Meeting Operational Demand – Determining Output for the Royal Canadian Navy

Dr. Van Fong **Jean-Denis Caron**
Defence Research & Development Canada
Centre for Operational Research and Analysis
101 Colonel By Drive, Ottawa, Ontario K1A 0K2
CANADA

van.fong@forces.gc.ca jean-denis.caron2@forces.gc.ca

ABSTRACT

The Royal Canadian Navy is currently in the midst of modernizing its fleet and procuring new classes of ships. Defence Research & Development Canada Centre for Operational Research and Analysis supports the Royal Canadian Navy and has been asked to provide scientific evidence and insight into fleet size, specifically how it affects the operational output of the fleet vis-à-vis the ability to meet operational demand. In this paper, we present a method for determining the fleet output by analyzing the supply and demand of a fleet. The demand is modeled using a set of vignettes where each vignette is characterized by a frequency of occurrence, its duration and an expected response in terms of number of ships required. The steady-state supply of ships is a function of the fleet size and the operational cycle, which depends on the maintenance profile of the ships, and how much time the ships are at high readiness. The method is illustrated using fictitious vignettes on a notional single class fleet. The intent in the future is to adapt and expand the methodology to analyze multiple classes of ships simultaneously with vignettes requiring more than one ship class.

1.0 INTRODUCTION

The Royal Canadian Navy (RCN) is currently in the midst of modernizing its fleet and procuring new classes of ships. The Director General of Naval Force Development (DGNFD) tasked the Defence Research & Development Canada (DRDC) Centre for Operational Research and Analysis (CORA) to provide scientific evidence and analysis to validate new concepts related to fleet size and how it affects the operational output of Canada's future fleet and its ability to meet demand.

Fundamentally, the question of fleet sizing and determining its operational output is a question of supply and demand: Given a set of constraints and conditions, how many ships are needed in order to meet a specified level of operational demand?

1.1 Previous Work on Fleet Sizing

There have been many efforts by DRDC in the past addressing this question, each with its unique set of strengths and shortcomings. Beginning in 1999, Massel et al [1] modeled fleet structures and their outputs using simulation studies to explore alternative fleet mixes for the RCN. In 2005, Allen et al [2] conducted the first in a series of 'Fleet Mix Studies' for the RCN, using a capability-based approach to matching fleet supply to operational demand in a stochastically driven model called Tyche. Between 2005 and 2010, the Fleet Mix studies continued by Allen and Eisler [3], Eisler [4-6], Bourque and Eisler [7-9] and culminated in 2010's Fleet Mix Study II by Bourque and Eisler [10], that provided the most detailed and comprehensive capability-based look across the RCN's future fleet options.

Between 2010 and 2015, in the face of fleet modernization and planning the introduction of new ship classes, the RCN based many of their decisions on the outcomes from the Fleet Mix Study II. However, there was a desire within DGNFD to look at individual ship classes again, to have focused answers to the specific problems they faced with each ship class rather than a whole-of-Navy look. DRDC began producing class-specific analyses, starting with Bourque [11], who initially continued with the Tyche model, and eventually developed new tools and tailored approaches to the class-specific questions at hand (Bourque and Mirshak [12-13]).

1.2 Current Approach

This paper illustrates the methodology used to produce the early results of the Future Fleet Analysis commissioned by DGNFD in 2016. The present method conducts a detailed demand analysis using a set of vignettes where each vignette is characterized by a frequency of occurrence, its duration and an expected response in terms of number of ships required. From the demand point of view, it approaches the characterization of operational demand for ships in a similar approach to Tyche. From the ship supply perspective, a steady-state supply was modeled as a function of the fleet size and the operational cycle, which depends on the maintenance profile of the ships and how much time the ships are at high readiness (HR), the most demanding readiness posture for RCN ships as defined in [5]. This approach is very similar to what was conducted in the recent years to provide the class-specific look at fleet sizing. In addition, this initial effort considered only vignettes that required ships that are at HR, which focuses the analysis to the case where the specific ship class being addressed becomes the limiting factor to providing the required response by the RCN.

2.0 METHODOLOGY

This section of the paper will present an overview of the methodology used to establish operational demand and supply levels, and determine the likelihood for a given configuration of fleet to meet the expected demand. The methodology will be presented using fictitious data on a notional single class fleet.

2.1 Establishing Operational Demand

As mentioned previously, the current methodology uses a stochastically driven approach to establishing operational demand for ships. The way in which this is accomplished is through the use of maritime vignettes – a set of hypothetical scenarios that aim to capture the spectrum of activities that the RCN is prepared to respond to. The usage of vignettes and scenarios is widely used within the Department of National Defence for Force Development projects using Capability-Based Planning [6] and thus, a well-understood mechanism for generating expected levels of demand.

