SIMULATION OF A READINESS-BASED SPARING OPTIMIZATION

John Wray[(a)], Arnold Buss[(b)], Javier Salmeron[(c)]

[(a),(b)] Modeling, Virtual Environments, and Simulation, Naval Postgraduate School, Monterey, CA 93940, U.S.A.
[(c)] Operations Research Department, Naval Postgraduate School, Monterey, CA 93940, U.S.A.

[(a)] jdwray1@nps.edu, [(b)]abuss@nps.edu, [(c)]jsalmero@nps.edu

ABSTRACT

We develop a discrete event simulation to complement a new optimization tool that establishes inventory levels for aviation weapon systems (WS) in the U.S. Navy. The optimization seeks cost minimization while achieving required readiness rates for hundreds of WS, each comprising thousands of indentured parts. Based on work in similar realms, the optimization employs the Vari-Metric model to estimate overall WS readiness and a variant of a greedy heuristic algorithm to set stock levels for all parts. Our simulation tests the assumptions and provides additional metrics for decision makers. We find that the estimates for readiness yielded by the optimization tool (a) have no systemic bias, and (b) remain within 5% in 53 of 64 WS (with an 8% worst-case). We also test two legacy optimization tools currently used by the Navy and find they have larger errors in expected readiness. We also identify factors correlated to these differences.

Keywords: Readiness based sparing; discrete event simulation; optimization; multi-indenture

1. INTRODUCTION

The United States Chief of Naval Operations (2011) requires the use of "readiness-based sparing (RBS) methodology to spares and repair parts allowance determination to ensure that prescribed readiness thresholds and objectives are achieved at the lowest possible cost." RBS uses advanced analytics to set inventory levels for most U.S. Navy parts and sub-parts at different locations.

To guarantee required combat power for the combatant commanders of the U.S. Navy (USN), all naval aviation Weapon Systems (WS) must maintain specified readiness (i.e. availability) rates. The term WS here identifies platforms such as the F/A-18 (Hornet) attack aircraft, or the MH60 (Seahawk) helicopter, among others. While reliability and maintainability are primarily set in the design phase of a WS, supportability is a crucial aspect of readiness that can be adjusted throughout the lifecycle of the system to achieve desired readiness rates. Supportability is affected by several factors; one of the key controllable elements is stock levels for spare parts at different echelons of supply.

Selecting the right mixtures of parts to stock at any given site in the USN is a very challenging task in a budget-constrained environment. A naval site contains numerous WS of different types and each WS may contain thousands of parts each failing at different rates. While it may not be possible to identify a provable optimal inventory for every site, our goal is to design and implement optimization and simulation tools that approximate such solutions and provide inventory policies that result in significant cost savings and improved fleet readiness over alternative solutions.

Although fill rate is a popular choice for evaluating inventory policies, it is problematic in a military setting where the ultimate goal is sufficient availability of WS. Although improving fill rates or reducing backorders will in fact improve readiness, policies developed with these metrics alone will be inefficient (Moulder et al. 2011). Looking solely at fill rates will inadvertently punish more complex WS. With all other factors such as failure rates and mean time to repair (MTTR) being equal, a WS with more parts will be requesting more parts from supply. If 95% of the parts are available upon request, a WS with more parts will be unavailable more often while awaiting parts than a WS with fewer parts.

In order to assist Naval Supply Systems Command (NAVSUP) with RBS planning, we develop an RBS Simulation (RSIM) to verify the recently developed Navy Aviation RBS Model (NAVARM) estimates and also compare its performance to the legacy Service Planning Optimization (SPO) and Aviation Readiness Requirements Oriented to Weapons Replaceable Assemblies (ARROWS) tools.

2. NAVARM AND RSIM

2.1. NAVARM Overview

NAVARM (Salmeron 2016) embeds a heuristic algorithm that approximates the optimal inventory quantities for a single-site, multi-indenture problem. Specifically, NAVARM recommends reorder quantities that minimize the cost of inventory held while maintaining pre-specified target availability rates for all WS.

