Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

MODELING OPERATIONAL DEMAND FOR CANADA'S FUTURE NAVAL FLEET: A CASE STUDY ON MAINTAINING EXPECTED FREQUENCIES OF MILITARY VIGNETTES

Lynne Serré[1] and Lise Arseneau[1]

[1]Defence Research and Development Canada Centre for Operational Research and Analysis, Department of National Defence, Ottawa, ON, CANADA

ABSTRACT

The Royal Canadian Navy is currently undertaking a fleet mix study to determine the optimal composition of its fleet to meet future operational requirements. The operational demand is modeled by generating possible future timelines of vignettes that can occur concurrently using a Monte Carlo discrete event simulation model. In this paper, we examine two aspects related to modeling operational demand stochastically: the number of replications required to ensure that the input frequencies of the vignettes correspond to a certain level of accuracy in the output frequencies from the simulation; and how the application of concurrency constraints to individual vignettes may lower the output frequencies if events are scheduled purely at random. We propose a modified scheduling algorithm that better maintains the input frequencies. Both aspects need to be taken into consideration to ensure that the modeled operational demands are representative of possible and probable futures for the Navy.

1 INTRODUCTION

The Royal Canadian Navy (RCN) is currently undergoing the largest recapitalization of its fleet since the Second World War. To determine the optimal composition of its fleet to meet the future operational requirements, the RCN is undertaking a fleet mix study. Fleet mix studies are often described as a supply and demand problem where, for the RCN, the supply consists of the number and type of platforms being considered in the fleet and the demand consists of several tasks and/or scenarios where the RCN will be expected to provide a response. Military fleet mix problems can be extremely difficult to solve due, in part, to a set of potentially conflicting objectives, such as performance, deployability, availability, cost, and risk (Baykasoğlu et al. 2019).

The determination of operational requirements (i.e., the demand) is a persistent challenge for defense policy makers and planners due in large part to the uncertainty of what missions will require military assets in the future. Furthermore, the military can be involved in a wide range of tasks, including combat, patrol, search and rescue, and surveillance, each requiring the use of a variety of assets to provide a response (Wojtaszek and Wesolkowski 2012). These tasks must also be considered under an ever-evolving security environment.

Several military fleet mix studies have modeled operational demand as being either deterministic or stochastic in nature (Fee and Caron 2021). For the deterministic case, operational demand can be expressed as a fixed demand based on the level of ambition or by examining specific scenarios. For example, the effective mix of the United States (US) destroyer fleet was assessed based on a particular planning scenario for a conflict occurring on the Korean Peninsula (Crary et al. 2002). The scenario was defined by four distinct phases, each with five primary missions for surface combatants. The problem was to then determine the appropriate mix of Aegis cruisers (CG-52), Aegis guided military destroyers (DDG-51) and the new destroyers (DD-21). For the US army, military force planning prior to September 11th was based upon military forces being sized to win two major regional conflicts nearly simultaneously, usually with scenarios describing wars against Iraq and North Korea. Although this approach focused on clear, concrete, and specific requirements to determine force size, it failed to consider the inherent uncertainty of international relations and potential threats. Furthermore, it did not explicitly include many small-scale US force commitments, which tend to decrease military readiness and can slow deployment times (Lane et al. 2022).

Stochastic modeling of operational demand, on the other hand, can support military planners in imagining and creating a wide range of possible, probable and plausible futures. In this paper, we will focus on the use of stochastic simulation to determine demand where RCN operational requirements are represented by possible future timelines of vignettes (or scenarios), which can occur concurrently. The list of hypothetical vignettes represents the full scale of activities that would require the use of a naval platform, where each vignette can be characterized by type (e.g., peacetime or wartime), frequency, and duration and are all distinct from one another. The description and characteristics of a scenario or vignette are typically derived from a combination of historical analyses and professional military judgement (Dobias 2019a).

In this paper, we will describe two aspects related to modeling operational demand using a Monte Carlo discrete event simulation model. First, in the generation of timelines of vignette combinations, we will examine the number of replications required so that the input frequencies of each vignette correspond to a certain level of accuracy in the output demand. Second, we will demonstrate how applying a concurrency constraint to individual vignettes in a timeline may result in unexpected frequencies observed in the output. We then propose a modified scheduling algorithm that better maintains the input frequencies.

2 BACKGROUND

2.1 Previous Military Studies

There have been several studies that have estimated operational demand using a stochastic approach. A Monte Carlo discrete event simulation model known as Tyche was developed for the RCN's fleet mix structure analysis (Eisler and Allen 2012). Within Tyche, the demand is constructed stochastically from scenarios using frequency, start date, and duration inputs. The scenarios can be randomly generated using a Poisson process or scheduled at known intervals. Then a force structure is selected to test a capability supply from the set of assets against the demand. Tyche is computationally intensive where a single simulation run can take hours to complete (Eisler et al. 2014). In (Caron et al. 2019), a two-fold approach was developed to determine the optimal number and types of platform modules required by the RCN to meet ambitions and mandate, where a Monte Carlo discrete event simulation is used to generate the operational demand and a mixed-integer linear programming (MILP) model is used to determine the optimal mix of modules. The RCN demand was generated from 54 vignettes representing both scheduled and random events.