2.1.1 Maritime Vignettes

The Directorate of Naval Strategy (DNavStrat) defined a set of 54 distinct maritime vignettes that span the spectrum of naval activities from peace-time routine patrols to major military conflicts. For the purpose of this paper, the nine notional vignettes shown in Table 1 provide a basic understanding of the structure and usage of the vignettes.

Each vignette in Table 1 carries the same type of information in the database: a unique internal ID number, a Description text field, a set of fields to establish the likelihood of occurrence (i.e. Type, Frequency, Duration, Time of Year and Co-occurrence limits) and finally, the Expected Response which identifies the number of ships required.

Table 1 Set of Nine Maritime Vignettes (with notional data)

ID	Description	Type	Frequency of Occurrence	Duration	Time of Year	Co-Occur	Expected Response (Number of Ships)
1	Single ship - Short deployment	Random	3 / 2 years	1-4 months	All year	Yes	1
2	Single ship - Medium deployment	Random	2 / 3 years	6-8 months	All year	Yes	1
3	Single ship - Long deployment	Random	1 / 3 years	2 x 6 months	All year	No	1
4	Two ships - Medium Deployment	Random	1 / year	4-6 months	All year	No	2
5	Task Group - Large Scale Commitment	Random	1 / 6 years	3 x 6 months	All year	No	3
6	Standing NATO Maritime Group 1	Scheduled	2 / 3 years	8 months	Feb to Nov	N/A	1
7	Flagship for Standing NATO Maritime Group 1	Scheduled	1 / 3 years	8 months	Feb to Nov	N/A	1
8	Spring / Fall scheduled operations	Scheduled	2 / years	1 month	May and Sep	N/A	1
9	Deployment w/ related exercises	Scheduled	1 year	5 months	Apr to Sep	N/A	2

In order for the model to stochastically generate the demand, it is necessary to first separate the vignettes into two different Types: Scheduled and Random. Scheduled vignettes represent maritime commitments that are pre-planned and occur with regularity over time. An example of such activity is Canada's ongoing commitment to sending a surface combatant ship to the Standing NATO Maritime Group 1 [7] (i.e., Vignette #6). Random vignettes represent maritime operations that are activated in response to an incident occurring, either domestically or internationally. One example of a maritime response to a random incident is Operation HESTIA, which was a humanitarian assistance operation in Haiti following a catastrophic earthquake in 2010 [8].

Based on the assigned vignette's frequency and duration, the model generates a timeline of *events* (i.e., specific instantiations of vignettes) that determines the operational demand. Deriving realistic numbers for the frequency and duration is one of the more difficult and time-consuming parts of the process, requiring as much art as science. DNavStrat is responsible for establishing these parameters for the maritime vignettes, which used a combination of historical analysis over decades of Canadian operations and strategic forecasting. In this paper, only notional values have been assigned to each vignette to provide the reader with an appreciation for the real data set.

In addition to the frequency and duration of occurrence, there are two additional fields to further refine the likelihood of events occurring in the model: Time of Year and Co-occurrence. To provide added realism to the Demand timeline, the Time of Year field is used to strictly limit the times in a calendar year when certain vignette events can be generated. For example, RIMPAC [9] is an international maritime exercise that typically takes place in the summer months (e.g., June – July). The Co-occurrence field is a second limit that is placed on random events to make sure that multiple events of the same vignette do not occur at the same time.

2.1.2 Stochastic Modelling of Operational Demand

Modelling operational demand is performed within *R* [10], which stochastically generates a list of events that occur over a user-defined length of time, which for the purpose of illustration, was chosen to be five years – the typical Operational Cycle (OPCYCLE) of a naval ship. The number of events for a given random vignette within the timeframe is assumed to be Poisson distributed with frequencies of occurrence corresponding to Table 1. The start time of both scheduled and random vignettes are drawn from uniform distributions with parameters as specified in Table 1 (e.g., Vignette #9 occurs between April and September). Finally, triangular distributions are used to determine the duration of a vignette (e.g., Vignette #2 has a distribution with $min=6$, $mode=7$, $max=8$). The end result is a timeline of events, where each event is identified by its vignette ID, a start time, a duration and the number of ships required. Figure 1 illustrates an example of an event timeline.

