The size of MARS: quantifying requirements for the Royal Navy's future afloat support fleet

Yvan Gauthier[1,2], William Wardrop[1], Adam Price[1] and Alan Lawrence[1]*

[1]*Dstl Naval Systems Department, Dstl Portsdown West, Fareham, Hants, PO17 6AD, United Kingdom*
[2]*DRDC Centre for Operational Research and Analysis, NDHQ, 101 Colonel By Drive, Ottawa, Ontario, K1A 0K2, Canada*

ABSTRACT

Most of the ships providing afloat support to the Royal Navy will go out of service during the next decade. The Military Afloat Reach and Sustainability (MARS) programme aims to deliver a new fleet of ships for supporting forces engaged in expeditionary operations. This paper describes some of the operational research methods that were developed to quantify MARS requirements. Fleet size is assessed with a scheduling model that attempts to meet concurrent demands from predictable and unpredictable deployments, whilst taking into account force generation priorities and readiness. Cargo capacities are optimised by considering consumption rates for various types of bulk consumables, transfer rates, transit speed, and other operational factors. The influence of systems' availability and survivability in theatre is also examined.

1 INTRODUCTION

All the ships currently providing afloat support to the Royal Navy (RN), with the exception of the Wave class tankers, are expected to leave service over the next 10 to 15 years. In similar timescales, the logistic requirements of the RN will change with the introduction of new platforms, notably the future aircraft carrier. The Military Afloat Reach and Sustainability (MARS) programme [1] aims to fill the capability gap by delivering a new fleet of ships to support Maritime Forces deployed at sea and landed Joint Forces operating ashore. The MARS capability will be delivered by three classes of auxiliary ships, as described in Table 1.

Ship class	Primary role (and consumables)
MARS Fleet Tanker	Liquid support (bulk ship fuel, aviation fuel, oil, fresh water)
MARS Fleet Solid Support Ship	Solid support to carrier groups (bulk ammunition, food, stores)
MARS Joint Sea-Based Logistics Ship	Solid support to amphibious groups (bulk ammunition, food, stores)

Table 1. MARS ship classes and roles.

* Yvan Gauthier is completing a two-year secondment from DRDC to Dstl. Correspondence regarding this work should be addressed to Alan Lawrence at ajlawrence@dstl.gov.uk.

Dstl has been providing operational research (OR) support during the formulation of the MARS programme. A key element of the OR work was to quantify requirements in terms of the number of hulls and their individual cargo capacities. This entails a comprehensive assessment of the minimum capability needed to adequately support forces and enable campaign success.

This paper gives an overview of the main OR tools and methods developed by Dstl to derive MARS requirements. The emphasis is placed on the toolset used for setting requirements of the Fleet Tanker (FT) [2], which is the immediate priority of the MARS programme and the first class to be delivered.

Section 2 describes a scheduling model used for assessing the required number of tankers in the fleet. Two other models used for setting cargo capacity requirements are then presented in Section 3. The first estimates requirements at the task group (TG) level, whereas the second is a more detailed simulation of replenishments at the platform level. Factors influencing the requirements, such as systems' availability and survivability, are examined more closely in this section. Lessons learned while using the OR toolset are discussed in Section 4.

2 FLEET SIZE REQUIREMENTS

2.1 Background

Six MARS FTs are currently expected to enter service during the next decade. They will replace old single-hulled tankers being phased out to comply with the International Maritime Organisation requirement for double-hulled tanker operation. Considering the two existing Waves, a total of eight tankers will form the future tanker fleet. Dstl contributed to determine if this number is adequate for deploying enough tankers to operations when needed.

Traditional methods for assessing fleet size often assume fixed demand levels and fixed readiness cycles. In practice, the demand for ships and their availability both vary over time. The fleet scheduling process is also flexible; ship activities can be delayed or interrupted to support higher-priority deployments. The necessity to model this flexibility and capture the dynamic nature of force generation motivated the development of an analysis tool called Ship Scheduler.