Another military application where vignette combinations were used to estimate the demand is in exploring ammunition stockpiles (Caron et al. 2023). The vignettes describe activities from several types of training and military missions that require ammunition. For each vignette, an annual frequency is set, a date range for the initial start, and a time range over which demand can occur which allows a demand to apply for part of a simulated time frame. The Stochastic Fleet Evaluation (SaFE) model is a Monte Carlo-based simulation to determine fleet sizes based on task frequency, duration and assignments for multiple platform types (Wesolkowski and Billyard 2008). The tool was applied to a set of 127 air mobility tasks using five different platforms over a one-year period. The work was extended with the development of the Stochastic Fleet Estimation - Robust (SaFER) model (Wesolkowski and Wojtaszek 2012) to find the minimal cost fleet by searching for task start times that result in minimal cost of platforms required to accomplish all tasks. As in the initial study, the modeling of air mobility requirements is done by stochastically generating many possible future scenarios (i.e., task requirements over a one-year period).

While fleet mix problems generally involve platforms (e.g., vehicles, ships, or aircraft), a similar problem arises in terms of personnel (e.g., infantry, sailors, or pilots), known as force mix problems. In a Canadian study of its armed forces personnel, operational demand was estimated using a set of 17 scenarios covering a full range of missions mandated by Canadian defense policy, with approximately 80 variants developed specifically for the study (Dobias et al. 2019b). A Monte Carlo-based simulation tool was developed that modeled scenario combinations over a predefined time period (e.g., five years) using the frequency and duration inputs, with imposed constraints on concurrency (e.g., a limited number of scenarios with certain characteristics can occur at the same time). The tool is a scheduling program that is similar to the Tyche simulation but requires much simpler inputs, thus taking less computation time. While the approach in modeling operational demand is similar to the above-mentioned fleet mix studies, it differs in that it imposed a constraint of not having more than seven lines of operation occurring concurrently (Dobias et al. 2019b). Constraints like this may reflect policy or other requirements imposing a maximum number of expected commitments at one time. However, the study's approach to generating timelines extends a Poisson point process to a finite duration of length equal to each operation's duration, which essentially results in an expansion of the time axis (Bryce 2024). This extension causes the observed frequencies in the discrete event simulation results to be lower than the set of input frequencies, which may artificially reduce the expected demand. Whether this is appropriate is dependent on the interpretation of the input frequencies.

The premise of this paper is that the input frequencies represent the expected number of events over a given time period to which the RCN would need to have the capacity to provide a response. Therefore, any conclusions drawn regarding the mix of naval platforms and its capacity is viewed by decision makers as being representative of a demand that respects these input frequencies. In other words, the input and output frequencies are assumed to align. In this paper, we consider a simplified concurrency constraint applied to each individual vignette rather than across vignettes to understand its impact on the output frequencies.

2.2 Simulation Output Analysis

To obtain accurate measures of model performance from a Monte Carlo discrete event simulation, decisions on three key areas are required: warm-up, run length and number of replications (Hoad et al. 2007; Law 2010). Since the RCN operational demand is being generated by a Monte Carlo simulation, one question that arose was how many replications are needed. The limiting factors on how many replications to run are computing time and expense. If performing N replications achieves satisfactory estimates of performance as required by the user, performing more than N replications may be unnecessary but performing fewer than N replications could lead to inaccurate results.

For the ammunition stockpile study (Caron et al. 2023), a total of 1,000 replications with timelines of 10-year duration were executed for each variation consisting of a specific combination of an order frequency and an annualized quantity for an ammunition type. The selection of the number of replications is not discussed, however, the computational speed to run 1,000 replications was a few minutes for a single type of ammunition. In (Dobias et al. 2019b), there are 10,000 possible futures representing trials of a five-year period with two-month sampling frequency. The paper does not provide the reason for selecting the number of replications nor the computing time. The RCN operational demand in (Caron et al. 2019) was based on generating a total of 2,500 replications of five-year timelines, with the results being aggregated to obtain unique combinations of daily demand. The selection of 2,500 was based on a related study where a good balance was obtained between the necessary accuracy (based on using the Law of Large Numbers (Law and Kelton 2000)) and the required computation time that took less than 30 minutes.

In this paper, we examine the number of replications needed to ensure the output (observed) frequencies are consistent with the input (expected) frequencies. An insufficient number of replications has the potential to bias the observed frequencies before any concurrency constraints are applied, which could result in a compounding effect.

3 SIMULATION MODEL

To address fleet mix questions from the RCN, Defence Research and Development Canada (DRDC)'s Centre for Operational Research and Analysis has developed a fleet capacity evaluation tool, which is a Monte Carlo discrete event simulation model. This tool builds on previous work, most notably Tyche (Eisler and Allen 2012) and the Platform Capacity Tool (Fee and Caron 2021). The latest version was implemented in a Python-based discrete event simulation graphical coding interface (DRDC 2023). An overview of an earlier version of the model, implemented in the Arena software, can be found in (Fee and Caron 2021). Here, we focus on how the model simulates operational demand. Like its predecessors, timelines of future events are stochastically generated based on a set of vignettes characterised by their frequency and duration.

3.1 Input Parameters

The model requires a list of vignettes to generate operational demand. For each vignette, an analyst provides inputs for the following parameters.

- Frequency: the average number of events per year.
- Duration: the minimum, maximum, and most likely number of days.
- Date range for event: earliest and latest day in a year where a vignette can begin.
- Concurrency: maximum number of events of the same vignette that can occur simultaneously (optional).

