time step), and exogenous information (in this case, the unscheduled maintenance).

The objective function is

$$\max_{\pi \in \Pi} \mathbb{E} \sum_{t=0}^{|T|} C(S_t, X_t^\pi(S_t)),$$

where $C(S_t, X_t^\pi(S_t))$ is the contribution function and $X_t^\pi(S_t)$ is the policy (decision function) indexed by π in the admissible set of policies for state S_t . There are many different policy types (choices of admissible policies), including policy function approximations, value function approximations, myopic policies, and deterministic and stochastic look-ahead policies (Powell, 2014).

3.2. Formulation

The model is described by (for sets of aircraft $a \in A$ and time $t \in T$ in days):

- **Decision variables**, which take two possible forms. Firstly, $x_{a,t}^e$ is used for continuous variables, where the ‘ e ’ superscript can be: f when representing aircraft able to fly; pm , plm,ri and plm,um for aircraft in phased maintenance (undergoing PM, RI or UM respectively); and flm,ri and flm,um for aircraft in flight line maintenance (undergoing either RI or UM). Secondly, $y_{a,t}^e$ represents binary variables, e.g., $y_{a,t}^{en,pm}$ is 1 when an aircraft enters phased maintenance and 0 otherwise. The caret symbol (\wedge) indicates the decision variable takes affect at the end of the day; otherwise it takes affect at the beginning of the day.
- **State variables**, which encode the information required at a particular time t to model a system from that time onwards. These also take two possible forms: $R_{a,t}^e$ for continuous variables, e.g., $R_{a,t}^{f,pm}$ for the flying hours accrued since the last PM; and $S_{a,t}^e$ for binary variables, e.g., $S_{a,t}^{pm}$ is 1 if an aircraft requires PM and 0 otherwise. A caret symbol indicates the state variable is measured (or observed) at the end of the day, otherwise it is measured at the beginning of the day.
- **Parameters**, where c^i represents continuous values with units, and n^i discrete dimensionless values.

Due to space limitations, we are unable to show the full formulation in this paper. However, we will endeavour to describe the key elements. The contribution function is

$$C(S_t, x_t) = \frac{1}{|T|} \sum_{a \in A} \left(\frac{w^f}{c^f} x_{a,t}^f + \frac{1}{c_{\max}^{plm}} (x_{a,t}^{pm} + x_{a,t}^{plm,ri} + x_{a,t}^{plm,um}) + \frac{1}{c_{\max}^{flm}} (x_{a,t}^{flm,ri} + x_{a,t}^{flm,um}) - \frac{\mathcal{E}}{2} (B_{a,t}^{bl,plm} + B_{a,t}^{bl,um}) \right)$$

The contribution function is a weighted sum of flying hours (first term in parentheses with weighting w^f), MMH (second and third terms), and a penalty for not performing PLM and UM on an aircraft on consecutive days. All terms are scaled such that they are dimensionless, e.g., c^f is the maximum flying hours per aircraft per day. Hence the objective function maximizes both the flying hours and MMH achieved. The MMH terms remove a symmetry that arises when aircraft are unserviceable beyond the solution time horizon.

The model contains $13|A||T|$ decision variables, $28|A||T|$ state variables, and $O(75|A||T|)$ constraints.

These constraints may be divided into three main categories:

- **Flying hour and MMH allocations** that limit the allocated flying hours or MMH to within specified ranges. For phased maintenance, these are $c_{\min}^{plm} y_{a,t}^{pm} \leq x_{a,t}^{pm} \leq c_{\max}^{plm} y_{a,t}^{pm}$, where $y_{a,t}^{pm}$ is a binary decision variable that takes the value 1 when phased maintenance is being undertaken. Extra constraints model the combinations of maintenance types at the allowed maintenance facilities.
- **Maintenance facility allocations** that determine when and where maintenance can be undertaken. The constraints are generally of the form $y_{a,t}^{pm} \leq S_{a,t}^{plm}$ that state an aircraft undergoing PM must be on a phased maintenance line (PLM). Due to the allowable combinations of maintenance at different facilities, more complicated constraints are used: e.g., $S_{a,t}^{flm} \leq S_{a,t}^{ri} + S_{a,t}^{um}$ states that an aircraft must require at least one of

an RI or UM to be in FLM, and $S_{a,t}^{flm} + S_{a,t}^{plm} + S_{a,t}^s \leq 1$ forbids an aircraft from being in multiple facilities (and serviceable) simultaneously. Additional constraints track maintenance capacity to ensure this is not violated at the different facilities.

