



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

STATISTICAL SENSITIVITY ANALYSIS OF THE REPLENISHMENT AT SEA PLANNER

by

Jared R. Deiter

September 2022

Thesis Advisor:
Second Reader:

Johannes O. Royset
Gerald G. Brown

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 2022	3. REPORT TYPE AND DATES COVERED Master's thesis		
4. TITLE AND SUBTITLE STATISTICAL SENSITIVITY ANALYSIS OF THE REPLENISHMENT AT SEA PLANNER			5. FUNDING NUMBERS	
6. AUTHOR(S) Jared R. Deiter				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) Underway replenishment is required for ships to operate at sea without port calls. The Replenishment At-Sea Planner (RASP) provides optimized schedules while considering a myriad of factors. We develop a statistical sensitivity analysis of the effect changes to RASP inputs have on outputs such as Combat Logistics Force (CLF) fuel consumption, CLF ship underway percentage, and combatant supply safety stock level. The resulting statistical models are useful for logistical planners if RASP is unavailable, yet decisions regarding the schedule must be made and avoid needing to re-solve RASP. Models of western Pacific scenarios schedule the replenishment of Carrier Strike Groups (CSGs) (e.g., one Aircraft Carrier, one Cruiser, and two Destroyers) and CLF ships. In a one-CSG scenario, we develop a statistical model that predicts CLF fuel consumption and percent of time CLF ships are underway with an average error of 4.6% and 13.7% respectively and these predictions are consistently below the actual values. In a two-CSG scenario, a statistical model either over or under-predicts CLF fuel consumption based on regional boundary constraints on CLF operations. Predictions are consistently between -26% and -14% under and 19% and 27% over. In order of importance, the number of days in the CSG sustainment cycle, regional boundary limitations imposed on CLF ships, and the number of CLF ships available are the most influential to RASP outputs.				
14. SUBJECT TERMS RASP, logistics at-sea planning, replenishment at-sea, optimization, Combat Logistics Force, CLF, at-sea replenishment scheduling			15. NUMBER OF PAGES 105	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**STATISTICAL SENSITIVITY ANALYSIS OF THE REPLENISHMENT AT SEA
PLANNER**

Jared R. Deiter
Lieutenant Commander, United States Navy
BA, San Diego State University, 2008

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2022**

Approved by: Johannes O. Royset
Advisor

Gerald G. Brown
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Underway replenishment is required for ships to operate at sea without port calls. The Replenishment At-Sea Planner (RASP) provides optimized schedules while considering a myriad of factors. We develop a statistical sensitivity analysis of the effect changes to RASP inputs have on outputs such as Combat Logistics Force (CLF) fuel consumption, CLF ship underway percentage, and combatant supply safety stock level. The resulting statistical models are useful for logistical planners if RASP is unavailable, yet decisions regarding the schedule must be made and avoid needing to re-solve RASP. Models of western Pacific scenarios schedule the replenishment of Carrier Strike Groups (CSGs) (e.g., one Aircraft Carrier, one Cruiser, and two Destroyers) and CLF ships. In a one-CSG scenario, we develop a statistical model that predicts CLF fuel consumption and percent of time CLF ships are underway with an average error of 4.6% and 13.7% respectively and these predictions are consistently below the actual values. In a two-CSG scenario, a statistical model either over or under-predicts CLF fuel consumption based on regional boundary constraints on CLF operations. Predictions are consistently between -26% and -14% under and 19% and 27% over. In order of importance, the number of days in the CSG sustainment cycle, regional boundary limitations imposed on CLF ships, and the number of CLF ships available are the most influential to RASP outputs.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	Background	1
1.1	Optimizing Naval Logistics	1
1.2	Underway Replenishment Scheduling Before Optimization.	1
1.3	CLF Optimization and Planner Development	2
1.4	Motivation and Goals	5
2	Key Logistics Concepts	7
2.1	Endurance	7
2.2	Replenishment-At-Sea	8
2.3	The Combat Logistics Force	9
2.4	Logistics Supply and Demand	10
2.5	Fleet Composition and Density	10
2.6	Sea Routes	11
2.7	Fleet Areas of Responsibility	11
2.8	RASP Model Formulation.	13
3	Statistical Sensitivity Analysis	21
3.1	Introduction	21
3.2	Statistical Sensitivity Analysis	21
3.3	Datasets: Scenario CSG-1 and Scenario CSG-2	22
3.4	Statistical Models	25
3.5	One-Hot Encoding.	26
3.6	Assessing Statistical Model Accuracy	29
3.7	Scenarios Constructed	29
3.8	Customary Discussion of Optimization Model.	31
4	Statistical Model Predictive Capability	33
4.1	Scenario CSG-1 Results	34

4.2	Scenario CSG-2 Results	42
4.3	Improving Error Through Subsets of Case B	46
4.4	Case C: Classification of Maximum Percent Below Observations	68
5	Conclusions and Future Research	73
5.1	Scenario Results	73
5.2	Key Take Aways.	75
5.3	Future Research	76
	List of References	79
	Initial Distribution List	81