NAVARM uses an (S-1, S) inventory model for all parts and sites. That is, S is the (maximum) stock level at a site determined by NAVARM and an order is placed as soon as that level decreases by one (i.e., the reorder point is S-1). This means that each time a part fails, it is turned into the system for repair. If the part cannot be repaired, a new part is ordered to replace it.
Assuming every part i is given a stock level S_i, every WS has an estimated availability that is calculated as a function of the expected backorders (EBOs) of the highest indenture parts in the WS. Naturally, backorders for any part in the system are a random variable which depends on: (a) the part's stock level; (b) its (possibly different) failure probability distributions for all common parts in the same or different WS; and (c) the backorder distribution for sub-indentured parts to all common parts. The underlying theory to calculate EBOs for a given set of inventory levels S_i is known as the Vari-Metric model, see Sherbrooke (1986, 2004, pp. 101-125).
The Vari-Metric model estimates EBOs under the assumption that, even though the number of failures for a given part can be approximated using a Poisson distribution, the actual number of failures after accounting for sub-indentured parts' failures is distributed as a Negative Binomial.
The multi-indenture structure used to describe WS repair with more fidelity complicates the problem significantly. Figure 1 illustrates this idea for three hypothetical WS at a site. For simplicity each WS has only one first-indenture part. If we are interested, for example, in improving the availability of WS 3, we can look at ways to decrease backorders of sub-parts "R" and "L". But, noting that part "L" is common to WS 2, its backorders are impacted by parts "M" and "N", and therefore by "G" in WS 2. Moreover, since this is common to WS 1, stocks of parts "H" and "I" in WS 1 will affect backorders of "L" in WS 3. The fact that WS 3 can be influenced by WS 1's parts (which have no direct commonality with parts in WS 3) is a challenging aspect of RBS optimization.

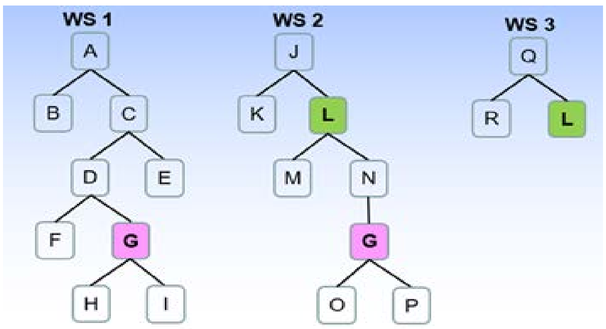

Figure 1: The Chain of Influence in a Multi-Indenture Part Structure

Sherbrooke (2004) points out that while the multi-indenture structure and the likelihood of common parts across WS types "does complicate the computer programs substantially … the basic logic is the same."

The use of heuristics to approximate the problem of satisfying a certain availability at minimum cost is justified due to the lack of a closed-form expression for expected readiness rates for a given reorder policy.
The Vari-Metric model suggests using a greedy heuristic based on an "effectiveness ratio" that measures improvement in EBOs with respect to cost. Parts with higher ratios are chosen until the desired availability is met. Even though this greedy heuristic is not provably optimal in a discrete setting (where we cannot order a fraction of a part), and counterexamples can be easily built, the method appears to work well in practice.
The matter becomes more complex when there are multiple WS with common parts. This is because if we follow the greedy algorithm for one WS at a time, we will achieve the desired availability at (approximately) minimum cost for, say, WS 1. But then, we will need other parts when working on WS 2. If some of those parts are common to WS 1, we will increase its availability unnecessarily above its target. Thus, refinements are needed, and NAVARM implements some of those, which basically consist of revisiting all of the WS above target in order to remove parts and reduce cost.

2.2. RSIM Introduction and Scope

RSIM (Wray 2017) simulates failures at the individual part level and then aggregates up to the individual WS level and WS type to help assess the accuracy of expected backorders and WS availability. To simulate the system of interest, three major classes of entities are created: parts, WS and part positions. Each part has attributes that include:

- Status (i.e., functioning or down for maintenance),
- Planned failure time (detailed below), and
- Position (specifying where the part is installed if currently in use)

Each WS has attributes that include:

- Type (e.g., CH-53 helicopter),
- Availability status (i.e., up or down), and
- A list of part positions that comprise the WS (e.g., hydraulic pump).

