

UNCLASSIFIED



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Fleet Sizing Analysis Methodologies for the Royal Australian Navy's Combat Helicopter Replacement Project

David Marlow and Ana Novak

Joint and Operations Analysis Division
Defence Science and Technology Organisation

DSTO-TR-2886

ABSTRACT

This report describes methodologies used to conduct fleet sizing analysis for the Royal Australian Navy's combat helicopter replacement project, known as AIR 9000 Phase 8. An initial analysis utilises a simple approach to predict fleet size based on the binomial distribution. The sole aim is to predict the likelihood that a certain fleet size will provide a required number of embarked aircraft. A more complex analysis requires the development of a discrete-event simulation that can also incorporate the hours flown by the fleet. This simulation model includes the representation of individual aircraft moving between embarked and ashore states and through various types of maintenance, including unscheduled maintenance. This methodology is robust and easily allows for sensitivity and trade-off analysis.

RELEASE LIMITATION

Approved for public release

UNCLASSIFIED

UNCLASSIFIED

Published by

*Joint and Operations Analysis Division
DSTO Defence Science and Technology Organisation
506 Lorimer St
Fishermans Bend, Victoria 3207 Australia*

*Telephone: 1300 DEFENCE (1300 333 363)
Fax: (03) 9626 7999*

*© Commonwealth of Australia 2013
AR-015-695
August 2013*

APPROVED FOR PUBLIC RELEASE

UNCLASSIFIED

UNCLASSIFIED

Fleet Sizing Analysis Methodologies for the Royal Australian Navy's Combat Helicopter Replacement Project

Executive Summary

This report describes the approaches used to analyse the fleet-sizing requirements for the Royal Australian Navy's combat helicopter replacement project, known as project AIR 9000 Phase 8. This project seeks to provide the Royal Australian Navy with a ship-based capability in the roles of anti-submarine and anti-surface warfare. This will replace the capability currently provided by the S-70B-2 Seahawk helicopter fleet. It will also fill the capability gap which emerged due to the cancellation of the SH-2G Super Seasprite. Two approaches were developed for use during the project.

Initially, the sole requirement from the project was to determine the minimum fleet size that would be able to embark at least the minimum number of aircraft on ships at sea. A simple method using the binomial distribution was employed to address this question. This allows the user to determine the probability that a certain fleet size would provide the minimum number of embarked aircraft, subject to a certain level of serviceability and availability in the fleet. This approach was suitable but only for the question posed and made many assumptions. One of the key assumptions was that it treated embarked and ashore aircraft as having the same level of serviceability, which is not the case in practice. Other questions could not be addressed using this method, including the number of flying hours that a given fleet could achieve.

Consequently, a more robust methodology was required. This method needed to explicitly determine the minimum fleet size that could embark the minimum aircraft throughout the fleet life, while simultaneously meeting annual requirements for a minimum number of embarked and ashore flying hours.

The method chosen was a discrete-event simulation. This is able to represent each individual aircraft as being in a particular state at a given time step. State transitions occur when certain criteria are met, such as the number of flying hours until a regular maintenance check is due, or the number of maintenance man hours until that service is completed. This method can explicitly handle random events, particularly unscheduled maintenance, which distinguishes it from other fleet-sizing approaches such as integer programming which generally assume fixed conditions. Being able to

UNCLASSIFIED

UNCLASSIFIED

track the state of each aircraft every day allows a check to be performed on whether the minimum number of aircraft are embarked each day across the fleet life. Once daily flying rates are determined for embarked and ashore aircraft, the method can also test whether the annual requirements are being met.

The model initialises the individual aircraft (called tails) in the fleet into various states: embarked and serviceable, ashore and serviceable, and in various types of scheduled maintenance ashore. For each day of the fleet life, it then updates the tails in their various states and tests for state changes: e.g. from serviceable to scheduled maintenance; from ashore serviceable to embarked serviceable, etc. Distinctions are made between embarked and ashore crews, with embarked personnel working every day of the year, while ashore personnel do not work on weekends or during holiday periods. Serviceable aircraft ashore may be detached to exercises at various stages of the year.

Tails may enter regular, phased or deep maintenance once a certain number of flying hours or elapsed time has been achieved. Unscheduled maintenance may occur on serviceable aircraft or those undergoing phased maintenance. Unscheduled maintenance frequency and duration were both found to follow a log-normal distribution following analysis of historical data. Therefore unscheduled maintenance occurs randomly in the model based on the parameters of this distribution, with the duration also randomly determined. Logistics delays are added to the unscheduled maintenance duration. Each type of maintenance has a fixed number of maintenance lines. If capacity is exceeded, tails must wait for a line to become free in a first-in first-out queue. Where applicable, modifications may be undertaken following completion of scheduled maintenance before the aircraft becomes serviceable again.

The daily flying rates for embarked aircraft depends on the operational tempo of the tail and progress against the annual requirement. Embarked aircraft fly at different rates depending on the operational tempo. Those deployed to an area of operations will fly at a higher rate. Other tails will fly at a lower rate to offset this higher rate of effort. The remainder will fly at a *pro rata* rate based on the annual requirement and time embarked. The daily flying rate is pre-determined for each type based on the required hours per ship and the expected amount of scheduled and unscheduled maintenance. Progress is checked on a monthly basis against the annual requirement. Higher tempo tails fly based on the pre-set rate, but can only fly when serviceable. The flying rate of tails designed to fly at a lower rate of effort can be adjusted, up or down as required, in order to ensure that the overall rate remains on track to meet, and not exceed or fall below, the annual requirement.

The daily flying rates for ashore aircraft are determined based on the desire to maximise the serviceability of ashore aircraft. This is done by matching ashore serviceable tails to phased maintenance lines. The flying rate is then the ratio of flying hours remaining until a phased service, to the number of days until the corresponding maintenance line becomes free. This ensures a constant flow of tails into maintenance and aims to keep maintenance lines (and personnel) occupied. It also means that individual aircraft flying rates may be increased or decreased as required to ensure this

UNCLASSIFIED

UNCLASSIFIED

flow. Tails that have been idle for longer than a set period are forced to fly to ensure tail rotation. The ashore fleet must meet a daily rate within these constraints which is determined *pro rata* from the annual rate. This method easily accommodates changes to the number of ashore serviceable aircraft (through tails entering or returning from unscheduled maintenance or embarkation) and phased maintenance lines (through unscheduled maintenance extending the time that a tail occupies a maintenance line).

Aircraft disembark from a ship under certain criteria, predominantly relating to achieving their required flying hours. Replacement aircraft are chosen for embarkation from the pool of ashore serviceable aircraft. Aircraft are ineligible if they fall due for deep maintenance prior to the end of the embarkation. Of the remaining aircraft, those with the least number of embarked hours are chosen in order to balance the effects in embarked environments (through achieved flying hours, exposure to corrosive environments and hard deck landings) across the fleet. The model prioritises the immediate replacement of disembarked aircraft and thus the ability to meet the embarked requirements first. Therefore any deficiency in fleet size is evident first through a deficiency in the ashore hours achieved.

The model procedures have been verified by key stakeholders in naval combat aviation following extensive consultation during model development. The model has also been verified to an acceptable level against flying hours produced by the current S-70B-2 fleet. This verification and validation process was essential to producing confidence in the robustness of the model in generating reliable results.

Indicative results against a generic requirement indicate the ability of the model to answer the fleet-sizing questions. Interesting effects are observed when the impact of attrition and a modification program are included. When attrition is included, smaller fleets show a steady decline in annual ashore flying hours over the fleet life. When a modification program is included, a large decline in annual ashore hours is seen, particularly during a mid-life upgrade program. The model has shown that this must be addressed by increasing the maintenance capacity.

The model is capable of conducting sensitivity analysis on a range of parameters, such as logistics delays, the number of attrition aircraft, and the maximum daily flying rate per tail. When sensitivity analysis is conducted on the required number of annual ashore flying hours, the results indicate that there can be an optimal daily flying rate for a given fleet based on the maintenance capacity for a given annual flying hour requirement. There is also a maximum requirement beyond which the hours achieved by the fleet cannot be exceeded due to these maintenance capacity constraints.

The model is also capable of trade-off analysis, with fleet size traded off with the number of lines of various types of maintenance. The results indicate that increasing the maintenance capacity can produce similar ashore hours as an increase in fleet size. This poses a question for decision-makers as to whether it is more prudent to purchase an extra aircraft or additional maintenance infrastructure and personnel.

UNCLASSIFIED

UNCLASSIFIED

The two methodologies can be compared by using the availability and serviceability outputs for a given fleet size from the discrete-event simulation as inputs to the binomial distribution. The fleet sizes predicted by the binomial distribution with high probability match those produced by the discrete-event simulation. This suggests that, if availability and serviceability data are known, using the binomial method may provide a good first estimate of the fleet size. However, if other factors such as achieved flying hours need to be determined, or variables like attrition effects or modification programs need to be included, a more robust approach such as the simulation method used here is required.

While both are useful tools, the discrete-event simulation methodology in particular is robust and potentially extensible to other fleet-sizing problems, given its flexibility, speed and explicit treatment of required aspects of the problem.

UNCLASSIFIED

Authors

Dr David Marlow

Joint and Operations Analysis Division

Dr David Marlow holds a Bachelor of Science in Mathematical Sciences from the University of Adelaide and a PhD in Applied Mathematics from the University of Wollongong. He commenced employment with Air Operations Division in DSTO in April 1996. Since that time he has worked on many of the major air acquisition projects from the late 1990s to the present time, including AEW&C, AIR 7000, AIR 6000 and AIR 9000 Phase 8. He has also published research into maritime surveillance modelling techniques in the open literature. The work described in this report earned him and the team a Defence Support Services Gold Level Commendation from the Chief Defence Scientist. His current interests are in undertaking fleet management studies for the MH-60R project. He also initiated and runs the Fleet Management theme in the OR Hub.

Dr Ana Novak

Joint and Operations Analysis Division

Dr Ana Novak holds a Bachelor of Information Technology and Telecommunications from the University of Adelaide and a PhD in Applied Mathematics from the University of Melbourne. She started work with Air Operations Division in DSTO in January 2010. Since then she has worked on major air acquisition projects including AIR 9000 Phase 8 and Phase 7. The work described in this report earned her and the team a Defence Support Services Gold Level Commendation from the Chief Defence Scientist.

UNCLASSIFIED

This page is intentionally blank

UNCLASSIFIED

Contents

ACRONYMS

GLOSSARY

1. INTRODUCTION.....	1
1.1 Project synopsis.....	1
1.2 Project problems and structure of report.....	1
1.3 Definition of terms	2
2. PRELIMINARY FLEET-SIZING ANALYSIS	6
2.1 Problem description	6
2.2 Methodology.....	6
2.3 Sample results	7
2.4 Limitations of preliminary analysis.....	10
3. DETAILED FLEET-SIZING ANALYSIS.....	12
3.1 Problem description	12
3.2 Review of fleet-sizing literature	14
3.3 Choice of discrete-event simulation method.....	15
3.4 Choice of modelling tool.....	16
4. DISCRETE-EVENT SIMULATION MODEL DESCRIPTION	18
4.1 Introduction	18
4.2 Inputs	18
4.2.1 General inputs.....	18
4.2.2 Unscheduled maintenance inputs.....	20
4.3 Initialisation.....	20
4.4 Main procedure	22
4.4.1 Test for work day.....	22
4.4.2 Embarked flying rate	22
4.4.2.1 Determining the embarked hours required for a given tail	23
4.4.2.2 Daily flying rate calculation.....	23
4.4.2.3 Monthly checks against requirement	25
4.4.3 Ashore flying rate.....	26
4.4.3.1 Problem description.....	26
4.4.3.2 Flying hour management literature.....	27
4.4.3.3 Methodology	29
4.4.3.4 Tracking against the annual requirement.....	32
4.4.3.5 Summary.....	33
4.4.4 Updating tails.....	34
4.4.5 State transitions.....	34
4.4.6 Maintenance transitions	35
4.4.6.1 Scheduled maintenance.....	35
4.4.6.2 Unscheduled maintenance	36

4.4.6.3	Logistics delays.....	36
4.4.7	Embarking and disembarking aircraft	37
4.4.8	Other transitions	38
4.4.8.1	Exercise periods.....	38
4.4.8.2	Attrition and life of type.....	39
4.5	Outputs	39
4.6	Summary.....	40
5.	SIMULATION MODEL VERIFICATION AND VALIDATION.....	41
6.	SIMULATION MODEL INDICATIVE RESULTS.....	44
6.1	Input data	44
6.2	Results excluding attrition and modifications	46
6.3	Results including attrition	51
6.4	Including attrition and a modification program.....	53
6.5	Sensitivity analysis.....	57
6.5.1	Varying the number of aircraft lost through attrition.....	57
6.5.2	Varying the maximum daily flying rate.....	60
6.5.3	Varying the logistics delay	61
6.5.4	Varying the ashore requirement.....	62
6.6	Trade-off analysis	65
7.	COMPARISON OF THE METHODOLOGIES.....	68
8.	SUMMARY	71
9.	FURTHER WORK.....	73
10.	ACKNOWLEDGEMENTS	74
11.	REFERENCES	75
APPENDIX A:	SAMPLE IMPLEMENTATION OF THE ASHORE DAILY FLYING PROGRAM.....	77
APPENDIX B:	LIST OF STATE TRANSITIONS.....	80

Acronyms

ADF	Australian Defence Force
ADO	Australian Defence Organisation
ASW	Anti-Submarine Warfare
ASuW	Anti-Surface Warfare
CDG	Capability Development Group
CFU	Carried Forward Unserviceability
CPU	Central Processing Unit
DES	Discrete Event Simulation
DM	Deep Maintenance
DMO	Defence Materiel Organisation
DSTO	Defence Science and Technology Organisation
FOC	Full Operating Capability
IOC	Initial Operating Capability
MFMR	Monthly Flying hour Maintenance Reports
MMH	Maintenance Man-Hours
MOE	Measure of Effectiveness
NASPO	Naval Aviation Systems Program Office
NetMAARS	Networked Maintenance Activity Analysis and Reporting System
RAAF	Royal Australian Air Force
RAM	Random Access Memory
RAN	Royal Australian Navy
ROE	Rate Of Effort

Glossary of Terms

Airframe hours	Terminology used by maintainers for tracking hours accumulated when flying
Ashore	A helicopter that is under the control of a squadron
Available	A helicopter that is not in deep maintenance, or undergoing upgrade or modification
Availability	The number of available helicopters divided by the number of helicopters in the fleet, expressed as a percentage
Deep maintenance	A major maintenance event which a helicopter routinely enters after achieving a certain number of flying hours or after a certain duration of time
Embarked/flight	A helicopter that is under the control of a ship
Flying hours	Terminology used by aircrew for tracking hours accumulated when flying
Maintenance frequency	The time between a service (scheduled maintenance) or failure (unscheduled maintenance), in units of elapsed time or flying hours
Maintenance duration	The time to complete a service (scheduled maintenance) or to repair a failure (unscheduled maintenance), in units of elapsed time or flying hours
Phased maintenance	A maintenance event which a helicopter routinely enters after achieving a certain number of flying hours
Scheduled maintenance	A regular planned maintenance event
Serviceable	A helicopter that is under the control of a ship or a squadron that is capable of flying
Serviceability	The number of serviceable helicopters divided by the number of available helicopters, expressed as a percentage
Tail	An individual helicopter platform
Unavailable	A helicopter that is in deep maintenance, undergoing upgrade or modification, or in reserve
Unscheduled maintenance	A maintenance event which arises unexpectedly, either during a regular service or when a tail is serviceable
Unserviceable	A helicopter that is undergoing scheduled or unscheduled maintenance while under the control of a ship or a squadron

1. Introduction

1.1 Project synopsis

Project AIR 9000 Phase 8 seeks to provide the Royal Australian Navy's surface combatant fleet with a combat aviation capability to undertake Anti-Submarine Warfare (ASW) and Anti-Surface Warfare (ASuW) roles. The AIR 9000 Phase 8 capability will replace that currently provided by the fleet of 16 S-70B-2 Seahawk helicopters. It will also fulfil the capability that would have been provided by the SH-2G Super Seasprite aircraft prior to its cancellation. The 2009 Defence White Paper [1] states that acquiring this naval combat aviation capability is "a matter of urgency". On 16 June 2011, the Minister for Defence announced that the Australian Government would purchase 24 MH-60R Seahawks to fulfil the AIR 9000 Phase 8 requirements.

1.2 Project problems and structure of report

This report describes the development of fleet-sizing analysis in support of AIR 9000 Phase 8. Two different methods were applied during the project as different questions were posed at each stage. Both methods are described in this report.

A preliminary analysis, using a probabilistic approach, was undertaken initially and is described in Section 2. The question to be addressed at this point was the minimum fleet size required to embark a certain number of helicopters on ships. Because this early stage of the analysis is solution-independent (i.e., no specific helicopter options are considered), this simple approach was justified.

The question was subsequently expanded to find the minimum fleet size that would also enable the fleet to meet annual requirements for flying hours for both embarked and ashore aircraft. To address these questions, a more detailed methodology was constructed using a discrete-event simulation approach. This approach was refined and updated throughout the project.

The focus of the report therefore is on this more detailed method. Section 3 describes the problem to be addressed, the solution method chosen and the reasons for this choice. Section 4 provides a description of the model structure including inputs and outputs and the procedures therein. Section 5 gives a brief description of how model verification and validation were achieved through engagement with stakeholders and comparison with Defence data. Section 6 then provides indicative results against a hypothetical requirement that demonstrates the applicability of the model in addressing the problem.

Section 7 provides a comparison between the two methods described in this report. Section 8 then summarises the body of the report, while Section 9 outlines the plans for future work.

1.3 Definition of terms

Various maintenance terms and states that apply to naval aviation assets need to be defined as they are used extensively throughout this report. These terms are applied consistently throughout the analysis.

The maintenance state hierarchy begins, at the highest level, with a number representing the **total fleet**. This fleet is divided into **available** and **unavailable** aircraft. **Unavailable** aircraft have been taken out of service to undergo **Deep Maintenance** (DM), or **modifications** and **upgrades** as part of a follow-on project or phase to extend or upgrade the capability of the original system. Major upgrades and modifications will be undertaken during a DM event. DM is a significant and thorough maintenance event that represents the most extensive maintenance tasking. During DM an aircraft is completely overhauled, requiring extensive disassembly and repair. The cycle and duration of a DM event varies between different aircraft types, but will generally be several years apart and may take anywhere from a few weeks to several months. For the Navy's S-70B-2 Seahawk fleet, DM is undertaken by contractors rather than Navy personnel. As such, they are not available for squadron usage.

Aircraft that are not unavailable due to undergoing DM or modifications are therefore considered **available**. However, an aircraft that is available may not necessarily be in a flying state and hence the available fleet is divided between **serviceable** and **unserviceable** aircraft. Unserviceable aircraft are those undergoing either **scheduled** or **unscheduled maintenance**.

Scheduled maintenance is the term given to the set of preventative maintenance tasks that are undertaken after a specified number of flight hours or elapsed time on an airframe. These usually occur in the form of **phased maintenance** or regular inspections. Regular inspections, such as safety or special inspections, occur on a frequent¹ basis (i.e., tens of hours), based on elapsed time or airframe hours consumed since the previous service. Phased maintenance (or a phased service) is a maintenance event that occurs with a lower frequency, e.g. after hundreds of airframe hours. The duration of these services in Maintenance Man-Hours (MMH) will be higher for phased services than for regular inspections. The aim of such maintenance is to ensure that the ongoing flight safety and operational reliability of the aircraft is preserved.