2.2 Naval force generation modelling

A prototype tool was first built as an MSc project [3] and further developed by Dstl to better represent some real-world factors. Different approaches to fleet scheduling were examined. A construction heuristic emulating the decisions made by fleet programmers was deemed the most suitable and was implemented in a spreadsheet tool using Visual Basic® for Applications.

The types of activities represented in Ship Scheduler are described in Table 2. They fall into four generic categories: deployments to contingent operations, predictable deployments, pre-scheduled activities, and routine tasks.

Activity types	Description
Contingent operations	Unpredictable deployments (e.g. response to a conflict or crisis)
Predictable deployments	Standing commitments and other predictable deployments (e.g. patrol task)
Pre-scheduled activities	Activities being planned long in advance and associated with specific ships (e.g. refit)
Routine tasks	Routine peacetime tasks that do not require long-term planning (e.g. routine training)

Table 2. Generic activity types represented in Ship Scheduler.

Deployments. For the first two types of activities, ships must be deployed on a specific date. Deployments can be of different types, each type being defined by:

- a required set of ships;
- a priority ranking indicating its importance relative to other deployment types;
- a warning time (for contingent ops); and
- a readiness time for re-deployment.

Deployment instances are then defined by stating the type, the start time, and the duration of the deployment. One or more streams of deployments can be scheduled by the model.

Pre-scheduled activities. These are predictable activities that are typically planned far in advance and associated with specific ships, such as upkeep activities. Pre-scheduled activities may have a single phase or successive phases, and a flexible start date, as shown in Figure 1.

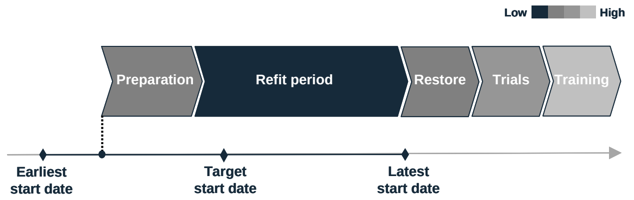

Figure 1. Example of a pre-scheduled activity with five phases.

Specifically, each pre-scheduled activity is defined by:

- the ship it applies to;
- the earliest start time, latest start time, and target start time of the activity;
- the duration of each phase;
- the ship's readiness during each phase;
- the priority ranking required by a deployment to justify an interruption of the activity; and
- the re-scheduling action to be taken in case of an interruption.

The re-scheduling action is selected from a list of options. For example, the user can request that an activity, if interrupted, has to be repeated upon return from the deployment that caused its interruption. Alternatively, the user can request not to re-schedule the activity, or only re-schedule particular phases.

Routine tasks. These tasks are defined in a similar way as deployments and require a specific set of ships to be generated. Routine tasks are assumed not to require long-term planning and no start times must be specified. Tasks may be prioritised between themselves, but are assumed to have a lower priority than deployments and pre-scheduled activities.

Scheduling method. Given a set of activities to perform and a set of ships of possibly different types, the model allocates ships to activities and determines how many of them can be adequately supported. A simplified description of the construction heuristic used to schedule activities is depicted in Figure 2. It works as follows:

(1) Pre-scheduled activities are loaded into the empty schedule. They are temporarily assumed to start as late as possible.

(2) The first deployment to start is considered. If more than one deployment starts at the same time, they are sorted by decreasing priority and duration.

(3) One of the ship types required by the deployment is selected for scheduling.

(4) For the selected ship type, the model first considers ships that are available on the start date.

(5) Assuming a ship is available on the start date, it is flagged as a candidate for deployment. If more than one ship is available, they are sorted by:
   - ships not forced to miss the target date(s) of upcoming pre-scheduled activities in order to perform the entire deployment; then by
   - ships already assigned to a similar deployment and already in theatre when deployment starts; then by
   - ships at highest readiness; then by
   - ships available the soonest.