A part position has attributes that include:

- The WS (if currently in use),
- Parameters describing expected failure times, and
- Parameters describing the time for a working part to return to inventory after breaking.

Modeling failures in a manner that closely mirrors reality is crucial to attaining realistic outputs. Expected failure rates can be derived from existing databases and are broken down into failures that can be repaired at the site and failures that cannot. Some parts have only one type of failure or the other while some have both.

The type of failure is tracked in RSIM to later develop an expected time the part will return to inventory in a working status. To handle the difference in types of failures, RSIM first combines the failure rates and assigns a failure time based on the rate. When the failure occurs, a random number draw is compared to the ratio of repairable and non-repairable failures to assign the type.

While multi-parameter distributions such as the Weibull that allow specification of mean and variance are generally preferred for detailed modeling of failure rates, the databases used for RBS currently provide only the mean failure rates. As a result, the exponential distribution is employed by RSIM to generate the next failure. This is also consistent with the assumptions established for NAVARM, SPO and ARROWS.

Failure rates in the database and RSIM are specific to a part position on a WS type; a hydraulic pump used on a CH-53E utility system may have a different failure rate than the same type of hydraulic pump used on an SH-60S utility system. In fact, the same pump may be installed on different WS types between failures and thus have different failure rates assigned based on where it is installed. When a part fails, RSIM removes the part from the usable pool for a specified period of time until repaired or replaced.

Although RSIM is best described with an event graph (see Figure 2), the basic steps are as follows:

- Read data in from database and instantiate all entities specified in the data.
- Assign parts to fill each WS and assign a first failure time stochastically for each part based on its specified distribution.
- When a failure occurs:
  - Assign a time the part will return to service.
  - If a part of the correct type is available in inventory, place WS in down status for the specified MTTR. If a part is not available, add the WS to a first-in, first-out (FIFO) queue for that part type.
- When a part returns to a ready-for-issue status at the site, use it to repair the first WS in the FIFO queue awaiting that part type. If no WS are awaiting that part type, return the part to inventory.

To manage complexity, simplify the verification and validation process, and ensure acceptable simulation run-time, RSIM tightly scopes the factors considered in the simulation while maintaining flexibility to add new factors as desired to closer mirror reality or support study objectives.

In its current state, RSIM ingests summary level data on flight hours, failure rates, repair times and shipping times and most of these factors are treated deterministically. While RSIM could simulate actual flight sorties and assign failures based on WS flight times, the effect of this added fidelity would likely be nominal when considering inventory policies and thus is not included. Likewise, scheduling the repair process at intermediate and depot level and including manpower and part availability consideration here would also have minimal effect on the metrics currently of interest; instead, expected values are substituted in lieu of this detailed analysis.

Finally, RSIM could consider the phasing of required repairs and how they may coincide with required periodic inspections and planned maintenance to minimize downtime. Part failures that render the WS partial mission capable could remain on the WS until an optimal time to complete the repair. Again, the effect of including this would not shed light on the objectives at hand though it may be a worthwhile future enhancement to provide decision makers with a fuller sense of what to expect if the given NAVARM solution is implemented.

2.3. RSIM Assumptions

A number of assumptions are made in the RSIM implementation, some of which could significantly impact the results. These assumptions are made for a variety of reasons to include limited data availability, code simplicity, and reduced run-time. The inherent flexibility in RSIM implementation makes these assumptions fairly easy to modify or eliminate through code manipulation. The following are significant assumptions currently made in RSIM:

- Failure rates are accurately represented by an exponential distribution: As stated early, failure rates would likely be better represented with a Gamma or Weibull distribution, but the limited failure data provided does not allow for such implementation. The exponential distribution is not very well suited to represent wear out failures that occur at fairly predictable intervals as compared to random failures.
- Failures are independent: Because failure times are scheduled into the future on a continuous timeline and there are no dependencies in our program, simultaneous failures will not occur despite real world experiences that suggest otherwise.
- Failures in the simulation should continue to happen when the WS is down: While failure rates in the database are given per flight hour, this data along with average flight hours is used to develop expected mean time between failures. Although parts are much less likely to fail when the WS is out of service, scheduled failures continue to occur in the simulation to ensure the expected failure rate is maintained. This may result in overlap of delay times for backordered parts. With a higher fidelity data set, this could be improved by developing conditional probabilities for failures that better reflect the empirical data.

