

ON THE AVAILABILITY OF THE CH149 CORMORANT FLEET

Raman Pall

Directorate Materiel Group Operational Research
Defence Research and Development Canada – Center for Operational Research and Analysis
National Defence Headquarters
Ottawa, ON K1A 0K2, CANADA

ABSTRACT

The CH149 Cormorant is the Canadian Forces (CF) designation for the AgustaWestland EH101, the Canadian Air Force's only dedicated search and rescue (SAR) helicopter. Since its procurement, the availability for operations of the CH149 fleet has been less than what was initially predicted. This study was undertaken to determine if the low serviceability of the fleet was due to its maintenance programme. A discrete-event simulation model was created to determine the number of aircraft available at any given time during a simulation run, assuming an ideal sparing situation. Analysis indicates that the current fleet cannot operate from four (or even three) Main Operating Bases (MOBs) in an ideal sparing situation, implying that the fleet's availability problem cannot theoretically be solved simply by addressing the logistical problem of the spares.

1 INTRODUCTION

The CH149 Cormorant is the Canadian Forces (CF) designation for the AgustaWestland EH101. It is the Canadian Air Force's only dedicated search and rescue (SAR) helicopter, and is designed to operate in the most severe conditions, making it ideal for Canada's challenging geography and climate.

Since its procurement, the availability for operations of the CH149 fleet has been less than what was initially predicted. Before the CF acquired the CH149, a study was done to determine the number of aircraft required to operate out of four Main Operating Bases (MOBs). This study assumed that all aircraft would have an availability of 75%, resulting in a conclusion that 15 aircraft would satisfy the minimum requirement as listed in the Statement of Operational Requirements (SOR) for the CH149 fleet (DND 1995).

Now that the CH149 has operated for 7 years and can, in reality, only operate from three MOBs, it must be acknowledged that the original study assumed that the



Figure 1: A Canadian CH149 Cormorant helicopter.

CH149 fleet would have a higher availability than was actually the case. An obvious question arose: Why is there such a low serviceability in the CH149 fleet?

It was hypothesized that the low serviceability of the CH149 fleet was due to its maintenance programme. This study was undertaken to determine if the current fleet of CH149 aircraft could possibly operate from four MOBs under an ideal sparing situation, and to determine if the low serviceability of the fleet is due to its maintenance programme.

Studying the performance measure of the fleet relating to SAR standby aircraft availability (see Section 2.4) in an ideal sparing situation provides an answer as to whether the fleet's poor availability is a result of its maintenance programme (which includes its scheduled inspections and maintenance, corrective maintenance, etc.), and whether or not the fleet's availability problem could possibly be solved by solving the logistical problem of its spares. If the measure was within its target range, then the availability problem could theoretically be solved by solving the logistical problem. If, however, this measure's target range was not met, the availability problem cannot be solved without reducing the

maintenance demands of the fleet, or possibly augmenting the size of the fleet.

This paper is structured as follows: Section 2 presents the methodology and the data, as well as the assumptions inherent in the data. The discrete-event simulation model is also described in this section, and the performance measure used to evaluate the results is given. The results of the simulation are presented in Section 3. The report concludes in Section 4 with a summary of the results and recommendations that were given to the CF concerning the size of the CH149 fleet and possible actions that can be taken to improve its availability.

2 METHODOLOGY

To tackle the problem of the fleet's availability in an ideal sparing situation, all maintenance activities that are performed on the aircraft were determined, along with their durations and frequencies. A discrete-event simulation model was created to determine the number of aircraft available at any given time during a simulation run, assuming an ideal sparing situation, and the performance measure pertaining to SAR standby aircraft availability in the model was compared to its desired levels.

2.1 The Data

In July 2000, the Aerospace Division of IMP Group Limited was awarded a major contract by the Government of Canada for the In-Service Support (ISS) of the CH149 Cormorant Search and Rescue Helicopter. This contract has since been renewed for an additional period of seven years (DND 2007). IMP Aerospace, per the terms of the contract, is responsible for the provision and management of total maintenance, logistics and life-cycle engineering support for the CH149 fleet. All data discussed here has been provided by IMP Aerospace.

