

Using discrete-event simulation to predict the size of a fleet of naval combat helicopters

D. O. Marlow^a and A. Novak^a

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

Email: david.marlow@dsto.defence.gov.au

Abstract: This paper describes the methodology developed to determine the size of a new fleet of combat helicopters for the Royal Australian Navy. The Australian Government's 2009 Defence White Paper described "as a matter of urgency" a requirement for a fleet of helicopters "to provide eight or more aircraft concurrently embarked on ships at sea". The objective therefore is to find the minimum fleet size that enables the fleet to meet both this minimum daily embarked requirement, as well as annual requirements for a specified number of flying hours for both embarked helicopters and the remaining ashore-based fleet.

The fleet sizing problem incorporates the helicopter fleet, ships, personnel and a home base with maintenance facilities. Individual helicopters (or *tails*) pass through various states throughout their life. These include the serviceable state when they are able to fly, or a range of maintenance states which will make them unable to fly. Maintenance types include regular inspections, phased maintenance and deep maintenance. Many of these states can occur when a tail is either embarked or ashore. Unscheduled maintenance adds a random element to the problem and influences which aircraft may be serviceable on a day-to-day basis. A discrete-event simulation approach was chosen to address this problem, due to the requirements to track tails in various states, test for state transitions and incorporate random effects.

The fleet sizing model begins using the nominated fleet size as an input, along with the required number of embarked aircraft and the annual embarked and ashore flying hours. Other inputs include maintenance data regarding the time between services and their duration. The fleet is initialized with the required number of tails embarked and the remaining ashore tails in various serviceability and maintenance states.

The model runs for the expected fleet life of 30 years, stepping through one day at a time. At the beginning of each day, it is determined whether it is a work day for ashore flying crews or maintenance crews. Embarked crews work every day of the year. Embarked serviceable aircraft will fly a certain number of hours depending on their operational tempo. A procedure is used to balance the flying hours of the embarked fleet while accounting for their expected and realised maintenance, in order to meet the annual requirement.

If it is an ashore work day, the next step is to allocate flying hours to aircraft, and maintenance crews to maintenance lines. As a baseline, a flying allocation method is used that seeks to minimise queuing of aircraft for phased maintenance. This is compared with a greedy technique that seeks to meet the daily requirement first, but at the expense of lower serviceability. Similar methods are used to allocate maintenance resources. The baseline in this case is allocating the same number of personnel to each maintenance line for both flight line maintenance (which handles regular inspections and unscheduled maintenance) and phased maintenance. An alternative is a greedy technique applied to flight line maintenance, where personnel are pooled and the aircraft with the least number of maintenance manhours is serviced first. This will increase the serviceability of the fleet, but may leave aircraft requiring a large amount of maintenance unserviceable for a long time.

Once the allocation of resources is completed, aircraft will fly or be maintained accordingly. At the end of the day, tests for state transitions occur. Aircraft will move in and out of scheduled or unscheduled maintenance, or in and out of a serviceable state, subject to criteria being met. Disembarking aircraft need to be replaced immediately. The model can also include tails going on exercises or being lost from the fleet.

To demonstrate the capability of the model in sensitivity and trade-off analysis, results are shown for a generic helicopter against an indicative requirement. The greedy maintenance allocation heuristic is shown to be superior, even superseding the differences between the flying allocation methods. When a modification program is included, a large decrease in ashore hours occurs due to queuing for deep maintenance during the mid-life upgrade, which can be overcome by increasing maintenance capacity. These results clearly show the influence that resource allocation and maintenance capacity can have when predicting fleet size.

Keywords: *Discrete-event simulation, military application, fleet size, resource allocation, heuristics*

1. INTRODUCTION

Determining the fleet size is an essential component of any acquisition project. This paper describes the methodology developed to determine the size of a new fleet of combat helicopters for the Royal Australian Navy. The Australian Government’s 2009 Defence White Paper (Commonwealth of Australia, 2009) described a requirement for a fleet of helicopters “to provide eight or more aircraft concurrently embarked on ships at sea”. Due to the strategic situation at the time, the decreasing relative capability of the current S-70B-2 Seahawk fleet and the decision to cancel the purchase of SH-2G Super Seasprites, the White Paper described this new acquisition as “a matter of urgency”.