The frequency and duration input parameters may be estimated from historical data or may be adjusted by subject matter experts to represent expectations for the future. It is possible that some vignettes may only occur in certain months, or seasons, each year. As an example, some search and rescue activities are concentrated in the peak recreational sailing months. Providing a date range for the earliest and latest day in a year where a vignette can begin allows the analyst to impose these types of restrictions. Further, in some cases, multiple events of the same vignette may not occur simultaneously, which can be controlled by the concurrency input parameter. As an example, public engagement events would not be scheduled simultaneously. This vignette would have a maximum concurrency of one.

3.2 Algorithm

An overview of the timeline generation algorithm is shown in Figure 1. Each replication generates a single timeline of a fixed duration, which will be set to five years in the case study. For simplicity, we assume a year has 30 days per month (i.e., 360 days per year). Each event in a timeline is scheduled by first determining its duration. This is done by sampling from a triangular distribution using the minimum, maximum, and mostly likely number of days. Next, a start date is selected uniformly at random from the days of the year (i.e., from 1 to 360 if there are no restrictions on the start date). A year is then also selected uniformly at random. If the randomly determined start results in the event ending after the five-year timeline, then the remaining portion is assigned to the beginning of the timeline as done in (Caron et al. 2019). This was done so that no start-up period is needed in the simulation.

```
for Each replication do
  for Each vignette do
    Draw number of events for the vignette from a Poisson distribution
    for Each event do
      Draw event duration from a triangular distribution
      Draw event start date within a year from a uniform distribution
      Draw event year from a uniform distribution
      if Number of events is not greater than the allowed concurrency do
        Schedule the event in the timeline
        if Event finishes after the end of the timeline do
          Assign remaining portion of the event to the beginning of the timeline
```

Figure 1: Pseudocode of the timeline generation algorithm.

If scheduling the current event would exceed the maximum concurrency level for its corresponding vignette, then it is skipped and not scheduled in the timeline. In this paper, we propose a second scheduling approach that attempts to minimize the number of skipped events by restricting possible start dates based on what has already been scheduled. This process is illustrated in Figure 2. In Row A, the first event (in green) has been scheduled from days 65 to 70. In Row B, the second event (in orange) has a duration of 5 days. Therefore, days 61 to 70 are removed from the possible start dates, and a start date is selected uniformly at random from the remaining days. Since this scheduling approach is attempting to minimize the number of skipped events, we will refer to it as the optimized scheduler, and the original scheduler as the random scheduler.

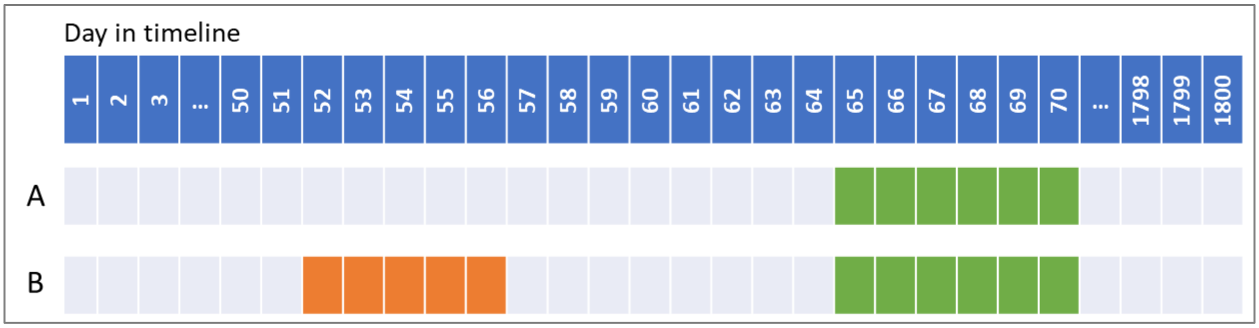

Figure 2: The green event in Row A restricts the possible start dates for the orange event added in Row B.

3.3 Output Analysis

As a Monte Carlo discrete event simulation model, many replications of the model are run. The case study focuses on two key metrics that summarize the timelines of operational demand.

1. The output (or observed) frequency of each vignette.
2. The most frequent combinations of vignettes and their average duration.

The first of the metrics above should match the input (or expected) frequency of each vignette within a prescribed level of accuracy, which is the premise for our case study. That is, the input frequencies represent the expected number of events, on average, for each vignette that the Navy would anticipate being asked to respond to each year. Therefore, it is important that the observed frequencies align with the input frequencies as much as possible.

The second of the metrics above is important for understanding concurrent operational demand. It is straightforward to assess whether the demand of a single vignette can be met by the supply of platforms; however, it becomes less straightforward to assess which combinations of vignettes can be met by the supply of platforms. As an example, with 56 vignettes, there are nearly 30,000 different possible combinations of two or three distinct vignettes.

4 CASE STUDY

To imagine the future operational demand of the RCN, a set of vignettes must be created that captures the full range of tasks. This may include peacetime tasks, such as public engagements or assistance in natural disasters; wartime tasks, such as support to a multinational coalition; discretionary tasks, such as assistance to international event security; and non-discretionary tasks, such as search and rescue. Here, we use a set of 56 anonymized vignettes that have a range of frequencies and durations. A subset of the vignettes is presented in Table 1. The type of vignette is provided as a high-level descriptor and is used in the analysis to visualize results but is not used in the generation of timelines. No restrictions are imposed on the start date for the vignettes and, initially, no concurrency limits on the vignettes are imposed.

Table 1: Subset of vignettes and their frequency and duration parameters.