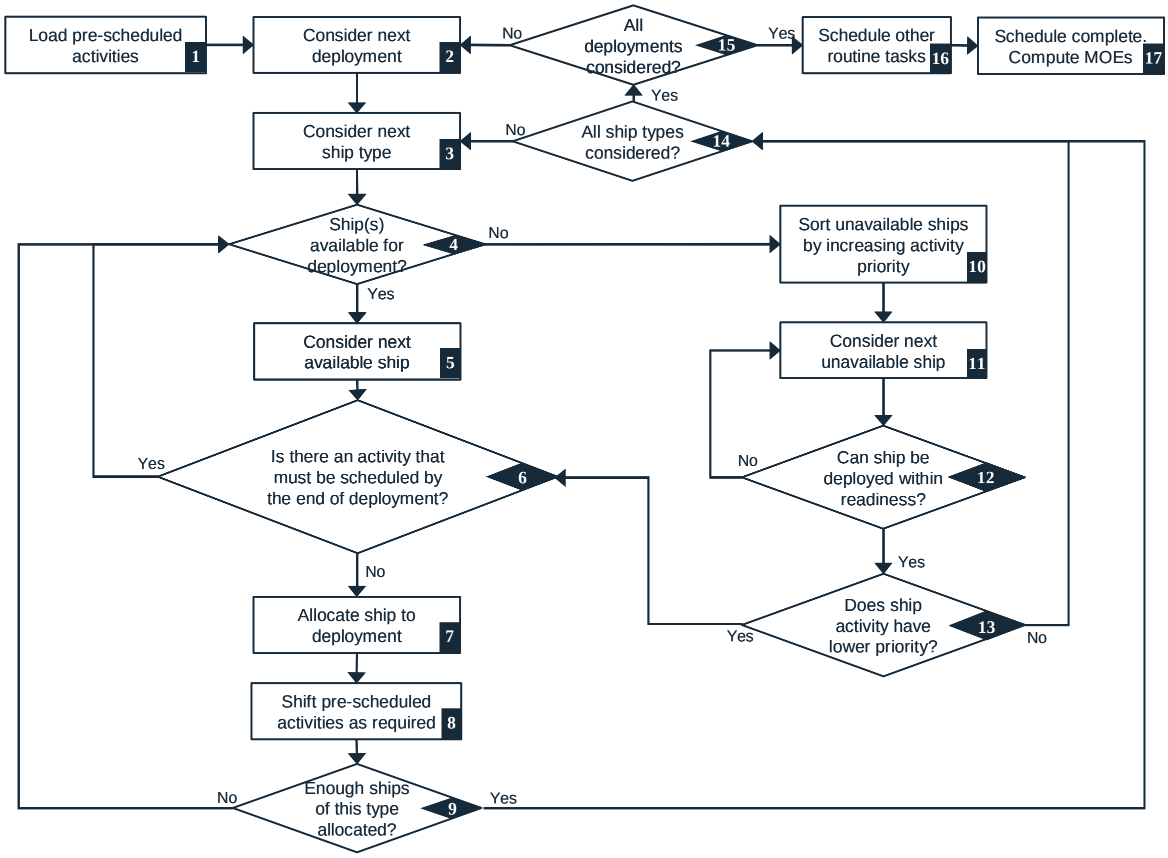

Figure 2. Simplified view of the Ship Scheduler heuristic.

(6) The model verifies whether the ship is available for the entire deployment. If the ship must start a pre-scheduled activity by the end of the deployment, the model selects the next available ship (if any).

(7) Assuming the ship is available for the entire deployment, it is allocated to it.

(8) A decision whether to start some of the pre-scheduled activities earlier is taken. If the model foresees deployments (based on their warning times) for which delaying a pre-scheduled activity may be necessary, the model temporarily keeps the start time of this activity as late as possible. Otherwise, the model tries to start the activity earlier. In doing so, the model attempts to meet the target date, unless:

- previous or upcoming deployments require deviation from the target date; or
- the number of ships being simultaneously at low readiness on target date exceeds a pre-set limit for that ship type.