List of Figures

Figure 2.1	US Numbered Fleet Areas of Responsibility (AORs) Worldwide .	12
Figure 3.1	Nimitz Class Aircraft Carrier	22
Figure 3.2	Ticonderoga-Class Cruiser	23
Figure 3.3	Arleigh Burke-Class Destroyer	23
Figure 3.4	Henry J. Kaiser Class Fleet Replenishment Oiler	24
Figure 3.5	Royal Australian Navy Replenishment Oiler	24
Figure 3.6	One-Hot Encoding Representation	26
Figure 3.7	Support Vector Machine	28
Figure 4.1	Scenario CSG-1 Fuel Consumed Prediction	36
Figure 4.2	Scenario CSG-1 Fuel Consumed Prediction Error	37
Figure 4.3	Scenario CSG-1 Supply Ship Underway Percentage Prediction . .	39
Figure 4.4	Scenario CSG-1 Supply Ship Underway Percentage Absolute Error	40
Figure 4.5	Scenario CSG-1 Supply Ship Underway Percentage Error	41
Figure 4.6	Scenario CSG-2 (Case B) Supply Ship Fuel Consumed Prediction	44
Figure 4.7	Scenario CSG-2 (Case B) Supply Ship Fuel Consumed Error . .	45
Figure 4.8	Scenario CSG-2 (Case B) Supply Ship Fuel Consumed Absolute Error	45
Figure 4.9	Scenario CSG-2 (Case B.1) Principal Component Analysis (PCA) (CLF AVAIL)	47
Figure 4.10	Scenario CSG-2 (Case B.1) Supply Ship Fuel Consumed Prediction	48
Figure 4.11	Scenario CSG-2 (Case B.1) Supply Ship Fuel Consumed Prediction Error	49

Figure 4.12	Scenario CSG-2 (Case B.1) Supply Ship Fuel Consumed Prediction Absolute Error	50
Figure 4.13	Scenario CSG-2 (Case B.2) Principal Component Analysis (PCA) (SDO DELAY)	51
Figure 4.14	Scenario CSG-2 (Case B.2) Supply Ship Fuel Consumed Prediction	52
Figure 4.15	Scenario CSG-2 (Case B.2) Supply Ship Fuel Consumed Prediction Error	53
Figure 4.16	Scenario CSG-2 (Case B.3) Principal Component Analysis (PCA) (PLH DELAY)	54
Figure 4.17	Scenario CSG-2 (Case B.3) Supply Ship Fuel Consumed Prediction	55
Figure 4.18	Scenario CSG-2 (Case B.3) Supply Ship Fuel Consumed Prediction Error	56
Figure 4.19	Scenario CSG-2 (Case B.4) Principal Component Analysis (PCA) (PORT AVAIL)	57
Figure 4.20	Scenario CSG-2 (Case B.4) Supply Ship Fuel Consumed Prediction	58
Figure 4.21	Scenario CSG-2 (Case B.4) Supply Ship Fuel Consumed Prediction Error	59
Figure 4.22	Scenario CSG-2 (Case B.5) Supply Ship Fuel Consumed Prediction	60
Figure 4.23	Scenario CSG-2 (Case B.5) Supply Ship Fuel Consumed Prediction Error	61
Figure 4.24	Scenario CSG-2 (Case B.5.1) Supply Ship Fuel Consumed Prediction	62
Figure 4.25	Scenario CSG-2 (Case B.5.1) Supply Ship Fuel Consumed Prediction Error	63
Figure 4.26	Scenario CSG-2 (Case B.5.2) Supply Ship Fuel Consumed Prediction	64
Figure 4.27	Scenario CSG-2 (Case B.5.2) Supply Ship Fuel Consumed Prediction Error	66
Figure 4.28	Scenario CSG-2 Error Summary	67
Figure 4.29	Correct Classification Percentage	69

Figure 4.30	Classification Optimal Value	69
Figure 4.31	False Positive and False Negative Classifications	70

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

Table 3.1	Overview of Inputs Used in Scenario CSG-1	30
Table 3.2	Overview of Inputs Used in Scenario CSG-2	31
Table 4.1	Overview of Cases Explored	34

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

AOR	Area of Responsibility; earth is divided into different areas based on responsible entity (e.g., Fleets, Combatant Commands)
CLF	Combat Logistics Force; fleet of supply ships used to conduct underway replenishment
CONSOL	Consolidation event; an underway replenishment event when two supply ships meet, perhaps underway, and transfer fuel and/or supplies before continuing to rendezvous with a customer ship
CONREP	Connected Replenishment; a method of underway replenishment used to transfer fuel between logistics ships and combatants
CPLEX	C-Language Simplex Method; method for solving optimization problems
CSG	Carrier Strike Group; a group of five to six combatants deployed supporting national objectives
DFM	Diesel Fuel Marine; fuel used by maritime surface vessels not nuclear powered
DOD	Department of Defense
GAMS	General Algebraic Modeling System; mathematical computer-based language used for solving optimization problems
INREP	In-port Replenishment; an event in which supply or customer ships receive fuel or other commodities pierside
JP5	Jet Propellant-5; fuel used by U.S. naval aircraft
MIP	Mixed Integer Program
MSC	Military Sea-lift Command; controlling command for all Combat Logistics Force ships