It was assumed in this data that the Yearly Flying Rate (YFR) for the fleet is 7000 Running Hours (RH). Note that this document makes reference to the number of *Running Hours (RH)* an aircraft operates for, as opposed to the more standard term *Flying Hours (FH)* as all documents pertaining to the in-service support contract from which data was derived specified the inspections of the aircraft in terms of their RH, and not in terms of the FH (IMP Group 2007). One can think of an aircraft's RH as the number of hours in which the aircraft's engine is operating, whereas its FH is the number of hours in which the aircraft is actually in flight.

It was also assumed that all major logistic impediments, save light *rob* requirements for heavy maintenance of the aircraft, have been properly addressed (note that the term *rob* in this context should be taken to mean the cannibalization of other aircraft for spares). This was done by IMP to provide

“the most optimistic projection possible” concerning the availability of aircraft spares (IMP Group 2007). The list of all activities to which aircraft can be assigned that render them not serviceable is provided in what follows.

- **Major and minor inspections.** Each aircraft undergoes two types of scheduled maintenance inspections each calendar year, known as major and minor inspections, which take place every 500 RH and every 250 RH, respectively. Moreover, at least one major and one minor inspection must take place each year for each aircraft, regardless of its running hours. These inspections are staggered throughout the year.
- **Out-of-sequence inspections (OSI).** Out-of-sequence inspections are inspections that the aircraft undergo that are not part of the regular scheduled maintenance for the aircraft handled in their major and minor inspections.
- **Depot-level inspections and repairs (DLIR).** The depot-level inspections and repairs that the CH149 aircraft undergo are a higher level of maintenance than the other scheduled maintenance activities, and are used to carry out activities such as corrosion surveys and repairs, aircraft paint activities, fuel bay repairs, and other complex repairs. Data provided by IMP showed that every six months, aircraft are rotated through a DLIR line for a period of approximately six months, resulting in all aircraft being rotated through the DLIR line in a period of approximately five years (if it is assumed that on average 1.4 aircraft are in DLIR at any given time).
- **Corrective maintenance.** Corrective maintenance results in aircraft not able to conduct missions due to ramp snags. Due to the high reliability of the CH149 Cormorant, the frequency and duration of corrective maintenance is quite small.

Note that the durations and frequencies of the different maintenance activities differ from MOB to MOB. These durations and frequencies may also vary depending on the age of the aircraft.

2.2 The Model

The model that was created to determine the number of aircraft available at each base at any given time was implemented as a discrete-event simulation model in the *Arena* simulation language (Kelton, Sadowski, and Sturrock 2004). An *Arena* model is a computer program containing components called modules that represent processes or logic. Connector lines are used to join these modules together and specify the flow of entities. While modules have specific

actions relative to entities, flow, and timing, the precise representation of each module and of each entity relative to real-life objects is subject to the modeller. Statistical data can be recorded and outputted as reports.

In this case, entities within the model represent the individual CH149 aircraft at each of the bases in operation. The entities move between various stations governing the activities to which they are assigned.

At the start of the simulation the aircraft are created and placed at the MOBs. During the course of the simulation, the aircraft move from station to station and follow a logic tree to determine what they are doing, which is reflected in an animated portion of the model.

In the beginning of every simulated year, the major and minor inspections are randomly staggered throughout the simulated year at each base and assigned to the aircraft at each base. The durations of the major and minor inspections are determined using triangular distributions.

During each simulated day the model determines to which activity the aircraft at each base is assigned. Some activities are scheduled (such as the scheduled major and minor inspections, whose start times are modelled as being determined at the start of the year, as described above); others types of activities are determined stochastically (such as corrective maintenance).

It was assumed that the scheduling of unscheduled activities for the aircraft is independent, in the following sense: the chance of an aircraft being assigned to an unscheduled activity is not affected by the number of other aircraft currently assigned to that type of activity. For example, if two aircraft are each supposed to undergo corrective maintenance 10% of the time, one can expect that both will be undergoing corrective maintenance approximately 1% of the time during the simulation.