Unlike scheduled maintenance, **unscheduled maintenance** is undertaken to rectify issues affecting the airworthiness or reliability of the aircraft which were not foreseen, or that have occurred outside regular scheduled maintenance intervals. This can occur on aircraft that are serviceable and programmed to fly on a given day, forcing their removal from the flying program. It may also occur during a phased service, as unexpected defects are discovered and repaired.

¹ Throughout this document, maintenance "frequency" is taken to be equivalent to the time between a service or a failure, rather than the mathematical definition of the failure rate per unit time. Similarly maintenance "duration" is taken to be equivalent to the time to complete a service or the time to repair. These are defined in the Glossary.

All remaining aircraft in a fleet are considered **serviceable**. This means they are able to be used for flight operations. As defined in NetMAARS², the aircraft is assumed to be “Mission Ready” at this point and has been cleared for release by maintenance personnel. The analysis makes no further distinction between whether serviceable aircraft are fully mission capable or are somehow limited in capability for particular missions.

Figure 1-1 illustrates the various maintenance and serviceability states. The Glossary of Terms provides further information on the definitions used throughout this report.




Figure 1-1 : Maintenance and serviceability states

Maintenance capacity for each type of maintenance is represented by a number of **maintenance lines**. If all lines for a particular type of maintenance are occupied, any additional helicopters requiring maintenance wait until another helicopter completes maintenance in one of the occupied lines. In this report, **flight-line maintenance lines** handle short-duration regular inspections, as well as unscheduled maintenance on serviceable helicopters. **Phased maintenance lines** and **deep maintenance lines** handle their respective maintenance types.

Figure 1-2 illustrates the allocation of various maintenance types to maintenance lines. If modifications are required, they will always follow a scheduled service. Different types of

² NetMAARS (Networked Maintenance Activity Analysis and Reporting System) is a reporting tool for analysis of Australian Defence Force (ADF) aviation and related data.

scheduled maintenance will be combined where possible. When deep maintenance occurs, maintenance planners will always try to add a phased service to realise efficiency savings.




Figure 1-2: Maintenance types and lines

Another distinction to be made is that aircraft may either be **embarked** or **ashore**. An embarked aircraft is commonly known as a “flight”, and one of the stated fleet-sizing requirements is to have a certain number of “flights at sea”. For the purposes of this analysis, the terms “flight” and “embarked” aircraft are used interchangeably. In reality, a flight is an aircraft that belongs to a ship until it is released back to the home squadron. Consequently, this aircraft can be ashore undergoing maintenance while still being part of the “flight”. Similarly an ashore aircraft is considered to be one that belongs to the home squadron, although it may be detached to a ship for the duration of a training exercise.

As described in Section 3 and beyond, the model developed for analysis is tail-based, where a **tail** is an individual helicopter. The model tracks the status and location (embarked or ashore) of a given tail, rather than whether a ship or a squadron is in “ownership” of the tail. Therefore, stating that an aircraft is embarked or ashore in the work that follows is a statement of location rather than ownership.

Embarked and ashore aircraft may both undergo scheduled or unscheduled maintenance. However, in this work only ashore aircraft are allowed to undergo deep maintenance and major modifications. This is appropriate given the tail-based nature of the model. In reality, DM for flights may occur “alongside”, that is, while a ship and its associated flight are in port.

Figure 1-3 demonstrates the delineation between the embarked and ashore fleets.




Figure 1-3: Embarked and ashore states

By definition, given the assumptions above, all embarked aircraft are available. From this:

- **Availability** can be determined by dividing the number of available aircraft by the total fleet size;
- **Embarked serviceability** can be determined by dividing the number of embarked serviceable aircraft by the number of embarked aircraft (since all embarked aircraft are available);
- **Ashore serviceability** can be determined by dividing the number of ashore serviceable aircraft by the number of available aircraft ashore.

2. Preliminary fleet-sizing analysis

2.1 Problem description

The initial fleet-sizing analysis conducted for AIR 9000 Phase 8 examined the single question of the number of aircraft needed to embark a certain number of flights at sea. This work pre-dated the release of the 2009 Defence White Paper which first specified the requirement for eight flights [1].

2.2 Methodology

Initial analysis in an Australian Defence Force (ADF) project is solution-independent: i.e., no particular aircraft or system is being considered. Therefore, given the question and the status of the project at the time, the simple binomial distribution was chosen to address the question. The binomial distribution is used to predict the probability of a certain outcome, given a certain number of successes from a certain number of trials. The trials are considered to be independent, so there is the same probability of success each time the trial is performed. This simple approach has been used previously for the AIR 7000 force mix study [2].

The binomial distribution takes the form shown in Equation 2-1 below.

$$P(K = k) = \binom{n}{k} p^k (1 - p)^{n-k}$$

Equation 2-1: Binomial distribution

Here, n is the number of 'trials'; k is the number of 'successes' from the number of trials; p is the probability of a single event occurring; and P is the probability of achieving $K=k$ successes from n trials.

The requirement for AIR 9000 Phase 8 is to achieve *at least* the required number of aircraft embarked. Therefore, summing the above equation from k to n , gives the probability of having at least k successes, as shown in Equation 2-2.

$$P(K \geq k) = \sum_{m=k}^n \binom{n}{m} p^m (1 - p)^{n-m}$$

Equation 2-2: Binomial distribution for at least a certain number of successes

For the fleet-sizing analysis, the aim was to use contemporary data on availability and serviceability from the current S-70B-2 Seahawk fleet, and apply this data to make platform-independent predictions about the size of the future AIR 9000 Phase 8 fleet. In

terms of the definitions of availability and serviceability, the binomial distribution formula definitions are applied as follows:

- n is the number of available aircraft;
- k is the number of serviceable aircraft;
- p is the probability that a single aircraft is serviceable; and
- P is the probability of achieving *at least* k serviceable aircraft from a fleet of n available aircraft.

The total fleet size N is then determined by adding the number of unavailable aircraft to the number of available aircraft, based on the percentage of the total fleet that is available a . This is given by $N = \text{ceiling}(n + n(1-a)/a)$, which becomes $N = \text{ceiling}(n/a)$. The 'ceiling' statement indicates that the value of N is always rounded up to the nearest whole number.

To generate results for AIR 9000 Phase 8, the following procedure was used. Data were obtained for the S-70B-2 Seahawk from the electronic Monthly Flying hour Maintenance Reports (MFMR) for the period incorporating financial years 2005-2007 inclusive. This provided information on average aircraft availability and serviceability as well as the number of embarked aircraft, which led to values for n , p and k respectively. The fleet size N is known to be 16, so a value for P can then be calculated. For the AIR 9000 Phase 8 fleet size, $k = 8$ and the fleet size N is unknown. The value previously calculated for P and the current levels of availability and serviceability for the S-70B-2 can be used as a baseline for determining the required AIR 9000 Phase 8 fleet size.

This method determines the probability that a number of serviceable aircraft will be provided from a number of available aircraft. When applied to the fleet-sizing problem here, it is implicitly assumed in the above definitions that all serviceable aircraft will be embarked, and thus all embarked aircraft are 100% serviceable. It therefore makes no explicit allowances for serviceable aircraft ashore. However, embarked aircraft are not always serviceable. An alternative way to interpret results is to assume a certain level of embarked serviceability and a specified number of ashore serviceable aircraft that would be utilised for flying training. For example, if $k = 8$, it may be assumed that 5 serviceable aircraft will be embarked (e.g. 62.5% embarked serviceability) and 3 serviceable aircraft will be ashore for training. Alternatively, higher values of k could be tested using this methodology. However, since embarked serviceability and ashore serviceable aircraft numbers were not specified, the assumption of 100% embarked serviceability is applied.

2.3 Sample results

Indicative results follow for the fleet size required to generate eight embarked serviceable aircraft. Varying values of availability and serviceability are used to show the ability of this method to undertake sensitivity analysis.

In the graphs that follow the curves are not always 'smooth'. The binomial distribution produces real numbers for the probability of a certain outcome based on the integer number of available aircraft and embarked aircraft. The chosen number of available aircraft is the lowest number with a probability that exceeds the chosen threshold. For

example, with $p = 0.5$ and $k = 8$, $n = 18$ gives $P = 0.760$, while $n = 19$ gives $P = 0.820$. Therefore 19 is chosen as the number of available aircraft required with a probability of 80%. If the availability is 70%, then this requires a fleet size $N = \text{ceiling}(19/0.7) = \text{ceiling}(27.1) = 28$ aircraft. These rounding effects combine to make the curves less smooth than may be expected.




Figure 2-1: Sample Results Output from Binomial Analysis of Fleet Numbers, assuming availability of 70% for various serviceabilities

Figure 2-1 shows the various fleet sizes that can be generated for the requirement of eight flights at sea with an indicative availability of 70%. Sensitivity analysis on the serviceability is also shown in this figure, with values ranging from 40% to 80%. The specified Defence White Paper fleet size of 24 is also included. It is clear that the fleet size is highly variable depending on the serviceability level (a function of maintenance) and probability (a user choice). In this case, to have a 90% probability that 24 helicopters is sufficient, a serviceability of slightly over 60% is required.

The same effect can be seen if availability is increased to 80% as in Figure 2-2. Here, a serviceability of around 55% would be enough to provide a probability of 90% that 24 helicopters would generate eight flights at sea.




Figure 2-2: Sample Results Output from Binomial Analysis of Fleet Numbers, assuming availability of 80% for various serviceabilities

Increasing the availability to 90% further reduces the required fleet size for the same serviceability and probability values as shown in Figure 2-3. In this case, to have a 90% probability of the fleet size being sufficient, a serviceability level of 50% is required.




Figure 2-3: Sample Results Output from Binomial Analysis of Fleet Numbers, assuming availability of 90% for various serviceabilities

Table 2-1 and Table 2-2 show summarised results of the required fleet size for various user-chosen probability levels, for different values of serviceability and availability.

Table 2-1: Fleet size required with 80% probability

		Availability		
		70%	80%	90%
Serviceability	50%	28	24	22
	60%	22	19	17
	70%	19	17	15

Table 2-2: Fleet size required with 90% probability

		Availability		
		70%	80%	90%
Serviceability	50%	30	27	24
	60%	25	22	19
	70%	20	18	16

It is clear that the same fleet size can be generated using different combinations of the three variables. For example, a fleet size of 22 can be derived if the user is satisfied with a probability of 80%, if availability is 70% and serviceability is 60%. However, if a probability of 90% is required, availability must be 80%.

2.4 Limitations of preliminary analysis

Using the binomial distribution to estimate fleet size is a simple technique designed to give indicative, solution-independent results, rather than definitive guidance. This technique was only utilised to answer the single question regarding the fleet size needed to embark a certain number of aircraft.

Given this aim, there were many limitations on the conclusions that could be drawn from the outcomes. These included the following:

- All embarked aircraft are assumed to be 100% serviceable, which is not correct;
- The binomial distribution assumes that the 'trials' are independent. For this analysis, it means that each available aircraft is assumed to have the same serviceability level, which is not the case;
- There was no distinction between embarked and ashore serviceability rates. In reality the priority is to maximise embarked serviceability for operations. Ashore serviceability is lower, sometimes deliberately so, e.g. to enable maintenance crews to be trained on unserviceable aircraft;
- The various types of maintenance were not explicitly considered in this analysis;
- Most notably, there was no consideration of any requirements for the ashore fleet. Ashore aircraft are used to train aircrew to be deployed on ships, so a consideration of ashore flying is essential for a more robust analysis of the problem; and

- Other effects, such as personnel numbers or logistics issues, were not included.

These acknowledged limitations led to the development of a detailed methodology that could more realistically model the scenario and address additional questions.

3. Detailed fleet-sizing analysis

3.1 Problem description

The 2009 Defence White Paper [1] stated that the Phase 8 capability would provide for eight or more concurrently embarked flights at sea. From this guidance it was assumed that this requirement exists for every day of the fleet life. Further guidance from the AIR 9000 Phase 8 project stated that the fleet was also required to achieve a certain number of flying hours each year, both for the embarked flights and the remaining ashore aircraft.

The fleet-sizing problem can thus be formulated as a multi-objective problem, or a single objective problem with multiple constraints. In the latter, the objective function is to minimise the fleet size such that the following constraints are satisfied:

- On any given day, a set minimum number of embarked aircraft is achieved;
- A set minimum amount of flying hours by the embarked fleet per year is achieved; and
- A set minimum amount of flying hours ashore per year is achieved.

This particular problem requires detailed tracking of both the number of aircraft and the number of hours flown, over a period of days and a period of years respectively. Additionally, the random events that occur throughout the life of the fleet – such as unscheduled maintenance and the duration of an embarkation for a single helicopter – also need to be represented.

The fleet-sizing system incorporates the fleet of aircraft, ships, personnel and a home base with various maintenance facilities. Within this system, individual helicopters, referred to as *tails*, pass through various ‘states’ throughout their life. The system is illustrated in Figure 3-1. The states shown in this example are as defined in Section 1.3.




Figure 3-1: Illustration of the fleet-sizing problem

At any given point in time, a tail may be in the state of being ashore and serviceable. At another time it may be in deep maintenance, while at some other time it may be embarked and in scheduled maintenance. Even though many state transitions are possible, there are some transitions that are not allowed. For example, an embarked helicopter that is serviceable may go into maintenance on board a ship, but it cannot go directly into scheduled maintenance ashore until its embarkation is completed.

The duration that a particular tail is in a given state will vary depending on factors such as the duration that a tail is embarked, the frequency of deep maintenance events, the duration of a service in work days, or the number of airframe hours between phased maintenance services. For some of these states, the duration that a tail is in a particular state is deterministic, such as the number of airframe hours between phased maintenance. Therefore the time between phased maintenance services will depend on the daily flying hours achieved for each tail (the daily flying rate). However, other parameters are stochastic in nature, such as the frequency and duration of unscheduled maintenance. It is unknown precisely when an unscheduled maintenance event will occur, or for how long a tail will be in unscheduled maintenance. The duration of an embarkation is also not known in advance.

Any method to address the fleet-sizing problem must be able to represent each individual tail in the overall fleet and track its movement between specific states on a daily basis. On each day, tests for a change of state need to be conducted, based on achieving a particular criterion or set of criteria. For example, a tail that has reached the number of airframe hours between phased maintenance services will move from a serviceable state to a scheduled maintenance state. A tail that has reached the number of days required to complete a service will move from a maintenance state (be it scheduled or unscheduled

maintenance) to a serviceable state. Tails that do not meet the criteria for a change of state have their information updated: serviceable aircraft update their flying hours by the daily flying rate; aircraft in maintenance update their time in maintenance by one day, etc.

By tracking the state of each individual helicopter, the total number of tails in each state can be determined on each day. This can then be analysed to provide a range of information, such as the number of available aircraft and the number of serviceable aircraft. It can also track the percentage of available and serviceable aircraft on each day and provide values averaged over the specified duration (run-time) of a particular model. The total time in maintenance for embarked and ashore aircraft can also be tallied. Similarly the number of flying hours achieved for all embarked and ashore aircraft can be obtained and converted into annual rates.

Therefore, this method can be used to test that the chosen minimum fleet size meets all of the problem constraints. If it does not, the fleet size needs to be increased. Alternatively, other parameters (such as those related to maintenance capacity) may need to be changed. The method also allows for sensitivity analysis by the variation of any number of input parameters. Any model needs to be run multiple times to generate statistically significant results to incorporate the randomness of various parameters, such as the frequency of unscheduled maintenance and the duration of the embarkation of a single tail to a ship.

3.2 Review of fleet-sizing literature

Fleet-sizing work in the scientific literature is quite extensive but primarily relates to commercial operations in the transport sector regarding entities such as trucks, buses, freight cars or ships. A comprehensive review of this literature has been undertaken by Hoff *et al.* (2010) [3].

Literature was sought that included fleet-sizing problems dealing with uncertainty because of the inherent uncertainty in unscheduled maintenance in this problem. One example is Lui *et al.* (2008) [4], who considered a network-based model of the transportation system in Hong Kong under different potential conditions (monopoly, oligopoly) while incorporating both transit operators and passengers. A disutility function incorporates travel time, fare, discomfort and the perceived cost of unreliability. Uncertainty is included through an unreliability level in the transit services. Turnquist and Jordan (1986) [5] examined freight transportation with containers shipping parts from a single component production plant to various assembly plants (a “one-to-many” model). The deterministic model is fairly straightforward, where fleet size incorporates the containers at the assembly plant and those in transit. Uncertainty comes in travel times, so extra containers are added at the component plant and the assembly plants. The model provides conditional results: i.e., “if the probability of running short of containers is to be no greater than x , the required fleet size is y ”. Du and Hall (1997) [6] examine a hub-and-spoke transportation system for trucks and seek to find the minimum fleet size that meets a given allowed stock-out probability (number of backorders) in the long term. A five-step decomposition approach is used to determine stock-out probabilities. Transportation demands are represented as independent Poisson processes. The results are found to be robust under travel time uncertainty.

In the military context, Wesolkowski and Billyard (2008) [7] introduced the stochastic fleet estimation model (SaFE) to determine fleet sizes based on task frequency, duration and assignments for multiple platform types. This Monte Carlo-based simulation tool was applied to a set of 127 air mobility tasks using five different platforms over a one year period. This work was further extended by Wesolkowski and Zhu (2008) [8], where a genetic algorithm was applied to the SaFE model in order to find the best overall cost for the fleet. Crary *et al.* (2002) [9] combined military judgement with integer programming to determine the size of the US destroyer fleet. Military officers were asked to compare pairs of missions in each phase of a scenario and the Analytic Hierarchy Process (AHP) was then used to determine mission weights. These weights were then aggregated over all participants using a probability distribution, and were then incorporated into the objective function of an integer program designed to maximise the probability of “winning” the scenario.

Simulation methods including discrete-event simulation (DES) have also been used in previous fleet-sizing work. Lesyna [10] described how a DES model is applied to sizing a fleet of industrial rail cars using a tool called ProModel. Although no explicit results are provided, the paper does describe how the model was able to prevent unnecessary planned expenses on additional rail cars. Godwin *et al.* [11] used simulation to consider the fleet size of locomotives for the Indian Railways. A well-defined deadheading policy (where empty locomotives are re-positioned) is sought. Locomotives respond to order generation, which is stochastic. A Petri Net model is constructed for the network, with heuristics used to choose how long to hold locomotives at stations before allowing deadheading. Results show that increasing the fleet size beyond a certain point is counter-productive due to the capacity constraints of the network.

Shyshou *et al.* [12] used DES to examine a fleet-sizing problem in anchor handling operations for a Norwegian oil company. Anchor handling tug vessels move rigs within or between various oil fields, where the movement can only occur during acceptable weather conditions. The tug vessels can be on long-term hire, or spot hire which is substantially more expensive. The problem is highly stochastic with all parameters represented as various probability distributions, which are determined based on historical data. The aim is to identify the number of tug vessels needed on long-term hire to minimize total hiring costs, based on various future contingencies for the spot-hire price.

3.3 Choice of discrete-event simulation method

In order to properly represent this helicopter fleet-sizing problem, any method must:

- Account for daily variations in tail state and be able to evaluate criteria for state changes;
- Track the number of tails in each state every day and the hours flown every day of the year, in order to test that the three main constraints of Section 3.1 are met; and
- Handle both deterministic and random state changes.

Therefore, based on the type and structure of the problem and the questions to be considered, a discrete-event simulation method was chosen for this work.