MVLR	Multi-Variable Linear Regression
NPS	Naval Postgraduate School
OPREP	Operational Report; daily informational report sent by an individual unit to a controlling command
OTTER	Optimized Transit Tool and Easy Reference
PCA	Principal Component Analysis; method of identifying most influential factors to a response quantity
PYOMO	Python-based, open-source optimization modeling
RAS	Replenishment-At-Sea; an event when fuel and other commodities are transferred between underway ships
RASP	Replenishment-At-Sea Planner
RMSE	Root Mean Square Error
SVM	Support Vector Machine
TRANSCOM	United States Transportation Command
UNREP	Underway Replenishment; a supply ship rendezvous with an underway customer ship to transfer to it fuel and other commodities
USN	United States Navy
USNS	United States Naval Ship; supply ship employed by the Military Sealift Command and used to replenish customer ships
USS	United States Ship; combatant ship employed by the United States Navy
VBA	Visual Basic for Applications; computer programming language used in Microsoft applications
VERTREP	Vertical Replenishment; a rendezvous between supply ships and customer ship where supplies are transferred by helicopter

Executive Summary

Dynamically and globally employed ships require recurring underway and in-port replenishment of many commodities to maintain operational availability and maximize on-station time. The Replenishment-At-Sea Planner (RASP) provides optimized schedules, considers myriad factors in support of combatant employment. Current schedules have a certain degree of vulnerability, often needing revisions when changes occur to combatant commodity levels or schedule requirements. The implementation of RASP alleviates many time-consuming steps that fleet schedulers used to perform by hand, but certain RASP solves preclude timely force employment and sustainment decisions due to complexity and time involved. To compound complexity, sustaining logistics in contested environments (either by weather or adversary) is an evolving concept moving at an increasingly rapid pace and warrants a change to business practices. Having the capability to solve these complex problems must come in advance of the need to produce solutions.

Statistical sensitivity analysis of the effect changes to input parameters have on corresponding outputs develops intuition regarding the ability to execute replenishment schedules. Every schedule has an associated degree of confidence associated with its feasibility after changes to initial conditions occur. Insight gained through analysis of predicting Combat Logistics Force (CLF) fuel consumption, CLF ship underway percentages, and any breach of safety stock thresholds is pivotal to informing a fleet commander's decisions. While considering overhead cost, planners make daily decisions about feasibility of mission execution, fleet-replenishment timing based on mission tasking requirements and force structure required to maximize on-station time. It is important to have quick and reasonably accurate assessments about these factors.

This thesis quantifies changes to RASP output (i.e., CLF fuel consumption, percent of days underway for CLF ships, and maximum percent below safety stock threshold) when its input parameters are perturbed. Perturbations are changes to the status quo; shortening the number of days between replenishment, limiting which CLF ships are available, or even limiting port availability for CLF ship resupply. Using statistical models, we determine which perturbations have the greatest effect on outputs. In particular, we predict whether any ship's inventory falls below a specified threshold. This provides significant insight about

schedules, especially their vulnerability to adversarial interdiction. Statistical models that predict outputs using specified inputs also enable quick what-if analysis without relying on potentially time-consuming RASP runs. Regardless of time constraints, many analysts may not be able to run RASP. The statistical models provide a simple and transparent tool for analyzing the effect of inputs deviating from those used to develop an incumbent plan.

For a scenario in the western Pacific involving a single Carrier Strike Group (CSG) (i.e., one Aircraft Carrier, one Guided-Missile Cruiser, and two Guided-Missile Destroyers) and four logistics replenishment oilers, we develop a statistical model that predicts CLF fuel consumption and percent of time CLF ships are underway with an average error (difference between predicted value and actual RASP output) of 4.6 and 13.7% respectively and these predictions are consistently below the actual values. A 4.6% error equates to 51,000 gallons of fuel, which is approximately 6,000 gallons of fuel, per ship, over the entire time horizon; the time horizon considered in these scenarios is 60 days. A 13.7% error equates to approximately 8 days, a significant error in planning.

In a second western Pacific scenario involving two CSGs and three CLF ships, a statistical model either over- or under-predicts CLF fuel consumption based on regional boundary constraints on CLF operations. Predictions are consistently between 26% and 14% under and between 19% and 27% over. Under predictions occur when CLF ships are required to stay within assigned geographic regions, and over predictions occur when ships are allowed to transit outside regional boundaries by a designated number of nautical miles.

Using the second scenario dataset again, a classification model (support vector machine) identifies an alternative way to determine schedule feasibility. Instead of focusing on prediction of output values directly, a classification of whether an observation exceeds a threshold is pursued. We develop a classification model for predicting whether two CSGs fall below the 30% safety stock threshold – a threshold controlled by the fleet commander. Using CLF ship and port availability as predictors, we classify with 82% accuracy whether such critical events occur.

We predict CLF fuel consumption decreases by 67,962 gallons per customer ship sustainment cycle day increase. It decreases by 117,916 gallons when the supply ship assigned to Guam is unavailable as a supply port. It decreases by 109,778 gallons when Pearl Harbor is unavailable as a resupply port. It also decreases when supply ships are restricted to their

assigned regions in support of the schedule. These reductions equate to savings of \$187,497, \$318,373, \$296,401, and \$528,938 respectively. We additionally predict that CLF underway percentage decreases by 2 per sustainment cycle increase. This reduction equates to 12 fewer days spent underway. Of course CLF fuel consumption savings typically come at the expense of other metrics planners consider (e.g., fuel inventory levels onboard combatants). Trade-offs occur where fuels savings are not always the leading metrics determining supply and customer ship schedules. The respective statistical model for each scenario is easily implemented in an Excel spreadsheet for convenient use by the planner to predict some outputs of RASP from various inputs.