The logic used to determine what the aircraft are doing in any given day, noted here as its *status*, is now explained. For each day, each aircraft is examined in sequence. If the aircraft is undergoing a major or minor inspection, its status is updated to reflect this fact.

Every six months, the model randomly chooses a number of aircraft to place on DLIR status, given by a Poisson distribution (Pfeiffer and Schum 1973) whose mean is the expected number of aircraft that are in DLIR at any given time. The duration of the DLIR activity is chosen from a triangular distribution.

If the aircraft is not in DLIR, then it may be placed in an out-of-sequence inspection. Aircraft reaching this point in the model have a chance of undergoing an out-of-sequence inspection given by the quotient of the number of days in which they are expected to undergo OSIs and the number of days they can possibly undergo this type of activity. Each such OSI activity is assumed to last one day.

If the aircraft are not placed in an out-of-sequence inspection, they are similarly tested to see if they are to

be placed in corrective maintenance. As for the chance of being placed in an OSI, the chance of an aircraft being placed in corrective maintenance is given by the quotient of the number of days in which they are expected to undergo corrective maintenance and the number of days they can possibly undergo this type of activity. Each such corrective maintenance activity is also assumed to last one day.

If the aircraft are not placed in any of these activities (major or minor inspections, DLIR, OSI, corrective maintenance), then they have their status assigned to being available. At the end of each day, the current status of each aircraft is written to an output file, which is then used to aggregate information about the number of aircraft available at any given moment.

During the model run, the statuses of the aircraft in the simulation are represented in an animated section of the model. For example, in Figure 2 is shown the statuses of the aircraft at one point of a run in which 19 aircraft are distributed amongst the four MOBs. (Note that the animation section of the model shows the statuses of up to 28 aircraft. As only 19 aircraft were used in the run shown, 9 of the aircraft are grayed out to show that they were not included in the model run.) Aircraft on scheduled major or minor inspections are shown in light blue, those undergoing OSIs are shown in dark blue, those in DLIR are shown in orange, those undergoing corrective maintenance are shown in pink, and the available aircraft are shown in green. At another point in the simulation, the statuses of these aircraft have changed to those shown in Figure 3.




Figure 2: The statuses of the aircraft shown in the animated portion of the model at one point during a run of the simulation.

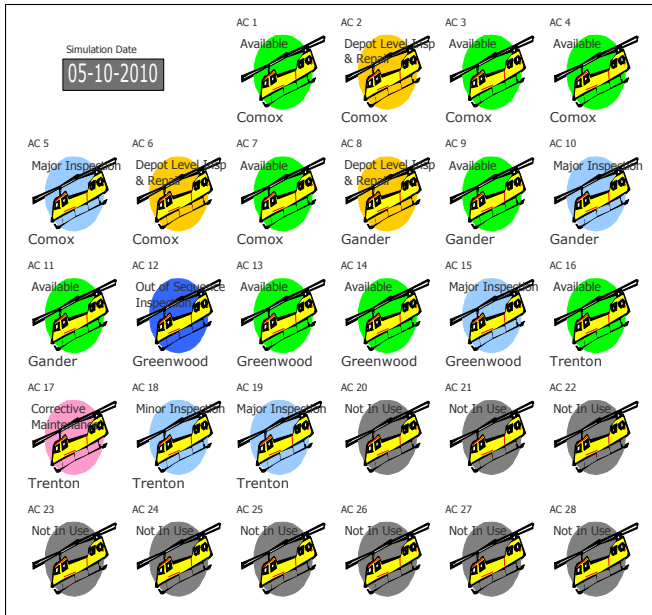


Figure 3: The statuses of the aircraft shown in the animated portion of the model at another point during a run of the simulation.

In this model, numerous stochastic processes were involved, including the scheduling of the activities to which the aircraft are assigned (such as the minor and major inspections), the aircraft selected to undergo DLIR or corrective maintenance, and the durations of all activities.

These distributions encompass a large number of distinct possible outcomes in the availability of the aircraft, even for a single set of circumstances (meaning the number of aircraft at each base, the number of bases in operation, etc.). Hence, the simulation was run repeatedly as a Monte Carlo simulation to consider all of the possible outcomes of the model at once for each set of model inputs, by using random sampling on the stochastic elements of the model. The simulation was repeated until a large enough number of possible future scenarios were explored, and a representative sample of the overall outcome emerged.