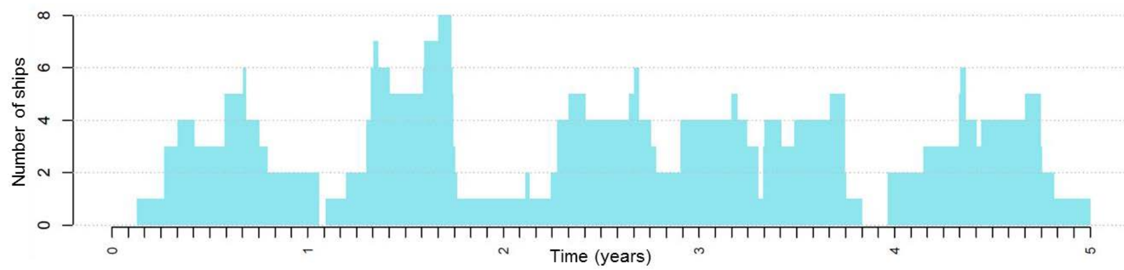

Figure 1 Example of a Single Event Timeline

In Figure 1, the horizontal axis represents time while the vertical axis represents the total number of ships required based on the Expected Response field for the vignettes. Often the events generated in the model overlap or co-occur in time, therefore, the total number of ships required per day is made up of any number of events. There are a number of days in this example where no ships are required, as in the beginning of Year 0 and the end of Year 3, which corresponds to the recognized fact that scheduled events are kept from occurring over the holiday season. More often, however, there are multiple events occurring simultaneously, which is most strongly illustrated in Figure 1 by the demand for eight ships in the late summer of Year 1.

In order to build up statistics to better characterize the operational demand, the model is run many times and the resulting individual timelines are aggregated together.

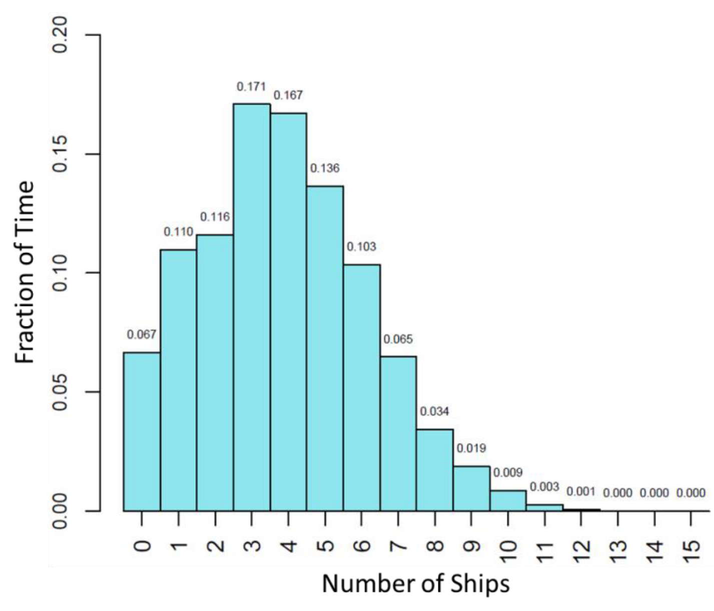

Figure 2 Histogram showing Fraction of Time when *x* ships were required

The discrete probability distribution function (PDF) in Figure 2, denoted by $Pr(X = x)$, summarizes the result of 1000 runs using the notional vignette data provided in Table 1. The histogram shows the Fraction of Time when x ships were required, which was computed by counting the number of days across 1000 timelines when the demand for x ships appeared divided by the grand total of days (365 days × 5 years × 1000 runs). It can be seen in Figure 2 that 6.7% of the time, no ships were required (i.e. no vignettes occurred). Also, the most likely demand (17.1% of the time) called for three ships, while the maximum simultaneous demand was for 15 ships (0.0005% of the time).

2.2 Establishing Operational Supply

In this paper, a steady-state supply of ships is assumed, which is driven by two main factors:

1. Fleet size – Number of ships in the fleet; and
2. OPCYCLE – The expected maintenance profile of the ships and how much time each ship spends at HR.

Figure 3 shows an example of a 5-year OPCYCLE, which will be used for illustration in this paper. The red area (18 months) represents the time when the ship is not available because it is either in maintenance, in tiered-readiness program or being worked up. The light green (total of 30 months) corresponds to periods of time when the ship is available only for force generation activities or missions not requiring the full-spectrum of capabilities.

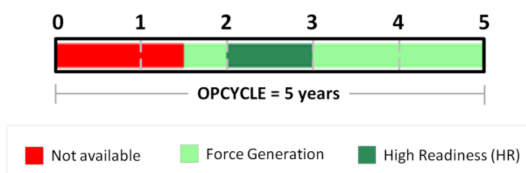

Figure 3 Example of a 5-Year OPCYCLE

The dark green, 12 months in this case, represents when the ship is at HR. In reality, to account for factors such as force generation requirements, scheduling conflicts and operational deployment patterns, ships may be ramped up to HR earlier (or later) in the OPCYCLE, and the HR period can even be broken down into more than one block. A generic OPCYCLE like the one in Figure 3 is used to build an idealized fleet schedule.