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgments

First and foremost, I would like to thank my beautiful wife, Samantha. Without your constant support as my grounding rod through the turbulent thunderstorm called “Operations Research” at Naval Postgraduate School (NPS) and keeping my focus True, my compass would be spinning in search of North today. There are no words to express my gratitude for you, filling in the gaps created by the military life and ensuring our children, and our collective livelihood, remain unforgotten.

Second, thank you Chloe and Jamesson, my children, for forcing much-needed distractions from coursework when clearly my patience waned, and focus could not be maintained. For your dedication to seeking my attention, I commend your unrelenting efforts.

Third, I thank my parents. Mom, if you only knew how much support you provided without even saying a single word. Your support through all endeavors persists without pause and will never tire. Thank you! Dad, your ability to help keep perspective is sometimes unmatched. I will never forget your repeated words: “If it was easy, everyone would do it.” Thank you!

Fourth, a hearty thanks to Dr. Johannes Royset for your guidance, both outside and inside the classroom, is in order. Your teaching ability, dedication to every student’s understanding of concepts, and expert knowledge are among the best campus wide.

Fifth, to Mr. Anton Rowe. Your superb dedication to implementing batch processing to the Replenishment-at-Sea Planner interactive model housed within Microsoft Excel. Without your efforts, gathering data for analysis would not have been possible in a reasonable amount of time. You are a tremendous asset, and I thank you.

Last, but certainly not least, I would like to thank the friends I have made through this educational journey. Our conversations, both in and out of OR context, have helped immensely in my pursuit of this degree. Without your help, I would not have gotten through this program in the same end-state, alive.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:

Background

1.1 Optimizing Naval Logistics

Sustaining operations, whether in wartime or peace, has always been predicated on the strength of the supporting logistics chain. Until after World War II, the impact logistics has on force sustainment was viewed as less important than tactics applied to objectives. Even more recently, optimizing logistics was not considered as tacticians focused on shorter planning horizons, meeting urgent and emergent requirements as they arise, instead of applying demand projection techniques; projecting demand has only been a focus in the past 30 years. To prevent wasting scarce resources, i.e., food supplies and energy sources among others, constraints help minimize waste and abuse. Improved efficiencies through technology and deliberate scheduling of ship underway time – training, qualifications, and deployments – are methods to minimize resource use. This is not easy; logistics at-sea is difficult and complex, thus optimizing its execution becomes paramount as increasing flexibility in the dynamic operational environments becomes ever more prevalent.

1.2 Underway Replenishment Scheduling Before Optimization

Before optimization-based planning tools, only non-analytical methods were used to schedule and track fleet replenishment requirements within an Area of Responsibility (AOR). Use of hand-written maps projecting plans of intended ship's movement, charts and tables depicting fuel curves predicting ship fuel consumption for each class of ship at different speeds, manually updated tracking boards and computer-based spreadsheets were required to plan and schedule Replenishment-At-Sea (RAS) and Consolidation events (CONSOLs); this was tedious, time consuming and highly prone to error and inefficiencies in schedules (Diaz 2010). A CONSOL event takes place when two supply ships meet, perhaps underway, and transfer fuel and/or supplies before continuing to rendezvous a customer ship. CONSOL events may involve smaller Military Sea-lift Command (MSC) in-charter deliveries of spare parts, mail, or other commodities intended for a customer the receiving supply

ship is scheduled to visit underway. We refer to a “supply ship” as any Combat Logistics Force (CLF) vessel or MSC in-charter vessel that is under our scheduler’s control. We refer to a “customer ship” as any ship, including US Navy ships and those of coalition partners, that we serve from our supply ships.

To support schedule development, information was made available to logistics planners, both at the tactical and operational levels, by fleet schedulers who did not necessarily consider the logistics portions of the problem. In some scenarios, especially those increasing in degree of complexity, time required to produce a complete and viable logistics schedule was not afforded; inefficient and ineffective combatant employment resulted. Therefore, scheduling, from the fleet tactics and the logistics perspectives, needed to be integrated, and then optimized, to ensure ships could sustain combat operations at-sea.

Information becomes available sooner through formalized communication formats with the institution of the Operational Report (OPREP) message, a daily requirement from each combatant and supply ships. OPREPs provide detailed information for operational logistics planning, such as geographic position, fuel states, stores and ordnance inventory levels, and other important information critical to mission sustainment (Diaz 2010). Although the information flow improved, Diaz claims that methods for developing logistics schedules did not, remaining time consuming, labor intensive and prone to human error; a need for optimized planning tools was obvious and thus began the effort to automate scheduling efforts, using computer software.