- **State transitions** that determine when aircraft enter (*en*) and exit (*ex*) the state of requiring different types of maintenance. For PM, $\hat{R}_{a,t}^{f,pm} \leq \hat{R}_{a,t-1}^{f,pm} + x_{a,t}^f$ and $\hat{R}_{a,t-1}^{f,pm} + x_{a,t}^f \leq \hat{R}_{a,t}^{f,pm} + c^P \hat{y}_{a,t-1}^{en,pm}$ track the flying hours since the last PM. The constraints $\hat{y}_{a,t}^{en,pm} \leq \hat{R}_{a,t}^{f,pm} / c^P$ and $\hat{R}_{a,t}^{f,pm} \leq c^P + c^f (\hat{y}_{a,t}^{en,pm} - 1)$ then work together to allow an aircraft to begin PM when $c^P \leq \hat{R}_{a,t}^{f,pm} \leq c^P$, where c^f is the minimum flight duration immediately before entering PM. The constraints take a similar form for ending PM and beginning and ending RI. Other constraints of the form $S_{a,t}^{pm} = S_{a,t-1}^{pm} + \hat{y}_{a,t-1}^{en,pm} - \hat{S}_{a,t-1}^{ex,pm}$ ensure aircraft can only be in allowable states at allowable times, such as when they begin and end maintenance.

A significant feature of the formulation is the model of UM. Three representative UM constraints are:

$$\hat{R}_{a,t}^{um} = \hat{R}_{a,t-1}^{um} + \hat{W}_{a,t-1}^{um} - x_{a,t}^{plm,um} - x_{a,t}^{flm,um} \quad \forall a \in A, t \in T$$

$$\hat{R}_{a,t}^{um} \geq c^{cfu} S_{a,t}^{en,um} \quad \forall a \in A, t \in T$$

$$\hat{R}_{a,t}^{um} - c^{cfu} \leq (M_{a,t}^{um} - c^{cfu}) S_{a,t}^{um} \quad \forall a \in A, t \in T$$

The first set of constraints tracks the amount of UM required for each aircraft each day, where $\hat{W}_{a,t-1}^{um}$ represents the new UM MMH added to an aircraft the previous day. This is generated from a probability distribution and is thus exogenous information to the model. The second and third sets of constraints prevent an aircraft from beginning UM if $\hat{R}_{a,t}^{um} < c^{cfu}$. The third set of constraints use a ‘‘big-M’’ parameter, in part, because $\hat{R}_{a,t}^{um}$ may be very large given the right-tailed distribution used to generate the values of $\hat{W}_{a,t-1}^{um}$, although we have derived a time and aircraft-dependent tight bound. For an aircraft to end UM, $\hat{R}_{a,t}^{um} \leq c^{um}$; these requirements are enforced with constraints similar to those above.

4. INDICATIVE RESULTS

4.1. Policies and inputs for computational experiments

To demonstrate the utility of the model, we consider a deterministic look-ahead policy model (Powell, 2014), which has the form

$$X_t^\pi(S_t) = \arg \min_{x_t} \left(C(S_t, x_t) + \sum_{t'=t+1}^{t+H} \lambda^{t'-t} C(S_{t'}, x_{t'}) \right),$$

where the index t indicates a base-model variable (these are implemented at the end of day t), and t' indicates a look-ahead variable (these are discarded at the end of day t). Look-ahead policies optimize over a time horizon $(t, \dots, t+H)$. The model implements the decision for the first day, then solves again over the time horizon, and so on. Deterministic look-ahead policies use a point estimate to forecast new UM events. We assume no new events in the forecast of UM, since this performed well and had improved numerical stability, due to smaller ‘‘big-M’’ values, in comparison with using a running mean during initial testing.