This requirement for eight helicopters embarked on ships is applied to every day of the fleet life, which is projected to be 30 years. Along with this requirement is the need for those helicopters to achieve a specified number of flying hours on an annual basis. This applies to both the embarked helicopters and those ashore located at a home base. Ashore flying is required to train aircrews for embarked missions, as well as participation in exercises and potential civil emergencies during natural disasters.

The problem therefore is to find the minimum fleet size given the following constraints:

- The minimum number of embarked aircraft is achieved every day;
- The minimum amount of embarked flying hours is achieved every year; and
- The minimum amount of ashore flying hours is achieved every year.

2. PROBLEM DESCRIPTION

2.1. Problem illustration

The fleet sizing problem incorporates the helicopter fleet, ships, personnel and a squadron home base with maintenance facilities. Individual helicopters (or tails) pass through various states throughout their life. Many states are possible, based on the fleet’s operating and maintenance regimes.




Figure 1. Illustration of helicopters changing state over time

Figure 1 indicates how tails in a sample fleet change states over time. The states shown here are:

- *Serviceable embarked* or *serviceable ashore*. Serviceable tails are defined as those that are able to fly.
- *Scheduled maintenance*: Scheduled maintenance can include regular inspections which are frequent and may last hours or days, and phased maintenance which is less frequent but may last days or weeks.
- *Unscheduled maintenance*. Unscheduled maintenance is stochastic, and the frequency (i.e., time between services) and duration follow probability distributions. It can occur on serviceable aircraft (preventing them from flying), or on aircraft in scheduled maintenance (prolonging the maintenance duration).
- *Deep maintenance (DM)* or *Modifications/Upgrades*: This is another type of scheduled maintenance. The frequency of these services is usually measured in years and the duration may be weeks or months. Modifications or upgrades may be appended to deep maintenance, thus increasing the duration.

Deep maintenance and modifications are undertaken by contractors, so aircraft in these maintenance types are unavailable to the squadron. Maintenance on available aircraft is performed by military personnel at the squadron, so aircraft in scheduled or unscheduled maintenance are defined as available but unserviceable.

Each type of maintenance has a capacity associated with it, represented by a number of maintenance lines. Only one tail can occupy one line of maintenance at any time. In this case there are three maintenance types. Flight line maintenance lines handle special inspections, and unscheduled maintenance occurring on serviceable tails. Phased maintenance lines and deep maintenance lines handle their respective maintenance types. Therefore squadron-level maintenance includes personnel on flight line and phased maintenance lines.

It is evident from Figure 1 that the number of tails in each state at any time (such as the daily number embarked) can be tracked and that flying hours for each year (from serviceable aircraft) can be determined.

2.2. Literature review and choice of method

The fleet sizing literature is comprehensive (Hoff et al., 2010). For problems involving uncertainty, as in this work, integer programs in complex liner shipping systems (Dong and Song, 2009) and deterministic models in freight transportation (Turnquist and Jordan, 1986) have been developed. In military applications, a simulation model was used to determine the fleet size applied to air mobility tasks (Wesolkowski and Billyard, 2008), and a combination of military judgement and integer programming to determine the size of the US destroyer fleet (Crary et al., 2002). Discrete-event simulation methods have been used in fleet sizing for train locomotives (Godwin et al., 2008) and anchor handling operations (Shyshou et al., 2010).

In this problem instance, to properly represent the helicopter fleet sizing problem, the chosen 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 criteria posed by the question are met; and
- Handle both deterministic and random state changes.

Discrete-event simulation can capture these requirements and thus was chosen to address the problem.

3. MODEL DESCRIPTION

3.1. Inputs

The fleet size, daily minimum number of embarked aircraft, and the required annual and embarked flying hours are all inputs to the model. Other inputs relate to the timing and details of exercise periods, standdown periods (where aircrew and maintenance personnel are on leave) and the development of the daily embarked and ashore flying programs described in Section 3.3.

Maintenance inputs include maintenance frequency and duration, maintenance capacity and manpower rates of effort. In this work both the frequency and duration of unscheduled maintenance are represented as log-normal distributions, based on statistical analysis of the current S-70B-2 Seahawk fleet. Unscheduled maintenance durations include an allowance for logistics delays in obtaining the relevant spare parts.