2.3 Assumptions

The full list of this study’s major assumptions inherent in the data and the model are described in detail below.

- Running Hours and Scheduled Inspections.** It was assumed that the fleet would fly for 7000 RH each year. Given a fleet of at least 14 aircraft (which is the current number of CH149 aircraft in service), on average an aircraft would operate for 500 RH each year. As major inspections are to occur every 500 RH or at least once a year, and

minor inspections are to occur every 250 RH, or at least once a year, it was assumed that each aircraft undergoes 1 major and 1 minor inspection each calendar year.

- Higher-Order Inspections.** In addition to the major and minor inspections, there are higher-order inspections that occur after the aircraft has flown for a longer period of time (for example, after 1000, 1500, and 2000 RH). In discussions with military staff, it was stated that these higher-order inspections are minor extensions of the major inspections, and activities that the aircraft undergo in these inspections can be encompassed in the major inspections. Hence, in the modelling of the availability of the fleet, it was assumed that there are only two types of scheduled inspections that the aircraft undergo – the major and minor inspections.
- Staggering of the Scheduled Inspections.** It was assumed that the inspections of the aircraft would be staggered whenever possible, to ensure that as little overlap in the scheduled inspections would take place at each MOB as possible.
- Aircraft Attrition and Distribution.** It was assumed that no attrition in the fleet would take place during the period simulated. Moreover, it was assumed that the distribution of the aircraft amongst the MOBs would stay constant throughout the simulation – that is to say that no aircraft would ever be transferred between the bases.
- Simulation Resolution.** The model built has a simulation resolution of 1 day – meaning that it was assumed that all activities take some integer quantity of days to perform (e.g., no activity takes 17.5 days, it would last either 17 or 18 days).

2.4 Performance Measure

The measure used to evaluate the results of the simulation, known as the *SAR Standby Aircraft Availability Rate (SARstbyA)*, is defined as:

“The extent to which mission capable SAR Standby aircraft are provided to support SAR mission operations on a 24 hour, 7 day a week basis.” (DND 2005)

In other words, this measure is the fraction of time that at least one aircraft is serviceable and available for SAR operations at each base. The target range for this measure is 99.0% to 99.5% at each MOB (DND 2005).

2.5 Model Inputs and Outputs

The model has several runtime inputs and options in order to set the model to behave as desired by the analyst. These inputs can be used to model situations known as “what-if” analyses – situations studied to determine the effects of changes in the operating environment of the CH149 fleet (such as the number of MOBs in operation, or the distribution of the aircraft amongst the bases) on their availability. A summary of the model inputs are listed below.

- **Number of MOBs in operation.** The analyst has the ability to specify the number of bases in operation (generally three or four bases are chosen).
- **Aircraft Distribution and Quantities.** At the start of the simulation, the analyst must specify how many aircraft are to be placed at each base in operation. This distribution is sometimes referred to in the form (w,x,y,z) , where this expression is taken to mean that the distribution of aircraft consists of w aircraft at Comox, x aircraft at Gander, y aircraft at Greenwood, and z aircraft at Trenton.
- **Reductions in the Durations and Frequencies of the Activities.** It may be desired by the analyst to see what effect a reduction in the time taken to perform any activities or in their frequencies would have on the availability of the fleet. In order to facilitate this type of analysis, there are model options specifying the factors by which the durations and frequencies of the different types of activities are to be reduced.
- **Number of Aircraft to be Placed in DLIR at a Time.** The number of aircraft that are to go in DLIR at a time (every 6 months) is determined through the use of a Poisson distribution whose mean is the expected number of aircraft that are to be in DLIR at any given time. This number can be specified by the analyst prior to running the model.
- **All Data on Activity Durations and Frequencies.** If desired, all data in the model concerning the durations and frequencies of the activities can be modified. For example, if it is desired to assume that the number of OSI days per aircraft at one base is substantially different than the corresponding number at the other bases, it may be modified to any value that the analyst desires.