- Expected sub-indentured part failure times are not reset when a parent part is changed: This assumes that all parts are repaired and that when they undergo repair, it does not affect reliability of the separate sub components. This assumption will fail if the part is replaced and sub components are not salvaged, but the available data does not delineate how often parts are repaired when they go off-site and what happens to sub-indentured parts when a replacement is necessary for the parent part. Of note, this assumption will lead to a conservative estimate of availability, though the extent of the impact is unknown with the data currently available.
- Demands are FIFO: This assumes that no priority will be given to WS of types that are below their availability goal or some other prioritization scheme.
- No lateral resupply: There is no cross-leveling between sites that have high and low inventories (or backorders) for a particular part.
- Cannibalization: While moving working parts from a down WS to one that can be returned to an up status is practiced in the real world, RSIM is conservative in not assuming so.
- Repair times are independent: RSIM does not attempt to simulate backlogs in the repair pipeline that would likely occur if multiple parts of the same type were in the repair pipeline simultaneously.

2.4. RSIM Event Graph and Implementation

Figure 2 depicts an event graph describing the overall model of part failures and subsequent repairs in RSIM. This version of the event graph is simplified and intended to provide a broad understanding of the flow of parts through the system.

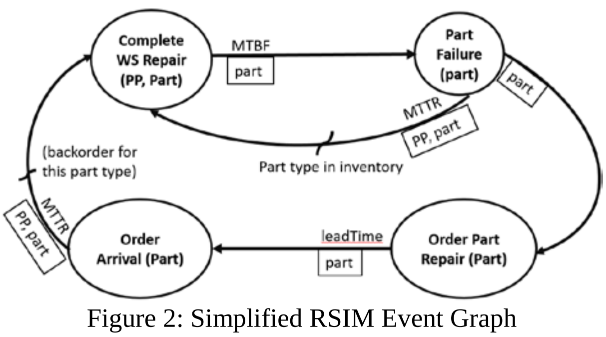

Figure 2: Simplified RSIM Event Graph

The RSIM event graph depicted in Figure 2 has been implemented in the Java programming language using the Simkit library (Buss 2002, 2004). Simkit provides the necessary support for easily converting the event graph into working code. As an open source library, Simkit is free of the encumbrances of commercial licensing, while providing excellent support for the model's features. An additional open source library, UCanAccess (2017), has been used to interact with the MS Access database inputs.

3. RESULTS

3.1. Introduction

In its current configuration, RSIM outputs several metrics by WS and part type to allow comparison to other RBS optimization software (i.e., NAVARM, SPO, and ARROWS). Additionally, it provides decision makers with a more comprehensive understanding of what to expect if a certain inventory policy is implemented. For each WS type, RSIM provides the mean number of WS available, the corresponding readiness rate, and the percent of simulated time the WS type was at or above its given readiness goal. For each part type, RSIM outputs the mean inventory level, mean number of backorders, and the fill rate.

The primary metric of interest is the readiness rate by WS type. Given the crucial nature of having required force levels available at any given time, NAVSUP must ensure the inventory quantities selected will enable this objective. RSIM, NAVARM, SPO, and ARROWS each have assumptions built in that may not be accurate, but a comparison of the outputs can be helpful in assessing the validity of the readiness estimates.

The data used for the analysis below is generated using Dell Inspiron I5378 laptop running Windows 10 with an Intel Core i7-7500U 2.7 GHz CPU and 8 GB of RAM. RSIM is implemented in JDK 1.8 and utilizes the 64-bit Simkit version 1.4.6 and UCanAccess version 4.0.1. Based on steady state analysis conducted for several sites, we use a warmup period of 3,000 simulated days before collecting 7,000 simulated days of data for 30 replications at each site analyzed. These settings result in less than 1% margin of error for readiness levels of all WS tested. Run times range between 2.5 and 59 minutes for the seven sites analyzed. NAVARM runs are conducted on the same laptop described above using the 32-bit version of Microsoft Excel 2016. While there are several NAVARM setting that can significantly affect the run time, standard settings result in run times ranging between 30 seconds and 18 minutes.