As an example, for a fleet of 20 ships following the OPCYCLE shown in Figure 3, there are exactly four ships at HR at any given time. This is calculated by taking the product of the number of ships (i.e., 20) and the number of months at HR (i.e., 12), and dividing it by the length (in months) of the OPCYCLE (i.e., 60). The idealized fleet schedule for this example is provided in Figure 4, which shows that there are exactly four ships at HR (i.e., dark green) at any given time in the OPCYCLE. To illustrate, at the beginning of Year 0, Ships 10, 11, 12 and 13 are all at their highest state of readiness, while at the beginning of Year 3, the ships at HR are 2, 3, 4 and 5.

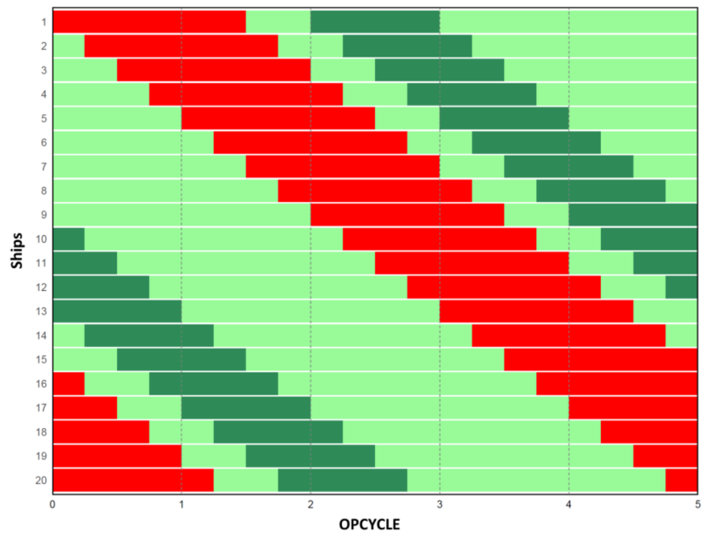

Figure 4 Idealized Fleet Schedule for 20 ships

Early in the process of procuring a new class of ship, there are many unknowns, starting with the number of ships to be acquired. Assumed operational and maintenance practices that were in place with the previous ship class can also be modified. This includes the way the ships are going to be employed in the future, in particular, how long the ships are expected to be kept at HR, their highest state of readiness.

These are the two main variables considered in this paper. The number of ships in the fleet was varied between 18 and 30 while the number of months a ship is at HR ranged between 9 and 21. This results in a total of 169 options for ship supply and are shown in Figure 5.

		Months at HR / OPCYCLE												
		9	10	11	12	13	14	15	16	17	18	19	20	21
Number of Ships in Fleet	30	4.50	5.00	5.50	6.00	6.50	7.00	7.50	8.00	8.50	9.00	9.50	10.00	10.50
	29	4.35	4.83	5.32	5.80	6.28	6.77	7.25	7.73	8.22	8.70	9.18	9.67	10.15
	28	4.20	4.67	5.13	5.60	6.07	6.53	7.00	7.47	7.93	8.40	8.87	9.33	9.80
	27	4.05	4.50	4.95	5.40	5.85	6.30	6.75	7.20	7.65	8.10	8.55	9.00	9.45
	26	3.90	4.33	4.77	5.20	5.63	6.07	6.50	6.93	7.37	7.80	8.23	8.67	9.10
	25	3.75	4.17	4.58	5.00	5.42	5.83	6.25	6.67	7.08	7.50	7.92	8.33	8.75
	24	3.60	4.00	4.40	4.80	5.20	5.60	6.00	6.40	6.80	7.20	7.60	8.00	8.40
	23	3.45	3.83	4.22	4.60	4.98	5.37	5.75	6.13	6.52	6.90	7.28	7.67	8.05
	22	3.30	3.67	4.03	4.40	4.77	5.13	5.50	5.87	6.23	6.60	6.97	7.33	7.70
	21	3.15	3.50	3.85	4.20	4.55	4.90	5.25	5.60	5.95	6.30	6.65	7.00	7.35
	20	3.00	3.33	3.67	4.00	4.33	4.67	5.00	5.33	5.67	6.00	6.33	6.67	7.00
	19	2.85	3.17	3.48	3.80	4.12	4.43	4.75	5.07	5.38	5.70	6.02	6.33	6.65
	18	2.70	3.00	3.30	3.60	3.90	4.20	4.50	4.80	5.10	5.40	5.70	6.00	6.30