1.3 CLF Optimization and Planner Development

Optimizing logistics planning first took form using Mixed Integer Programming, in 2001, when Borden implemented them for scheduling CLF CONSOL events and evaluating their force levels and capabilities. More specifically, Diaz states that Borden’s optimization model tries to determine whether the CLF fleet composition in the AOR is capable of sustaining [battle groups] in various, logistically demanding scenarios. This initial look at applying optimization modeling to logistics scheduling sparked interest in developing more robust models capable of identifying not only solutions to singular replenishment events, but multiple events over short and long time horizons. Using predictions of demand, to the highest fidelity possible, helps identify potential pitfalls associated with sustaining planned

operations, over the specified time horizon, so issues or limitations can be mitigated before they take effect.

A second study, performed by J. Cardillo in 2004, analyzed CLF support capabilities in response to a global deployment of all available United States Navy (USN) combatants (Cardillo 2004). Cardillo wanted to illustrate the number and capacities of supply ships that were required to sustain a large-scale maritime contingency operation, while reacting to demands of a second one. His analysis “demonstrates the advantages of planning CLF commodity load outs based on supporting a [battle group’s] forecasted daily requirements vice using the average daily demand data, as the traditional basis for determining fleet [logistics] requirements” (Cardillo 2004; Diaz 2010). Improving plan fidelity, using anticipated commodity consumption rates from projected operating environment factors, strengthens viability of predictions and fidelity of supply ship requirements, since variability in commodity demand is better captured and helps prevent depleting on-hand replenishment stock levels that would cause a myriad of cascading effects. For example, using historical average fuel consumption rates for a ship conducting both submarine search efforts and counter-piracy operations over the course of a two-month period will inadequately represent demand for spikes in fuel consumed when conducting counter-piracy operations; misrepresenting demand at the appropriate times could lead to fuel tanks running dry because on average enough fuel is available, but not the right quantities are available at the right times.

1.3.1 Formal Logistics Optimization Planning Tools

Like many other efforts to develop optimization models, initial efforts leading to Replenishment-At-Sea Planner (RASP) began with exploratory analysis (Brown and Carlyle 2008). Focus started with trying to answer the following questions, among others:

1. What should the logistics support force be to support an operation of size X with mission set Y?
2. How many of each type of logistics support asset will be necessary to support an operation of size X with mission set Y?
3. How much fuel will be required to support an operation of size X with mission set Y?
4. What quantities of supply commodities will be required to support an operation of size X with mission set Y?

CLF Planner and Optimized Transit Tool and Easy Reference (OTTER) were among the first attempts to answer the questions above.

CLF Planner

Brown and Carlyle (2008), at Naval Postgraduate School (NPS), developed the CLF Planner, an optimization model that determines supply ship requirements to meet customer ship replenishment needs. CLF Planner answers questions one and two above and is primarily used to evaluate new supply ship design proposals to see what restrictions to combatant force employment exist (e.g., CLF owned by fleet). CLF Planner was used as a guide to support operational planning.

OTTER

OTTER is a planning tool used by schedulers to calculate fleet fuel requirements. Accounting for fuel consumption by ship class, fuel quantity needs are projected for planned operations scheduled over a given time horizon. OTTER's insights improve fuel savings, allowing analysts to determine optimal transit speeds based on engineering data received from ships in operation. Fuel curves developed are the best prediction tools available, aside from being able to gather necessary information first-hand, in the moment, or relying on other optimization models available.

RASP

RASP is a Mixed Integer Program (MIP) optimization model, implemented in the General Algebraic Modeling System (GAMS) optimization software and uses the C-Language Simplex Method (CPLEX) solver, to generate logistics support schedules which meet customer RAS and/or In-port Replenishment (INREP) requirements. Microsoft Excel[®], via a Visual Basic for Applications (VBA) supported interface, manages user inputs to manipulate RASP functionality. RASP was developed at NPS in 2010, by Dr. G. Brown, Dr. W. Carlyle, and CAPT P. Burson, SC, USN (Brown et al. 2017). Since customer ship schedules are fixed, schedules can only influence supply ship fuel consumption. Therefore, RASP aims to produce schedules that are as cost effective and efficient as possible, i.e., it minimizes supply ship fuel consumption while meeting the operational needs of customer ships as best possible. RASP takes as input all customer employment plans, including locations over time

and anticipated consumption rates for fuel and all consumable commodities. RASP chooses locations in the operating area from which to load supplies required for sustaining customer ships as they execute tasking. RASP considers several real-world constraints, such as ship storage capacities or fuel consumption rates based on type of ship employed. RASP inputs include projected supply commodity consumption rates and supply ships available for tasking. RASP outputs include estimates of supply ship fuel consumed and supply ship underway percentage.

RASP takes as input availability of ports for supply ship replenishments, supply ship availability, and supply commodity replenishment cycles. These data elements may in practice change between the time RASP creates a schedule and when that schedule is executed. RASP, like all optimization models, may amplify even small changes of such inputs to wholesale revisions of its scheduling advice (Brown et al. 1997). Even small changes of such inputs may render a RASP schedule either inefficient or completely infeasible.