After the simulation has run to completion, several outputs are collected. This output includes the inputs used in the model run, such as the number of aircraft at each base, the number of bases in operation, and whether or not the durations of the activities were reduced.

However, the main output of the model concerns the activities to which the aircraft were assigned during the simulation. After each simulated day, every aircraft in the simulation records the type of activity to which it was assigned in a file that is later used in several database queries to determine the number of aircraft available at each MOB in each day. The proportion of the simulated time that at least a given number of aircraft are available at each base can then be determined.

3 RESULTS

3.1 Basic Results

Results obtained using inputs that describe the current situation of the fleet – i.e., where there are fourteen aircraft in operation – are presented here. These results, which will sometimes be referred to as the *basic results*, are comprised of two specific cases. The following model inputs were used for these two cases, respectively:

1. Four MOBs in operation with an aircraft distribution of (5,3,3,3); and
2. Three MOBs in operation with an aircraft distribution of (6,4,4,0).

Note that the model used the same inputs for all MOBs other than Comox. In both cases, the model was run for a total of 50 replications. The results of the simulation concerning the availability of the first aircraft at each MOB for these two cases are presented in Tables 1 and 2, respectively.

Table 1: Simulation results on aircraft availability with standard inputs and an aircraft distribution of (5,3,3,3), ordered as listed below.

MOB	Median	5 th Perc.	95 th Perc.
Comox	96.9%	93.7%	98.4%
Gander	90.1%	84.3%	93.2%
Greenwood	89.6%	84.2%	93.4%
Trenton	90.5%	85.9%	93.4%

Table 2: Simulation results on aircraft availability with standard inputs and an aircraft distribution of (6,4,4,0), ordered as listed below.

MOB	Median	5 th Perc.	95 th Perc.
Comox	98.7%	96.9%	99.4%
Gander	95.6%	91.2%	97.0%
Greenwood	94.9%	91.8%	97.0%
Trenton	–	–	–

The results show that in both cases, the availability of the first aircraft at all MOBs is lower than the desired level of

99.0%, indicating that under the assumptions of this study, the aircraft cannot be expected to meet the performance measure's target range if they are expected to operate from three or four MOB, even under the assumption of an ideal sparing situation.

3.2 Effects of Changing the Size of the Fleet

As the availability of the first aircraft did not meet its desired target range when fourteen aircraft were included in the simulation (which is the current number in the fleet), the obvious question that arose concerned the number of aircraft necessary to meet the measure's target range. The simulation was run for several other cases, in which the number of aircraft at MOB Comox ranged from 5 to 8 and the number of aircraft at the other MOB's ranged from 3 to 7. Note that the model used the same inputs for all MOB's other than Comox. Simulations were run for a total of 50 replications in all cases.

A summary of the results is presented in graphical form in Figures 4 and 5, where the availability of the first aircraft is shown as a function of the number of aircraft at the MOB.




Figure 4: The results obtained concerning the availability of the first aircraft at MOB Comox.

The representation of the results shown in the figures is a variation of the well-known statistical plots known as box-and-whisker plots (Hogg and Tanis 1996). For each ordinate in each chart, five pieces of data are shown: the median value of all replications (shown as the dots in each chart), the 25th percentile (the top point of the lower bars), the 5th percentile (the bottom point of the lower bars), the 75th percentile (the bottom point of the upper bars), and the 95th percentile (the top point of the upper bars). In other words, one can think of the bars as illustrating the spread found in the Monte Carlo simulation runs of the model.

It was found that the measure is within the target range only when at least 7 aircraft are located at MOB Comox




Figure 5: The results obtained concerning the availability of the first aircraft at all MOB's other than Comox.

and at least 6 aircraft are located at the other three MOB's, resulting in a total requirement of 25 or 19 aircraft if four or three MOB's, respectively, are desired. Again, these results concern only the minimum number of aircraft required in an ideal sparing situation, under the conditions and assumptions of this study.