Figure 5 Average number of ships at HR as a function of the number of ships in fleet and the number of months in that readiness state per OPCYCLE

2.3 Determining likelihood for Supply to meet Demand

The next step in the methodology is to determine how well the supply can meet the demand defined by the PDF $Pr(X = x)$ shown in Figure 2. In the PDF, X is the discrete random variable on the number of ships at HR on any given day, and x represents all the possible values, i.e. 0, 1, 2, …, 15 in this case. For example, $Pr(X = 2) = 11.0\%$ meaning that on any given day in the OPCYCLE, there is an 11.0% chance of requiring 2 ships at HR.

The corresponding cumulative distribution function (CDF) of the random variable X, expressed as $F(X) = Pr(X \leq x)$, is shown in Figure 6. The objective is to use the CDF to determine how well a steady-state supply of ships at HR can meet the demand. For example, as discussed earlier, a fleet of 20 ships with each ship spending 12 months at HR in an OPCYCLE, provides an average of 4 ships at HR at any given time. In this case, we can expect that the demand would be met 63.1% of the time, which is calculated as follows:

$$\begin{aligned} F(4) = &\ Pr(X = 0) + Pr(X = 1) + Pr(X = 2) + Pr(X = 3) + Pr(X = 4) \\ = &\ 6.7\% + 11.0\% + 11.6\% + 17.1\% + 16.7\% = 63.1\% \end{aligned} \quad (1)$$

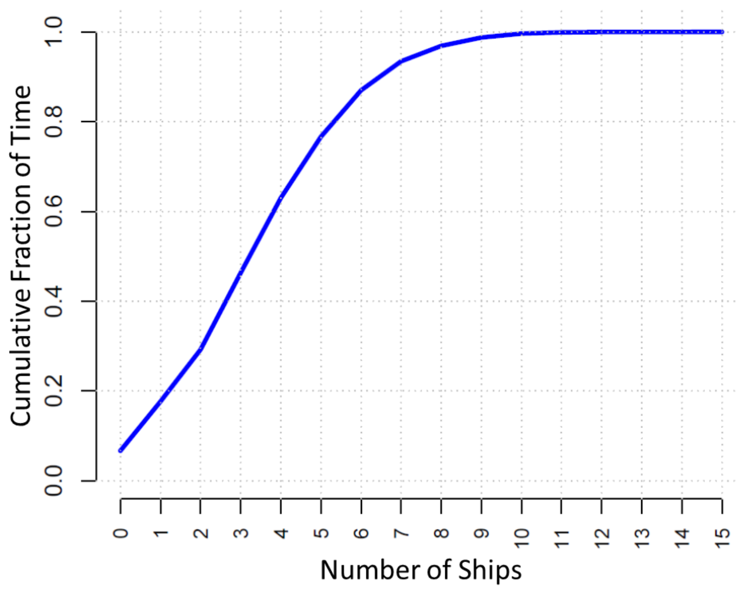

Figure 6 Cumulative Distribution Function for the Demand

Typically, the CDF of a discrete random variable is not a continuous function like the one in Figure 6. Rather, it is a discrete, piece-wise, function. However, as shown in Figure 5, the steady-state supply of ships is not always an integer value. For example, a fleet of 18 ships with each ship spending 15 months at HR leads to a supply of 4.5 ships at HR at any given time. This means that half of the time there are 4 ships at HR and the other half, 5 ships at HR. A linear interpolation between the x values of the original PDF (i.e., 0, 1, 2, …, 15) was done to make the CDF a continuous function (instead of piece-wise) in order to provide realistic probabilities for non-integer values of x. Section 3.0 will present the results of how well each of the 169 options for ship supply at HR enumerated in Figure 5 can meet the demand.

2.4 Assumptions and Limitations of the Methodology

This methodology provides a convenient way to consider operational supply and demand. This section lists the assumptions and limitations inherent with this approach.

- Steady-State ship supply: This assumption does not look in detail at ship scheduling, as opposed to previous efforts [2-10]. The result is that the model only assesses the long-term average ability of fleet to meet operational demand, and cannot predict individual scheduling conflicts that may arise. Over long time horizons, such as the lifetime of a fleet, one can expect the actual performance to approach the steady-state average behaviour predicted by this model.

- No Attrition: Following on from the first assumption, there is no stochastic attrition rate built into this model. The steady-state supply model assumes that all ships in the fleet will operate in a constant fashion over the entire time period being modeled.