3.2. RSIM Compared to NAVARM

We have run RSIM on seven representative USN sites to compare expected WS availability rates and expected backorder rates for a given allowancing to those anticipated by NAVARM. The number of WS types at these Naval sites ranges from 3 to 23 with a mean of just over 9 WS types per site and a mean of 111 individual WS per site.

Figure 3 shows the summary histogram of the differences in expected availability for the 64 WS types tested. Out of the 64 WS types analyzed, 53 have expected readiness levels within 5% and the mean difference for all WS types in this sample is .2% with no systemic bias to over or under estimate readiness noted.

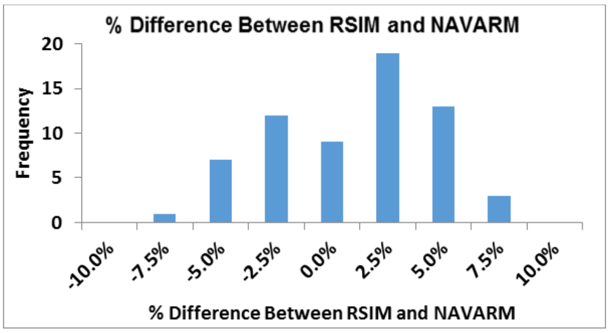

Figure 3: Differences in Expected Availability between NAVARM and RSIM

3.3. Comparison to Legacy RBS Tools

As NAVSUP considers whether to switch RBS optimization tools, it is crucial for the decision makers to assess the accuracy of the NAVARM expected readiness calculations and compare its accuracy to the SPO and ARROWS tools currently in use. Because RSIM models the system at the part and WS level, its method of observing readiness rates through the course of a simulation provides an independent observation to compare against the optimization tool estimates available. SPO, ARROWS, and NAVARM have been run at a representative site with seven different WS types and a total of 62 WS. Their recommended inventory policies have been simulated in RSIM in order to compare the expected readiness rates for each WS type. Figure 4 shows a summary of the resulting differences in estimates. In this case, it becomes clear that NAVARM's estimated readiness rates are much closer to RSIM than SPO or ARROWS estimates are.

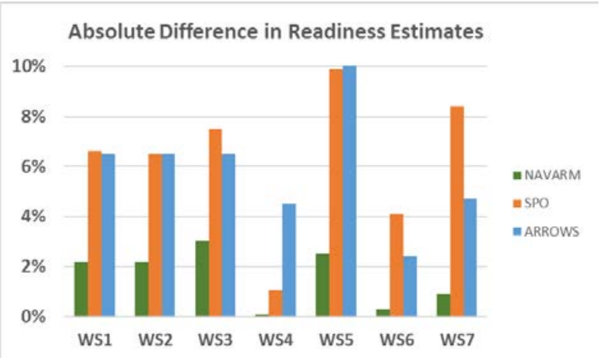

Figure 4: Absolute % Difference between Available Optimization Tools and RSIM Readiness Estimates by WS Type at a Single Site

3.4. Analysis of Attributes Affecting Performance

While the difference between expected readiness given by RSIM and NAVARM is likely acceptable in the current versions, we have tried to identify key drivers of any differences found in the hope of further explaining the differences and ideally reducing the errors. First, we consider attributes of the WS type that may complicate calculations for availability in the models. The factors of interest by WS type are:

- Quantity per application (QPA): Some part types are found on a WS multiple times at the same indenture level. In these cases, the part type is represented with a single line in the database and the number of parts used is given as QPA. While QPA implementation is straightforward in simulation, its use in the readiness estimation for NAVARM has been debated.
- Commonality: This is a measure of how often the same part type is used throughout the site. This adds a layer of complexity in the optimization model as changes in inventory can affect numerous WS types.
- Number of parts: The number of parts tracked on a given WS type in our data sets ranged from 80 to over 8,000.
- Indenture depth level: Assumptions are made in both RSIM and NAVARM regarding the impact of indenture depth level on part failure rates and readiness.

Each of the above factors has been examined compared to the difference in readiness estimates produced by NAVARM and RSIM. None of the correlations are significantly strong, with total number of parts being the highest (0.49), indenture depth level and mean number of common parts slightly lower (0.43 and 0.40, respectively), and QPA being clearly non-significant (0.02).