3.3 Effects of Reduced Inspection Times

As the availability of the first aircraft at all MOB's did not reach the target range under the assumptions of this study with the fleet comprised of 14 aircraft, meeting the performance measure's target range would require at least one change to the model's options. In the previous section, the size of the fleet was changed, and the effects of these changes on the fleet's availability were studied. In this section, the effects of reducing the durations of specific activities on the fleet's availability are studied.

The durations of the major, minor, and out-of-sequence inspections were reduced by a fixed percentage to present a study of the effects of reducing the durations of the inspections on the fleet's availability. The durations of other activities, such as the depot-level inspections and repairs as well as the corrective maintenance, were not altered. Discussions with military staff led the author to conclude that the depot-level inspections and repairs would take the same amount of time regardless of any possible changes in the time taken for the other inspections. The corrective maintenance remained unchanged as reductions in inspection times were not expected to reduce the chance or durations of ramp snags or any other type of corrective maintenance.

As in Section 3.1, two cases will be described. The following model inputs were used for these two cases, respectively:

1. Four MOBs in operation with an aircraft distribution of (5,3,3,3); and
2. Three MOBs in operation with and an aircraft distribution of (6,4,4,0).

In both cases, the model was run for a total of 50 replications, and the durations of the major, minor, and out-of-sequence inspections were reduced by 25%. The results of the simulation for these two cases are presented in Tables 3 and 4, respectively.

Table 3: Simulation results on aircraft availability with selected inspection durations reduced by 25% and an aircraft distribution of (5,3,3,3), ordered as listed below.

MOB	Median	5 th Perc.	95 th Perc.
Comox	99.0%	97.5%	99.7%
Gander	94.5%	90.5%	97.2%
Greenwood	94.6%	91.4%	96.7%
Trenton	94.1%	90.6%	97.5%

Table 4: Simulation results on aircraft availability with selected inspection durations reduced by 25% and an aircraft distribution of (6,4,4,0), ordered as listed below.

MOB	Median	5 th Perc.	95 th Perc.
Comox	99.6%	98.9%	99.9%
Gander	97.7%	94.5%	98.9%
Greenwood	97.8%	94.5%	99.0%
Trenton	–	–	–

In both the first and second cases, it was found that the only MOB where availability of the first aircraft met the measure’s target range was Comox. Hence, these results indicate that under the assumptions of this study, with a 25% reduction in the durations of selected inspections, and operating under the assumption of an ideal sparing situation, the fleet of 14 aircraft cannot be expected to meet the target range of the performance measure if they are expected to operate from three or four MOBs.

As was done for the cases using the basic inputs, the simulation was run when the size of the fleet was increased (while still using the reduced inspection durations described above) to study the effects of changing the size of the fleet under these conditions. In these additional runs of the model, the number of aircraft at MOB Comox ranged from 5 to 8 and the number of aircraft at the other MOBs ranged from 3 to 7. A summary of the results is presented in graphical form in Figures 6 and 7.

It was found that the availability of the first aircraft is within its target range only when at least 5 aircraft are located at each of the bases, resulting in a total requirement




Figure 6: The results obtained concerning the availability of the first aircraft at MOB Comox, when the durations of the inspections were reduced by 25%.

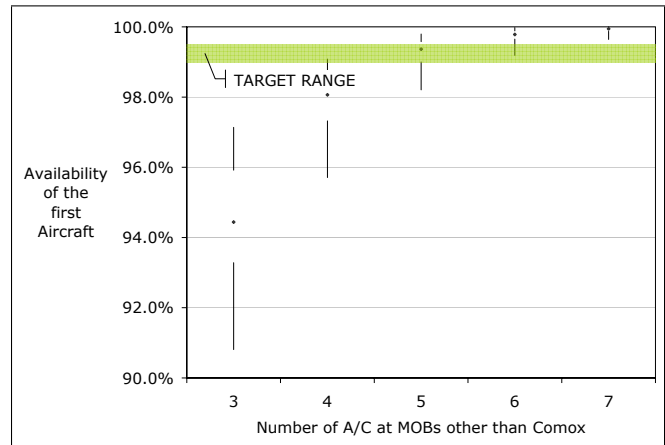


Figure 7: The results obtained concerning the availability of the first aircraft at all MOBs other than Comox, when the durations of the inspections were reduced by 25%.

of 20 or 15 aircraft if four or three MOBs are desired, respectively (contrast these values to the requirements of 25 or 19 aircraft when the durations of the inspections were not reduced).