(9) The model verifies the allocation of the required number of ships. If more ships of that type are required, the model looks for more available ships.

(10) If there are not enough ships available to meet the requirement, the model sorts, by increasing order of priority, ships that are already deployed or allocated to pre-scheduled activities.

(11) The model then attempts to pull a ship from the lowest-priority activity.

(12) The ship is pulled from the activity only if it can be ready on time for the deployment date. Otherwise, the next unavailable ship is considered.

(13) The deployment requiring support must have a higher priority, and no other activities must be assigned to the ship before the end of the deployment. If there are still not enough ships allocated to the scenario, the model notes the shortfall and moves on to the next ship type.

(14) The process described above is repeated for all ship types requiring deployment.

(15) The process is then repeated for all deployments.

(16) Once all deployments have been considered, the model attempts to fit other routine tasks into the parts of the schedule remaining empty. Before proceeding, the model sorts routine tasks by

- priority; then by
- the total number of ships required; then by
- duration.

(17) Once the schedule is complete, the model reviews the list of all deployments, pre-scheduled activities and routine tasks. It then computes different measures of effectiveness (MOEs) on how well activity requirements are met.

2.3 Assessing the fleet size

Demand streams including various types of deployment potentially undertaken by UK forces are generated by an external, stochastic model [4]. The streams are sequences of standing commitments and contingent operations of different types and scales. The contingent operations are initiated based on probabilities derived from historical analysis. One-off operations have a fixed duration, whereas enduring operations are periodically extended based on probabilities also derived from historical analysis. Constraints are applied such that the number of concurrent deployments and recuperation intervals between deployments are within limits set by defence policy.

The demand streams are then imported into Ship Scheduler. Following the method previously described, the model allocates tankers to deployments, taking into account force generation priorities, upkeep requirements, and ship readiness.

Based on the resulting schedules, the model produces detailed ship-to-deployment allocation tables. These tables are used to calculate various MOEs (e.g. percentage of activities supported with the required number of tankers). The allocation tables are also used to infer the likelihood of supporting deployments in specific contexts, for example when multiple operations must be concurrently supported in different theatres.

Different MARS solutions were compared by repeating this analysis with different number of tankers in the fleet. The risk of not being able to support concurrent operations was the primary MOE used. Operations being the most at risk of not receiving adequate tanker support were also identified.

3 CAPACITY REQUIREMENTS

3.1 Background and definitions

Once deployed, tankers must carry enough consumables to meet the demand. In general, afloat support ship activities can be divided into two main roles. Ships in the *station role* stay with a TG at all times to provide replenishment at sea (RAS) and extra holding capacity to the group. Ships in the *consolidation role* support the TG by:

(1) transiting to a well-found port;

(2) loading commodities;

(3) transiting back to the TG; and

(4) replenishing station ships.

This process is repeated as many times as necessary over the course of a deployment. The time required for completing steps (1) and (3) is a function of the transit speed and the distance to port, which varies depending on the scenario of interest. The time required for completing step (2) or step (4) is the *turnaround time*. The total time required to complete all steps is the *consolidation period*.

The consolidation process does not start immediately as a TG arrives in theatre. There is an *unsupported period* during which the TG must operate independently, without any logistic pipeline in place. This period extends from the last consolidation performed during transit to the first time a consolidator returns to the TG after entry into theatre.

The *reserve level* is the minimum fuel holding allowed by command to maintain operational flexibility across the TG.

3.2 Group-level replenishment modelling

Cargo capacity requirements are first estimated at the group level using the Support Ship Optimiser (SupShOp) [5]. SupShOp is a spreadsheet-based model optimising the cargo capacities of support ships within a TG. The objective is to ensure there are enough bulk consumables within the group to sustain it over the entire duration of a given campaign.