3.5. Comparison of EBO Outputs

Because EBOs play an integral part in the NAVARM calculations of expected readiness, we configure NAVARM and RSIM to output EBOs for every part position to compare expected EBO levels. At the sample site described above, there are over 11,000 part positions tracked. Of these, approximately 6,400 part positions are expected to have some backorders based on their failure rate and stock level. The magnitude of the difference is generally negligible with only a 2% difference in the sum of EBOs for NAVARM and RSIM and an average difference of less than .0003 per part. Figure 5 shows a histogram comparison of EBOs. Parts with extremely low EBO levels will not significantly impact overall WS readiness. Of note, the counts are nearly identical for EBOs above .001.

We consider the same abovementioned factors as potential drivers in the difference in EBO levels between NAVARM and RSIM, but again neither of these factors demonstrates strong correlation with the differences in EBO levels. Specifically, the fact that indenture depth does not strongly influence availability or EBO differences suggests that the Vari-Metric assumption of negative binomial distribution for modeling EBOs of sub-indentured parts is reasonable.

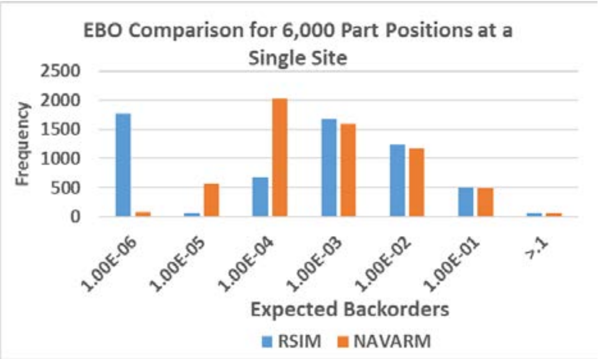

Figure 5: Observed RSIM EBO Levels Compared to NAVARM Expected EBO Levels at a Single Site

3.6. Additional RSIM Insights

In addition to verifying NAVARM outputs and providing an independent comparison of the three available RBS optimization tools, RSIM can provide additional insights not available with the optimization output. One example of this is the expected readiness levels. While the optimization tools only provide the expected readiness levels overall, RSIM provides metrics that include percentage of time above the stated readiness goal and the readiness levels observed at the beginning of each simulated day.

For example, Figure 6 shows a histogram of observed daily readiness levels over a period of 7,000 simulated days for a single WS type with 22 WS at a single site. Even though the most important output of RSIM is the expected readiness achieved (in this case 60.7%, slightly below the 63% goal), additional valuable information can be gleaned: In this simulation run, 48.2% of the simulated time had readiness rates above the goal. A decision maker may be more interested in worst-case scenarios to ensure that assumptions made for contingency planning are realistic. The fact that we expect less than 50% readiness during 11% of the time may be of interest.

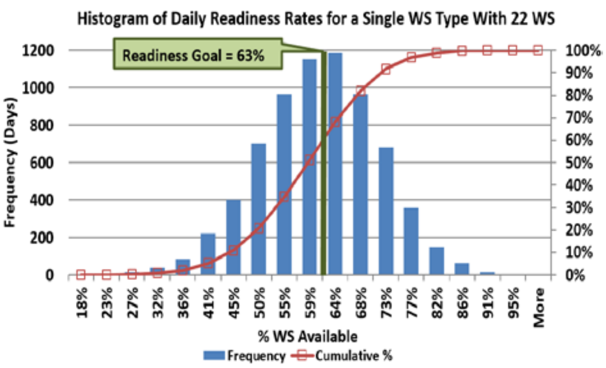

Figure 6: Histogram of Daily Readiness Rates in RSIM for 7,000 Simulated Days of a Single WS Type

In addition to providing useful metrics for decision makers, RSIM provides metrics that may be used to improve recommended inventory levels calculated by NAVARM. Table 1 shows a small sample of output from RSIM for six part types at a single site. The mean backorder level, mean inventory level, number of failures over the period of the simulation and the fill rate are output for each part type. The full output will contain thousands of entries, but an analyst could sort this list to help identify areas where inventory levels could be manually adjusted to incorporate other factors not accounted for in the NAVARM optimization, such as those due to modeling assumptions described in Section 2.3. This process could be automated and take advantage of both NAVARM and RSIM to evaluate the changes. Moreover, RSIM could be extended to implement its own adjustments and become a complement to NAVAIR's optimization.