Other applications of DES relevant to this work are to investigate aircraft availability and flight and maintenance processes. There have been a number of recent papers applying the ARENA® DES model to maintenance processes of the Finnish Air Force. Initially, Raivio *et al.* [13] represented various levels of maintenance with different distributions. This work was extended by Mattila *et al.* [14] who considered operations in wartime, including the effect of discarding periodic maintenance at various stages of higher-intensity operations. A later paper (Mattila *et al.* [15]) captured the essence of the aforementioned papers, with an emphasis on the interactions with the Finnish military in the process of undertaking the work. The organisational-level maintenance, intermediate-level maintenance and depot-level maintenance represented in their work are analogous to the regular, phased and deep maintenance types used here.

Similar work has been undertaken by Cook & DiNicola [16] for the US Army Black Hawk fleet. They developed a computer model for a fleet undergoing combat operations. The model uses a range of probability distributions for mission scheduling, battle damage, system failure and repair times. Aircraft experience failures during the mission causing the mission to be aborted. The aim of this work is to reduce the maintenance workload on combat helicopters.

Interestingly, Bender *et al.* [17] noted some potential dangers in examining whole-of-life issues. Using a DES that represented random processes, as in the fleet-sizing study, they discovered that bifurcations occur after a very large number of time steps (around 100,000). In their work, maintenance durations were represented as probability distributions. However, in the fleet-sizing study, while concerned with representing a fleet over its entire life, all input parameters are deterministic and the number of time steps is less at around 10,000 (i.e. for each day of a 30 year fleet life). No such effects were observed during the course of this analysis, and no testing was undertaken to see if such effects may be observed with more time steps.

3.4 Choice of modelling tool

The discrete event simulation model for this analysis was built in the MATLAB R2010a® environment developed by the MathWorks. The primary reasons for using MATLAB as opposed to specific DES tools such as ARENA® were a) purchase costs; b) project time constraints, given the time required to learn how to use new software; and c) existing familiarity with MATLAB®. The MathWorks has other arguably more appropriate tools such as Simulink® for visual-based design, Stateflow® for logic-driven system modelling and SimEvents® specifically for DES modelling. However, these other tools were not utilised for the same reasons.

Other potential tools such as systems dynamics models were considered but rejected for their lack of flexibility. Goodwin [18] developed a generic systems dynamic model of an aircraft fleet using the iThink system dynamics software. This model included many of the elements described in the fleet-sizing model, such as maintenance types and flying rates. However, the report notes that system dynamics modelling is not naturally tailored to handle events, such as discrete numbers of aircraft entering maintenance as they

accumulate flying hours. Therefore a separate model for maintenance triggers was subsequently required in the overall system dynamics model. Discrete-event simulations are far more naturally capable of handling these state changes.

Undertaking all of the model development in MATLAB also eliminated any issues with using “black box” software, where details on how particular elements of the software work are hidden from the user. For a high-profile project, as AIR 9000 Phase 8 quickly became, transparency was essential, both for the model developer and the stakeholders. Translating military procedures into algorithms and code, and explaining their implementation, was a relatively simple process within the MATLAB environment. This environment was also able to easily handle some of the more complex algorithms for calculating daily flying rates, as well as the more traditional state-tracking components of discrete-event simulation.

A description of the model follows in the next section.

4. Discrete-event simulation model description

4.1 Introduction

DES provides a means to effectively deal with complex dependencies and variability. It utilises mathematical and logical models of a physical system that portrays state changes at precise points in simulated time, all of which are significant features of the fleet-sizing problem as formulated for this study. The components of a discrete-event simulation include:

- Clock – The simulation keeps track of current simulation time: in this case the measurement unit is days, with a time step of one day.
- Set of possible states – Each state change triggers a code that is used to simulate that event.
- Random number generators – The simulation needs to generate random variables of various kinds, such as to create unbiased initial conditions as well as to model stochastic variables.
- Statistics – The simulation keeps track of relevant model statistics, which quantify the properties of interest.
- Ending condition – here defined as the time when the fleet reaches end-of-life.

4.2 Inputs

4.2.1 General inputs

The model inputs for the fleet-sizing model can be grouped into three categories, with a number of subcategories. These are:

- Fleet size;
- Parameters for sub-procedures:
 - Maintenance man-hours per day (embarked and ashore);
 - Maintenance frequency and durations for regular services;
 - Maintenance parameters for representing unscheduled maintenance;
 - Maintenance capacity (i.e., number of maintenance lines for each type of maintenance);
 - Proportion of maintenance delays due to logistics;
 - Expected number of attrited (permanently lost) aircraft;
 - Life of type per tail in airframe hours;
 - Timing/frequency and duration of planned modification and upgrade program; etc.
- Project requirements:
 - Requirement for minimum number of embarked aircraft;
 - Annual flying hours requirement per embarked flight;
 - Average annual flying hours requirement (both embarked and ashore).

A diagram showing the model flowchart is presented in Figure 4-1.

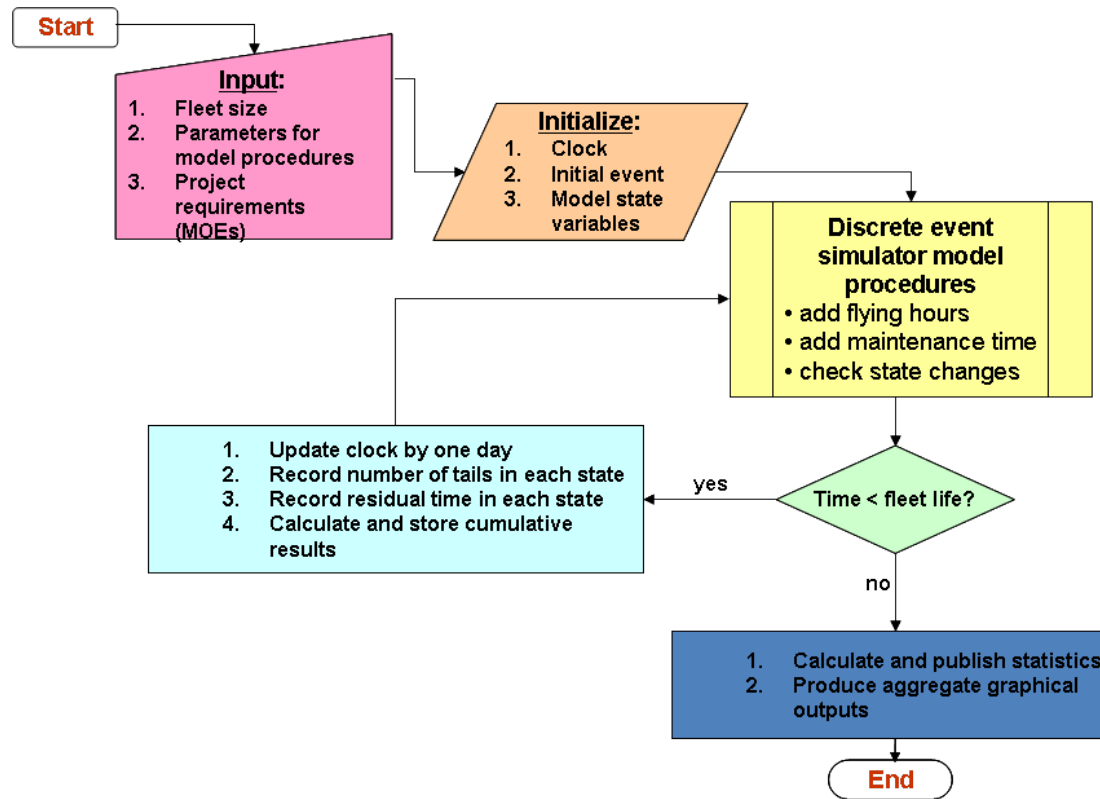


Figure 4-1: Discrete-event simulation model flowchart

4.2.2 Unscheduled maintenance inputs

Unscheduled maintenance data were analysed for the Royal Australia Navy's existing Sikorsky S-70B-2 Seahawk fleet from 2002-2010 to derive appropriate probability distributions. A statistical analysis of the random data showed that both unscheduled maintenance frequency and duration followed a log-normal distribution. The chi-squared goodness of fit test was applied to confirm the appropriateness of the distribution. For both frequency and duration, the log-normal distribution provided a superior fit than others tested such as the exponential and Weibull distributions.

The log-normal distribution is commonly found to apply for the time to repair (duration) in corrective maintenance [19], so this was as expected. It was also expected that the time to failure (frequency) would follow a Poisson process and therefore an exponential distribution, but this was not the case.

No concrete explanation is proffered for this unexpected finding. One possibility is that the period analysed is approximately in the mid-life of the capability. The data may be showing the combined effect of an exponential distribution in the early part of that period corresponding to a constant failure rate, and perhaps a normal distribution in the later part corresponding to an increasing failure rate. These would be expected if unscheduled maintenance failure rates followed the typical 'bathtub curve' [20]. There is no comparable data from 1992 (when Full Operational Capability was achieved) to 2002 to test this proposition. Moreover, the maintenance regime for the Seahawk fleet was overhauled during this period.

Based on these findings, both unscheduled maintenance and frequency follow log-normal distributions in the model. It is eminently reasonable to expect that improvements in maintenance processes and technology will reduce both the frequency and duration of unscheduled maintenance (and indeed scheduled maintenance) during the fleet life. However, in lieu of any reliable information quantifying these improvements, it is assumed that these parameters are unchanged over the fleet life.

4.3 Initialisation

All tails are initialised with a random number of total flying hours, based on the fleet flying hours achieved over this introductory period. The model 'starts' once the full fleet of aircraft has been received, with the start 'day' being any random day of the year.

Usually, the introduction of a fleet of aircraft is staggered from the Initial Operating Capability (IOC), when the first batch of aircraft is received, through to the Full Operating Capability (FOC), when the full fleet has been received and accepted. Therefore, it is assumed that the one-fifth of the fleet is introduced from the year of IOC to the year of FOC four years later. These aircraft fly an average of one-fifth of the ashore requirement in each year. A random variation using a normal distribution around the average value is used to spread these initial values. The final aircraft delivered at FOC will have zero flying hours. This means that the total fleet hours achieved by FOC will be twice the ashore

requirement, i.e., $4/5 + 3/5 + 2/5 + 1/5 + 0/5 = 2$ times the requirement. The split between embarked and ashore hours is randomly determined from a normal distribution such that it falls between 25/75 and 75/25 for each tail.

The model initialises all tails into various states. The initial number of embarked aircraft is set to the minimum requirement. Embarked aircraft are further allocated into three categories with different flying rates: 'surge', 'normal' and 'low' rate of effort (ROE) explained further in Section 4.4.2. A mix of aircraft is allocated to deep maintenance and phased maintenance, depending on the frequency and duration of the maintenance parameters. The numbers in the various maintenance states are then determined based on the frequency and duration of the maintenance processes, as given by Equation 4-1.

$$N_{maint} = N \frac{T_{maint}}{v_{maint}}$$

Equation 4-1 : Average number of helicopters in maintenance

Here N is the fleet size, T_{maint} is the duration of a type of maintenance and v_{maint} is the frequency. As an example, if there are 16 tails in the fleet and deep maintenance is every 4 years and lasts for 3 months, then on average there would be exactly 1 tail in deep maintenance at a time. This case would constitute a perfect stagger, with tails leaving maintenance as a new one arrives.

If the service frequency is flying-hours based, the time between services in days will then depend on the daily tail flying rate. An initial estimate of this can be derived from the fleet size and the combined annual requirement for embarked tails G_{emb} and ashore tails G_{ash} , i.e.:

$$f_{tail} = \frac{G_{ash} + G_{emb}}{365N}$$

Equation 4-2: Estimate of maintenance frequency (in days) when frequency is given in flying hours

From this, v_{maint} in days is determined by dividing the flying hours-based maintenance frequency by the estimated value for f_{tail} in flying hours per day.

The remaining aircraft are placed in an ashore serviceable state. Initially there are no tails in unscheduled maintenance and all embarked tails are considered to be serviceable.

For tails initialised into maintenance, the number of days into maintenance is randomly allocated. The previous random allocation of flying hours will place each tail to have a certain number of flying hours until the next service is due. Embarked tails are allocated an embarked duration and are initialised at a certain time during that embarkation and with a certain number of accumulated flying hours while embarked. All of these random allocations use a uniform distribution from zero to the appropriate value: e.g. for serviceable aircraft, the airframe hours until the next service.

Model outputs for the first three years are regarded as the initialisation period and are not counted towards the overall results. This allows time for annual ashore hours to reach a 'steady-state' and for an improved alignment of the maintenance stagger. Some previous runs revealed that a steady-state was not achieved until after the first few years of the fleet life in certain cases. Following this initialisation period, the model 'officially' begins to record the results for the originally specified fleet life. This reflects the practice and findings of others such as Matilla *et al.* [15], who used a six-month initialisation period.

4.4 Main procedure

The main procedure is the core of the fleet-sizing model. This component calculates the embarked and ashore flying rates for the given day. It then updates tails in their current state, and finally tests for state changes.

4.4.1 Test for work day

First, the current day is tested to see if it is a working day. Embarked crews work every day of the year, whereas ashore crews do not work on weekends or during holiday periods. These periods also differ between flying and maintenance crews. Ashore personnel have a specified "standdown" period over the Christmas/New Year period, while flying crews have an additional shorter standdown period in the middle of the year.

Consequently, flying hours and maintenance activities only accrue during work days. However other parameters, such as time between deep maintenance and embarkation durations, are tracked in calendar days.

It also checks to see whether the current day is within an exercise period, where some ashore tails are detached and fly at different rates as explained in Section 4.4.8.1.

Following this, flying rates are determined for embarked and ashore aircraft.

4.4.2 Embarked flying rate

The embarked fleet is required to achieve a certain number of flying hours every year. The annual requirement is split evenly between each of the flights. Each embarked tail has a requirement to achieve a certain number of flying hours per embarkation in line with the annual requirements of the particular flight.

The procedure for determining the embarked flying rate is three-fold. First, the number of hours to be achieved for a newly-embarked tail needs to be determined depending on the type of embarked tail. Second, the flying rate for the embarkation needs to be calculated based on the hours to be achieved and the expected amount of maintenance. Third, a dynamic procedure needs to be implemented to track the progress of all embarked tails against the annual requirement. The flying rate can vary depending on which tails are serviceable to ensure that the hours achieved do not greatly exceed or fall below the annual requirement.

4.4.2.1 Determining the embarked hours required for a given tail

When tails are embarked, the duration of the embarkation and the number of hours that the tail is expected to fly is pre-determined, based on the embarked fleet reaching the annual requirement. In accordance with current naval combat aviation practice, the model represents three levels of flying rate for embarked aircraft:

- A surge rate of effort, where the tail consumes all of the flight's annual hours during embarkation (e.g., as with high-tempo operations);
- A normal rate of effort, where a flight consumes its annual hours on a *pro rata* basis; and
- A low rate of effort, where remaining embarked flights offset the rate of effort generated by tails flying at the surge and normal rates.

In the model, of the minimum number of flights at sea at any one time, it is assumed that one-quarter will be experiencing the surge rate of effort and another one-quarter will be flying at a low rate of effort. Therefore half of the fleet will be flying at the *pro rata* normal rate. If the annual requirement per flight is 400 hours, the surge tails will fly this amount during their embarkation duration (which is somewhere between 6-12 months). The normal tails will fly, for example, 300 hours for a 9 month embarkation, or 200 hours for a 6 month embarkation. The low tails will then fly whatever amount is required to ensure that the annual embarked requirement (in this case, 8 flights * 400 hours/flight/year = 3200 hours/year) is met. The initial calculation for the embarked hours required for a low tail is given by:

$$F_{emb_low} = F_{flight} \left(\frac{T_{emb_low} + T_{emb_surge}}{365} - 1 \right)$$

Equation 4-3: Initial calculation for required embarked hours for a low rate-of-effort tail

In Equation 4-3, F_{flight} is the required hours for a flight over a given year and T_{emb_low} and T_{emb_surge} are the time embarked for low and surge tails respectively. Given that the hours required for a surge tail is F_{flight} , the average of the hours for the surge and low tails should over time equate to the hours required for a normal tail.

4.4.2.2 Daily flying rate calculation

A tail does not fly every day that it is embarked. There will be designated periods of scheduled maintenance and unpredictable occurrences of unscheduled maintenance during which a tail cannot fly. The more unserviceable a tail is, the greater the rate it will have to fly on the days on which it is serviceable.

The daily embarked flying rate for the duration of an embarkation is:

$$f_{emb} = \frac{1}{s_{emb}} \frac{F_{emb}}{T_{emb}}$$

Equation 4-4 : Top-level daily embarked flying rate equation

In Equation 4-4, F_{emb} is the required flying hours while embarked (set prior to embarkation), T_{emb} is the duration of the embarkation, and s_{emb} is the embarked serviceability.

Serviceability is defined as the time serviceable divided by the time available. By definition, all embarked tails are available, and tails are unserviceable if they are in maintenance. So,

$$s_{emb} = 1 - \frac{T_{emb_unserv}}{T_{emb}}$$

Equation 4-5 : Formula for embarked serviceability

In Equation 4-5, T_{emb_unserv} is the time unserviceable while embarked.

Time unserviceable consists of scheduled and unscheduled maintenance. Scheduled maintenance is deterministic and unscheduled maintenance is stochastic, so an expected amount of time in both scheduled and unscheduled maintenance needs to be established in order to estimate the time unserviceable while embarked.

Scheduled maintenance while embarked consists of phased services and other regular inspections, such as safety inspections and special inspections. The number of airframe hours between phased services is known, as is the frequency of other inspections, so the average time in phased maintenance during an embarkation can be estimated.

The average amount of unscheduled maintenance while embarked can be determined using the mathematical expectation of the frequency and duration. The mathematical expectation of the log-normal distribution for a random variable X , where μ is the mean and σ the standard deviation, is as given in Equation 4-6:

$$E[X] = e^{\mu + \frac{1}{2}\sigma^2}$$

Equation 4-6: Expectation of a log-normal distribution for a random variable X with mean μ and standard deviation σ

These expectations can then be generated for frequency and duration and used to calculate the average embarked downtime due to unscheduled maintenance. For example, for a tail that is planned to be embarked for 300 days, with an average frequency of unscheduled

maintenance every 20 days and an average repair duration of 5 days, a total of $(300/20)*5 = 75$ days of unscheduled maintenance is expected.

Thus, the expected amount of unserviceability for an aircraft during its embarkation is as shown in Equation 4-7:

$$T_{emb,unserv} = \frac{F_{emb}}{R_{reg}} T_{reg} + \frac{F_{emb}}{R_{ph}} T_{ph} + \frac{T_{emb}}{E[U_{frequency}]} E[U_{duration}]$$

Equation 4-7 : Equation for embarked unserviceability

The first two expressions cover the calculated amount of maintenance downtime from regular inspections and phased maintenance, with R representing the maintenance frequency and T the duration. If maintenance frequency is determined based on elapsed time rather than achieved flying hours, then the embarked duration should be used in the calculations rather than the flying hours. The final term is the expected downtime due to unscheduled maintenance.

Substituting the value of $T_{emb,unserv}$ into Equation 4-5 gives the value for embarked serviceability S_{emb} , and subsequently inserting this into Equation 4-4 gives the initial embarked flying rate.

4.4.2.3 Monthly checks against requirement

In order to achieve the annual requirement, the model breaks any given year into monthly periods with a specified 'target' of hours to be achieved based on the monthly *pro rata* rate. If the fleet is running ahead or falling behind the target, it can be adjusted as necessary. A daily flying rate is then calculated based on the procedures defined previously.