Vignette ID	Vignette type	Annual frequency	Duration (in days)		
			Minimum	Maximum	Mode
1	Peacetime	5	7	90	25
9	Non-discretionary	8.8	1	3	2
18	Peacetime	0.2	3	30	17
45	Wartime	0.5	18	1,000	250
51	Discretionary	0.03	30	180	105

4.1 Replication Analysis

Since the premise of the study is that the input and output frequencies should match (as much as possible) for each vignette, we first analyze the number of replications required. Within the simulation, the number of events for each vignette is generated from a Poisson distribution with a mean, λ, that is equal to the annual input frequency times the number of years in the timeline. Therefore, the analysis of the input and output frequencies was done based on the expected and observed frequencies over five years.

Figure 3 shows the observed and expected frequencies over five years for three subsets of vignettes that each share the same expected frequency of either 0.15, 0.20, or 0.25. When the number of replications is smaller, there is enough variability in the observed frequencies that vignettes with the same expected frequency have different observed frequencies. Furthermore, the observed frequencies are the same for some vignettes where the expected frequencies are different. For instance, with 1,000 replications, Vignette 14 has the same observed frequency as Vignette 37 of 0.24; however, Vignette 14 is expected to occur 0.20 times over five years while Vignette 37 is expected to occur 0.25 times. This highlights the importance of having sufficient replications in our context. However, it also raises a question of desired accuracy. In Figure 3, the observed frequencies were rounded to two decimal places based on the assumed precision in the input frequencies.

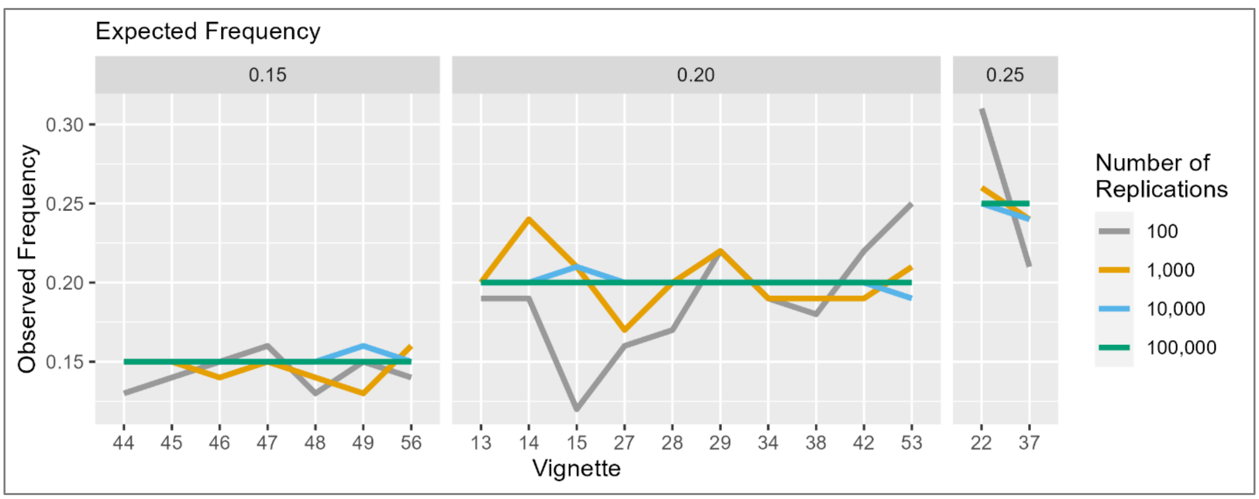

Figure 3: Comparison of the observed and expected frequencies for selected vignettes in a five-year time period.

The desired accuracy level in the output frequencies must be balanced against computational time. Modeling operational demand is only the first part of the fleet capacity evaluation tool. The remainder of the simulation model focuses on the supply side of the fleet mix problem by attempting to assign platforms to the generated timelines of operational demand. This task has several complexities, including accounting for maintenance schedules (i.e., platform availability), the location and priority of events (which may lead to platform relocation or reassignment), and choosing from multiple response options for an event. Therefore, high computational time for demand generation has the potential to increase the overall runtime of the simulation to an unreasonable amount of time.

To determine the number of replications, suppose we want a $100(1-\alpha)\%$ confidence interval for λ to be no wider than that given by $\bar{\lambda} \pm \varepsilon$. Then, as shown in (Hogg and Tanis 2006), the sample size n is given by

$$n = \frac{z_{\alpha/2}^2 \sigma^2}{\varepsilon^2}$$

Since the number of events is drawn from a Poisson distribution, the variance σ^2 is known (and equal to λ), which is why the z distribution rather than the t distribution can be used in the sample size calculation. It remains to determine the desired accuracy level, ε, which may vary depending on the magnitude of the input frequency for a vignette. Here, to balance computational time and desired accuracy, we allow for some deviation in the output frequencies that are small enough to not alter which vignettes are more frequent than others. Among the 56 vignettes, there are 26 unique expected frequencies. Table 2 shows the identified accuracy level and resulting sample size for the unique set of expected frequencies. As an example, the vignettes in Figure 3 have an accuracy level of 0.01. This allows for some variability but ensures that vignettes with a frequency of 0.15 occur less often than vignettes with a frequency of 0.20. Allowing some deviation also signals that the input frequencies are not exact numbers, which is important to communicate and discuss with stakeholders. As shown in Table 2, the largest sample size to maintain the desired accuracy level is 28,812. Therefore, in the remainder of the paper, results are based on 30,000 replications.

Table 2: Sample sizes (n) for the unique set of five-year expected frequencies (λ) based on the desired accuracy level (ε).