Table 1: Sample RSIM Output by Part Type

part type	mean backorders	mean inventory	# failures	fill rate
Part 1	0.01	0.89	11	0.91
Part 2	0.37	1.23	1592	0.62
Part 3	0	0.93	7	0.86
Part 4	0.11	0.55	65	0.63
Part 5	0	4.96	6	1
Part 6	0	0.93	6	1

Simulation lends itself well to the collection of numerous metrics. Here we have presented several metrics that may be useful to the decision makers or analysts. As our understanding of the problem continues to develop, we expect to modify the RSIM assumptions and metrics accordingly.

ACKNOWLEDGMENTS

The authors thank Naval Supply Systems Command, Weapon Systems Support for their support of this research. Special thanks to Ms. Karla Lizzo for her continued assistance.

REFERENCES

Buss, A., 2002. Component Based Simulation Modeling with Simkit. Proceedings of the Winter Simulation Conference, pp. 244-249. December 8-11, San Diego, California.

Buss, A., 2004. Simkit Analysis Workbench for Rapid Construction of Modeling and Simulation Components. Paper Presented at the Fall Simulation Interoperability Workshop. April 18-23, Arlington, Virginia.

Moulder, R.D., O'Malley, T.J. and Peterson, D. K., 2011. Sizing Inventory with Readiness-Based Sparing. Air Force Journal of Logistics 35: 18-29.

Salmeron, J., 2016. "Readiness-Based Sparing Inventory Optimization Model." Presentation to NAVSUP. Naval Postgraduate School, Monterey, CA. (Available from author.)

Sherbrooke, C. C., 1986. VARI-METRIC: Improved Approximations for Multi-indenture, Multi-

echelon Availability Models. Operations Research 34 (2): 311-319.
Sherbrooke, C. C., 2004. Optimal Inventory Modeling of Systems: Multi-Echelon Techniques Second Edition. New York: Kluwer Academic Publishers.
UCanAccess, 2017. A pure java JDBC Driver for MS Access. Available from: https://sourceforge.net/projects/ucanaccess/ [accessed 29 June 2017]
United States Chief of Naval Operations, 2011. Readiness Based Sparing. OPNAV Instruction 4442.5A.
Wray, J.D., 2017. A Simulation of Readiness Based Sparing Policies. M.S. Thesis in Modeling, Virtual Environments and Simulation. Naval Postgraduate School, Monterey, CA.

AUTHORS BIOGRAPHY

Major John Wray is a helicopter pilot in the United States Marine Corps. He earned a B.S. degree in Systems Engineering from the United States Naval Academy, and two M.S. degrees in (a) Operations Research and (b) Modeling, Virtual Environments, and Simulation, from the Naval Postgraduate School in Monterey, CA.

Arnold Buss received his B.A. from Rutgers University, M.S. from University of Arizona, and Ph.D. from Cornell University. He is currently a Research Associate Professor in the Modeling, Virtual Environments, and Simulation Institute at the Naval Postgraduate School. His research interests are in simulation methodology, applied simulation modeling, and agent-based modeling. His previous positions include Assistant Professor at Washington University and Visiting Assistant Professor in Operations Research at the Naval Postgraduate School. His simulation software, Simkit, has been used to teach discrete event simulation for the past 20 years. http://faculty.nps.edu/vitae/cgi-bin/vita.cgi?p=display_vita&id=1023568056

Javier Salmeron received the M.S. and Ph.D. degrees in mathematics from Complutense University and Polytechnic University of Madrid, Spain, respectively. He is currently an Associate Professor in the Operations Research Department at the Naval Postgraduate School, Monterey, CA, USA, where he teaches courses in optimization and its applications. His research focuses is in the area of applied modeling and optimization, which he has actively used in multiple civilian and military projects. http://faculty.nps.edu/jsalmero/