Tails that are surge tails or normal tails have their daily flying rates calculated first. This rate is constrained by a user-defined minimum and maximum value. The calculated rate flown is identical to the rate calculated using Equation 4-4, except that F_{emb} is the hours *remaining* to be flown and T_{emb} is the time remaining to be embarked. This constant update allows a given tail to be disembarked as closely as possible to the originally planned date. If the tail is behind schedule for disembarkation, most likely due to a long-duration unexpected unscheduled maintenance event, the "time remaining" may be zero or negative. In this case, "time remaining" is re-set to one, allowing the tail to fly as fast as possible in order to disembark as close as possible to the original date.

The actual flying rate of the low tails is determined dynamically, depending on current circumstances. If F_{emb_month} is the monthly requirement, the model calculates the combined daily flying rate for the low rate of effort tails based on the amount flown by the surge and normal tails, given by:

$$f_{emb_low} = \max(0, \frac{F_{emb_month}}{30} - (f_{emb_surge} + f_{emb_normal}))$$

Equation 4-8 : Updated daily flying rate for 'low' tails

Sometimes, all of the surge and normal tails will be serviceable and capable of flying at their determined rates. This will result in a very low (and sometimes zero) rate of effort for the low tails. Conversely, not all embarked tails will be serviceable on any given day. If an unexpected number of these tails is unserviceable for a given month, the low tails will fly at a greater rate in that month to address the shortfall. In some instances there may also be less than the minimum number of embarked flights for certain periods. This procedure will also allow any low tails to compensate for a shortfall in fleet embarked hours.

4.4.3 Ashore flying rate

Within the context of the overall model to determine the minimum fleet size are the requirements for this fleet to achieve a particular number of flying hours every year. On each ashore flying day over the fleet life, which ashore serviceable tails fly and how much they fly needs to be determined. This is a non-trivial problem.

4.4.3.1 Problem description

The aim in this work is to achieve a certain daily flying rate, based on progress against the annual requirement. Tails are chosen to participate in the flying program based on the number of tails, the number of phased maintenance lines, and the hours remaining before each serviceable tail is due for phased maintenance. Failing to meet the annual ashore flying requirement may be an indication that insufficient ashore serviceable aircraft and/or insufficient maintenance lines are available.

Input parameters for the ashore flying program are as shown in Table 4-1.

Table 4-1: Input parameters used in determining the daily flying program

Parameter	Units	Notation
Number of ashore serviceable tails	number	n
Number of phased maintenance lines	number	p
Airframe hours between phased services	airframe hours	R
Phased maintenance duration	work days	$m_1, \dots, m_n$
<i>Pro rata</i> daily flying rate for the ashore fleet	hours/day	F_{pro}
Required daily flying rate for the ashore fleet	hours/day	F_{req}
Maximum daily flying rate for the ashore fleet	hours/day	F_{max}
Maximum daily flying rate per tail	hours/day	f_{max}
Minimum daily flying rate per tail	hours/day	f_{min}
Maximum days idle per tail	work days	I

The number of ashore serviceable tails n and the required daily flying rate F_{req} are potentially variable on a daily basis. All other parameters are fixed for the duration of the fleet life.

Table 4-2 gives the variables and output parameters for this problem.

Table 4-2: Variables and outputs in the determination of the daily flying program

Parameter	Units	Notation
Airframe hours remaining until next phased service	airframe hours	$r_1, \dots, r_n$
Days until current phased maintenance is completed	work days	$d_1, \dots, d_p$
Consecutive days idle	work days	$s_1, \dots, s_n$
Calculated daily flying rate per tail	hours/day	$f_1, \dots, f_n$

Figure 4-2 illustrates the problem. In this instance there are five serviceable tails $h_1, \dots, h_5$ ashore and three phased maintenance lines $l_1, \dots, l_3$. The question is how to allocate flying hours on any given day in order to ultimately meet the flying hour requirement over the course of a full year. The problem is dynamic, with the number of ashore serviceable tails varying day to day. Tails may enter and exit phased maintenance, enter and exit unscheduled maintenance, and be embarked or return from being embarked. While the number of phased maintenance lines is constant, the status of the lines change, with some tails possibly experiencing extensions to their original maintenance durations due to unscheduled maintenance occurring.




Figure 4-2 : Illustration of the daily flying program problem

4.4.3.2 Flying hour management literature

Research into fleet management and allocating flying hours is a small but growing field with regards to military aircraft with various techniques observed in the literature. Integer programming techniques are commonly employed in these types of scheduling and allocation problems.

Examples include two Masters of Science theses from the US Naval Postgraduate School. Sgaslik [21] described a decision support tool for the German Army which included a mixed integer linear program to manage their fleet of 45 Bell UH-1D helicopters. The tool can develop an annual plan which then feeds into a short-term plan. These helicopters have two types of phased maintenance that are considered. The annual plan seeks to balance priorities across the fleet, such as minimising the variation in hours, meeting planned hours each month and complying with maintenance capacity, and generates a maintenance schedule. This acts as input to the short term plan for mission assignment. The mixed integer linear program also uses two heuristics to increase the computation speed and provides results that, while sub-optimal, were found to be acceptable to the military users.

Pippin [22] also developed a mixed integer linear program to ensure the steady-state flow of individual tails into maintenance for the US Army. He begins by describing the “sliding scale method” used by the US Army that plots remaining flight hours of a fleet in increasing order. By plotting aircraft on a chart from least to most hours remaining until a phased service, and overlaying with a line from the origin to the phased maintenance frequency (the “goal line”), planners can organise the aircraft to enter maintenance at regular intervals. They then check as to whether an aircraft is ‘above’ or ‘below’ the line, and adjust the flying rate accordingly to maintain a regular arrival of tails into maintenance. He uses the “sliding scale method” as a baseline and seeks to minimise the overall deviation of the fleet from the “goal line”. A similar method using quadratics to minimise the square of the deviation provides similar results. However, his thesis does not consider types of maintenance.

Work by Greek researchers builds on that described by Pippin [22] in determining a flying and maintenance plan for the Greek Air Force. Kozanidis [23] and Kozanidis *et al.* [24] also described the sliding scale method or aircraft flowchart as in Pippin [22] as it is also used among the world’s Air Forces. However, his primary objective is to maximise the fleet availability over a given period, rather than minimise the deviation from the “goal line”. He then constructs a mixed-integer linear program that seeks to achieve this over a period of six months with monthly intervals. The first paper [23] is set up as a multi-objective model which seeks to maximise the minimum availability and the minimum residual flight time (i.e., the time remaining until a phased service) across an aircraft wing and corresponding squadrons with various weightings applied. He then compares these results with various heuristics, including one similar to the quadratic method used by Pippin [22] based on minimising variations from the flowchart goal line. In a follow-up paper, Kozanidis *et al.* [24] re-formulated the previous work as a single objective problem to maximise the minimum availability of the wing over the planning period. The other three previous objectives are included as constraints. These papers both consider various types of maintenance as in Sgaslik [21].

Hahn & Newman [25] developed a mixed integer linear program for managing helicopter maintenance and deployment for the US Coast Guard’s Sikosky HH60J helicopters. The planning period is 12 weeks with weekly blocks, with the aim to keep tails deployed as long as possible and for each tail to achieve a certain number of flying hours. Winata *et al.* [26] uses the same planning duration but included separate deployment areas as well as a

home base for the fleet. The objective in this case was to minimise the number of deployments while meeting an overall weekly objective.

Each paper included different types of maintenance regimes. Pippin [22] and Winata *et al.* [26] considered only one type of maintenance, Sgaslik [21] included two types of phased maintenance, while Kozanidis [23] and Kozanidis *et al.* [24] both considered various types of maintenance analogous to those considered in this work.

The primary difference in the fleet-sizing work described here compared with the papers described above is in the daily variability of the problem. None of the above papers considers the effects of unscheduled maintenance on fleet availability. They also all assume a constant number of serviceable tails over the entire planning period, along with constant numbers in maintenance. The removal of tails from the serviceable pool is not considered, nor is the addition of tails to the serviceable pool, both of which may change the relative positions of individual tails in the flowchart. Furthermore, the maintenance durations are effectively assumed to be constant for each type of maintenance, implied by the steady-state arrival of tails into maintenance.

An alternative approach is simulation-based optimization used by Mattila and Virtanen ([27], [28]). Their aim is to maximise fleet availability over a one-year period. Firstly they calculate the ratio of remaining flight hours until periodic maintenance to the remaining calendar time until the scheduled commencement of maintenance. Those aircraft with the highest ratio are chosen for flight duty. They then applied heuristics to improve the maintenance schedule, such as genetic algorithms [27] and simulated annealing [28], in order to minimise queuing of tails for maintenance. Random failures of aircraft are also included but are not fully described, although from Mattila *et al.* [15] it would appear to follow a Poisson process. This type of approach has the inherent flexibility required for the daily ashore flying program here.

4.4.3.3 Methodology

The overall aim of the procedure implemented in this model is to maximise serviceability by minimising maintenance queuing as with Mattila & Virtanen ([27], [28]). It is in the course of achieving this objective that the daily requirements should be met if there are sufficient serviceable aircraft and maintenance capacity.

This aim reflects the intent of flying management procedures currently practiced by RAN engineers in managing the S-70B-2 Seahawk fleet. Under these circumstances, if a tail enters unscheduled maintenance, there is a higher probability that a replacement serviceable tail can be sourced. It minimises queuing of tails for phased maintenance while also keeping phased maintenance lines occupied. This is ideal in maximising the utility of the maintenance workforce. If achieving the daily flying requirement were the objective, this may lead to inefficient choices of which aircraft to fly and the resultant queuing of aircraft for maintenance. In some cases these tails can become 'hangar queens': i.e., tails that have achieved their hours and have been awaiting maintenance for a long time. These tails are called 'dead' as they cannot be flown until they are serviced.

Initially, tails are prioritised from least to most remaining hours until their next phased service. Similarly, maintenance lines are prioritised from least to most days until the current service is completed. The model then seeks to match each serviceable tail to a corresponding maintenance line. This is done in ascending order of airframe hours remaining for serviceable tails, and days remaining for maintenance lines. Tail h_1 is matched to line l_1 , tail h_2 to line l_2 etc.

Assume that all phased maintenance lines are occupied, so $d_j \neq 0$ for all $j=1, p$. If $n \geq p$, for tail i , where $i = 1, \dots, p$, then the daily flying rate per tail is simply the ratio of flying hours remaining to days remaining as in Equation 4-9:

$$f_i = \frac{r_i}{d_i}, \quad i = 1, \dots, p$$

Equation 4-9: Basic equation for calculating the daily flying rate per ashore tail

This approach is similar to that employed by Mattila and Virtanen [27]. However, in their method, they select aircraft with the highest ratio of remaining flying hours to remaining calendar time and then employ heuristics to improve the maintenance schedule and minimise queuing. Without the heuristics, it would be expected that a bottleneck of tails would form using that procedure as these ratios would converge.

If there is an excess of serviceable tails to maintenance lines, then these additional tails will be matched to maintenance lines in the same order. Therefore, tail h_{p+1} will be matched to line l_1 , tail h_{p+2} will be matched to line l_2 , etc. If $n \gg p$, then tail h_{2p+1} , $h_{3p+1} \dots$ will be matched to line l_1 , and tail h_{2p+2} , $h_{3p+2} \dots$ will be matched to line l_2 etc., until all tails are matched. For tail i where $i = \{p+1, 2p\}$, the aim is to fly the remaining serviceable tails at a rate that plans their arrival in maintenance when both the current tail occupying the line *and* the one ahead of it in the queue completes maintenance. Therefore, these tails will fly at a rate given by Equation 4-10:

$$f_i = \frac{r_i}{(d_{i-p} + m_{i-p})}, \quad i = p+1, \dots, 2p$$

Equation 4-10: Flying rate for ashore tails where there are tails ahead which will fall due for maintenance first

In general terms, for $n \geq p$ with tail $i \geq p$, until $i = n$, the flying rate can be calculated using Equation 4-11:

$$f_i = \frac{r_i}{d_{i-kp} + \sum_{j=1}^k m_{i-jp}}, \quad i = kp+1, \dots, (k+1)p, \quad k = 1, \dots, \text{div}(n/p)$$

Equation 4-11: General expression for calculating flying rate for ashore tails

For neatness, this calculated flying rate is rounded, for example to the nearest half-hour, to generate the actual flying rate for that tail on that day.

The calculated ratio must be such that the flying rate is within the minimum and maximum daily flying rate as in Equation 4-12, i.e.:

$$f_{\min} \leq f_i \leq f_{\max}$$

Equation 4-12: Upper and lower bounds for the daily flying rate per ashore tail

If $f_i \leq f_{\min}$, then tail i does not fly on that day, so $f_i = 0$. This is indicative of a tail that is 'ahead of schedule' in hours flown and would be due for phased maintenance well before a phased maintenance line becomes available. The only exception is when $r_i \leq f_{\min}$ and $d_i = 0$, so the tail flies its remaining hours and immediately enters phased maintenance. Conversely, if $f_i \geq f_{\max}$, then the model re-sets the rate for tail i such that $f_i = f_{\max}$. This is indicative of a tail that is 'behind schedule' in hours flown and needs to make up remaining hours as quickly as possible to avoid maintenance lines remaining unoccupied. In both cases, these corrections are made before allocating the flying hours for the next tail.

If one or more phased maintenance lines are unoccupied (i.e., $d_j = 0$ for some of $j = 1, \dots, p$), these maintenance lines are placed at the head of the queue as before and matched to serviceable tails. The flying rate in these instances is then calculated to fill this line as quickly as possible, up to the maximum rate, as in Equation 4-13:

$$\forall i : d_i = 0, \quad f_i = \min(f_{\max}, r_i)$$

Equation 4-13: Flying rate for a tail matched to an empty maintenance line

Tails may remain idle: i.e., capable of flying but not required to fly, for long periods under certain circumstances. This may be when maintenance turnaround times are rapid, or when there are a large number of serviceable tails. This is because these few tails flying at or near their daily maximum rate may achieve the daily requirement. The idle tails in these circumstances are usually those that have recently completed maintenance, so they have a large amount of flying hours (close to R) until their next service is due.

The number of days that a tail is idle s_i is tracked by the model. The model ensures tail rotation by forcing tails to fly that have been idle for longer than a chosen period: i.e. when $s_i \geq I$. When this occurs, these tails are prioritised above all others on that particular day and matched to the maintenance lines as before, and fly at the rate as determined by the above equations. If there are still flying hours to be allocated, any remaining tails are matched to lines and their flying rates calculated. For example, if of the n tails, k tails have been idle for longer than I days (often tails $n-k+1$ to n), those will fly first, and the number of days idle is re-set so $s_i = 0$. The tail h_1 that actually has the lowest number of hours remaining will be matched to phased line $k+1$, and so on. Note that this procedure does not apply to tails that are 'ahead of schedule' as described above.

4.4.3.4 Tracking against the annual requirement

The *pro rata* daily flying rate can be calculated from the ashore requirement and the number of flying days. It must also allow for exercise periods, during which the total flying rate for the ashore and detached aircraft is increased by 50% over the normal rate. Therefore the *pro rata* daily flying rate is as shown in Equation 4-14:

$$F_{pro} = \frac{F_{ash}}{T_{ash} + \frac{1}{2} T_{ex}}$$

Equation 4-14: Equation to calculate the *pro rata* daily flying rate

Here F_{ash} is the annual ashore requirement, T_{ash} is the number of ashore flying days per year and T_{ex} is the number of exercise days per year.

The model tallies the total flying hours per day for the ashore fleet in order from tail 1 upwards. The progressive tally may be such that applying the calculated rate to the $(j+1)^{th}$ tail would cause the required daily rate F_{req} to be exceeded as in Equation 4-15, i.e.:

$$\begin{aligned} \sum_{i=1}^j f_i &< F_{req}, \quad j < n \\ \sum_{i=1}^{j+1} f_i &> F_{req}, \quad j+1 \leq n \end{aligned}$$

Equation 4-15: Example where the number of ashore serviceable tails may exceed the daily flying hours requirement

In this case, that tail will only fly enough hours to reach the maximum, as shown in Equation 4-16:

$$f_{j+1} = F_{req} - \sum_{i=1}^j f_i$$

Equation 4-16: Correction technique applied to ensure that the daily flying requirement is not exceeded

On the first day of a year, $F_{req} = F_{pro}$. Every day the progress against the *pro rata* requirement is checked based on the accumulated flying hours and the number of flying days remaining. Then F_{req} can be adjusted upwards over the course of the year if it is falling behind the *pro rata* rate up to a maximum value F_{max} . The flying rate per tail can correspondingly be increased as follows:

- If $F_{req} > 1.2 * F_{pro}$ then $f_{max} = 2 * f_{max}$ (i.e. double the maximum daily flying rate per tail if the requirement is now more than 20% higher than the *pro rata* rate);

- If $1.1 * F_{pro} < F_{req} \leq 1.2 * F_{pro}$, then $f_{max} = 1.5 * f_{max}$ (i.e. increase the maximum daily flying rate per tail by 50% if the requirement is now between 10% and 20% higher than the *pro rata* rate);
- If the requirement is less than 10% higher than the *pro rata* rate, then the maximum daily flying rate per tail is unchanged.

These adjustments allow the daily flying program, while ‘local’ (daily) in substance, to still aim to reach the ‘global’ (annual) goal. Note that a fleet is not allowed to “bank” hours by flying more than required earlier in the year, even if it has the capacity to do so.

While this correction technique can allow the fleet to make up shortfalls in the flying rate, it may also be counter-productive. Increasing the flying rate increases the speed at which tails approach maintenance, so if the maintenance burden is too great, the methodology will counter-act the effectiveness of this correction technique.

4.4.3.5 Summary

This procedure for determining the ashore daily flying rate has the flexibility and speed required to deal with the day-to-day changes in the ashore serviceable pool, given that it is a component of a larger life of type model. These requirements preclude the use of typical optimisation techniques, such as mixed integer programming as in Kozanidis *et al.* [24], for this type of problem. Here, tails can be added or removed from the ashore serviceable pool every day, by embarking or disembarking, entering or exiting deep maintenance, or experiencing unscheduled maintenance.

This ashore flying rate procedure can also easily accommodate changes in maintenance durations when unscheduled maintenance is discovered on tails in phased maintenance. In this case, the additional duration of unscheduled maintenance is added to the days until that line becomes free. This means that this maintenance line is now in a different position relative to the other lines and may therefore be matched to a different tail.

In these circumstances, the model simply applies the same procedure every day regardless of the number of serviceable tails that day and how many airframe hours remaining on each tail. The model does not ‘remember’ the placing of those tails in unscheduled maintenance, since others may ‘move past’ their original position during their unserviceability. When the unscheduled maintenance tails return from maintenance, they will be added back into the pool and matched to a phased line as before.

This procedure can also manage tails that are due for DM when DM servicing is flying-hours based. When their next service is a DM (which will be combined with phased services), they are managed separately from the rest of the fleet. However, DM frequency may be tracked in elapsed time. This contrasts with phased maintenance, which is generally tracked in airframe hours. Where possible, a phased service is combined with a DM, in accordance with general practice of RAN engineers. If a tail is within a certain duration of a DM, the tail will fly at a rate consistent with the above description, except that the denominator will be days until DM is due. If r is close to R when this period commences, the tail may be unable to fly sufficient hours to combine the two services

(particularly if affected by unscheduled maintenance). Conversely if r is close to zero, the tail's flying rate will be very slow until DM. Tails due for DM are prioritised first when calculating the daily flying rate for the fleet, so tails due for phased maintenance will fly the remaining hours on a given day following allocations to the DM-due tails. A more considered treatment of this situation is a potential area for further work.