- Coastal disposition: Canada is a tri-coastal nation, being able to operate in the Atlantic, Pacific and Artic Oceans. However, this analysis does not consider the coastal disposition of ships within the fleet. The effect of this is that sometimes, an event may require a longer transit time than foreseen in the vignettes or at worst the fleet is not able to meet the demand at all since in reality some events require all ships to be sent from the same coast.

- Other ship classes: Some missions require a mix of ship classes to operate as a single naval task group. This current model only looks at the requirement for and the availability of a single ship class within a vignette and does not consider any other ship classes. The result is that the reported performance can be interpreted as the expected values when the vignettes are *not* limited by other ship classes.

Taking into consideration these assumptions and limitations, the net effect is that the reported performance sets an upper bound to the operational output of a single ship class. The RCN can expect that the real performance of the fleet will be less than what is reported, depending on a host of real-world variables that are outside the scope of this simple model.

3.0 RESULTS AND DISCUSSION

The main result of the analysis using the illustrative data for vignettes and fleet configurations are presented in Figure 7. It combines the results of the estimated operational demand from Section 2.1.2 with the operational supply from Section 2.2, and shows the performance levels for 169 possible fleet configurations. It can be seen that there are many ways to achieve the same output, such as a fleet of 20 ships with ships deploying 15 months out of the OPCYCLE providing the same level of output as a fleet of 25 ships each deploying 12 months out of the OPCYCLE. Both achieve the average steady-state ship supply of five ships per day, which equates to meeting 76.6% of the expected operational demand; it is simply a matter of having more ships running ‘lightly’ or less ships running ‘hard’.

Number of Ships in Fleet	Months at HR / OPCYCLE												
	9	10	11	12	13	14	15	16	17	18	19	20	21
30	69.8%	76.6%	81.8%	87.0%	90.2%	93.5%	95.2%	96.9%	97.8%	98.8%	99.2%	99.6%	99.8%
29	67.8%	74.4%	79.9%	84.9%	88.8%	92.0%	94.3%	96.0%	97.3%	98.2%	98.9%	99.3%	99.7%
28	65.7%	72.1%	78.0%	82.9%	87.4%	90.4%	93.5%	95.1%	96.7%	97.6%	98.5%	99.1%	99.5%
27	63.7%	69.8%	76.0%	80.8%	85.4%	88.9%	91.8%	94.2%	95.7%	97.1%	97.9%	98.8%	99.2%
26	61.3%	67.6%	73.5%	78.7%	83.2%	87.4%	90.2%	93.0%	94.7%	96.2%	97.3%	98.1%	98.8%
25	58.8%	65.3%	71.0%	76.6%	81.0%	85.3%	88.6%	91.3%	93.8%	95.2%	96.6%	97.5%	98.3%
24	56.3%	63.0%	68.5%	73.9%	78.7%	82.9%	87.0%	89.6%	92.2%	94.2%	95.5%	96.9%	97.6%
23	53.8%	60.2%	66.0%	71.2%	76.4%	80.4%	84.4%	87.9%	90.3%	92.8%	94.4%	95.7%	97.0%
22	51.3%	57.4%	63.5%	68.5%	73.5%	78.0%	81.8%	85.6%	88.5%	90.9%	93.3%	94.6%	95.9%
21	48.8%	54.7%	60.5%	65.7%	70.5%	75.3%	79.2%	82.9%	86.5%	88.9%	91.2%	93.5%	94.7%
20	46.3%	51.9%	57.4%	63.0%	67.6%	72.1%	76.6%	80.1%	83.5%	87.0%	89.1%	91.3%	93.5%
19	43.8%	49.1%	54.4%	59.7%	64.6%	68.9%	73.2%	77.3%	80.6%	83.9%	87.1%	89.1%	91.2%
18	41.2%	46.3%	51.3%	56.3%	61.3%	65.7%	69.8%	73.9%	77.7%	80.8%	83.9%	87.0%	88.9%

Figure 7 Percent of Expected Demand met by Supply

It is worth noting that the four corners of the matrix in Figure 7 represent the extreme cases of all the possible options being considered. The bottom-left quadrant provides the weakest solutions in terms of meeting operational demand. A minimum performance level can be placed if a requirement to meet all scheduled vignettes is enforced, which would represent Canada keeping her standing naval commitments. In this example, a minimum of four ships are required to meet all scheduled vignettes, and therefore a minimum of 63% performance must be achieved. All cells in Figure 7 have been coloured red to represent this lower-bound.

Similarly, the top-right quadrant represents a very 'rich' solution, where the fleet size is maximized and each ship is deployed for the longest amount of time in the OPCYCLE. While meeting all or nearly all of the expected demand, solutions in this quadrant come at a cost, both fiscally and in terms of wear and tear to the ships.