3.2. Initialisation

All tails are initialised into various states. The initial number of embarked aircraft is the minimum required. Aircraft are allocated to deep and phased maintenance depending on the frequency and duration of these services. All other aircraft begin as ashore serviceable. The model 'starts' on day one of the full fleet and assumes a progressive introduction of aircraft to that stage. Aircraft are randomly initialised with a number of flying hours, with those most recently 'introduced' having zero hours. Aircraft in maintenance are randomly allocated a duration and an accrued time in maintenance. The model includes an initialisation period where no results are recorded, similar to Mattila et al. (2008), allowing it to come to a 'steady-state'.

3.3. Main procedure

The model has a one-day time step and runs until the specified end of the fleet life. Firstly, each day is tested to see whether it is a work day for embarked crews (who work every day) and ashore flying and maintenance crews (who do not work weekends and have different holiday periods). Then, if required, flying rates and maintenance allocations are determined, tail data are updated, and tests occur for state transitions.

Determining the embarked flying rate

The embarked fleet must achieve a certain number of flying hours every year. The annual requirement is split between each of the flights, where a flight represents a particular aircraft allocated to a ship. Each newly-embarked tail is designated to fly a pre-determined number of hours. The initial daily flying rate is set based on the required flying hours and the expected maintenance that will occur during the embarkation.

The model includes 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 deployment (e.g. during 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 the remaining embarked flights offset the rate of effort generated by tails flying at the surge and normal rates.

The fleet’s progress against the overall annual requirement is checked against the *pro rata* rate on a monthly basis. The monthly requirement is then adjusted accordingly if found to be tracking ahead or behind schedule. On a daily basis, the tails flying at the surge and normal rates of effort will fly at their calculated daily rate when they are serviceable. The flying rates for the low tails can then act as a ‘lever’, being adjusted down or up based on whether the fleet’s progress is above or below the *pro rata* rate respectively for a given month. This may mean that the ‘low’ flights accumulate only a small number of hours for their embarkation, or they may make up any shortfall in the annual embarked hours. This shortfall may be caused by low serviceability for the surge or normal tails, or by a deficiency in the number of embarked aircraft.

Determining the daily ashore flying rate

The ashore fleet must also achieve an annual requirement for hours flown. This is broken down into a daily requirement which is determined on a *pro rata* basis. The problem can be visualized as shown in Figure 2. Each day there are a number of ashore serviceable aircraft. Hours flown accumulate towards their next phased service, and there are a fixed number of phased maintenance lines.




Figure 2. Visualisation of the daily ashore flying program problem

The model needs to predict the ashore daily flying rate every day for 30 years, while accounting for the daily variation in the number of ashore serviceable tails. Tails can go into and return from regular inspections or unscheduled maintenance, or be deployed to or return from embarkation or exercises. Also, the maintenance duration can increase if unscheduled maintenance occurs on a tail already in maintenance. These factors preclude using optimisation techniques typically employed for these problems (e.g. Kozanidis et al. (2010)).

This paper tests two possible approaches to this problem. One procedure seeks to maximize the serviceability of the ashore fleet by minimising queuing for phased maintenance. However, it is not the explicit objective of this method to meet the daily requirement.

Firstly, all serviceable tails are matched to a phased maintenance line. The tail with the lowest number of flying hours until its next phased service is matched to the phased maintenance line with the least number of maintenance man hours (converted to days) to being free, and so on. If the number of tails exceeds the number of lines, the phased maintenance duration of the tails ahead must be included. The flying rate is then determined from the ratio of flying hours remaining to maintenance days remaining. Additionally, if this ratio is below the minimum daily rate, an aircraft will not fly on that day. Conversely an aircraft will not exceed the maximum daily rate if the ratio is above it. To ensure tail rotation, if serviceable tails are idle for longer than a specified period, they will be prioritised and their flying rates determined as above. This approach is similar to that of Mattila and Virtanen (2006), where they select aircraft with the highest ratio of remaining flying hours to remaining calendar time and employ heuristics to minimise queuing.

Alternatively, a ‘greedy’ heuristic can be used. This simply allocates the flying hours, starting with the aircraft with the least number of hours remaining until a phased service, until all hours are satisfied. However, this method may result in a larger number of unserviceable tails queued for phased maintenance.