One shortfall of the current procedure is that it does not take account of any short-term planning horizon, other than the described adjustments to the maximum fleet and tail rates when falling behind the annual *pro rata* requirement. In this respect it differs from the literature (e.g. Winata *et al.* [26] & Kozanidis *et al.* [24]). There is also no limit as to how much a tail can fly, other than the daily limit. This may lead to unrealistic allocations of hours to particular tails over short-term periods of weeks or months, although the tail rotation procedure will reduce this effect somewhat. Further work will seek to add a short-term planning capability to the daily flying program methodology.

While this procedure has been designed to maximise serviceability, this procedure is not necessarily optimal. Other procedures may seek to maximise flying hours (as in Winata *et al.* [26]) or maximise the minimum ashore serviceability over a given period (as in Mattila & Virtanen [27]). Future work will seek to examine these techniques in the dynamic environment described here. This will be done in the context of balancing flying requirements, not just over the short or medium term, but over the fleet life, such that no tail(s) reaches life of type well before others.

A sample implementation of the methodology is found in Appendix A.

4.4.4 Updating tails

Once flying rates are determined and the type of day has been established, tails are updated in their current state during a given day. Serviceable tails accumulate flying hours based on the amounts determined by the embarked and ashore daily flying rates. They will therefore accrue flying hours towards the total annual requirement for the fleet. Individually they will accrue airframe hours towards the commencement of the next phased service, as well as time towards their next DM. Tails on exercise will accumulate hours towards the ashore annual requirement. If tails are embarked, they will individually accrue hours towards the completion of the embarkation, and collectively accumulate hours towards the embarked annual requirement. Tails in maintenance will accumulate an extra day in maintenance (if it is a work day). Parameters recorded in calendar days, such as time embarked, will increment every day.

4.4.5 State transitions

Testing for transitions of state occurs at the end of the day. If certain criteria are met, tails will move to a new state. These sub-procedures include tests for:

- Tails entering and exiting various types of maintenance;
- Tails embarking to and disembarking from ships; and

- Other tests, such as being deployed on exercises, suffering attrition, or reaching life of type.

A list of all possible state transitions allowed by the model is given in Appendix B.

4.4.6 Maintenance transitions

4.4.6.1 *Scheduled maintenance*

Maintenance transitions involve serviceable tails either entering or exiting maintenance. For deterministic maintenance types such as regular, phased or deep maintenance, this occurs after reaching a certain elapsed time or flying hours. Extensions to these limits may be considered in practice but only in exceptional circumstances, so they are not considered here.

If maintenance capacity is exceeded for a particular maintenance line, tails may be forced to wait in a first-in first-out queue. A tail can only then enter maintenance once a maintenance line is vacated.

Ideally, scheduled maintenance will be staggered to minimise queuing. For time-based deep maintenance, the maintenance stagger is managed by programming tails to enter deep maintenance on a prescribed date. Where plausible, when deep maintenance is due, it is combined with a phased service. For flying hours-based deep maintenance, a phased service will always occur in conjunction with deep maintenance. In this case, the daily ashore flying program methodology automatically manages the deep maintenance stagger. For phased maintenance, the daily ashore program also dynamically manages the phased maintenance stagger, as explained in Section 4.4.3. For flight-line maintenance, a maintenance stagger for tails entering regular maintenance is not attempted. This is because serviceable tails that randomly fall due for unscheduled maintenance also use flight-line maintenance, which may continually disrupt the stagger.

Tests are performed when deterministic maintenance occurs to see whether the timing coincides with a designated modification type (e.g. an annual modification). Modifications are due if the tail is in the appropriate maintenance type and the specified time to the next modification has elapsed. If due, once the original maintenance is undertaken, the aircraft remains in its current maintenance state (and maintenance line) until the modifications are also completed. Examples of upgrade programs may be:

- Regular (e.g. annual) to make miscellaneous improvements, that may take days;
- Frequent (e.g. every few years) to make more significant improvements, e.g. to the sensor suite to update the system to accommodate the latest software improvements, that may take days or weeks; and/or
- At the mid-life of an aircraft. This may involve an extensive re-fit of the entire aircraft, such as replacing parts of the fuselage, rotor blades and other work, that may take weeks or months.

All modification work is completed ashore. Regular modifications may be undertaken following a phased service. The other modification types are more detailed and can only be prosecuted following deep maintenance. These modification types may be combined if their timing coincides.

Once the designated time in maintenance has been reached, the aircraft can return to a serviceable state.

4.4.6.2 *Unscheduled maintenance*

Tails that are eligible for unscheduled maintenance are those that are embarked serviceable, ashore serviceable, embarked in phased maintenance or ashore in phased maintenance. Tails undergoing regular inspections, deep maintenance, or awaiting maintenance are not eligible.

The criteria for an unscheduled maintenance event are determined from the parameters of a log-normal distribution as discussed in Section 4.2. The duration of an unscheduled maintenance event, and the length of time till the next event (frequency) are determined as follows. For a given tail, a random number X is computed from the log-normal distribution with mean μ and standard deviation σ as in Equation 4-17:

$$X = e^{\mu + \sigma R}$$

Equation 4-17 : Equation for determining frequency or duration of unscheduled maintenance event in the model

Here, R is a seed from a normal distribution with a mean of zero and a standard deviation of one.

If the aircraft experiencing unscheduled maintenance was serviceable, it moves to a flight-line maintenance line (if capacity exists). If it was already in scheduled maintenance, it remains in its current maintenance line until the unscheduled maintenance is completed. Following this, any remaining scheduled maintenance must also be completed before it becomes serviceable.

4.4.6.3 *Logistics delays*

Logistics delays are added to the duration of an unscheduled maintenance event to allow for the unavailability of the required spare parts. Logistics delays are measured in terms of the percentage of calendar days: i.e., a logistics delay of 5% means that an average of 5 days in every 100 is spent awaiting parts due to logistics delays. However, unscheduled maintenance events are random and must occur for logistics delays to be applied. When an unscheduled maintenance event occurs, the actual unscheduled maintenance duration is determined by Equation 4-18:

$$u_{dur,l} = u_{dur} \left(1 + l \frac{E[U_{freq}]}{E[U_{dur}]} \right)$$

Equation 4-18: Calculation of unscheduled maintenance duration when including logistics effects

Here, u_{dur} is the initial unscheduled maintenance duration (determined probabilistically), l is the logistics delay, and $E[U_{freq}]$ and $E[U_{dur}]$ are the expected values of unscheduled maintenance frequency and duration respectively. So if on average an unscheduled maintenance event happens every 33.3 days and requires 10 days to repair, there will be 30 days out of every 100 where an aircraft will be unserviceable. Applying the above factor for a logistics rate of 5% would add 16.7% to the unscheduled maintenance duration for each event, which equates to a 1 day logistics increase for every 6 days in unscheduled maintenance. Therefore, an additional $30 * 16.7\% = 5$ days in every 100 would result in unserviceable aircraft due to logistics delays.

Logistics delays are assumed not to apply for deterministic maintenance. This is because these are regular planned services, so any items to be replaced are known in advance.

4.4.7 Embarking and disembarking aircraft

The model always seeks to keep the number of embarked aircraft at the minimum required. If the number of embarked aircraft is detected to be below the minimum, replacement tails are immediately sought. Tails must be capable of flying immediately, so only tails from the ashore serviceable pool are eligible for consideration. Multiple tails can be replaced on a given day if required. It is assumed that tails stay with their ship for the duration of their embarkation and thus do not transfer between ships.

Tails must be able to be embarked for the randomly-determined required number of flying hours and must not fall due for DM while embarked. The required flying hours will also depend on the type of embarked tail being replaced (surge, normal or low rate of effort). Of the tails that meet these criteria, the one(s) with the lowest number of total embarked hours will be chosen, so as to balance the total embarked hours across the fleet over the fleet life. Chosen tail(s) would be embarked for between six and twelve months. The actual duration is randomly determined from a uniform distribution.

If no tails meet all of the criteria, the model will record that day as not meeting the minimum embarked requirement and will use the same procedure to attempt to find replacements at the next time step under the same criteria. If further replacement tails are required but cannot be found immediately, a backlog will ensue. When this occurs, replacements will be sought in a first-in first-out queue. For example, consider that a surge tail needs to be replaced and needs to fly 300 hours, but there are either no ashore serviceable tails, or all that are will be due for DM in that time. Several days later a normal tail needs to be replaced that requires 200 hours, but again no ashore serviceable tail meets the criteria. Any tails that exit maintenance and become ashore serviceable following that time will be tested to replace the surge tail requiring 300 hours first.

In reality, tails will be pre-prepared for embarkation, usually months in advance. However, for simplicity it is assumed that the tail to be replaced is chosen on the day. Complications would arise in the model if the chosen tail suddenly experienced a major unscheduled maintenance event, thus requiring a replacement. Developing procedures for examining how to choose tails in advance is a potential area for further research.

Disembarkation occurs when tails that are flying at the surge or normal rate have achieved their pre-determined number of flying hours. Therefore they are disembarked in a serviceable state. These tails can also be automatically disembarked if they have been unable to achieve their hours within the 12 month maximum embarkation period. This usually occurs when a tail has experienced an unexpectedly large amount of unscheduled maintenance during its embarkation.

For low rate of effort tails, they are disembarked under two possible criteria. One is when a low tail has exceeded its initially planned embarked duration, as its hours flown depends on the rate flown by the surge and normal tails. The other is when a low tail has reached its flying hours until its next DM service. This only occurs for tails with flying hours-based DM when these low tails exceed their initially planned flying rate. This occurs when the surge or normal tails are unable to meet their rates due to high unserviceability, or due to an overall deficiency in embarked aircraft numbers.

Sometimes, a tail flying at the low rate of effort may accumulate a very small number of flying hours during its embarkation. In this case, it will disembark once its planned embarked duration has elapsed. When a replacement tail is sought, it is likely that this same tail will re-embark immediately as it may still have the lowest number of embarked hours of any ashore serviceable tail. To prevent the continued re-embarkation of the same tail, the model allows a maximum of three consecutive embarkations per tail.

Tails that embark to and disembark from ships undergo a changeover period of one to two weeks to allow for the transit time between the ship and the home base. For the purposes of tracking the number of embarked tails, they are still considered to be embarked while they are undergoing changeover.

4.4.8 Other transitions

4.4.8.1 *Exercise periods*

A designated number of ashore tails may be chosen for detachment to an exercise for certain periods of time throughout the year. The chosen tails will have the most hours remaining until their next phased service. During these periods they will fly at a higher rate of effort than their normal ashore rate.

No scheduled maintenance occurs during the exercise period. If tails do not have sufficient hours to fly all of those required during the exercise, they fly all of the hours remaining until their phased service is due. Unscheduled maintenance events may occur according to the rules in Section 4.4.6.2, but they do not affect the serviceability of the aircraft during

the exercise period. Instead, these events accumulate and so are “carried over”³ to be completed on return to the squadron.

The daily flying rate during exercise periods is increased by 50% by combining the flying rates of the tails on exercise and those remaining at the squadron. This usually means that the daily requirement for the remaining ashore aircraft is correspondingly reduced for the exercise period.

4.4.8.2 Attrition and life of type

Serviceable aircraft can be randomly lost through attrition: i.e. the permanent loss of the aircraft from the fleet, for example through a crash. If this occurs, the lost aircraft is immediately removed from the fleet and the fleet size is decremented. The user defines the number of attrited aircraft as an input, and the model randomly pre-determines the days on which aircraft can be lost based on that input number and the aircraft fleet life.

Serviceable aircraft may also reach the maximum designated number of flying hours for the aircraft during operations. As for attrition, when this occurs, the aircraft is removed from the fleet.

If in either instance the particular serviceable aircraft is embarked, the model records that the number of embarked aircraft is below the minimum and will seek to find a replacement using the procedure described in Section 4.4.7.

4.5 Outputs

The simulation chronologically steps through events in daily increments, and at each increment the following outputs are produced:

- Number of tails in each state on the current day;
- Residual time in each state on the current day (e.g. the hours since the last service, or hours achieved while embarked); and
- Cumulative results up to the current day.

These daily outputs can be collated into cumulative results from which relevant statistics and graphical outputs can be generated.

Once the clock has reached the end of the fleet’s life, the model publishes all the relevant statistics as well as numerous graphical results. The model can be tailored to produce a range of outputs depending on the focus of the study. For example, some of the current outputs include the primary Measures of Effectiveness (MOEs):

- The percentage of time the minimum requirement for embarked aircraft is met,

³ Such unscheduled maintenance events are known as Carried Forward Unserviceabilities (CFUs). CFUs are defects which are acknowledged but are deemed to be no threat to safety or to adversely affect mission capability.

- The percentage of time that the required minimum number of ashore serviceable aircraft is met, and
- The annual embarked and ashore hours achieved by the fleet.

Other parameters that are calculated include:

- Availability;
- Serviceability (embarked and ashore);
- Average daily hours flown (embarked and ashore);
- Average number of ashore serviceable tails;
- Average number of ashore serviceable tails that are idle;
- The percentage of time that the minimum requirement for ashore serviceable aircraft is met on all flying days (analogous to the embarked requirement);
- Annual days in maintenance (e.g., scheduled and unscheduled maintenance for embarked and ashore aircraft);
- Annual days queued awaiting various types of maintenance (flight-line, phased or deep maintenance).

Some of these parameters can be recorded for a single run, while others require multiple runs. Some can be recorded for both, such as the parameters recording average annual data. A single run allows a state plot to be generated, which allows the individual states of each tail to be tracked with time.

This range of outputs provides data for analysis that informs the fleet size. It also can inform the input values of other parameters, such as the workforce rate of effort and the number of maintenance lines. Changes to these or other parameters will also potentially influence fleet size. Trade-off analysis is therefore possible using this methodology: for example, if a particular fleet size is found to be insufficient, it may be more cost-effective to increase the manpower rate of effort rather than increase the fleet size, and the model allows the effects of varying both to be tested. As such the fleet-sizing model developed allows investigation into the sensitivity of factors other than aircraft numbers in determining the required fleet size.

4.6 Summary

The discrete-event simulation model provides a thorough treatment of the fleet system. This includes deterministic and stochastic effects and detailed procedures for calculating flying rates.

5. Simulation model verification and validation

The primary Defence stakeholders in the outcomes of the fleet-sizing model were Capability Development Group (CDG), Defence Materiel Organisation (DMO) and the Royal Australian Navy. Engagement with stakeholders was strong and continuous through the development of the model.

Stakeholder engagement occurred in numerous ways through visits, telephone calls, emails and meetings. CDG and DMO provided general information about procedures and processes during model development. The RAN provided guidance on the intent behind determining embarked and ashore flying rates. The RAN and DMO both assisted with the understanding of maintenance processes and provided data that was used to determine the appropriate treatment of unscheduled maintenance. Information provided by CDG, DMO and RAN personnel was then translated into algorithms and, ultimately, code for the discrete-event simulation model.

Once the model was sufficiently mature, endorsement and verification of model design was sought at a meeting of the stakeholders group. All aspects of the model were described step-by-step so that stakeholders could see how their inputs had been included. The stakeholders agreed that the model appropriately represented procedures, either as currently practiced by RAN personnel for the S-70B-2, or as would reasonably be expected for the Phase 8 aircraft. The meeting was minuted and any further changes required were subsequently implemented.

Validation of model results was undertaken by testing results against the annual hours flown by the S-70B-2 Seahawk fleet recorded in the Defence Annual Reports. Input data was taken from the current Seahawk fleet regarding maintenance information, embarked and ashore flying rates and other procedures.

Figure 5-1 shows the comparison of outputs from the fleet-sizing model with actual Defence Annual Reports data. Such Defence output data has only been recorded from the 1999-2000 financial year. The results indicate that the model outputs provide almost an average value of the Defence outputs.




Figure 5-1: Results of validation of fleet-sizing model against Defence Annual Reports data

While undertaking validation of model results was a useful exercise, obtaining an accurate validation was not possible as the quality of the input data was not appropriate. The reasons for this are specifically related to some significant one-off events during the life of the S-70B-2 fleet, such as:

- A maintenance regime overhaul several years into the life of the fleet. This led to the new implementation of a deep maintenance service which was not previously practiced; and
- A major review of all RAN aircraft maintenance procedures following the crash of an RAN Sea King helicopter in 2005. Aircraft availability was lower for the period following the implementation of the review's recommendations (as seen in the later part of the Defence Annual Reports results in Figure 5-1).

The model assumes that the inputs (such as maintenance information) are unchanged throughout the fleet life. Additionally, being a Monte Carlo simulation model, the comparison of results had to be undertaken with multiple runs of the simulation, with random factors such as unscheduled maintenance effectively averaged out. By contrast, the annual hours flown equate to the results from a 'single run', equivalent to a single fleet life, including the manifestation of random instances of unscheduled maintenance. Replication of the results would require precise knowledge of the frequency and duration of every unscheduled maintenance event.

Combined with the maintenance changes described above, this led to complications in obtaining an accurate validation. Nevertheless, when accounting for these difficulties, the results of the validation were deemed to be acceptable by the stakeholders.

Subsequent to the delivery of reports containing the results and analysis from the fleet-sizing model, the analysis team and others associated with the project received a Defence Support Services Gold Level Commendation from the Chief Defence Scientist for the fleet-sizing work. An excerpt from the award citation states that “the ultimate measure of your success has been the unconditional acceptance of the results of your analyses by the key stakeholders associated with this Project”. This acceptance was a direct result of the continual stakeholder engagement practiced throughout the verification and validation process. This process provided the stakeholders with confidence that the model outputs would provide reliable results when testing various fleet-sizing options.

6. Simulation model indicative results

This section provides results designed to demonstrate the capability of the fleet-sizing model in addressing the questions posed. Input data and the input requirements are indicative only and do not relate to any actual aircraft type or capability requirement.

6.1 Input data

Table 6-1 specifies the input requirements for this particular problem. The minimum embarked requirement is for eight aircraft as given in the 2009 Defence White Paper [1]. The embarked flying hour requirement is split between the flights, noting the different embarked flying rates as explained in Section 4.4.2. Aircraft may be embarked for between six and twelve months. The fleet life is 30 years. Life of type issues with aircraft are not considered.

Table 6-1: Sample input requirements for fleet-sizing model

Input	Units	Value
Minimum embarked	Number	8
Annual embarked requirement	Flying hours	3200
Annual ashore requirement	Flying hours	4400
Minimum embarked duration	Calendar days	180
Maximum additional embarked duration	Calendar days	185
Fleet life	Years	30

Table 6-2 provides the relevant parameters for the ashore and embarked flying program. As mentioned in Section 4.4.3, the daily ashore flying rate per tail may be increased if the achieved daily rate falls behind the *pro rata* rate.

Table 6-2: Sample daily ashore and embarked flying program parameters

Input	Units	Value
Maximum daily ashore flying rate	Flying hours/day	24
Maximum daily ashore flying rate per tail	Flying hours/day	6
Minimum sortie length ashore	Flying hours	1
Maximum allowable tail idle days	Work days	5
Maximum daily embarked flying rate per tail	Flying hours/day	10
Minimum sortie length embarked	Flying hours/day	0.5

Table 6-3 gives the input values for aircraft maintenance data covering maintenance capacity, types of maintenance and their frequency and duration. The unscheduled maintenance parameters are those for a log-normal distribution as described in Section 4.2. Logistics delays are calculated as described in Section 4.4.6.3.