In essence, SupShOp calculates the total consumption of a TG for each day of the campaign, based on the following scenarios assumptions:

- platforms deployed;
- expected activities of individual platforms on that day (e.g. transit, loiter, normal air ops, peak air ops, amphibious ops, maintenance);
- typical time spent at minimal, transit and maximum speed during a ‘standard day’ of the activity planned for each platform; and
- consumption rates for each platform at minimal, transit and maximum speed.

These assumptions are derived from workshops involving military experts and logisticians. They are based on representative missions, timelines and TG locations throughout the force planning scenarios of interest.

SupShOp then finds the minimum amounts of commodities that support ships must carry in order to maintain sufficient commodity levels within the TG, whilst satisfying four fundamental constraints:

- There must be enough bulk consumables within the TG (including any consolidators) to last the unsupported period.
- The station ship(s) must hold enough consumables to sustain the TG through each consolidation period.
- The consolidator(s) must bring enough consumables to sustain the TG through each consolidation period.
- The consolidator(s) must not bring more consumables than the amount of empty space within the TG.

Figure 3 illustrates how this is done for a commodity such as ship fuel. It shows the holding level of the entire TG over time, first during transit, then during the unsupported period (shaded), and then during each of the consolidation loops. The combined holding capacities of warships and support ships are just large enough to avoid the group’s stock going below the reserve level.

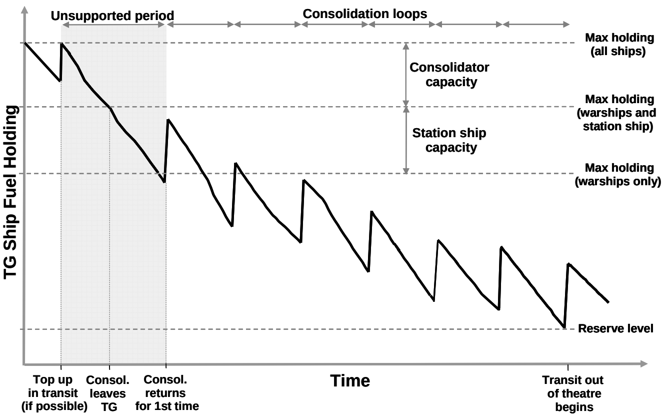

Figure 3. Example of TG fuel holding as a function of time with minimised tanker capacities.

In this example, the minimum is reached just before leaving theatre, which means that capacity requirements are driven by the consolidation process. If requirements were driven by the unsupported period, the minimum would occur at the end of it and stocks would be progressively re-built during consolidation.

The consolidation requirement is dependent on the station ship requirement. If the station ship capacity were to increase, the consolidation requirement would decrease accordingly. Figure 4 shows an approximation of this relationship for a particular scenario. The relationship is piecewise linear because time units in SupShOp are integer days.

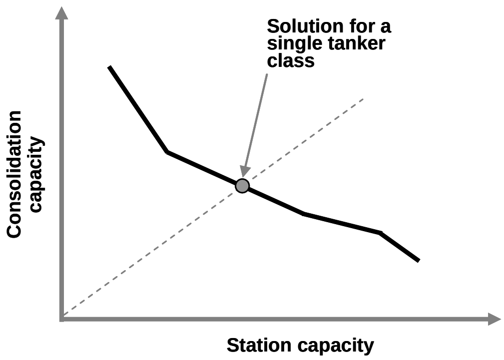

Figure 4. Relationship between station and consolidation requirements for a given scenario.

If there is a single station ship and a single consolidator in the TG, and if both ships are of the same class, SupShOp takes this into account and attempts to find a solution where the station and consolidation requirements are equal.

3.3 Deriving capacity requirements

Capacity requirements are determined by repeating the analysis described above for all consumables and TGs of interest. Multiple force planning scenarios with different characteristics (location, forces being deployed, combat intensity, etc.) must be covered to ensure that the solution will meet the demand from a variety of campaigns.