1.4 Motivation and Goals

Until recently, fleet operational planners relied upon running RASP, repeatedly, updating replenishment schedules for customer and supply ships when changes occur within the operational environments. This is not always swift. Sometimes changes are small enough, and do not warrant alterations to the current schedule; this is only learned after the fact. However, this is not always the case; certain changes require a complete RASP re-solve, spending time and allocating already constrained computational resources re-doing efforts. Even though solution times can be as short as ten seconds in simple scenarios, suggesting a re-solve might be the best course of action, most scenarios are not so simplistic. Most solutions can take two or more hours to calculate. In these circumstances, it will be useful to have a simple statistical model available to provide initial assessments of how changes may affect the overall solution and advise whether to re-run RASP. Being able to predict supply ship fuel consumption and supply ship utilization (i.e., percentage of days underway) enables planners and leaders to determine if current supply ship utilization can support more customers in their area, or even an increasingly dynamic environment, requiring longer travel distances between Carrier Strike Group (CSG) and INREP ports. With this additional information, leaders can determine potential stress points, or possible break points, for their employed fleet and better plan for changes to requirements.

Sensitivity analysis of the effect changes to RASP inputs have on RASP outputs leads to answering some underlying questions similar to:

- What would happen if a change to the plan occurs? (e.g., a supply ship breaks and is unable to conduct replenishment; or a replenishment port experiences change to diplomatic authorizations, and no longer allows ships under the United States flag to enter)
- Would a completely new solution be necessary? (Depending on the type, or magnitude of change, the easy answer might be: Yes.)
- What if the answer could be “No, and this is how our predictions could change...”? (e.g., If INREP events are not feasible (input), then supply ship fuel consumption could be reduced by “X” amount (output) since fuel is not consumed by multiple ships transiting from operational area and port of replenishment.)
- Without having to wait for a new RASP solution, would decision-makers then become more flexible, quickly altering plans, and capitalizing on windows of opportunity?

This thesis quantifies changes to RASP output (i.e., supply ship fuel consumption, supply ship underway percentage, and maximum percent below safety stock threshold) when its input parameters are perturbed. Perturbations are changes to the status quo; shortening the number of days between replenishment, limiting which supply ships are available, or even limiting port availability for supply ship resupply are all examples. Using linear regression, we determine which perturbations have the greatest effect on RASP outputs. This provides significant insight about schedules RASP produces, specifically their vulnerability to adversarial interdiction (e.g., destruction of an in-port refueling station). Statistical models that predict output quantities from input parameters also enable quick what-if analysis without relying on potentially time-consuming runs of RASP. Regardless of time constraints, most analysts may not be able to utilize RASP, so the statistical models provide a simple and transparent tool for analyzing the effect of deviating from an incumbent CLF plan.

CHAPTER 2:

Key Logistics Concepts

Logistically supporting vessels at-sea has been important since nations first became seaworthy. Sustaining efforts in the open ocean is no easy endeavor, especially when reasons behind ships being underway grow increasingly complex over time. What started as exploration and expansion of influence, morphed into full-fledged sea battles in a few thousand years. Today, food and essentials are not the only things required to keep ships underway; ammunition, fuel, and the tools and supplies necessary for conducting maintenance are also required, especially by a militarized ship. How are these needs met? What solutions exist to maximize mission endurance? How do we measure solution effectiveness to determine if one schedule is better than another? For the United States Navy, the most effective solution is conducting underway replenishment when possible; time spent transiting between the operating area and port of replenishment decreases a customer ship's mission availability. Delivering fuel and supplies underway minimizes off-station transit time and maximizes time on-station for employment. After all, a combat vessel is not designed to be moored to a pier.

Scheduling RAS requirements is simplified when the number of ships is small (five or less); modeling them as one unified group is typical practice, basing the level of support required on the individual customer ship in the group with the shortest endurance. However, when multiple, dynamically employed groups are present, requiring a wider range and increased quantity of supplies, the problem becomes too difficult for a human to solve well.

2.1 Endurance

A ship's endurance is measured as a function of supply commodity consumption rate, the associated storage capacities, and the number of personnel aboard to keep the ship operational (e.g., the watch team maintaining safe voyage cannot stay awake indefinitely, so reliefs are required to mitigate the effects of fatigue). Balancing trade-offs between storage and consumption rates has been an ongoing effort for every sea-going nation, especially those that are militarized. The smaller a ship's crew becomes limits the number of lives

that can be potentially lost; however, assuming the workload required to keep the ship operational remains static, and as the number of available personnel on-board decreases the effort required from each individual increases. Increased strain on personnel introduces risks associated to ship safety (e.g., a misjudged traffic situation by a ship's navigator could result in collision).

Food storage capacity becomes a limiting factor in a ship's endurance just as much as reducing crew size – a lower storage capacity requires a more frequent replenishment cycle, which reduces endurance. However, the primary factor is not food, it is fuel consumption. A ship requires fuel replenishment much more frequently than replenishment of any other commodity.

2.2 Replenishment-At-Sea

2.2.1 What Is Replenishment-At-Sea?

A RAS, defined as “all methods of transferring fuel, munitions, supplies, and personnel from one ship to another while the vessels are underway;” replenishment ships travel most of the distance required to transit between ports and operational environments, instead of the customer (Pike 1999). RAS is also known as Underway Replenishment (UNREP), which consists of a combination of Connected Replenishment (CONREP) (physical connection of hoses and high-tension steel cables to transfer fuel and supplies, and sometimes personnel, between ships underway) and/or Vertical Replenishment (VERTREP) (transfer of supplies and/or personnel via helicopters). Inherently, both evolutions are extremely dangerous, so minimizing time spent conducting them is always the goal. Minimizing the amount of time ships spend physically connected to each other minimizes risk for collision and/or vulnerability to enemy action.