Table 6-3: Sample maintenance input data for fleet-sizing model

Input	Units	Value
Frequency of regular inspections	Flying hours	50
Frequency of phased services	Flying hours	250
Frequency of deep maintenance services	Flying hours	1000
Duration of regular inspection	Man-hours	100
Duration of phased maintenance (1)	Man-hours	1000
Duration of phased maintenance (2)	Man-hours	900
Duration of phased maintenance (3)	Man-hours	1000
Duration of phased maintenance (4)	Man-hours	1200
Duration of deep maintenance	Work days	40
Maximum additional deep maintenance duration	Work days	20
Unscheduled maintenance frequency (mean/standard deviation)	Calendar days	3.10/0.69
Unscheduled maintenance duration (mean/standard deviation)	Calendar days	1.77/0.92
Logistics delays	% of Calendar days	3

Four phased services of varying duration are included, with the fourth phased service coinciding with deep maintenance. Deep maintenance takes at least two calendar months. Here a random period of up to an extra month of DM is included to allow for any unexpected unscheduled maintenance. This additional duration is drawn from a uniform distribution.

Table 6-4 shows the data regarding the maintenance workforce, including the number of maintenance lines for each maintenance type and the workforce rate of effort. There is no queuing of embarked aircraft for maintenance as each ship has its own maintenance team for its flight.

Table 6-4: Sample maintenance capacity and workforce rate of effort data for fleet-sizing model

Input	Units	Value
Number of flight-line maintenance lines	Number	5
Number of phased maintenance lines	Number	4
Number of deep maintenance lines	Number	3
Maintenance workrate (embarked)	Man-hours/day	80
Maintenance workrate (ashore - flight-line)	Man-hours/day	40
Maintenance workrate (ashore - phased)	Man-hours/day	60

Table 6-5 gives the details on the exercise periods included in the model. There are two exercise periods, each of one month duration of varying intensity of effort.

Table 6-5: Sample annual exercise regime for fleet-sizing model

Exercise time of year	Exercise duration	# aircraft	Required hours/tail
March	4 weeks	2	80
November	4 weeks	3	100

Finally, details regarding the work days for crews are provided in Table 6-6. Note that embarked crews work every day, whereas ashore aircrew and maintainers have different holiday periods. The summer flying standdown period begins one week before and ends one week after the maintenance standdown, allowing maintainers the opportunity to prepare serviceable tails before aircrew return from their holidays.

Table 6-6: Sample crew work day information for fleet-sizing model

Crew type	Work days/year	Non-work days
Embarked crews (all)	365	o None
Ashore aircrews	220	o Weekends o 2 weeks winter (Jun) o 6 weeks summer (Dec/Jan)
Ashore maintainers	240	o Weekends o 4 weeks summer (Dec/Jan)

Three sets of results are considered for an indicative analysis. The first set of results in Section 6.2 assumes that no aircraft are lost through attrition and that no modification or upgrade program is pursued, so the fleet conditions are effectively unchanged over its life. The second set of results in Section 6.3 includes attrition but excludes a modification program. The third set of results in Section 6.4 includes the effects of both.

Additional results are provided to demonstrate the utility of the model. Section 6.5 shows how the model can be used to undertake sensitivity analysis on various parameters. Section 6.6 gives an example of trade-off analysis, where results for different fleet sizes are compared with results for various maintenance regimes.

In all cases results are provided for 100 runs. On an Acer Veriton T661 computer with an Intel® Core™ 2 Duo CPU running at 3.16 GHz with 1.98GB of RAM, using MATLAB version R2012a, 100 runs takes around 35-40 minutes.

6.2 Results excluding attrition and modifications

Initial results are given where the effects of attrition and a modification program are excluded. Table 6-7 gives the results for the three primary MOEs, relating to the ability of the fleet to meet the requirements for a minimum number of embarked aircraft daily, and ashore and embarked flying hours annually. Mean results are shown over the 100 runs,

with '+/-' column showing the size of the 95% confidence interval. The small variability in these results indicates that 100 runs are an appropriate number for statistical significance.

Table 6-7: Primary MOE results using sample data and requirements for varying fleet sizes, ignoring attrition and modifications

Fleet size	%t = 8 flights		Annual ashore hours		Annual embarked hours	
	Mean	+/-	Mean	+/-	Mean	+/-
20	99.8	0	4254	6	3198	1
21	99.9	0	4329	4	3198	1
22	99.9	0	4366	3	3198	1
23	99.9	0	4378	2	3198	1
24	99.9	0	4385	2	3198	1
25	100	0	4386	2	3198	1

The MOEs for both the minimum number of embarked aircraft and the embarked hours achieved are both met. Meeting the requirement for eight embarked aircraft 99.8% of the time equates to less than one whole day every year. The embarked hours requirement is unlikely to be met precisely due to the vagaries of unscheduled maintenance. Tails may be assigned to fly on a given day, including the last day of the year, but may unexpectedly experience unscheduled maintenance, resulting in a deficiency.

The ashore hours achieved is more variable. The model places a priority on meeting the embarked requirements first - e.g., in immediately seeking to replace a tail that has disembarked. Consequently, the ashore fleet must meet its requirement with whichever ashore serviceable tails remain. Therefore a deficiency in fleet size will be evident when the ashore hours requirement is not met.

% ashore & embarked hours with fleet size: 4400 hr requirement, no att or mods, 5fl/4ph/3DM




Figure 6-1: Graph showing the variation of mean annual ashore and embarked hours with fleet size when attrition and modifications are excluded using sample data and requirements

Figure 6-1 shows the variation in mean annual ashore and embarked hours achieved with fleet size in graphical form. This graph plots the hours achieved as a percentage of the requirement. This shows that as fleet size increases, the ashore hours requirement is more likely to be met. Smaller fleet sizes are less likely to provide serviceable aircraft to meet the ashore flying requirements. This is exacerbated when many tails experience unscheduled maintenance, which affects the ability of the fleet to meet the requirement in a given year and thus the average results over 30 years. Larger fleets however have enough spare capacity to cover these losses. As the fleet size increases, there are more years of the fleet life in which the annual requirement is met and so the averages are higher.

These results pose interesting questions for a decision maker on fleet size. Clearly a fleet size of 25 provides statistically no advantage over a size of 24. On average, there are only 20 hours more achieved by a fleet of 25 over a fleet of 22 in this case, and a fleet size of 22 still provides more than 99% of the required ashore hours. The steeper decline in achieved ashore hours for fleet sizes of 20 and 21 indicates the greater difficulty that smaller fleets have in meeting the requirements.




Figure 6-2: Graph showing the number of ashore serviceable tails, including idle tails for varying fleet sizes when attrition and a modification program are excluded using sample data and requirements

Figure 6-2 shows the results for the average number of daily ashore serviceable aircraft. It also includes the average number of those tails that are idle on a given day.

Firstly, these results show that adding one extra tail to the fleet size does not result in one extra ashore serviceable aircraft. The extra tail will require maintenance within the confines of the existing maintenance capacity. Secondly, the number of idle tails increases with fleet size. This is due to the methodology behind the daily ashore flying program, where not all serviceable tails need to fly on a given day. The more ashore serviceable tails present, the more idle tails there will also be.

The number of idle tails may also be considered to be an indicator of the amount of 'slack' in the fleet size. The preferred number of idle tails depends on the amount of unscheduled maintenance and the views of the stakeholders. Too few idle tails may mean that there will not be enough to replace any that break on a given day, resulting in an inability to meet the day's flying requirements. Conversely, too many idle tails indicates an excessive fleet size with the potential for wastage of resources.

A question that may arise from these results is: if the ashore requirements are not met for smaller fleet sizes, why is there at least one idle tail every day for those cases? This is

because these results are the average over the whole 30 year fleet life. For some years of the fleet life the annual requirement *is* met, so in those years there are likely to be more idle aircraft. However, for smaller fleet sizes there are also more years when it is not met (shown in the average results), so there will be less idle aircraft. Therefore, as fleet size increases, the average number of idle aircraft will increase, as the ashore requirements will be met for more years of the fleet life.




Figure 6-3: Graph showing average annual days awaiting various types of maintenance for varying fleet sizes when attrition and a modification program are excluded using sample data and requirements

Figure 6-3 shows the variation in average annual days for tails queued for the three different types of maintenance. Flight-line maintenance queuing increases significantly with fleet size, which is expected since there is no increase in flight-line maintenance capacity as fleet size is increased. However, phased and deep maintenance queuing does not increase with fleet size in the same way. This is because the daily ashore flying program is designed to minimise queuing for phased maintenance by adjusting the daily flying rate for each ashore serviceable aircraft. In this case it acts in the same way for deep maintenance, since the frequency of deep maintenance is also based on achieved flying hours. These results serve to indicate that the daily ashore flying program methodology is achieving its objectives regardless of the fleet size.

Table 6-8 provides the results shown in Figure 6-2 and Figure 6-3 in tabular form.

Table 6-8: Additional results when excluding attrition or modifications using sample data and requirements

Fleet size	# Ash Serv		# Idle Ash serv		Flt line maint queuing		Phased maint queuing		DM queuing	
	Mean	+/-	Mean	+/-	Mean	+/-	Mean	+/-	Mean	+/-
20	4.8	0	0.8	0	18.7	1.3	4.4	0.7	23.3	2.3
21	5.3	0	1.1	0	37.7	1.9	5.8	1.1	29.2	3.0
22	5.9	0	1.4	0	71.2	3.0	6	0.9	35.7	3.9
23	6.4	0	1.7	0	122.7	3.9	8	1.5	32.5	2.9
24	6.9	0	2.1	0	192.2	6.3	8.5	1.5	34.7	3.7
25	7.4	0	2.4	0	295.8	6.7	9	1.3	35.9	4.3

These results considered in combination indicate that it is more than simply the notion of “what is the minimum fleet size” that needs to be considered. Incorporated into any decision should be questions such as:

- What improvement does an increase of one additional aircraft provide in terms of hours achieved?
- How much slack should there be in the fleet size to accommodate unscheduled maintenance on a given day?
- What threshold for maintenance queuing should be considered as prohibitive?

6.3 Results including attrition

Aircraft fleets may lose aircraft through accidents. Table 6-9 shows the results against the primary MOEs when three aircraft are randomly lost through attrition over the fleet life.

Table 6-9: Primary MOE results when attrition of three aircraft is included for varying fleet sizes using sample data and requirements

Fleet size	%t = 8 flights		Annual ashore hours		Annual embarked hours	
	Mean	+/-	Mean	+/-	Mean	+/-
20	99.2	0.1	3981	0.3	3199	1
21	99.6	0.1	4146	0.2	3198	1
22	99.8	0	4255	0.1	3198	1
23	99.8	0	4324	0.1	3198	1
24	99.9	0	4361	0.1	3198	1
25	99.9	0	4377	0.1	3197	1

The results in Table 6-9 show that there is a notable decline in performance of smaller fleets in meeting the ashore hours requirement. Figure 6-4 demonstrates this visually by plotting the average annual hours achieved for various fleet sizes for each year of the fleet life.




Figure 6-4: Variation of mean annual ashore hours achieved for each year of the fleet life when including the attrition of three aircraft using sample data and requirements

The impact of attrition on smaller fleets is clearly shown in Figure 6-4 by the steady decrease in ashore hours achieved each year of the fleet life. Given that these results are averaged over 100 runs, random attrition events are spread evenly over the fleet life, causing the steady decline. Conversely larger fleets (e.g. 25) are able to absorb the loss of aircraft and are relatively unaffected by attrition.

Table 6-10: Additional results when including attrition using sample data and requirements

Fleet size	# Ash Serv		# Idle Ash serv		Flt line maint queuing		Phased maint queuing		DM queuing	
	Mean	+/-	Mean	+/-	Mean	+/-	Mean	+/-	Mean	+/-
20	3.9	0.1	0.5	0	7.7	1.1	2.6	0.6	19.2	2.0
21	4.5	0.1	0.7	0	16.5	2.1	3.4	0.7	24.2	2.3
22	5.0	0.1	0.9	0	34.3	3.9	4.4	0.9	32.4	3.3
23	5.6	0.1	1.2	0	57	5.6	6.0	1.1	31.4	3.0
24	6.1	0.1	1.5	0.1	101	8.6	6.9	1.2	31.2	3.6
25	6.6	0.1	1.8	0.1	168.5	12.7	7.7	1.2	34	3.5

Table 6-10 shows the results for the average number of ashore serviceable aircraft, including idle aircraft, and the amount of maintenance queuing. When compared with the results for ashore serviceable aircraft in Table 6-8 it is evident that attrition leads to a smaller number of ashore serviceable aircraft, and thus less idle aircraft. Less aircraft in the fleet means that there will be less available, and therefore serviceable, and more of the serviceable aircraft will be required to fly on a given day to meet the flying rate. When compared with the results for maintenance queuing in Table 6-8, it is seen that the amount of queuing for flight-line maintenance has decreased significantly for larger fleets. As aircraft are lost from the fleet, maintenance demand is reduced and thus there is less maintenance queuing. Phased and deep maintenance are effectively unchanged, as is desired according to the daily flying program methodology.

6.4 Including attrition and a modification program

The final set of results adds the effect of a generic modification program. It also includes the effect of the loss of three aircraft as in the previous section. Table 6-11 provides information on the sample modification program to be included in the model to see what effect this has on the ability of the fleet to meet its requirements.

Table 6-11: Sample input data for modification program for the fleet-sizing model using sample data and requirements

Input	Units	Value
Minor modification frequency	Years	1
Major modification frequency	Years	5
Mid-life upgrade start time	Years from FOC	15
Minor modification duration	Work days	10
Major modification duration	Work days	30
Mid-life upgrade duration	Work days	120

Results are given in Table 6-12 for the primary MOEs.

Table 6-12: Results for the primary MOEs when attrition and the modification program are included using sample data and requirements

Fleet size	%t = 8 flights		Annual ashore hours		Annual embarked hours	
	Mean	+/-	Mean	+/-	Mean	+/-
20	85.7	0.7	2856	42	3187	2
21	84.8	0.8	2894	44	3186	2
22	84.1	0.9	2877	46	3186	2
23	81.7	0.7	2853	48	3180	3
24	80.6	0.9	2797	50	3179	3
25	78.5	0.7	2748	52	3172	3

Table 6-12 shows that the given fleet sizes are unable to meet either the requirement for eight flights at sea or for ashore hours. The embarked hours achieved is within 1% of the requirement. This demonstrates the capability of the embarked flying program, even in circumstances where the minimum number of embarked aircraft is not met. The embarked flying program methodology will increase the flying rate of the low rate-of-effort tail if the surge and normal tails are not meeting their hours – although in this case, it is because those tails may not even be embarked.

Table 6-13 shows the annual days awaiting various maintenance types for the given fleet sizes when attrition and modifications are included.

Table 6-13: Annual days queuing for maintenance lines when attrition and the modification program are included using sample data and requirements

Fleet size	Flight-line maintenance queuing		Phased maintenance queuing		Deep maintenance queuing	
	Mean	+/-	Mean	+/-	Mean	+/-
20	2.6	0.4	1.5	0.4	564.3	36.1
21	4.9	0.8	2.1	0.7	786.5	42.1
22	6.6	0.9	1.5	0.5	1081.0	56.4
23	9.3	1.1	2.2	0.7	1398.2	61.6
24	12.3	1.7	2.4	0.7	1771.5	68.9
25	17	1.9	3.8	1.2	2159.6	67.2

The results in Table 6-13 show that increased queuing for deep maintenance is the reason why the chosen fleet sizes can no longer meet the primary requirements. This queuing is caused by the increased deep maintenance burden on the fleet resulting from the modification program. Deep maintenance queuing causes a decrease in ashore availability, and therefore ashore serviceable aircraft to meet the ashore daily flying program.

Figure 6-5 shows the annual ashore hours achieved for varying fleet sizes for every year of the fleet life for this case.




Figure 6-5: Annual ashore hours achieved for varying fleet sizes for each year of the fleet life when attrition and a modification program are included using sample data and requirements

The results in Figure 6-5 starkly demonstrate the impact of the modification program on annual hours achieved for each year of the fleet life. The most significant impact is during the mid-life upgrade program where the maintenance burden is at its worst and there is a dramatic decrease in ashore hours. However, even the five-yearly major modification program has a cyclical impact on fleet performance. This level of ashore hours achieved would not be sufficient to train aircrews to participate in embarked operations.

Interestingly, smaller fleets cope better during the mid-life upgrade, as smaller fleets mean less aircraft to upgrade and therefore less queuing for maintenance. Consequently a fleet of 20 recovers more quickly than a fleet of 25. However it is also seen that the maximum hours that a smaller fleet can achieve are also significantly less than for a larger fleet when the effects of maintenance queuing are lower.

In response to this finding, a reasonable approach would be to increase deep maintenance capacity to reduce aircraft queuing.




Figure 6-6: Annual ashore hours achieved for a fleet size of 25 for varying numbers of DM lines when attrition and a modification program are included using sample data and requirements

Figure 6-6 shows the results for a fleet size of 25 when the number of DM lines is increased from three to seven. From the results with the loss of 3 aircraft in Figure 6-4, this fleet size is the most likely to still be able to achieve the annual ashore hours requirement throughout the fleet life.

These results show that one additional DM line is able to reduce the deep maintenance queuing sufficiently except during the mid-life upgrade program. To address the shortfall over the mid-life upgrade period would require significantly more investment in deep maintenance capacity. Even an additional four maintenance lines does not enable the fleet to fully meet the ashore requirements.

Table 6-14: Average annual ashore hours achieved for various fleet sizes and numbers of DM lines using sample data and requirements

Fleet size	Number of DM lines					
	3	4	5	6	7	8
20	2856	3526	3702	3742	3741	3744
21	2894	3679	3884	3929	3950	3951
22	2877	3790	4022	4073	4088	4095
23	2853	3843	4104	4173	4186	4197
24	2797	3892	4181	4242	4258	4266
25	2748	3921	4224	4291	4307	4309
26	2716	3921	4243	4322	4333	4338
27	2684	3919	4263	4343	4349	4352

Table 6-14 shows the average annual ashore hours achieved for a range of fleet sizes and DM lines, with two additional fleet sizes included. As the deep maintenance capacity increases, larger fleet sizes are able to achieve more flying hours. This result indicates that increasing fleet size simultaneously with maintenance capacity may be a way to overcome the deficiency in ashore hours achieved. However, the substantial extra costs involved over this period in infrastructure and personnel would have to be considered.

These results are an important finding in helping to identify potential bottlenecks in the system. While increasing the number of deep maintenance lines is one possible solution, reducing the maintenance duration by paying for more contractor staff is another. This would reduce the duration of the upgrade period and help to remove the backlog of tails.

6.5 Sensitivity analysis

This section demonstrates the ability of the model to conduct sensitivity analysis on some of the model input parameters. Sensitivity analysis is a method of testing how much influence a single parameter may have on the results.

6.5.1 Varying the number of aircraft lost through attrition

It was evident from the results in Section 6.3 that aircraft attrition may have an impact on the ability of smaller fleets to meet the requirements. In this sensitivity analysis the number of aircraft lost is varied from 0 (no attrition) to 4 to test the effect on various fleet sizes.

The first set of results is for a fleet size of 22. For the results when attrition was excluded in Table 6-7, 22 aircraft were able to achieve over 99% of the ashore hours requirements.