Several factors can drive the requirements up or down. The key ones are presented below and their influence must be well understood before setting the final ship specifications.

Consolidation distance and transit speed. Because the consolidation period is mainly driven by the distance to port and the consolidator's transit speed, both factors influence the requirements. Figure 5 shows how deviating from a baseline speed and distance in a given scenario changes the capacity requirements for ship fuel and aviation fuel. Note that deviating from the baseline has slightly different effects on capacities for the two types of fuel, mainly because their consumption rates vary differently over the course of the campaign.

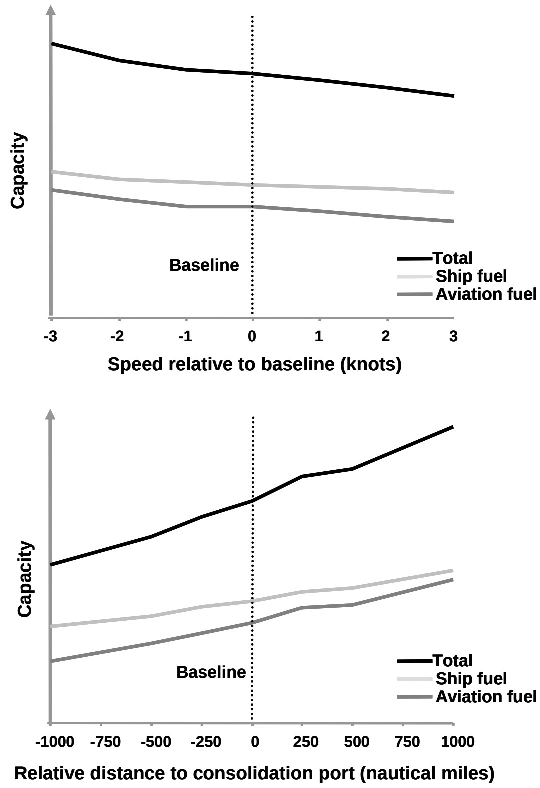

Figure 5. Relationship between capacity requirements and consolidator speed (top) and distance to port (bottom).

Scenario assumptions. The number of ships being deployed, the number of aircraft and joint forces embarked, as well as the expected platform activities obviously have a major influence on the requirements. Since these assumptions are mainly obtained from expert judgement, sensitivity analyses are beneficial for quantifying the potential implications of having underestimated the demand.

Reconfigurable cargo space. One way to set the requirements would be to identify, for each type of consumable, the largest capacity required amongst all of the TGs and scenarios analysed. However, the total ship capacity obtained by summing up the requirements would often be larger than necessary.

Consider two hypothetical scenarios with the capacity requirements shown in Figure 6. In Scenario A, each tanker must be able to carry at least 8,000 m^3 of aviation fuel, and in Scenario B each tanker must be able to carry at least 10,000 m^3 of ship fuel. Simply adding up these two figures would mean that a tanker with 18,000 m^3 of total capacity is needed for supporting either scenario A or B when it arises. However, if at least 2,000 m^3 of the capacity can be used to carry either ship fuel or aviation fuel, both scenarios can be supported with a tanker of 16,000 m^3 capacity.

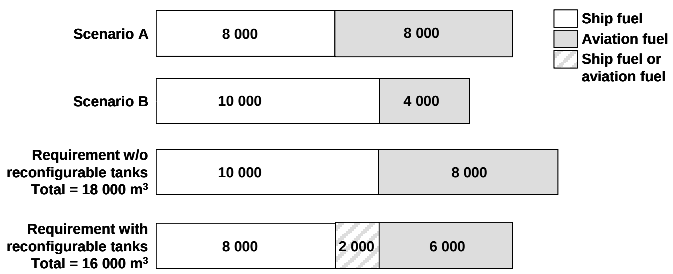

Figure 6. Fuel capacities (in m^3) required by two notional scenarios. A tanker with reconfigurable tanks can support either scenario with 16,000 m^3 of usable capacity.