λ	**ε**	**n**	**λ**	**ε**	**n**
0.005	0.001	19,208	1.80	0.05	2,766
0.05	0.01	1,921	2.50	0.05	3,842
0.15	0.01	5,762	2.85	0.05	4,379
0.20	0.01	7,683	3.35	0.05	5,148
0.25	0.01	9,604	3.75	0.05	5,762
0.35	0.01	13,446	5.00	0.05	7,683
0.55	0.01	21,129	6.55	0.05	10,065
0.60	0.01	23,050	12.3	0.1	4,725
0.70	0.01	26,891	15.0	0.1	5,762
0.75	0.01	28,812	15.5	0.1	5,954
1.00	0.05	1,537	25	0.5	384
1.45	0.05	2,228	44	0.5	676
1.60	0.05	2,459	120	0.5	1,844

4.2 No Concurrency Constraints

Figure 4 shows an example of a timeline generated in a single replication. In this case, no concurrency constraints have been imposed on the vignettes. That is, each vignette can occur with itself as many times as by random chance. Of the 56 vignettes, 28 appear in this timeline. Most vignettes only have one event occurring at a time, but some have up to three concurrent events.

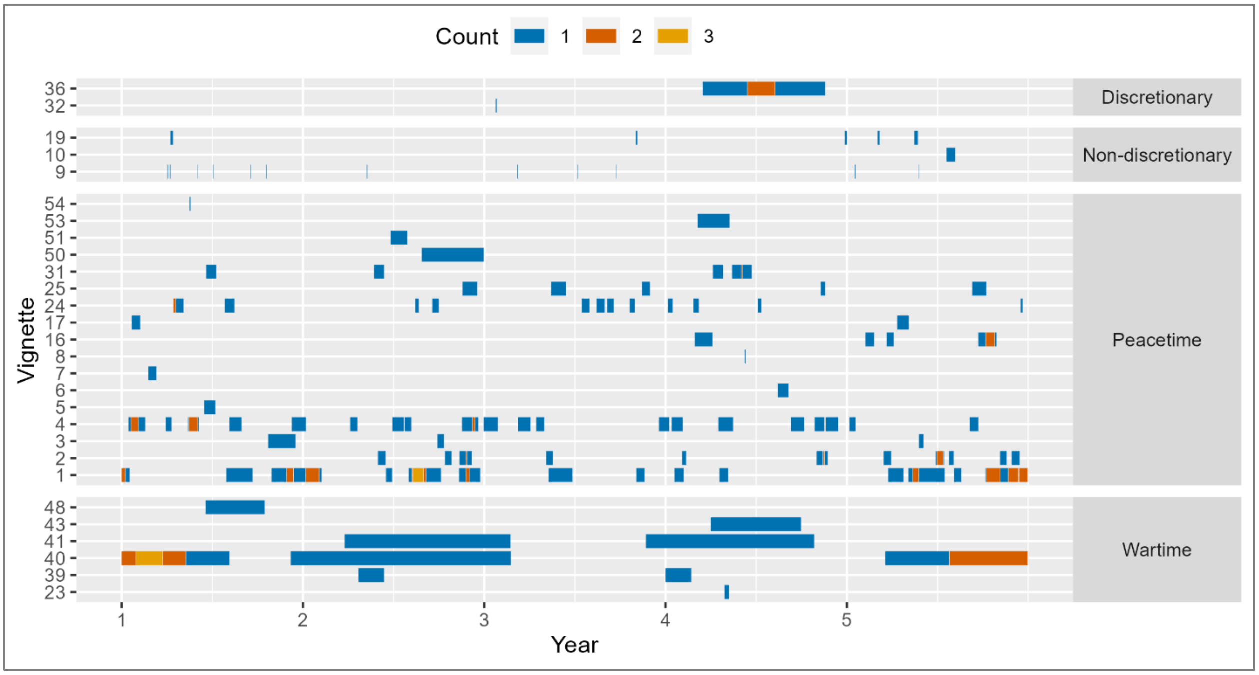

Figure 4: A single timeline generated by the simulation.

Figure 5 shows the 50 most frequent combination of vignettes by type under no concurrency constraints. The most frequent combination had one peacetime and one wartime vignette, and, on average, occurred for about 162 days (just over five months) over five years. There are also days with no events scheduled, this situation is in eighth place, occurring, on average, for about 75 days (two and a half months) over five years. Among the top 50 combinations, the maximum number of events at one time was seven.