Figure 6-7: Sensitivity analysis on the number of attrition aircraft for a fleet size of 22 showing average annual ashore hours for each year of the fleet life using sample data and requirements

The results in Figure 6-7 show how the mean annual ashore hours achieved is relatively constant when there are no aircraft lost. As the amount increases to 4, the decrease in ashore hours amounts to 10% by the end of the fleet life. This compares to a decrease of 2% when 2 aircraft are lost.




Figure 6-8: Sensitivity analysis on the number of attrition aircraft for a fleet size of 25 showing average annual ashore hours using sample data and requirements

Figure 6-8 shows the same set of results for a fleet size of 25. In this case even a loss of 3 aircraft can be absorbed, without a notable decrease in ashore hours.

Table 6-15: Average annual ashore hours achieved across fleet life for various fleet sizes and attrition rates using sample data and requirements

Fleet size	Attrition aircraft				
	0	1	2	3	4
20	4254	4191	4091	3981	3875
21	4329	4291	4230	4146	4051
22	4366	4346	4315	4255	4188
23	4378	4369	4355	4324	4284
24	4385	4382	4372	4361	4336
25	4386	4386	4381	4377	4366

Table 6-15 shows the results for the average annual hours achieved across the fleet life for the fleet sizes of 20 to 25 and the number of attrition aircraft of 0 to 4. Note that the average hours achieved for a fleet size of 21 with no attrition is similar to those for a fleet of 22 with 1 aircraft lost, 23 with 2 aircraft lost, 24 with 3 aircraft lost and 25 with 4 aircraft lost, with just over 1% difference between them. This is perhaps not surprising as the final fleet size

in all instances is 21. Similar trends are seen for final fleet sizes of 22, 23 and 24. However, it is important to note that the requirement is to meet the minimum annual ashore hours for *every year* of the fleet life, not just on average. Using the average result may mask a significant number of years whereby the annual requirement is not met.

6.5.2 Varying the maximum daily flying rate

Varying the maximum daily flying rate for an individual tail leads to changes in the number of aircraft that need to be used on a given day. Previous results used a value of 6 hours/day. Figure 6-9 shows the results from 4 hours/day through to 10 hours/day. Note that the maximum daily flying rate for the fleet is unchanged. Recall that the *actual* daily flying rate is determined based on the ashore requirement and the number of annual flying days.




Figure 6-9: Sensitivity analysis on the maximum ashore daily flying rate per tail for varying fleet sizes showing average annual ashore hours using sample data and requirements

Results from Figure 6-9 show that lowering the maximum daily flying rate requires more ashore serviceable tails, while an increase in the maximum rate means less are required. This has potential implications for the fleet size, as less aircraft are required to generate the same number of ashore flying hours if the daily rate is increased. A fleet of 21 flying at 10 hours per day generates the same number of hours as a fleet of 23 flying at 6 hours per day.

Table 6-16 shows the average number of ashore serviceable aircraft for the respective cases, including the number idle. As expected, an increase in the flying rate leads to an increase in the number of idle aircraft, as less aircraft are required to meet the daily requirement. This is despite the fact that the number of ashore serviceable aircraft is virtually unchanged when varying the maximum daily flying rate.

Table 6-16: Number of ashore serviceable and idle aircraft for varying fleet sizes and maximum daily ashore flying rates using sample data and requirements

Fleet size	4 hrs/day		6 hrs/day		8 hrs/day		10 hrs/day	
	# serv	# idle	# serv	# idle	# serv	# idle	# serv	# idle
20	4.9	0.4	4.8	0.8	4.7	1.1	4.7	1.5
21	5.5	0.6	5.3	1.1	5.3	1.5	5.3	1.9
22	6.0	0.8	5.9	1.4	5.9	1.9	5.8	2.4
23	6.5	1.0	6.4	1.7	6.4	2.3	6.4	2.8
24	7.0	1.3	6.9	2.1	6.9	2.7	6.9	3.2
25	7.4	1.5	7.4	2.4	7.4	3.0	7.4	3.5

Therefore the potential downside of increasing the maximum daily flying rate is increasing the number of idle tails ashore. Conversely, overly constricting the daily flying rate requires more serviceable tails than can be generated, given the constraints of the maintenance system. Fleet managers would need to balance these requirements.

6.5.3 Varying the logistics delay

Varying the logistics delay is a way of testing the impact of problems with the supply chain in providing spare parts when unscheduled maintenance occurs. Here the rates are increased from the 3% of calendar days utilised in the earlier runs through to 6% and 9%. The case where there are no logistics delays are also included. The impact on the average annual ashore hours achieved is shown in Figure 6-10.




Figure 6-10: Sensitivity analysis on logistics delays for varying fleet sizes showing average annual ashore hours using sample data and requirements

Increasing the logistics delays has a serious impact on the ability of a chosen fleet to meet the ashore requirements. A rate of 9% suggests that even 25 aircraft may not be sufficient. Conversely, no logistics delays leads to a decrease in unscheduled maintenance durations, providing more serviceable aircraft ashore. In this case a fleet size of 22 produces more ashore hours than a fleet size of 25 with a 3% delay. These results suggest that ensuring an adequate supply of spare parts is an important factor in meeting fleet requirements.

6.5.4 Varying the ashore requirement

Varying the ashore requirement is a means of testing the maximum flying hours that can be generated by the fleet. Table 6-17 shows the results when the annual ashore requirement varies from 4,000 hours to 6,000 hours.

Table 6-17: Mean ashore hours achieved for various fleet sizes and annual ashore hours requirements using sample data and requirements

Fleet size	Annual ashore hours requirement							
	4,000 hrs	4,400 hrs	4,800 hrs	5,000 hrs	5,200 hrs	5,400 hrs	5,600 hrs	6,000 hrs
20	3962	4254	4446	4425	4305	4330	4354	4352
21	3984	4329	4577	4586	4509	4552	4570	4577
22	3992	4366	4664	4710	4668	4719	4756	4766
23	3994	4378	4717	4798	4794	4854	4895	4910
24	3995	4385	4735	4851	4883	4947	4996	5013
25	3996	4386	4751	4879	4936	5015	5080	5100

As expected, reducing the requirement allows smaller fleet sizes to be feasible. An annual rate of around 4,800 hours appears to be close to the limit for the chosen fleet sizes.

Increasing the requirement leads to two interesting results. One is that, as expected, the maximum annual ashore hours achieved by the fleet increases with the requirement. This indicates that if the fleet were able to “bank” hours by flying faster than the *pro rata* rate, it could meet higher requirements. For example, a fleet size of 24 falls 65 hours short of the requirement for 4,800 hours per year, but achieves almost 5,000 hours when the target is 5,600 hour per year. This is because the *pro rata* rate is higher when the overall requirement is higher. If the *pro rata* rate constraint were relaxed, this would allow the fleet to fly as much as possible when capacity allowed, potentially enabling it to meet higher targets.

Figure 6-11 plots the results from Table 6-17 against the various annual ashore requirements.




Figure 6-11: Sensitivity analysis on the ashore requirement for varying fleet sizes showing average annual ashore hours using sample data and requirements

Figure 6-11 displays the other interesting result in that the hours achieved do not always keep increasing for all fleet sizes. Indeed, in some cases the hours achieved reach a local maximum before increasing again. This is particularly evident for smaller fleet sizes. Of the chosen ashore requirements values, a peak in ashore hours achieved is reached for a fleet size of 20 with a requirement of 4,800 hours (4,446) and for a fleet size of 21 at 5,000 hours (4,586). The other fleet sizes also show this characteristic. However for these fleet sizes, as the requirement increases, the hours achieved increase again before eventually plateauing between 5,600 and 6,000 hours. These results are shown in graphical form in Figure 6-12, where the hours achieved are plotted against the requirement for each fleet size.




Figure 6-12: Results for varying fleet sizes, showing the actual hours achieved against the varying ashore hours requirement using sample data and requirements

The reason for this shape can be explained by the required increase in the daily flying rate and its impact on the maintenance capacity, particularly for deep maintenance. The greater the annual requirement, the higher the *pro rata* flying rate, so aircraft will move more rapidly into scheduled maintenance. For smaller fleets this demand is spread across fewer tails, so as the requirement increases there are less tails available to meet the demand. Therefore there becomes a critical requirement at which point the fleet is capacity constrained and deep maintenance queuing increases. Any increase in the requirement beyond this point leads to more tails entering deep maintenance even more rapidly, leaving even less serviceable tails to meet these requirements. This suggests that there is an optimal flying rate for a fleet which is appropriate for the maintenance capacity. For larger fleets, these effects can still be seen, but the load can be spread more evenly, leading to a steadier upward trend. Eventually, as the requirement increases, these larger fleets reach maximum flying rates based on the maintenance capacity, leading to the maximum ashore hours output and the plateauing effect seen.

6.6 Trade-off analysis

Trade-off analysis can be used to explore the potential costs or benefits of changing the input parameters for comparable items. In the model, fleet size is an input, with an individual aircraft costing many millions of dollars in purchase and through-life support costs. Maintenance capacity is also an input. The cost of constructing a maintenance

facility, as well as the ongoing costs to support the infrastructure and the maintenance workforce, would also run into many millions of dollars. However, overall it may be a cheaper option than purchasing an aircraft. Here, an increase in maintenance capacity is tested using various fleet sizes to examine whether this gives any overall benefits.

The results from Case A given in Table 6-7 are used as the baseline for comparison. One extra line of maintenance of each maintenance type is then added and the results compared in Table 6-18. This approach follows from the results in Figure 6-3 for the average annual days awaiting maintenance for the three maintenance types. For flight-line maintenance there was a sharp increase in queuing with increases in fleet size, so an extra flight-line maintenance line may result in more rapid servicing times and therefore more serviceable aircraft. While the maintenance queuing for phased and deep maintenance were effectively constant, increasing maintenance capacity there will generate an increased flying rate according to the methodology behind the daily flying program, thus generating more annual hours.

Table 6-18: Mean results for annual ashore hours achieved when one extra maintenance line of each type is added to the input maintenance regime using sample data and requirements

Fleet size	Ashore hours achieved			
	Input maintenance regime	With one extra flight-line maintenance line	With one extra phased maintenance line	With one extra deep maintenance line
20	4254	4264	4274	4301
21	4329	4342	4345	4352
22	4366	4373	4374	4378
23	4378	4388	4386	4386
24	4385	4394	4389	4389
25	4386	4396	4392	4392

These results are shown graphically in Figure 6-13. The results for the input maintenance regime are shown in green. The results for adding an extra flight-line, phased and deep maintenance line are given by the red, blue and black colours respectively.




Figure 6-13: Average annual ashore hours for various maintenance regimes using sample data and requirements

In this example it is noted that a fleet size of 23 for the input maintenance regime produces statistically the same annual ashore hours as a fleet size of 22 when an extra maintenance line is added of any type. This also applied for a fleet size of 24 with the input maintenance regime compared to a fleet size of 23 for an extra maintenance line. Adding an extra maintenance line in each case significantly reduces queuing on that line: from 295.8 to 103.6 days for flight-line maintenance queuing, from 9 to 2.4 days for phased maintenance queuing and from 35.9 to 4.6 days for deep maintenance.

For a stakeholder considering various fleet-sizing options, these results suggest that adding extra maintenance capacity may be an alternative solution to purchasing additional aircraft. In this case the costs of purchasing aircraft, including the through-life support costs, may be weighed against the costs of providing additional maintenance capacity. This may be in the form of squadron-level maintenance provided by additional uniformed personnel, or by providing extra funding for a contractor to expand their deep maintenance capacity.

7. Comparison of the methodologies

This section compares the results of the simulation model with those that could be provided using the binomial distribution method (see Section 2.2).

The binomial distribution uses availability and serviceability data as inputs. The simulation model generates availability and serviceability as outputs, based on input data regarding maintenance frequency, duration and capacity. In order to compare the two, the availability and serviceability data from the simulation model are used as inputs to the binomial distribution.

Only results with no attrition or modification program are included. The binomial distribution method implicitly assumes that the fleet dynamics are unchanged over the fleet life, so this is the only appropriate case for comparison.

Table 7-1 shows the mean availability and serviceability results, along with the mean number of embarked and ashore serviceable aircraft, generated from the fleet-sizing model with no attrition or modification program. These are taken from the same results set as shown in Table 6-7. The 95% confidence levels from the simulation model outputs are very small (0.03 for availability, 0.09 for embarked serviceability and 0.16 for ashore serviceability), which indicates that the model outputs provide a reliable estimate of the actual mean.

Table 7-1: Mean availability and serviceability results from fleet-sizing model with no attrition or modification program using sample data and requirements

Fleet size	Availability (%)	Embarked serviceability (%)	Ashore serviceability (%)
20	87.3	71.3	50.2
21	87.7	71.2	51.0
22	88.1	71.4	51.7
23	88.7	71.3	51.7
24	89.1	71.4	51.8
25	89.5	71.4	51.4

To obtain a serviceability level for the fleet requires the total number of serviceable aircraft divided by the total number of available aircraft. The mean number of embarked serviceable and ashore serviceable aircraft are recorded as model outputs. The mean number available is calculated from the output availability and the fleet size. Table 7-2 provides the results. The 95% confidence levels for the number of embarked and ashore serviceable aircraft are less than 0.02 in both cases.

Table 7-2: Calculating mean fleet serviceability from fleet-sizing model using sample data and requirements

Fleet size	Number embarked serviceable	Number ashore serviceable	Number available	Fleet serviceability (%)
20	5.7	4.8	17.5	60.0
21	5.7	5.3	18.4	59.8
22	5.7	5.9	19.4	59.8
23	5.7	6.4	20.4	59.3
24	5.7	6.9	21.4	58.9
25	5.7	7.4	22.4	58.5

Figure 7-1 shows the results when the generated and calculated availability and serviceability numbers from Table 7-1 and Table 7-2 respectively are used as input into the binomial distribution using the method described in Section 2.2. Probability levels given are greater than 90%. The availability and serviceability inputs are the minimum and maximum calculated to cover the results from the full range of fleet sizes examined.




Figure 7-1: Results from binomial distribution using simulation model outputs for availability and serviceability using sample data and requirements

The binomial distribution results indicate that there is a 95% probability that a fleet size of 22 will be sufficient to embark 8 aircraft, using the availability and serviceability rates produced by the simulation model. For a fleet size of 25, using the availability and serviceability rates produced from the simulation (89.5% and 58.5%), the binomial distribution results say that there is a 99% probability that this is sufficient to embark 8 aircraft.

While the results are similar for both methods, clearly care must be taken when comparing them. The binomial distribution only answered the single question regarding a fleet size large enough to embark 8 aircraft. Moreover, this assumed that all of these embarked aircraft would be 100% serviceable and there was no allowance for ashore serviceable aircraft for training. The simulation model was used to answer the single question, plus additional questions requiring the fleet to achieve annual embarked and ashore flying hours. The simulation model implicitly provides a higher “probability” that a fleet size will be sufficient, given the level of detail represented and a statistically significant number of runs.

However, a more robust comparison between the results can be made. From Table 7-2, the average numbers of serviceable and available aircraft across the fleet are provided from the simulation results. It is seen that for a fleet size of 21, on average 11 are serviceable and 18.4 are available. The average serviceability is 59.8% and average availability is 87.7%. Feeding these numbers into the binomial distribution (i.e. $k = 11$ and $p = 0.598$ from Equation 2-2) gives $P = 0.557$ when $n = 18$. That is, there is a 55.7% probability that at least 11 aircraft will be serviceable from 18 available, which is as expected given that 11 is an average number of serviceable aircraft. Using an availability of 87.7% on $n = 18$ gives $N = 21$, which matches the results in Table 7-2. A similar check for other fleet sizes with close to integer numbers of average serviceable aircraft (23 and 25 in Table 7-2) provide similar matches. Although this check says nothing about meeting the minimum number of embarked aircraft, it does indicate consistency in the outputs produced by both methods.

In summary, if there is enough information provided regarding the availability and serviceability of a given fleet, estimating the required fleet size using a simple tool such as the binomial distribution is a useful first step. However, a more detailed simulation model of the type described here is necessary if the problem is more complex. This includes more requirements, such as annual hours flown. It also includes the impact of particular occurrences, such as attrition or modification programs, which affect the fleet capability at different stages of the fleet life. The binomial method is only applicable for average values of availability and therefore assumes no variation from these over the fleet life.

8. Summary

This report describes the methodologies developed in order to determine the fleet size required for the naval combat helicopter capability replacement project (AIR 9000 Phase 8). Two separate methodologies were developed as appropriate for the respective stages of the project, based on the questions that were posed.

Initially, the only question to be considered was the minimum fleet size that could continually embark a minimum number of helicopters on ships. The solution method applied the simple binomial distribution to the problem, by treating the number of embarked aircraft as the number of 'trials' and the fleet size as the number of 'successes'. The outputs then provided a probability of the fleet size being sufficient based on particular levels of serviceability (used as the 'probability of success') and availability (with unavailable aircraft added). The limitations of this method were noted, particularly regarding its inability to answer any questions regarding fleet rate of effort, and in implicitly applying the same serviceability rates to ashore and embarked aircraft.

Subsequently the requirements were changed to include the ability of the fleet to meet annual flying hour requirements for ashore and embarked aircraft. This required the development of a more detailed approach. Given the multiple objectives and the need to include random effects such as unscheduled maintenance, a discrete-event simulation method was adopted. This method is able to represent the state of each individual tail at each time step, including transitions between various states. Each tail will be in one of several possible states on a given day, such as being serviceable embarked or ashore or undergoing some type of maintenance. The number of embarked tails can be tallied to check that the minimum embarked requirement is being met on a given day. The embarked and ashore daily flying rates can be tallied over the course of a year to determine whether or not the annual requirements can be met.

The embarked and ashore daily flying rates are calculated using differing approaches regarding which tails are allocated to fly and how much they fly. Embarked aircraft fly at different tempos, and the daily flying rate for each is initially pre-determined based on their required hours and the expected amount of maintenance. Depending on the actual flying rate of the higher and normal tempo tails, the lower tempo tails adjust their flying rate to offset this to ensure that the total embarked hours are neither greatly exceeded nor underachieved. The ashore flying rate is determined based on the number and flying hours of ashore serviceable aircraft and the number and status of phased maintenance lines. Tails are flown at such a rate as to minimise queuing of tails for maintenance in order to maximise serviceability. Tails with the least number of hours remaining until a service are matched to maintenance lines with the least number of days until they become free. If there are sufficient serviceable tails and sufficient maintenance capacity, the daily requirement, determined *pro rata* from the annual requirement, should be met. This method is simple and fast, and has the necessary flexibility to handle variations caused by the effects of unscheduled maintenance and embarking/disembarking aircraft.

Indicative results show the ability of the simulation model to address the posed questions. It demonstrates the effects of attrition on the ability of smaller fleets to meet the requirements over the fleet life. It also shows the effects of a modification program on the ashore hours achieved, and how this may be somewhat addressed through increasing deep maintenance capacity. Sensitivity analysis provided insights into the importance of various input parameters, such as those for minimising logistics delays. Trade-off analysis indicates that various combinations of fleet size and maintenance capacity may provide similar outputs in ashore hours achieved. This may provide the stakeholder with potential options in whether or not to procure additional aircraft or alternatively resource the maintenance facilities.

The simulation model developed is a useful and detailed tool for representing complex fleet-sizing problems. The methodology is potentially extensible to other platforms, especially if issues such as flying rates and unscheduled maintenance need to be considered.

9. Further work

The Australian Government decided to purchase 24 MH-60R Seahawk helicopters in June 2011 to satisfy the requirements of AIR 9000 Phase 8. The next stage of the analysis described here is to extend some of the methodology developed during the fleet-sizing work and to apply it to fleet management.