The same strategy could be applied to other types of consumables, for instance on solid support ships carrying food, stores, and ammunition in reconfigurable cargo space.

3.4 Platform-level replenishment modelling

The main limitation of SupShOp is that it neglects how consumables are distributed within the TG. Although the group as a whole may have enough consumables, some ships may have large stocks whilst other may need replenishment. The factors driving how commodities are distributed can only be analysed at the platform level.

Fleetflow [6] simulates the RAS processes by which support ships supply consumables to individual warships at sea. As in SupShOp, ships have pre-defined capacities and consumption rates. But instead of setting a reserve level for the entire TG, each ship has its own reserve (minimum allowable) level of holdings below which it loses the operational flexibility needed to conduct missions.

On each day, many ships approaching their reserve levels may request replenishment. Fleetflow emulates the planning of a TG logistic coordinator and predicts when the ships will fall below their reserve levels. It also determines how many ships can be replenished on each day. In the process, it takes into account factors influencing the duration of individual replenishments, such as:

- the time to connect or disconnect RAS rigs;
- transfer rates (based on hose diameters, maximum flow rates, weight tolerance of RAS rigs);
- the amounts of consumables to transfer; and
- the number of RAS rigs that can be connected.

Fleetflow also considers constraints on the maximum number of replenishments that can be performed each day, based on manpower limits (in terms of maximum working hours per day) and the amount of time ships have to transit before each RAS.

Conducting the analysis at the platform level makes it possible to model the Availability, Reliability, and Maintainability (ARM) of individual ship systems, as well as their survivability to battle damage.

Fleetflow models ARM based on the Mean Time Between Failure (MTBF) and the Estimated Time for being Back Online (ETBO) of critical replenishment capabilities. The MTBF and ETBO values are derived from a historical database of system breakdowns and repairs. These values define probability distributions from which downtimes are stochastically generated and considered by Fleetflow when computing RAS plans. Fleetflow also models system downtimes due to battle damage. The times and severity of damage are obtained from external combat simulations and survivability modelling.

The impact of ARM and survivability on replenishments depends on the nature of the damage or defect. If a propulsion system (e.g. generator, engine, shaft) is broken, a support ship may still operate, but the lower speed will reduce the number of replenishments achievable per day. If replenishment systems (e.g. rigs, pumps) are broken, the impact on RAS plans may be minor if other rigs on the ship are still usable.

If the battle damage or defect is severe, the support ship may have to return to port for repair. This is likely to cause major delays in the re-supply process, especially if the ship remains unavailable for several days.

Fleetflow is run as a Monte-Carlo simulation for producing statistics on the consequences of defects and battle damage. The MOE is the number of missions during which warships remain above their minimum holding levels (a mission is defined here as a portion of a campaign where a TG is involved in a specific activity). A mission is considered inadequately supported if ship holdings go below their reserve levels.

Figure 7 shows how the MOE can vary depending on the scenario analysed. It gives the average percentage of missions adequately supported by supply ships for three notional scenarios, along with the 5th and 95th percentiles.

A combination of factors can make a scenario more challenging to support than another. For instance, if the consolidation distance is very long and the consolidator suffers a defect shortly before it is due to rejoin the group, it may have to return to port for repair. This greatly increases the time needed for new stocks to get to the group, causing many missions to fail. A similar situation can also arise if the consolidator must go for repair just before the end of the unsupported period.

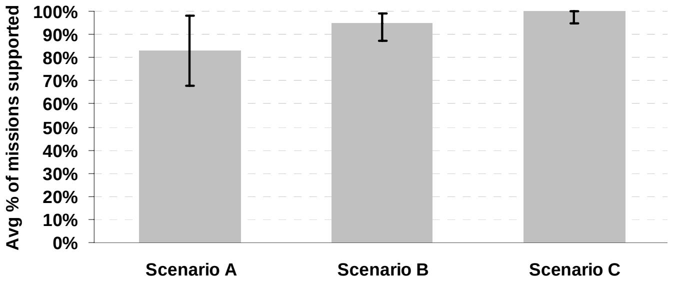

Figure 7. Examples of MOEs produced by Fleetflow for a same tanker design in three notional scenarios.