CONREP consists of steel cables, called span-wires, connected between customer and supply ships to transfer fuel and supplies. Span-wires must undergo high-tension to facilitate sending fuel hose and stores transfer assemblies across the 180-foot lateral separation ships maintain during UNREP; each assembly weighs more than 500 pounds, without respective load, so maintaining tension on these steel lines is required for preserving safe conditions and minimizing risk to losing life or damaging equipment.

VERTREP consists of transferring stores from supply to customer ships via helicopter; it is highly inefficient to transfer fuel via this method as helicopters cannot carry large enough amounts of fuel to replenish a ship efficiently. Stores, such as food and other supplies, are transferred using pallets, cargo nets, and rigging assemblies attached to the underside as external loads of the aircraft. VERTREP can sometimes be a quicker method of transferring materiel, depending on space available on the receiving ship.

CONSOL events between supply ships are another method to prepare a supply ship for underway replenishment of a customer ship. CONSOL events between supply ships occur prior to RAS with a customer to limit the number of events the customer must undergo to receive material; time spent conducting UNREP limits time available to conduct mission sets. CONSOL events may be required if customer materiel is loaded onto separate supply ships because required amounts exceed current capacity available for onload, customers transferring materiel between each other, or stores are delivered pier-side too late to be received by the customer prior to getting underway again.

2.3 The Combat Logistics Force

2.3.1 What Is the Combat Logistics Force?

The Military Sealift Command (MSC) (a supporting command to United States Transportation Command (TRANSCOM)) employs United States Department of Defense (DOD) civilian and military personnel to operate its CLF. The CLF executes RAS events required to maximize customer ship endurance. The MSC's CLF "consists of approximately 30 special transport ships that carry ship and aircraft fuel, ordnance, dry stores, and food, and deliver these to customer ships underway, making it possible for our naval forces to operate at sea for extended periods," and in some cases, indefinitely (Brown and Carlyle 2008). Methods and hardware used since they were developed in the 1930s have been continually improved and "permit our navy today to operate continuously for extended periods at sea without returning to any port" (Brown and Carlyle 2008). Today, the CLF fleet is comprised of United States Naval Ship (USNS) ships capable of supporting multi-national sea-going customer ships, but primarily serve United States Ship (USS) vessels.

The United States Navy, like only a couple of other countries, relies on underway replen-

ishment as its primary solution. MSC performs this capability.

2.3.2 Why Model It?

It is important to assess capabilities of the CLF based on its projected operating environment. Modeling its capabilities, in their entirety, allows leveraging machine processing power to identify shortfalls in operational plans and scheduling. Completely modeling many details gives logistics planners the ability to determine whether supporting the intended schedule is feasible. Using projected supply consumption rates as inputs, planners can provide replenishment schedules to decision-makers. It is noteworthy that these models experience decreasing fidelity as the time horizon lengthens. Ultimately, schedules produced by any model are projections, since describing every aspect of controlled and uncontrollable behavior such as weather or unexpected equipment failure is not possible.

Further, modeling supply ship support allows decision-makers to execute some “what-if?” analysis. For example, one can consider the effect of losing a supply ship to enemy action, or perhaps a logistics support hub. Conducting analysis of this manner enables planning for worst case scenarios, with respect to specific operational environments.

2.4 Logistics Supply and Demand

Supply, as defined in logistics, is “the quantities of goods or services offered for sale at a particular time or at one price” (Supply 2022). Conversely, demand is “the quantity of a commodity or service wanted at a specified price and time” (Demand 2022). In order to ensure customer ships remain at maximum mission readiness, commodity demand shall never be greater than supply; in other words, the amount of supply on-hand for a given commodity should be enough to accommodate demands from customer ships.

2.5 Fleet Composition and Density

When conducting underway operations, Carrier Strike Group (CSG) commanders concern themselves with the composition and density of their forces. Knowing types and numbers of ships available and where they are located is required for proper force employment as well as determining the type and amount of logistics support required to conduct operations.

Inadequate and late information, in this respect, can subject forces to unnecessary hardship, lead to equipment malfunction, and death of personnel.

2.5.1 Fleet Composition

Composition refers to the general makeup of the CSG; makeup refers to type and number of ships in the CSG, type and number of aircraft aboard each ship, and personnel aboard each ship who operate, maintain, and support ship systems and pilot aircraft.

2.5.2 Fleet Density

Formally, density is defined as “the average number of individuals or units per unit of space” (Density 2022). Density, in the context of RASP and United States Navy vessels, refers to how close CSG assets are grouped with respect to each another; more dense means ships are operating close together (within a few nautical miles of each other) and less dense means the opposite.

2.6 Sea Routes

Sea routing refers to the locations and paths used by ships to traverse the global waterways. Restrictions are identified with respect to land masses, underwater hazards, and restricted bodies of water. Generally speaking, geographic obstacles to navigation, allied nations, and international waters construct boundaries for the feasible set of way points and routes ships can use to execute their voyage between points in the ocean.