4 CONCLUSIONS

Studying the performance measure of the fleet relating to Search and Rescue (SAR) standby aircraft availability in an ideal sparing situation provides an answer as to whether the fleet’s poor availability is a result of its maintenance programme (which includes its scheduled inspections and maintenance, corrective maintenance, and so on), and whether

or not the problem of the fleet's availability could possibly be solved by solving the logistical problem.

It was found that in order to meet the target range, at least 7 aircraft are required at MOB Comox, and at least 6 aircraft are required at all other MOBs assuming an ideal access to spares. Hence at least 25 or 19 aircraft are required if four or three bases are desired, respectively.

Additionally, it was shown that a 25% reduction in the durations of the major, minor, and out-of-sequence inspections of the aircraft had significant impacts on the number of aircraft required. It was found that with these reductions in place, only 5 aircraft would be required at each base in an ideal sparing situation, for a total of 20 aircraft if 4 MOBs are in operation, or a total of 15 aircraft if 3 MOBs are in operation. It should not be expected that the current fleet of 14 aircraft can operate indefinitely from 4, or even 3, MOBs without significant changes to the fleet's maintenance program.

These results imply that the fleet's availability problem (of not meeting the target range for the performance measure) cannot theoretically be solved simply by addressing the logistical problem of the spares – the maintenance demands of the fleet need to be reduced or the size of the fleet would have to be increased in order to meet the performance measure's target range.

A number of recommendations regarding the availability of the CH149 fleet were made to military staff upon completion of this study. These recommendations have had a significant impact on the maintenance programmes of the fleet, as well as on future acquisitions of rotary wing search and rescue aircraft.

REFERENCES

- Department of National Defence. 1995, February. A2678 Canadian search and rescue helicopter statement of operational requirement.
- Department of National Defence. 2005. Performance management framework and plan for in-service support of the Canadian search and rescue helicopter CH149 Cormorant.
- Department of National Defence. 2007. DND NR - 07.082: Renewed contract to sustain flight-readiness of Canada's search and rescue helicopters, September 12, 2007. Available via http://www.forces.gc.ca/site/newsroom/view_news_e.asp?id=2451 [accessed November 15, 2007].
- Hogg, R. V., and E. A. Tanis. 1996. *Probability and statistical inference*. Fifth ed. Prentice Hall.
- IMP Group 2007. *Fleet un-availability analysis*. Halifax: IMP Group Aerospace.
- Kelton, W. D., R. P. Sadowski, and D. T. Sturrock. 2004. *Simulation with Arena*. Third ed. McGraw-Hill Series in Industrial Engineering and Management Science. McGraw-Hill.
- Pall, R. 2007. On Forecasting the Availability of the CH149 Cormorant Fleet – Determining the Number of Aircraft Necessary to Meet the Statement of Operational Requirements. Technical Report DRDC CORA TR 2007-20. DRDC – Centre for Operational Research and Analysis.
- Pall, R. 2008. On the Availability of the CH149 Cormorant Fleet in an Ideal Sparing Situation – Studying the SAR Standby Aircraft Availability Performance Measure at Different Fleet Sizes. Technical Memorandum DRDC CORA TM 2008-11. DRDC – Centre for Operational Research and Analysis.
- Pfeiffer, P. E., and D. A. Schum. 1973. *Introduction to applied probability*. New York: Academic Press.

AUTHOR BIOGRAPHY

RAMAN PALL received a B.Sc. and an M.Sc. in Mathematics from the University of Ottawa. He is currently a defence scientist with DRDC CORA. His research interests include military operations research, operational logistics, inventory management, transportation modeling, and simulation modeling. He has been published in a wide range of peer-reviewed journals and has authored several internal technical papers. Mr. Pall is also involved in the North Atlantic Treaty Organization. His e-mail address is <Raman.Pall@drdc-rddc.gc.ca>.