We consider a squadron of 12 aircraft operating over a 180 day period. Aircraft can fly up to 6 hours per day. Aircraft must begin PM between $c^P = 180$ and $c^P = 220$ flying hours (inclusive) and require 300 MMH to return to serviceability. RI are required every 7 days for 10 MMH. There is a total of 150 MMH available each to both FLM and PLM, with a maximum of 50 MMH per aircraft per day, and a capacity of 3 maintenance lines each 4.2.for FLM and PLM. UM time between failures are drawn from a lognormal distribution with $\mu = 3$ and $\sigma = 1$, while the time to repair lognormal parameter values are $\mu = 1$ and $\sigma = 1.25$.

Aircraft must undertake UM when $c^{cfu} \geq 20$ and return to serviceability when $c^{um} \leq 5$. In our contribution

function, we set $\varepsilon = 0.01 = 0.5(c_{\min}^{plm} / c_{\max}^{plm}) = 0.5(c_{\min}^{flm} / c_{\max}^{flm})$ and $w^f = 10$. We initialise the squadron with 7 aircraft serviceable, 3 in phased maintenance and 2 in regular inspections, with the aircraft ‘staggered’ in each state, e.g., the 9 aircraft not in PM are all separated by 20 flying hours.

We consider two parameters for our policy search: the time horizon H (from 1 to 14 days), and a discount factor λ (ranging from 0.8 to 1). We wish to determine which combination of time horizon and discount factor produce the maximum expected flying hours, as shown in the objective function and implied by the choice of $w^f = 10$. We use the given values of H and discretize the discount factor by steps of 0.05, giving a total of 70 design points for the experiments, which were conducted on a computer with an Intel® Xeon® 2.6 GHz CPU with 64 GB of RAM running Ubuntu 16.04LTS. The model was coded in Pyomo 5.0.1 and the solver was Gurobi® 7.5.0. Using 24 cores, all 96 runs of the 14 day rolling horizon case takes approximately 90 minutes to solve with an optimality gap of 1% (for each day); cases with a smaller discount factor run faster than those with a larger value. We performed 96 runs at each design point as this is a



Figure 1. Sample output from a single instance for a single design point showing (a) the accumulated flying hours since the last PM and (b) the remaining accumulated UM MMH

multiple of the 24 cores, and results in the widths of the 95% confidence intervals being less than 5% relative to the mean.

Figure 1 provides indicative schedules from a single run (days 15-29) for one design point to demonstrate how the model works in practice. The table on the left (a) shows the values for $\hat{R}_{a,t}^{f,pm}$, while the table on the right (b) is for $\hat{R}_{a,t}^{um}$. Green represents a serviceable aircraft, blue is an aircraft in PLM, orange is an aircraft in FLM, and red is an idle aircraft, i.e., neither flying nor being maintained. Given the objective, aircraft generally fly at the maximum rate.

Figure 1(a) shows how aircraft begin PM at different values of $\hat{R}_{a,t}^{f,pm}$: aircraft A10 after 218 flying hours; A9 at 216 flying hours. RI occur on A4 on day 17, A3 on day 18, etc. The impact of UM is seen in **Figure 1(b)**: aircraft A8 has 1763 MMH of UM remaining at the end of day 16 when previously serviceable; and aircraft A10 incurs at least 573 MMH during a phased service on day 21. The UM constraints are clearly evident for aircraft A1: on day 16 it has 64 MMH of UM remaining at the end of the day, then a further 100 MMH is remaining on day 19. Aircraft with large remaining UM MMH are given a lower maintenance priority compared with aircraft requiring services that can be completed more quickly, thus generating more flying hours; this follows from the form of the contribution function.

Table 1. A summary of results for the top 6 combinations of time horizon and discount factor, and the myopic policy ($H = 0$).