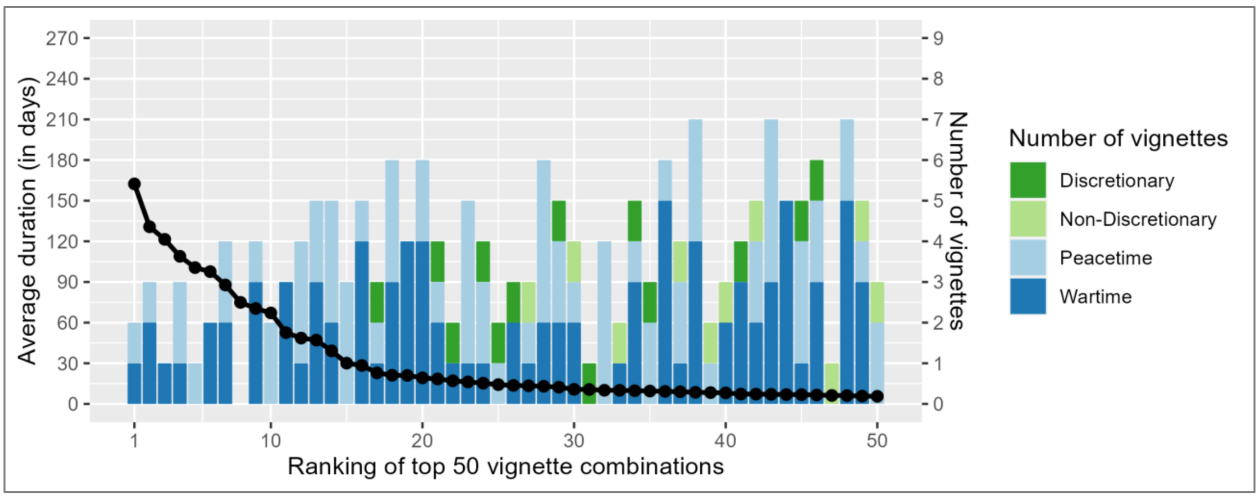

Figure 5: Most frequent combinations of vignettes by type under no concurrency constraints.

4.3 Concurrency Constraints with the Random and Optimized Schedulers

To illustrate the potential impact of concurrency constraints, we imposed a maximum concurrency of one for each vignette. That is, two events of the same vignette cannot occur simultaneously on any day (i.e., they cannot overlap). The probability that vignettes are randomly scheduled to overlap is a function of both the expected frequency and duration of the vignette. Therefore, the impact of this constraint will be greater on some vignettes than others.

Under the random scheduler, out of the 56 vignettes, 54 of them had skipped events. In contrast, under the optimized scheduler, only 7 of them had skipped events. The ten vignettes with the largest percentage of skipped events under the random scheduler are shown in Table 3. Vignette 41 had 42.8% of its sampled events skipped under the random scheduler, which resulted in an observed frequency of 1.6 compared to an expected frequency of 2.9. Under the optimized scheduler, the same vignette had only 14.5% of its sampled events skipped, which resulted in an observed frequency of 2.4. While an improvement, the optimized scheduler still had skipped events due to the relatively long duration of Vignette 41, which has a mode duration of 250 days. In contrast, Vignette 1 has a mode duration of 25 days. So, while it has a much higher expected frequency of 25 (compared to 2.9 for Vignette 41), the optimized scheduler skipped just 0.7% of the sampled events compared to 38.9% for the random scheduler. Under the random scheduler, Vignette 1 had an observed frequency of 15.3 while under the optimized scheduler it had an observed frequency of 24.9, which is much closer to the expected frequency of 25.0. Vignettes with high frequencies but short durations and vignettes with low frequencies but long durations tended to be most affected by the type of scheduler.

Table 3: Comparison of observed frequencies for selected vignettes under the random and optimized schedulers.

Vignette	Expected Frequency	Random Scheduler		Optimized Scheduler	
		% Skipped Events	Observed Frequency	% Skipped Events	Observed Frequency
41	2.9	42.8%	1.6	14.5%	2.4
40	2.5	40.7%	1.5	13.2%	2.2
1	25.0	38.9%	15.3	0.7%	24.9
43	1.8	16.1%	1.5	0.0%	1.8
4	15.5	14.7%	13.2	0.0%	15.5
36	1.5	11.8%	1.3	0.0%	1.4
2	15.0	10.4%	13.4	0.0%	15.0
24	12.3	7.7%	11.3	0.0%	12.3
16	6.6	7.2%	6.1	0.0%	6.5
50	1.0	5.7%	0.9	0.0%	1.0

Figure 6 and Figure 7 show the 50 most frequent combination of vignettes by type under the concurrency constraint for the random and optimized schedulers, respectively. In both cases, the most frequent combination remains unchanged from the unconstrained case (Figure 5). However, the average duration of this combination has increased to about 251 days (over eight months) under the random scheduler and to about 238 days (just under eight months) under the optimized scheduler. Under the random scheduler, there are more days with no events compared to the unconstrained case – this situation now ranks fourth compared to eighth because the random scheduler skips events when the concurrency constraint is met. Conversely, under the optimized scheduler, there are fewer days with no events compared to the unconstrained case – this situation is now in eleventh place because the optimized scheduler attempts to find available space in the timeline, thus minimizing the number of skipped events. This, however, changes the distribution of events in the timeline. Among the top 50 combinations, under both the random and optimized schedulers, the maximum number of events in a day has dropped to six from seven in the unconstrained case. Under the random scheduler, this is due to events being skipped while, under the optimized scheduler, this is due to events being scheduled to minimize overlap.