Another example is the sinking of a station ship in combat, forcing the consolidator to replenish warships by itself. The increase in the consolidation period, combined with the loss of holding capacity within the group, can cause many missions to be inadequately supported.

A particularly challenging situation is when defects or damage simultaneously occur on the station ship and the consolidator, as the entire re-supply process may then stall. When this occurs, warships can only rely on their individual holdings for a short period of time beyond which several missions will not be adequately supported and may fail.

4 LESSONS LEARNED

Optimising the size and capacities of a support fleet is complex. A considerable amount of analysis is required to ensure that the logistic requirements of expeditionary forces will be met, without recommending a solution that is unnecessarily big and costly. Many lessons can be drawn from the OR modelling carried out by Dstl for MARS. A few of them are discussed below.

A scheduling model is essential to correctly assess the size and force generation capability of a fleet. Although Ship Scheduler is still a relatively crude representation of fleet programming, it proved to be very useful in comparing candidate MARS solutions and estimating the risks of not being able to deploy enough tankers to operations.

A group-level analysis model such as SupShOp can produce reliable estimates of the minimum cargo capacities required to support a TG, as a cross-validation with Fleetflow demonstrated. Results can be quickly produced and are easy to interpret. Group-level modelling is also very useful for identifying the factors that drive requirements and perform 'what if' analyses around them.

There are however many factors that cannot be easily modelled at the group level, such as ARM, survivability, and the distribution of consumables within the group. A platform-level simulation such as Fleetflow is more appropriate for analysing these issues. More complex RAS strategies (e.g. strategies that involve partial replenishments or supply ships supporting more than one TG) are also easier to represent in a simulation. Modelling such strategies is particularly worthwhile to determine if specific warships can be adequately supported during peak consumption periods.

Another lesson learned from modelling is the importance of analysing many scenarios with different characteristics when setting the requirements. Because several factors have an influence on the requirements, the driving scenarios may not be intuitively the most demanding. Furthermore, different scenarios may drive capacity requirements for different commodities. Having ships with reconfigurable cargo space can greatly increase the flexibility of the support fleet and its ability to enable campaign success.

ACKNOWLEDGEMENTS

The authors would like to thank the MARS Integrated Project Team and Mr. Duncan Feveyear for their guidance during the development of the tools described in this paper.

REFERENCES

[1] DE&S MARS Integrated Project Team (2008), *Defence Fact Sheet – Military Afloat Reach and Sustainability (MARS)*, http://www.mod.uk/DefenceInternet/FactSheets/ProjectFactsheets/MilitaryAfloatReachSustainabilitymars.htm

[2] DE&S MARS Integrated Project Team (2008), *Fleet Tanker Schedule 3 Technical Specification – Unclassified version*, http://www.mod.uk/DefenceInternet/MicroSite/DES/OurPublications/FleetTankerTechnicalSpecificationUnclassified.htm

[3] L. Fildes (2006), *A Decision Support Tool for the Procurement and Planning of Royal Navy Ships*, MSc Dissertation, Faculty of Engineering, Science and Mathematics, University of Southampton

[4] A. Kashyap and S. Fagg (2006), *Demand in Deployed Operations (DIDO) User Guide*, Dstl TR15805 v1.0, Unclassified

[5] A.J. Price (2006), *Optimising the Support Ship Fleet*, MSc Dissertation, Faculty of Engineering, Science and Mathematics, University of Southampton

[6] W. Wardrop (2006), *Analysis of Maritime Force Sustainment*, Proceedings of the International Conference on Logistics and Supply Chain Management