For either method, the fleet is able to fly more than the *pro rata* rate (in this case up to 20%) in order to ‘stay ahead’ of the annual requirement if capacity exists to do so. This allows the fleet to ‘bank’ hours in case of major unserviceability or unavailability of aircraft, such as in simultaneous unscheduled maintenance events.

Determining the daily maintenance allocation

Each maintenance line has a specified amount of manpower allocated to it in terms of maintenance man hours (MMH) per day. This amount incorporates allowances for administrative duties and other non-maintenance tasks, so it represents hands-on time at the maintenance facility. For squadron-level maintenance, personnel are allocated to flight line and phased maintenance as explained in Section 2.1.

Nominally maintainers remain with their maintenance line and work on the individual aircraft occupying that line. However, the daily pooling and allocation of squadron maintenance personnel may also be considered. In this work, two approaches to maintenance allocation are tested:

- A ‘minimum’ approach, where crews remain with their allocated line and work at the daily rate; and
- A ‘greedy’ approach, where flight line maintenance personnel are pooled and allocated in order to return as many aircraft to a serviceable state on a given day, starting with the tail with least amount of maintenance remaining and so on. This means that tails with a small amount of remaining maintenance will be returned quickly, whereas tails with larger amounts may remain in maintenance for a long time.

Updating tails

Once these allocations are completed, tail data are updated if it is a work day. Using the allocation methods described above, serviceable tails accumulate flying hours, while unserviceable tails accrue time in maintenance. Other parameters which track time in calendar days (such as time between deep maintenance or time embarked) increment every day.

Testing for state transitions

The tests for state transitions occur at the end of each day. Firstly, tests occur to see if any embarked or ashore serviceable tails have been permanently removed from the fleet. This can be due to attrition (losing an aircraft) or life-of-type (reaching their maximum designated number of flying hours).

Secondly, tests are conducted to determine tails returning to a serviceable state. These can be through leaving maintenance, returning from exercises, or disembarking from a ship. Tails disembark once they have achieved their allocated number of flying hours. They may be forcibly disembarked in an unserviceable state if large unscheduled maintenance events cause the maximum time embarked to be exceeded. A ‘changeover period’ of up to one week allows for the transit time of tails embarking to or disembarking from a ship.

Thirdly, maintenance tests are undertaken. Scheduled maintenance falls due after reaching a certain elapsed time or achieved flying hours. Regular inspections of small duration can be addressed on the same day. Otherwise, tails will be unserviceable for at least one work day and may have to wait in a first-in-first-out queue if maintenance capacity is exceeded at the relevant facility. Where plausible, when DM is due, it is combined with a phased service. If modifications are also due, these are applied following the completion of the original service. Unscheduled maintenance can occur randomly on any available tail.

Finally, remaining serviceable tails are tested to see if they are eligible to be deployed for either embarkations or exercises. A disembarked tail triggers the need for a replacement, as the number of embarked aircraft will fall below the minimum. Only ashore serviceable tails are eligible and they must not fall due for DM while embarked. Then, the one(s) with the lowest number of embarked hours will be chosen in order to balance the embarked hours across the fleet. If no tails are eligible, the model records that the minimum embarked requirement was unmet on that day. Exercises occur at specified times throughout the year, with a number of ashore serviceable tails participating and flying at a higher rate of effort than normal.

3.4. Outputs

At the end of the fleet life, the model publishes relevant statistics as well as graphical results. Output parameters can include the primary Measures of Effectiveness (MOEs): i.e. the percentage of time the minimum requirement for embarked aircraft is met and the annual embarked and ashore hours achieved by the fleet; as well as availability, serviceability, and annual days in and awaiting various types of maintenance.

4. INDICATIVE RESULTS

The model was verified and validated through extensive consultations with stakeholders and by testing against the current S-70B-2 Seahawk fleet. This section demonstrates the types of outputs and insights that the model can provide using generic input data against indicative requirements.