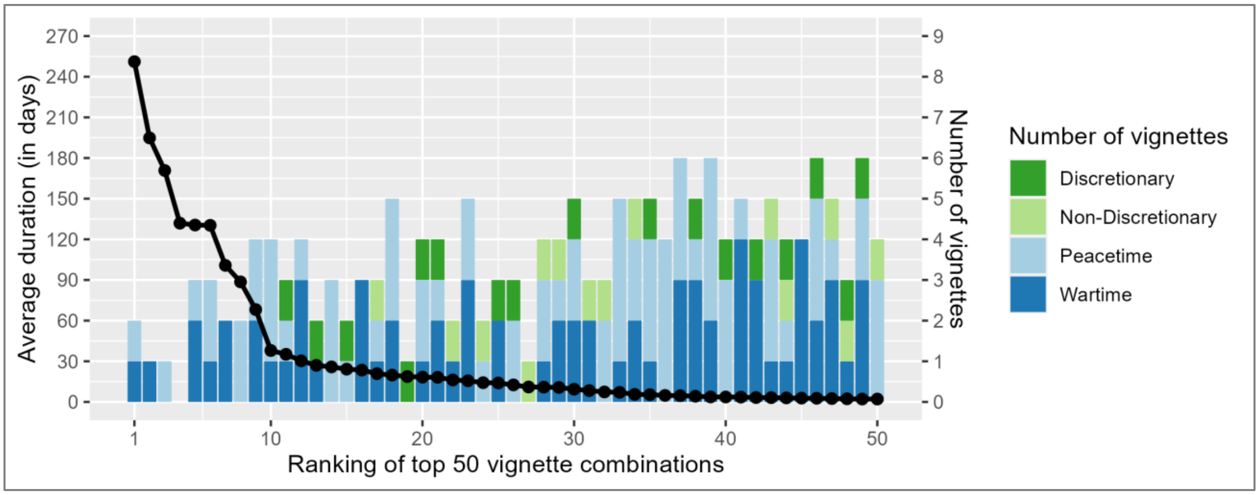

Figure 6: Most frequent combinations of vignettes by type under a maximum concurrency of one and the random scheduler.

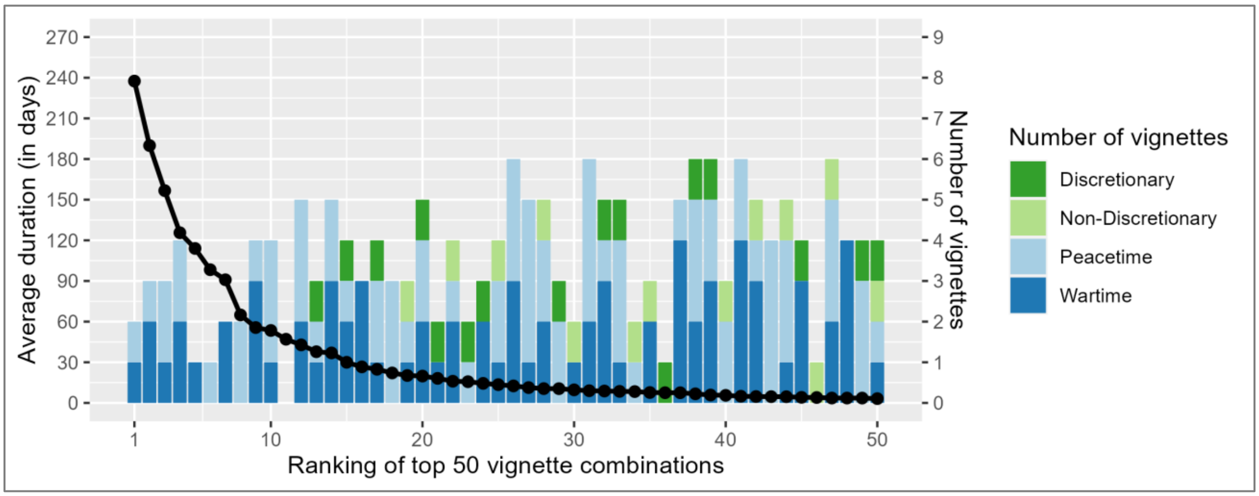

Figure 7: Most frequent combinations of vignettes by type under a maximum concurrency of one and the optimized scheduler.

5 CONCLUSION

For various military applications, there is a requirement to model operational demand representing future operational requirements to assess how the supply, which can be the number and type of platforms in a fleet mix study or the number of personnel in a force mix problem, meets the demand. When operational demand is modeled stochastically, it allows military planners to imagine and create a wide range of possible futures. For the RCN, the operational demand is generated by using a list of hypothetical vignettes representing the full scale of activities requiring a naval platform. Each vignette is characterized by its frequency, duration, date range and maximum concurrency, where these data are obtained from a combination of historical data and subject matter experts. Therefore, the input frequencies provided should match closely with the output (or observed) frequencies of each vignette. In this paper, we examined the number of replications required to ensure that the output frequencies maintain a certain level of accuracy. The resulting number of replications in the case study was 30,000, which is higher than that used in other related studies. However, the output metrics in these studies were different. Unlike previous related studies, this study focuses on understanding the modeled demand signal itself. It may be possible to reduce the number of replications, but this should be done as part of a discussion with stakeholders especially if it would result in the observed frequencies deviating significantly from the expected frequencies.

The case study also examined the impact of concurrency constraints on the modeled demand signal. We applied a simple concurrency constraint to each vignette so that no vignette occurs simultaneously with itself. Under a random scheduler, this constraint was shown to lower the observed frequencies, in some cases by as much as 45%. It was then shown that by attempting to minimize the number of skipped events using an optimized scheduler that the observed frequencies were much closer to the expected frequencies. However, in one case, the observed frequency was still lower than the expected frequency by about 18%. This highlights that depending on the frequency and duration of a vignette, it may not be possible to accommodate a concurrency constraint without impacting the observed frequencies. In these cases, discussions with stakeholders are important to resolve conflicting objectives. The impact on the demand signal of applying a concurrency constraint was also seen in the 50 most frequent combinations of vignettes, which further differed depending on the type of scheduler. In particular, the amount of time with no scheduled vignettes shifted from eighth to fourth place under the random scheduler and from eighth to eleventh place under the optimized scheduler.

While the optimized scheduler better maintained the observed frequencies compared to the input frequencies, it did increase the run time for 30,000 replications from the order of minutes to hours. Future work could develop a more efficient algorithm that attempts to schedule vignettes in the timeline subject to the concurrency constraint. Furthermore, the algorithm developed could incorporate other concurrency constraints such as having a maximum number of vignettes per day or not allowing certain vignettes to occur concurrently. This would also better represent what the RCN may be expected to do at any one point in time, but it would also be important to understand the impact on the modeled demand signal.