The top-left and bottom-right corners represent extreme trade-off situations, minimizing the months deployed by adding more ships to the fleet, or minimizing fleet size by maximizing months of deployment respectively.

In the absence of additional information, this approach cannot narrow the option-space down to a single solution, it merely provides an aid that decision-makers can refer to when discussing options for fleet sizing.

3.1 Characterizing Multi-ship Demand

It was found that presenting actual examples of *x*-ship demand scenarios was extremely enlightening to RCN staff. Based on the data derived from the model, it was possible to show the most likely scenarios that required 2, 3, 4, … ships. Figure 8 below illustrates this using the notional vignette data used throughout this paper.

(A) Two-ship demand

Single ship Short deployment [1]	Single ship Medium deployment [1]	Single ship Long deployment [1]	Two ships Medium Deployment [2]	Task Group Large Scale Commitment [3]	Standing NATO Maritime Group 1 [1]	Flagship for Standing NATO Maritime Group 1 [1]	Spring / Fall scheduled operations [1]	Deployment w/ related exercises [2]	Percentage of time
			1						2.87%
	1				1				1.07%
		1			1				0.89%
	1	1							0.88%
1					1				0.84%
1	1								0.79%
1		1							0.67%
					1		1		0.64%
	1					1			0.58%
	2								0.55%

(B) Three-ship demand

Single ship Short deployment [1]	Single ship Medium deployment [1]	Single ship Long deployment [1]	Two ships Medium Deployment [2]	Task Group Large Scale Commitment [3]	Standing NATO Maritime Group 1 [1]	Flagship for Standing NATO Maritime Group 1 [1]	Spring / Fall scheduled operations [1]	Deployment w/ related exercises [2]	Percentage of time
					1			1	4.18%
						1		1	2.13%
				1					1.67%
			1		1				1.19%
	1		1						1.14%
		1	1						0.98%
1			1						0.82%
			1			1			0.58%
	1	1			1				0.34%
1	1				1				0.32%

(C) Four-ship demand

Single ship Short deployment [1]	Single ship Medium deployment [1]	Single ship Long deployment [1]	Two ships Medium Deployment [2]	Task Group Large Scale Commitment [3]	Standing NATO Maritime Group 1 [1]	Flagship for Standing NATO Maritime Group 1 [1]	Spring / Fall scheduled operations [1]	Deployment w/ related exercises [2]	Percentage of time
					1		1	1	1.68%
	1				1			1	1.61%
		1			1			1	1.36%
1					1			1	1.19%
						1	1	1	0.88%
	1					1		1	0.79%
		1				1		1	0.72%
	1			1					0.66%
				1	1				0.64%
1						1		1	0.62%

(D) Five-ship demand

Single ship Short deployment [1]	Single ship Medium deployment [1]	Single ship Long deployment [1]	Two ships Medium Deployment [2]	Task Group Large Scale Commitment [3]	Standing NATO Maritime Group 1 [1]	Flagship for Standing NATO Maritime Group 1 [1]	Spring / Fall scheduled operations [1]	Deployment w/ related exercises [2]	Percentage of time
			1		1			1	1.84%
			1			1		1	0.90%
			1	1					0.74%
	1				1		1	1	0.64%
		1			1		1	1	0.56%
1					1		1	1	0.52%
	1	1			1			1	0.51%
1	1				1			1	0.45%
1		1			1			1	0.42%
	1					1	1	1	0.34%

Figure 8 Top ten most likely vignette combinations requiring *x* ships

In the figure, the four sub-charts (A), (B), (C) and (D) show the top ten most likely combinations of vignettes that required 2, 3, 4 and 5 ships respectively. For each sub-chart, the nine vignettes are listed across the top with the five Random vignettes to the left of the double line and the four Scheduled vignettes to the right. The rightmost column shows the Percentage of Time that each combination of vignettes occurred over the 1000 runs in this illustration. The columns are colour-coded so that blue corresponds to vignettes requiring one ship, orange corresponds to two ships and purple to three ships. Each row in the table represents a specific combination of vignettes, which when totalled, adds up to the required number of ships. The numbered cells within the table represent the number of times a particular vignette occurs within a specific combination of *x* ship demand.

For two-ship demand (A), it can be seen that the most likely two-ship scenario is a random two-ship medium duration event. The second and third most likely two-ship combinations of vignettes are one Random single-ship plus one Scheduled single-ship. Totalling the Percentage of Time column for all two-ship combinations equals 11.6%, which recovers the results first presented in Figure 2.