Table 1. Sample general input data for fleet sizing model

Input	Units	Value
Minimum embarked	Number	8
Annual embarked requirement	Flying hours	4000
Annual ashore requirement	Flying hours	4000
Required annual embarked hours	Flying hours	500
Maximum daily ashore flying rate	Flying hours/day	24
Maximum daily ashore flying rate per tail	Flying hours/day	6
Minimum daily ashore flying rate	Flying hours/day	1

4.1. Input data

Table 2. Sample maintenance input data for fleet sizing model

Input	Value
Embarked maintenance	1 line per aircraft at 80 MMH/day
Flight line maintenance	5 lines at 60 MMH/day
Phased maintenance	4 lines at 40 MMH/day
Deep maintenance	3 lines
Regular inspections	Every 50 flying hours/14 days for 30 MMH; Every 100 flying hours/28 days for 60 MMH.
Phased maintenance	Every 200 flying hours for 600 MMH
Deep maintenance	Every 4 years for 60 + (0,30)work days
Modifications	Annually for 200 MMH; Every 4 years for 20 work days; After 15 years for 90 work days

Table 1 provides the sample general input data for the fleet sizing model. As mentioned in Section 3.1, fleet size is an input that can be tested. In this example, the possible loss of aircraft from the fleet through attrition or reaching life-of-type is excluded.

Table 2 shows the maintenance data for this example. The workforce rate of effort varies depending on the type of service. Ashore maintenance is often used for training, while embarked crews must be suitably trained. Therefore embarked crews have a higher rate of

effort reflecting the higher tempo of operations and their greater competency. Deep maintenance lasts for a minimum of 60 days, with an extra random duration of up to 30 days included to allow for possible unscheduled maintenance.

4.2. Results

The first set of results excludes the effects of a modification program. In all instances the embarked requirements for a daily minimum number of aircraft and annual flying hours are satisfied, as the model prioritises meeting embarked requirements over ashore requirements. Therefore any deficiencies in fleet size are revealed in the number of ashore flying hours achieved.

The annual ashore hours achieved are given for fleet sizes from 20 to 25 using four combinations of flying and maintenance allocation. These are: 1) minimum maintenance allocation and maximise serviceability flying allocation (denoted as Mmin, Fserv); 2) minimum maintenance allocation and greedy flying allocation (Mmin, FphG); 3) greedy flight line maintenance allocation and maximise serviceability flying allocation (MfltG, Fserv); and 4) both greedy flight line maintenance and flying allocation (MfltG, FphG).

The results for the annual ashore hours achieved for this example are shown in Figure 3. The influence of resource allocation on fleet size is clear from these results. For a minimum allocation of maintenance manpower to each line of maintenance, using a greedy allocation of flying hours provides an improvement on maximising the serviceability. This is despite queuing for phased maintenance increasing from an average of 5 days per year to an average of 450 days per year for a fleet size of 25. When a greedy procedure for allocating flight line maintenance is used, the number of annual ashore hours improves significantly. This efficient use of resources leads to an extra two serviceable aircraft each day on average, allowing greater flexibility in achieving the daily flying program. It would be left to the stakeholder to decide as to whether the achieved ashore hours are acceptable: for the (MfltG, FphG case), the values range from 3954±6 (20 aircraft) to 3986±4 (25 aircraft), i.e. within 1.2% to 0.4% of the requirement.

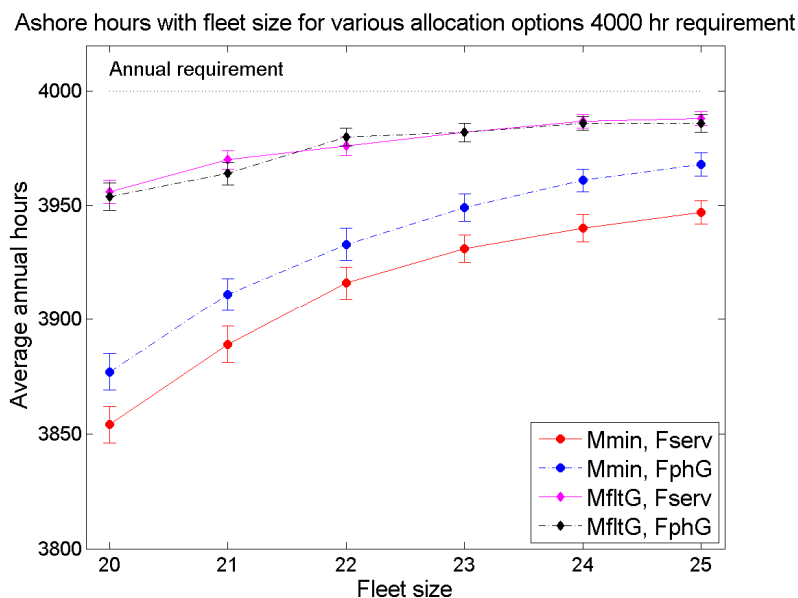


Figure 3. Annual ashore hours achieved for varying fleet sizes using possible flying and maintenance allocation heuristics