Time horizon	Discount factor	Flying hours (mean)	Flying hours confidence interval width (%)
11	0.90	7355	3.38
9	0.95	7325	3.57
13	0.90	7321	3.04
14	0.80	7306	3.26
7	0.90	7303	4.05
14	0.90	7303	3.69
0	NA	3934	6.28 (576 runs)

Table 1 shows the 6 best mean total flying hours results, together with the myopic policy. Observe that the best case with a time horizon of 11 and that with a horizon of 7 are not significantly different; this indicates the potential for savings in computation time. In fact, the top 6 results are not significantly different from each other. While longer time horizons will produce larger optimal objective values for a given parameter set and random data, simulation results may not significantly differ when compared with shorter time horizons due to random variation. Additional statistical analysis indicates the length of the time horizon has a greater influence on the total flying hours than the discount factor.

5. DISCUSSION AND FUTURE WORK

We have developed a new fleet management model for military aircraft that explicitly includes random UM. The model also includes aircraft flying and scheduled maintenance and can be used to search for policies that achieve specified objectives. Incorporating UM gives this model much greater fidelity than deterministic models in the highly dynamic and uncertain environment of military aircraft fleet management.

Future work will explore the use of additional policies, such as stochastic look-ahead policies (Powell, 2014), to determine which policies best meet fleet objectives. It will also include more comprehensive tests of policy parameters, such as longer time horizons and representing the discount factor as a continuous variable. Such extensions may warrant a more thorough treatment using simulation experimental design techniques.

ACKNOWLEDGMENTS

This research has been done under a Collaborative Project Agreement between DST and Deakin University.

REFERENCES

- Cho, P. (2011). Optimal Scheduling of Fighter Aircraft Maintenance. M. Sc. Thesis, Massachusetts Institute of Technology.
- Gavranis, A. and Kozanidis, G. (2015). 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*, 242, 631-643.
- Hahn, R. A. and Newman, A. M. (2008). Scheduling United States Coast Guard helicopter deployment and maintenance at Clearwater Air Station, Florida. *Computers & Operations Research*, 35, 1829-1843.
- Kozanidis, G., Gavranis, A. & Kostarelou, E. (2012). Mixed Integer Least Squares Optimization for Flight and Maintenance Planning of Mission Aircraft. *Naval Research Logistics*, 59, 212-229.
- Marlow, D. O. and Dell, R. (2017). Optimal short-term military aircraft fleet planning. Submitted to *Journal of Applied Operations Research*.
- Marlow, D. O. and Novak, A. (2013). Using discrete-event simulation to predict the size of a fleet of naval combat helicopters. In Piantadosi, J., Anderssen, R.S. and Boland J. (Eds), *MODSIM2013, 20th International Congress on Modelling and Simulation* Adelaide: Modelling and Simulation Society of Australia and New Zealand, 2506–2512.
- Mattila, V. & Virtanen, K. (2011). Scheduling Fighter Aircraft Maintenance with Reinforcement Learning. In S. Jain, R. R. Creasey, J. Himmelspach, K. P. White and M. Fu (Eds.), *Proceedings of the 2011 Winter Simulation Conference*. Phoenix: IEEE, 2540-2551.
- Mattila, V. & Virtanen, K. (2014). Maintenance scheduling of a fleet of fighter aircraft through multi-objective simulation-optimization. *Simulation: Transactions of the Society of Modeling and Simulation International*, 90(9), 1023-1040.
- Mattila, V., Virtanen, K. and Raivio, T. (2008). Improving maintenance decision-making in the Finnish Air Force through simulation. *Interfaces*, 38(3), 187-201.
- Pippin, B. (1998). Allocating Flight Hours to Army Helicopters. M. Sc. Thesis, Naval Postgraduate School.
- Powell, W. (2011). Approximate Dynamic Programming: Solving the Curse of Dimensionality. Wiley Series in Probability and Statistics, 2nd edition. Wiley, Hoboken, New Jersey.
- Powell, W. (2014). Clearing the jungle of stochastic optimization. *Bridging Data and Decisions*, INFORMS Tutorials on Operations Research, 109-137.