The overall aim of any future work will be to help to maximise the operational availability and serviceability of the fleet in order to undertake its required missions at any time. Therefore the main focus of this future work is likely to be the improved management of flying and maintenance activities. This will include:

- Development of short, medium and long-term schedules and plans for managing fleet flying hours;
- Studying the optimal spread of flying hours across the fleet (i.e. from the minimum to maximum at any one time); and
- Optimising maintenance activities to realise efficiencies.

These aspects will extend the methodology developed for the ashore daily flying program in particular. If scheduling is incorporated, some of the integer programming approaches described in the literature review in Section 4.4.3 may well be appropriate.

10. Acknowledgements

This work could not have been completed without the assistance of many personnel from across the Australian Defence Organisation (ADO) and DSTO.

Within the ADO, the authors would like to thank staff from Capability Development Group, Defence Materiel Organisation (particularly the AIR 9000 Phase 8 project and NASPO) and the Royal Australian Navy (particularly the Fleet Air Arm and 816 Squadron) for their experience and inputs. Assistance from RAAF personnel is also acknowledged.

Within DSTO, the authors would also like to thank staff and management from the former Air Operations Division, particularly within Air Operations Research Branch (now within Joint and Operations Analysis Division), for their assistance and support. Several staff from different areas within the former Air Vehicles Division (now within Aerospace Division) also made substantial contributions to the work and this is also gratefully acknowledged. Additional assistance from former Joint Operations Division staff (now within Joint and Operations Analysis Division) is also acknowledged.

11. References

1. "Defending Australia in the Asia Pacific Century: Force 2030", *Defence White Paper 2009*, Second Edition, Commonwealth of Australia, 2009.
2. Niem Tri and Stephen Bocquet, "Analysis methods used in the AIR 7000 Force Mix study," DSTO Technical Report, DSTO-TR-2286, May 2009.
3. Arild Hoff, Henrik Andersson, Marielle Christiansen, Geir Hasle and Arne Løkketangen, "Industrial aspects and literature survey: Fleet composition and routing," *Computers & Operations Research*, **37**, p2041-2061, 2010.
4. Zhi-Chun Lui, William H. K. Lam and A. Sumalee, "Modeling Impact of Transit Operator Fleet Size Under Various Market Regimes with Uncertainty in Network", *Transportation Research Record: Journal of the Transportation Research Board*, number 2063, p 18-27, 2008.
5. Mark A. Turnquist and William C. Jordan, "Fleet sizing under Production Cycles and Uncertain Travel Times", *Transportation Science*, **20**, 4, p227-236, November 1986.
6. Yafeng Du and Randolph Hall, "Fleet sizing and Empty Equipment Redistribution for Center-terminal Transportation Networks," *Management Science*, **43**, 2, p145-157, February 1996.
7. Slawo Wesolkowski and Andrew Billyard, "The Stochastic Fleet Estimation (SaFE) Model", *Proceedings of the 2008 Spring simulation Multiconference*, Ottawa, Canada, April 2008.
8. S. Wesolkowski and Z. Zhu, "Optimizing the Stochastic Fleet Estimation Model Using Evolutionary Computation", *Proceedings of IEEE 2008 World Congress on Computational Intelligence*, Hong Kong, June 2008.
9. Michael Crary, L. K. Nozick and L. R. Whitaker, "Sizing the US destroyer fleet," *European Journal of Operational Research*, **136**, p680-695, 2002.
10. William R. Lesyna, "Sizing Industrial Rail Car Fleets using Discrete-Event Simulation," in P. A. Farrington, H. B. Newbhard, D. T. Sturrock and G. W. Evans (eds.), *Proceedings of the 1999 Winter Simulation Conference*, Phoenix, 5-8 December 1999, p1258-1261.
11. T. Godwin, Ram Gopalan and T. T. Narendran, "Tactical locomotive fleet sizing for freight train operations," *Transportation Research Part E*, **44**, p440-454, 2008.
12. Aliaksandr Shyshou, Irina Gribkovskaia and Jaume Barceló, "A simulation study of the fleet sizing problem arising in offshore anchor handling operations," *European Journal of Operational Research*, **203**, 2010, p230-240.
13. Tuomas Raivio, Eemeli Kuumola, Ville Mattila, Kai Virtanen and Raimo Hämäläinen, "A Simulation Model for Military Aircraft Maintenance and Availability," *Proc. 15th European Simulation Multiconference, ESM2001*, 6-9 June 2001, Prague, Czech Republic, p190-194.
14. Ville Mattila, Kai Virtanen and Tuomas Raivio, "A Simulation Model for Aircraft Maintenance in an Uncertain Operational Environment," *Proc. 17th European Simulation Multiconference*, 9-11 June 2003, Nottingham, England, p456-461.
15. Ville Mattila, Kai Virtanen and Tuomas Raivio, "Improving maintenance decision-making in the Finnish Air Force through Simulation," *Interfaces*, **38**, 3, May-June 2008, p187-201.

16. Thomas N. Cook and Ralph C. DiNicola, "Modeling Combat Maintenance Operations", *Proceedings of the 1984 Annual Reliability and Maintainability Symposium*, San Francisco, USA, 24-26 January 1984, p390-395.
17. Axel Bender, Adrian H. Pincombe and Gregory D. Sherman, "Effects of decay uncertainty in the prediction of life-cycle costing for large scale military capability projects", In Anderssen, R.S., R.D. Braddock and L.T.H. Newham (eds), *18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation*. Cairns, July 2009, pp. 1573-1579.
18. Andrew Goodwin, "A Generic System Dynamics Aircraft Fleet Model," DSTO Technical Report DSTO-TR-2323, Defence Science and Technology Organisation, July 2008.
19. Linda C. Wolstenholme, *Reliability Modelling: A Statistical Approach*, Chapman & Hall/CRC, 1999.
20. Klutke, G-A., Kiessler, P. And Wortman, M., A Critical Look at the Bathtub Curve, *IEEE Transactions on Reliability*, **52**, 1, p125-129, 2003.
21. Achim Sgaslik, "Planning German Army Helicopter Maintenance and Mission Assignment," M. Sc. Thesis, Naval Postgraduate School, Monterey, California, USA, March 1994.
22. Bradley W. Pippin, "Allocating Flight Hours to Army Helicopters," M. Sc. Thesis, Naval Postgraduate School, Monterey, California, USA, June 1998.
23. George Kozanidis, "A Multiobjective Model for Maximising Fleet Availability under the Presence of Flight and Maintenance Requirements," *Journal of Advanced Transportation*, **43**, 2, p155-182, 2008.
24. George Kozanidis, George Liberopoulos and Christos Pitsilkas, "Flight and Maintenance Planning of Military Aircraft for Maximum Fleet Availability," *Military Operations Research*, **15**, 1, p53-73, 2010.
25. R. A. Hahn and Alexandra M. Newman, "Scheduling United States Coast Guard helicopter deployment and maintenance at Clearwater Air Station, Florida", *Computers & Operations Research*, **35**, p1829-1843, 2008.
26. Leonard Winata, Melih Ozlen & Chun H. Wang, "Heuristic approaches for flight and maintenance planning", in *AIAC14 Fourteenth Australian International Aerospace Congress*, Don Love (ed.), Engineers Australia & Royal Aeronautical Society, Australian Division, Melbourne, Australia pp. 1-17, 2011.
27. Ville Mattila and Kai Virtanen, "A Simulation-based Optimization Model to Schedule Periodic Maintenance of a Fleet of Aircraft," *Proceedings of the 2005 European Simulation and Modeling Conference*, 24-26 October 2005, Porto, Portugal, p479-483.
28. Ville Mattila and Kai Virtanen, "Scheduling Periodic Maintenance of Aircraft Through Simulation-based Optimization," *Proceedings of the 47th Conference on Simulation and Modeling, SIMS2006*, 27-29 September 2006, Helsinki, Finland, p38-43.

Appendix A: Sample implementation of the ashore daily flying program

This appendix provides a sample implementation of how the ashore daily flying rate is determined. It expands on the material covered in Section 4.4.3.

Table A-1 gives the values of the parameters used in this example.

Table A-1: Parameter values used in sample implementation of ashore daily flying program

Parameter	Units	Number
Number of ashore serviceable tails	number	variable
Number of phased maintenance lines	number	3
Airframe hours between phased services	airframe hours	200
Phased maintenance duration	work days	15
Required daily flying rate for the ashore fleet	hours/day	20
Maximum daily flying rate per tail	hours/day	6
Minimum daily flying rate per tail	hours/day	1
Maximum days idle per tail	work days	5

Table A-2 gives an example of how the ashore flying program works over a generic ten-day period. Against each tail number are the hours remaining until phased maintenance is due, while against each maintenance line number are the days remaining until the current service is completed and the line becomes free.

Table A-2: Sample results using methodology for determining daily flying program, showing the airframe hours remaining until the next phased service at the start of the given day

	Start of day ...									
	1	2	3	4	5	6	7	8	9	10
Tail no.	Airframe hours remaining									
1	180	180	U	U	180	180	180	180	175	169
2	0	0	U	U	U	0	0	200	200	200
3	154	154	154	148	148	145	145	143	U	143
4	75	71	67	64	58	58	52	46	40	34
5	42	36	30	27	21	15	9	3	0	0
6	0	0	0	0	0	200	200	200	200	198
7	E	E	E	73	69	63	59	U	U	U
8	123	118	113	109	105	99	95	89	83	77
9	3	2	1	1	1	0	0	0	0	0
10	48	44	40	38	E	E	E	E	E	E
Line no.	Days remaining									
1	3	2	5	4	3	2	1	0	15	14
2	5	4	3	2	1	15	14	13	12	11
3	12	11	10	9	8	7	6	5	4	3
FLYING HOURS	20	20	20	18	20	20	20	20	20	20

The number of available tails is 10 and there are three phased maintenance lines. Line 3 remains occupied for the duration of this period. The number of ashore serviceable tails varies from day to day depending on the frequency and duration of scheduled and unscheduled maintenance, with a range of five to seven in this example. Unscheduled maintenance can occur during a phased service, as demonstrated with tail 2. Tails can also be rotated for embarkation: at the end of day 3, tail 7 finishes its disembarkation (with hours remaining) and tail 10 replaces it.

On day 1, tail 9 is matched to line 1, tail 5 to line 2 and tail 10 to line 3. Therefore tail 9 flies $3/3 = 1$ hour/day; tail 5 flies at $\max(6, 42/5) = 6$ hours/day; and tail 10 flies at $48/12 = 4$ hours/day. Tail 4 is then also matched with line 1 and will fly at a rate that will allow it to enter maintenance when tail 9 completes maintenance, at $75/(3+15) \sim 4$ hours/day. Next is tail 8 which will be matched to line 2. It can fly at a rate of $\max(6, 123/(5+15)) = 6$ hours/day. However, the first four tails have already achieved 15 hours to this point. Therefore, tail 8 will fly the remaining 5 hours for the day to exactly achieve the 20 hours required.

Consider tails matched with line 3. When tail 10 is embarked on day 4, tail 4 is then matched to line 3 as it has the third-lowest remaining airframe hours. When tail 9 enters phased maintenance on day 6, the order of tails is now 5, 4, 7, 8, 3, 1 and 6, so tail 7 is now matched with line 3. When tail 7 enters unscheduled maintenance on day 7, tail 8 now has the third-lowest hours. When tail 5 enters maintenance on day 9, tail 1 has the third-lowest hours.

Tail 9 is initially only 3 hours from a phased service. The minimum flying rate is 1 hour/day, so on days 1 and 2 it flies at this rate. At the start of day 3 it has 1 hour left and line 2 will be free the next day. However, unscheduled maintenance unexpectedly occurs on tail 2 in line 1 on day 3. According to the methodology any tail must wait until the ratio of hours remaining over time remaining is at least this value. Therefore tail 9 remains idle for 2 days before flying at 1 hour/day and entering phased maintenance. This also causes a re-alignment of the tail/line matching: instead of entering line 1 on day 3, it now enters line 2 on day 6.

Tail 3 is marked as being idle for longer than the maximum duration on day 3. Subsequently it is placed at the head of the queue on day 4 and flies its maximum allocation. The re-matching of the other tails to different maintenance lines means that only 18 hours are achieved on this day rather than the required 20. This shortfall will be noted and the flying rate may be increased on the following day, depending on progress against the annual requirement. This re-matching also slows down tail number 5, which now enters an empty maintenance line 1 on day 8. If there are multiple idle tails, those that have been waiting the longest are flown first.

Appendix B: List of state transitions

The range of all possible state changes represented in the model is shown in the tables below. Since the time step for the model is one day, all state changes are assumed to occur at the end of each day.

In reality, events happen on an hour-by-hour or even minute-by-minute basis. This particularly applies to tails experiencing unscheduled maintenance. When this happens in practice, replacement tails need to be found in order to fulfil the daily flying program. Other tails may return to service during the day after having exited maintenance and be serviceable on the flight-line. Even though the model does not capture the dynamic nature to the level of granularity that is observed in the daily operations, it still captures sufficient information that is required to determine the fleet size over the fleet life. Thus the time step chosen is appropriate, given the balance between the amount of information to be captured and the 30-year time frame of the model.

Embarked flying and maintenance occurs every day. Ashore flying and maintenance occurs only on work days, which therefore excludes weekends, public holidays and standdown (specified holiday) periods. Some event triggers are based on elapsed time in calendar days: when this occurs ashore, the test for this will not be undertaken until the next work day.

Table B-1: Possible state changes with criteria for embarked tails under the model

Current State	Possible Future State	Criterion/criteria for state change
Embarked-serviceable	Embarked-scheduled maintenance	Airframe hours till next regular or phased service achieved
	Embarked-unscheduled maintenance	Unscheduled maintenance event is due as determined randomly from probability distribution
	Changeover embarked-ashore	Completion of assigned embarked hours, or maximum time embarked exceeded
	Attrited	Randomly determined attrition day reached
	Life of type	Flying hours achieved reaches maximum
Embarked-scheduled maintenance	Embarked-serviceable	Duration in work days in regular or phased maintenance completed
	Embarked-unscheduled maintenance	Unscheduled maintenance randomly determined to have occurred during scheduled service
Embarked-unscheduled maintenance	Embarked-serviceable	Randomly-determined duration in work days in unscheduled maintenance duration completed
	Embarked-scheduled maintenance	Work days in unscheduled maintenance completed for tail in phased maintenance
Changeover embarked-ashore	Ashore-serviceable	Calendar days in changeover completed

Table B-2: Possible state changes with criteria for ashore serviceable tails under the model

Current State	Possible Future State	Criterion/criteria for state change
Ashore-serviceable	Changeover ashore-embarked	Embarked tail numbers fall below the minimum. Tail(s) that will not be due for DM during embarkation and has the lowest total embarked hours
	Ashore-scheduled maintenance (regular)	Elapsed time or airframe hours till next regular service achieved and flight-line maintenance line is free
	Ashore-scheduled maintenance (phased)	Airframe hours till next phased service achieved and phased maintenance line is free
	Ashore-unscheduled maintenance	Unscheduled maintenance event due as determined randomly from probability distribution and flight-line maintenance line is free
	Awaiting flight-line maintenance (unsched)	Unscheduled maintenance event due as determined randomly from probability distribution but all flight-line maintenance lines occupied
	Awaiting flight-line maintenance (regular)	Elapsed time or airframe hours till next regular service achieved but all flight-line maintenance lines occupied
	Awaiting phased maintenance	Airframe hours till next phased service achieved but all phased lines occupied
	Awaiting deep maintenance	Elapsed time or airframe hours since previous deep maintenance achieved but all DM lines are occupied
	Deep maintenance	Elapsed time or airframe hours since previous deep maintenance achieved
	Detachment	Detachment starts and tail(s) with most hours until next phased service due
	Attrited	Randomly determined attrition day reached
	Life of type	Flying hours achieved reaches maximum
Detachment	Ashore-serviceable	Calendar day duration of exercise period completed
	Awaiting flight-line maintenance (unsched)	Calendar day duration of exercise period completed and unscheduled maintenance event(s) have randomly accrued during exercise
Changeover ashore-embarked	Embarked-serviceable	Calendar days in changeover completed

Table A-5: Possible state changes with criteria for ashore maintenance tails under the model

Current State	Possible Future State	Criterion/criteria for state change
Ashore-scheduled maintenance	Ashore-serviceable	Work days in scheduled maintenance completed
	Ashore-unscheduled maintenance	Unscheduled maintenance randomly determined to have occurred during scheduled service
	Modifications	Work days in scheduled maintenance completed and elapsed time for modification reached
Ashore-unscheduled maintenance	Ashore-serviceable	Work days in unscheduled maintenance achieved for tail in flight-line maintenance
	Ashore-scheduled maintenance	Work days in unscheduled maintenance completed for tail in phased maintenance
	Modifications (squadron-level)	Work days in unscheduled maintenance completed for tail undergoing modifications
Deep maintenance	Ashore-serviceable	Work days in deep maintenance completed
	Modifications	Work days in deep maintenance completed and modifications are due
Modifications	Ashore-serviceable	Work days in modifications completed
	Ashore-unscheduled maintenance	Unscheduled maintenance randomly determined to have occurred during scheduled service (only for modifications undertaken at squadron level)
Awaiting flight-line maintenance	Ashore-unscheduled maintenance	Flight-line maintenance line becomes free for tail experiencing unscheduled maintenance
	Ashore-scheduled maintenance	Flight-line maintenance line becomes free for tail due for regular service
Awaiting phased maintenance	Ashore-scheduled maintenance	Phased maintenance line becomes free
Awaiting deep maintenance	Deep maintenance	Deep maintenance line becomes free

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. DLM/CAVEAT (OF DOCUMENT)			
2. TITLE Fleet Sizing Analysis Methodologies for the Royal Australian Navy's Combat Helicopter Replacement Project			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)				
4. AUTHOR(S) David Marlow and Ana Novak			5. CORPORATE AUTHOR DSTO Defence Science and Technology Organisation 506 Lorimer St Fishermans Bend Victoria 3207 Australia				
6a. DSTO NUMBER DSTO-TR-2886		6b. AR NUMBER AR-015-695		6c. TYPE OF REPORT Technical Report		7. DOCUMENT DATE August 2013	
8. FILE NUMBER 2012/1132922/1	9. TASK NUMBER DMO 07/223	10. TASK SPONSOR DGNAS		11. NO. OF PAGES 83		12. NO. OF REFERENCES 28	
13. DSTO Publications Repository http://dspace.dsto.defence.gov.au/dspace/			14. RELEASE AUTHORITY Chief, Joint and Operations Analysis Division				
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>							
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111							
16. DELIBERATE ANNOUNCEMENT No Limitations							
17. CITATION IN OTHER DOCUMENTS				Yes			
18. DSTO RESEARCH LIBRARY THESAURUS fleet-sizing, naval combat helicopter analysis, discrete-event simulation, heuristics, sensitivity analysis							
19. ABSTRACT This report describes methodologies used to conduct fleet sizing analysis for the Royal Australian Navy's combat helicopter replacement project, known as AIR 9000 Phase 8. An initial analysis utilises a simple approach to predict fleet size based on the binomial distribution. The sole aim is to predict the likelihood that a certain fleet size will provide a required number of embarked aircraft. A more complex analysis requires the development of a discrete-event simulation that can also incorporate the hours flown by the fleet. This simulation model includes the representation of individual aircraft moving between embarked and ashore states and through various types of maintenance, including unscheduled maintenance. This methodology is robust and easily allows for sensitivity and trade-off analysis.							