Table 3. Results for ashore hours achieved when including a modification program for various numbers of DM lines

Average annual ashore hours achieved				
Flt size	3 DM/no mods	3 DM/mods	4 DM/mods	5 DM/mods
22	3976	3822	3917	3924
23	3982	3833	3928	3935
24	3987	3835	3942	3949
25	3988	3844	3952	3960

Table 3 shows the impact of adding a modification program, with the (MfltG, Fserv) results given for 22-25 aircraft. There is a considerable decrease in the number of ashore hours achieved with 3 DM lines, due to the impact of the mid-life upgrade program. This causes significant queuing for DM, leading to a reduction in availability and ashore

serviceability. Adding extra DM lines reduces this queuing and leads to an increase in ashore hours achieved as shown. The results also suggest that it may be more cost-effective to employ extra DM capacity than a larger fleet, as the ashore hours achieved with an extra DM line are similar to those for an extra aircraft.

5. CONCLUSION AND FURTHER WORK

This paper demonstrates the capability of this discrete-event simulation model in determining fleet sizes. It also shows the ability of the model in conducting trade-off and sensitivity analysis, and the importance of considering the allocation of maintenance and aircraft resources in the process. The methodology is extensible to other aircraft types and systems that require complicated representation of maintenance and other influences on aircraft availability. The model is now being adapted as a fleet management tool using a fixed fleet size. Future work therefore involves further developments in improving the flying and maintenance allocation programs in order to maximise fleet availability and the efficiency of resources.

ACKNOWLEDGMENTS

The authors would like to acknowledge the assistance of numerous personnel from Capability Development Group, the Defence Materiel Organisation, the Royal Australian Navy and from the Defence Science and Technology Organisation, particularly from Joint and Operations Analysis Division and Aerospace Division.

REFERENCES

- Crary, M., Nozick, L.K. and Whitaker, L.R. (2002). Sizing the US destroyer fleet. *European Journal of Operational Research*, 136, 680-695.
- Commonwealth of Australia (2009). Defending Australia in the Asia Pacific Century: Force 2030. *Defence White Paper 2009*, Australian Government Department of Defence.
- Dong, J.-X. and Song, D.-P., (2009). Container fleet sizing and empty repositioning in liner shipping systems. *Transportation Research Part E*, 45, 860-877.
- Godwin, T., Gopalan, R. and Narendran, T.T. (2008). Tactical locomotive fleet sizing for freight train operations. *Transportation Research Part E*, 44, 440-454.
- Hoff, A., Andersson, H., Christiansen, M., Hasle, G. and Løkketangen, A. (2010). Industrial aspects and literature survey: Fleet composition and routing. *Computers & Operations Research*, 37, 2041-2061.
- Kozanidis, G., Liberopoulos, G. and Pitsilkas, C. (2010). Flight and maintenance planning of military aircraft for maximum fleet availability. *Military Operations Research*, 15(1), 53-73.
- Mattila, V. and Virtanen, K. (2006). 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, 38-43.
- Mattila, V., Virtanen, K. and Raivio, T. (2008). Improving maintenance decision-making in the Finnish Air Force through simulation. *Interfaces*, 38(3), 187-201.
- Shyshou, A., Gribkovskaia, I. and Barceló, J. (2010). A simulation study of the fleet sizing problem arising in offshore anchor handling operations. *European Journal of Operational Research*, 203, 230-240.
- Turnquist, M.A. and Jordan, W.C. (1986). Fleet sizing under production cycles and uncertain travel times. *Transportation Science*, 20(4), 227-236.
- Wesolkowski, S. and Billyard, A. (2008). The Stochastic Fleet Estimation (SaFE) model. Proceedings of the 2008 Spring Simulation Multiconference, Ottawa, Canada, April 2008, 1-5.