REFERENCES

Baykasoğlu, A., K. Subulan, A.S.Taşan, and N. Dudaklı, N. 2019. "A review of fleet planning problems in single and multimodal transportation systems". *Transportmetrica A: Transport Science* 15(2):631–697.

Bryce, R.M. 2024. "Finite interval processes: Simulating military operations". In *Proceedings of the 13th International Conference on Operations Research and Enterprise Systems (ICORES)*, February 24-26, Rome, Italy.

Caron, J.-D., V. Fong, and V. Brion. 2019. "On the Use of Simulation and Optimization for Mission Modules Selection in a Maritime Context". *Military Operations Research Journal* 24(1):41–56.

Caron, J.-D., R.M. Bryce, and C. Young. 2023. "A Simulation Tool for Exploring Ammunition Stockpile Dynamics". In *Proceedings of the 12th International Conference on Operations Research and Enterprise Systems (ICORES)*, February 19-21, Lisbon, Portugal, Volume 1:38-49.

Crary, M., L.K. Nozick, and L.R. Whitaker. 2002. "Sizing the US destroyer fleet". *European Journal of Operational Research* 136(2002):680-695.

Defence Research and Development Canada. 2023. ORIGAME. https://github.com/DND-DRDC-RDDC/OS_ORIGAME, accessed March 28, 2024.

Dobias, P. 2019a. "Estimation of frequency and duration of future Canadian Armed Forces Operations". In *Proceedings from the 36th International Symposium on Military Operational Research*, July 23-26, London, United Kingdom.

Dobias, P., Maj D. Hotte, J. Kampman, and B. Laferriere. 2019b. "Modeling Future Force Demand: Force Mix Structure Design". In *Proceedings from the 36th International Symposium on Military Operational Research*, July 23-26, London, United Kingdom.

Eisler, C. and D. Allen. 2012. "A Strategic Simulation Tool for Capability-Based Joint Force Structure Analysis". In *Proceedings of the 1st International Conference on Operations Research and Enterprise Systems*, February 4-6, Vilamoura, Portugal, 21-30.

Eisler, C., S. Wesolkowski, and D. Wojtaszek. 2014. "The Tyche and SaFE Models: Comparting Two Military Force Structure Analyis Simulations". In *Proceedings 28th European Conference on Modelling and Simulation*, May 27-30, Brescia, Italy.

Fee, M., and J.-D. Caron. 2021. "A simulation model to evaluate naval force fleet mix". In *2021 Winter Simulation Conference (WSC)*, 1–12 https://www.informs-sim.org/wsc21papers/176.pdf.

Hoad, K., S. Robinson, and R. Davies. 2007. "Automating DES Output Analysis: How Many Replications to Run". In *2007 Winter Simulation Conference (WSC)*, 505-512 https://doi.org/10.1109/WSC.2007.4419641.

Hogg, R.V., and E.A. Tanis. 2006. *Probability and Statistical Inference*. 7th ed. Upper Saddle River, New Jersey: Prentice-Hall, Inc.

Lane, M., B. Frederick, J. Kavanagh, S. Watts, N. Chandler, and M.L. Smith. 2022. "Forecasting Demand for U.S. Ground Forces: Assessing Future Trends in Armed Conflict and U.S. Military Interventions". Santa Monica, CA: RAND Corporation https://www.rand.org/pubs/research_reports/RR2995.html.

Law, A., and W. Kelton. 2000. *Simulation Modeling and Analysis*. McGraw-Hill Series in Industrial Engineering and Management Science, McGraw-Hill.

Law, A. 2010. "Statistical Analysis of Simulation Output Data: The Practical State of the Art". In *2010 Winter Simulation Conference (WSC)*, 65-74 https://doi.org/10.1109/WSC.2010.5679174.

Wesolkowski, S., and A. Billyard. 2008. "The Stochastic Fleet Estimation (SaFE) model". In *Proceedings of the 2008 Spring Simulation Multiconference*, April 14-17, Ottawa, Ontario.

Wesolkowski, S., and D. Wojtaszek. 2012. "Multi-objective optimization of the fleet mix problem using the SaFER model". In *Proceedings of WCCI 2012 IEEE World Congress on Computational Intelligence*, June 10-15, Brisbane, Australia.

Wojtaszek, D., and S. Wesolkowski. 2012. "Military Fleet Mix Computation and Analysis". *IEEE Computational Intelligence Magazine* 7(3):53-61.

AUTHOR BIOGRAPHIES

LYNNE SERRÉ is a Defence Scientist in Defence Research and Development Canada's Centre for Operational Research and Analysis at the Department of National Defence in Ottawa, Canada. Through quantitative analysis, modeling, and simulation, her research has supported analysts and decision makers in the Military Personnel Command, the Canadian Army, and the Department's Materiel Group. She obtained her Master's degree in computational mathematics from the University of Waterloo, Canada. Her email address is lynne.serre@forces.gc.ca.

LISE ARSENEAU is a Defence Scientist in Defence Research and Development Canada's Centre for Operational Research and Analysis at the Department of National Defence in Ottawa, Canada. She currently works on a team providing decision support to the Royal Canadian Navy through quantiative analysis, modeling and simulation. She obtained her Master's degree and PhD from the University of Waterloo, Canda. Her email address is lise.arseneau@forces.gc.ca.