For three-ship demand (B), the most likely combination of vignettes is a scheduled single-ship plus a scheduled two-ship vignette. For four-ship demand (C), the most likely is two scheduled single-ship vignettes plus a scheduled two-ship vignette. Finally for five-ship demand (D), the most likely is a scheduled single-ship vignette, a scheduled two-ship vignette plus a random two-ship vignette.

While simplistic in its approach, this allowed the RCN staff to gain an appreciation of what output they can expect to achieve if they only had *x* ships available at the highest readiness state at any given time. In effect, they were able to gain an intuitive sense of what they were "buying" in terms of operational output for a given combination of ships in fleet running at a particular operational/maintenance profile.

4.0 CONCLUSION AND FUTURE WORK

This paper presented a method for determining the operational demand and assessing the ability of a given configuration fleet to meet this demand under the assumption of a steady-state ship supply model. The operational demand was estimated using a vignette-based stochastically-driven model, which has strong basis in the Canadian Force Development process as well as foundational work provided by previous DRDC studies. The ship supply was modeled to allow for quick analysis of the effectiveness of a fleet, using only two main variables to drive the model: the number of ships in fleet and the number of months at a specific readiness state. This resulted in presenting an option-space to the senior RCN leadership that was straight-forward and transparent.

Future work in this area naturally extends to enlarging the scope of the model to include multiple ship classes, requiring further definition and refinement of the demand vignettes, as well as modifying the steady-state ship supply model to accommodate the concept of task groups, sending mixed classes of ships to answer an event.

5.0 REFERENCES

[1] P. Massel, R. Burton, D. Mason and D. a. J. McLean, "The Canadian maritime forces 2015 study phase II: analysis of maritime force structure alternatives using the FLEETSIM model," Department of National Defence - Operational Research Division, 1999.

[2] D. Allen, D. Blakeney, R. Burton and L. D. Purcell, "Fleet mix study: determining the capacity and capability of the future naval force structure," Defence R&D Canada, 2005.

[3] D. Allen, C. Eisler and A. Forget, "A user guide to Tyche version 2.0: providing a joint flavour to Tyche," Defence R&D Canada, 2006.

[4] C. Eisler, "Fleet mix study: preliminary reporting from iteration II," Defence R&D Canada, 2007.

[5] C. Eisler, "Improving the ability to model future maritime force structures: from Tyche version 1.0 to 2.1," Defence R&D Canada, 2007.

[6] C. Eisler, "Fleet Mix Study: Linking the Maritime Vignettes to the 2006 DND/CF Force Development Scenarios," Defence R&D Canada, 2008.

[7] A. Bourque and C. Eisler, "Fleet Mix Study: Establishing the historical and policy basis for the maritime vignettes," Defence R&D Canada, 2008.

[8] A. Bourque and C. Eisler, "Fleet Mix Study: Measures of Performance and Sensitivity Analysis," Defence R&D Canada, 2008.

[9] A. Bourque, C. Eisler and L. W. Reive, "Fleet Mix Study: Capability Supply and Demand Requirements for Iteration II," Defence R&D Canada, 2009.

[10] A. Bourque and C. Eisler, "Fleet Mix Study Iteration II: Making the Case for the Capacity of the "Navy after Next"," Defence R&D Canada, 2010.

[11] A. Bourque, "Canadian Surface Combatant Variant Mix: Modeling the Quantitative Assessment using Tyche," Defence R&D Canada, 2012.

[12] A. Bourque and R. Mirshak, "Effect of platform numbers on the projected capacity of the Canadian Surface Combatant fleet," Defence R&D Canada, 2015.

[13] A. Bourque and R. Mirshak, "Impact of Canadian Surface Combatant Fleet Size on Delivering Naval Ambitions," Defence R&D Canada, 2015.

[14] Royal Canadian Navy, "CFCD 129 Readiness and Sustainment".

[15] Chief of Force Development, "Capability Based Planning Handbook," Department of National Defence, 2014.

[16] NATO, "NATO Standing Maritime Groups," [Online]. Available: http://www.mc.nato.int/org/smg/Pages/default.aspx. [Accessed 27 09 2016].

[17] Department of National Defence, "ARCHIVED - Operation HESTIA," 12 01 2015. [Online]. Available: http://www.forces.gc.ca/en/operations-abroad-past/op-hestia.page. [Accessed 27 09 2016].

[18] Department of National Defence, "Rim of the Pacific," 22 08 2016. [Online]. Available: http://www.forces.gc.ca/en/operations-exercises/rimpac.page. [Accessed 27 09 2016].

[19] The R Foundation, "R: The R Project for Statistcal Computing," [Online]. Available: https://www.r-project.org/. [Accessed 27 9 2016].