2.7 Fleet Areas of Responsibility

The world is divided by the US Navy into geographical areas to separate operational environments and provide opportunities for leaders to command forces located in the area to best fit the area’s needs; not every area has the same requirements and forces must have a level of oversight capable of affecting change, as necessary, with respect to specifics that may be unique to a given area. The biggest difference is the location of INREP ports depending on the region. The DOD has established an AOR for each of these geographic sections of the world; the USN “has [further] divided the world’s oceans into administrative divisions,” with each being a numbered fleet (Brown et al. 2017). Brown et al. (2017) explain that 2ND

Fleet refers to the Western Atlantic Ocean; 3RD Fleet the Eastern Pacific Ocean; 5TH Fleet the Red Sea, Indian Ocean, and Arabian Gulf; 6TH Fleet the Eastern Atlantic Ocean and Mediterranean Sea; and 7TH Fleet the Western Pacific Ocean. Figure 2.1 gives a pictorial representation of these divisions.

Our example is from US Navy 7TH Fleet, the largest geographic expanse of all our fleet areas for which distance is a key consideration. By contrast, other areas such as 5TH Fleet around the Arabian Gulf or 6TH Fleet around the Mediterranean provide much less geographic expanse (Brown et al. 2017).




Figure 2.1. US Navy Numbered Fleet Areas of Responsibility (AORs) World-wide: Commanders have areas of responsibility in which they are responsible for all ship operations. Replenishment of customer ships is a task assigned to the commander's staff; use of RASP is standard by almost all fleet planning staffs. Source: Brown et al. (2017).

2.7.1 How Does RASP Consider Composition and Density?

To determine feasibility of meeting demand, RASP considers the number of supply ships available in the AOR, their storage capacities and the respective customer ship demand by commodity group. RASP identifies closest replenishment port to each customer ship, and calculates distances between them, to construct constraints for identifying minimum fuel consumed by supply ships in the schedule produced by the optimization model. To RASP, composition of the CSG only determines the types of commodities that will be demanded

by customer ships; each of these commodities are predefined input parameters defined by consumption rates depending on the operational tempo. Density is slightly different; the relative proximity of ships to each other determines where and when RAS events are scheduled.

2.8 RASP Model Formulation

RASP is a Mixed Integer Linear Program, optimization problem that schedules supply ship utilization to execute customer ship replenishment events, while minimizing fuel costs. The following formulation is reproduced from (Brown et al. 2017).

Indices and Sets [~cardinality]

- $g \in G$ Combatant strike group [~30]
Each group is composed of one or more combatants traveling in company
- $d \in D$ Day of planning horizon, a contiguous ordinal set [~45 - 360] (alias δ)
- $t \in T$ Time period, a contiguous ordinal set (alias τ)
May be a day, or some fraction of a day. $|T| = k|D|$ for some integer $k \geq 1$
For example, with six four-hour watches per day, $k = 6$
- $d(t)$ Day of time period t , $d(1) = 1, d(k) = 1, d(k + 1) = 2, \dots$
- $t \in T_d$ Set of time periods during planning day d
- $s \in S$ Shuttle ship [~10]
- $c \in C$ Commodity group (DFM, JP5, DRY, FRZ, CHL, UHT, H2O, ORDN) [~8]
- $p \in P$ Port [~25]
- $a \in A$ Set of potential actions for a shuttle ship at any location. $A = \{POS, LOG\}$, indicating, respectively, that a shuttle must simply be in a given location on the start of a time period, or that the shuttle ship will have an opportunity for a logistic event with some strike group g .
- $l \in L$ Commodity level (e.g., SAFETY, EXTREMIS, NEGATIVE), an ordinal set

* Diesel Fuel Marine (DFM) is the fuel used by United States combatant and support vessels and Jet Propellant-5 (JP5) is fuel used in United States Naval Aviation. DRY, FRZ, CHL, UHT, and H2O refer to the different food and supply types a ship needs to support its crew. ORDN refers to the ordnance required for the ship to perform its duties (Brown et al. 2017).

Section 2.8 delineates the number of ships, by type and class, and the number of battle groups, including their composition, RASP considers. Additionally, the time horizon specified is divided into individual days for scheduling events. Identifying specific days individually within RASP enables the scheduler to establish windows for event planning (Brown et al. 2017). Lastly, the commodities to be replenished are modeled separately, keeping track of deliveries individually so that costs can be more accurately stratified.

Derived Indices and Sets [$\sim$ cardinality]

- $v \in V_s$ Voyages for shuttle s [$\sim 10,000$]
- $\{g, t\} \in GT_v$ For voyage v , two-tuples of combatant strike group g time period t RAS rendezvous
- $\{t, p\} \in TP_v$ For voyage v , two-tuples of time periods with port p visits

Provided Data [units]

- $lat_{g,d}, lon_{g,d}$ Coordinates of combatant strike group g at start of day d [degrees]
- $RAS_OK_{g,t}$ = 1 if RAS permissible during time period t , 0 otherwise [binary]
- $window_g$ Minimum number of days between RAS events (from any shuttle) for combatant g [days]
- s_lat_s, s_lon_s Initial coordinates of shuttle ship s [degrees]
- min_speed_s, max_speed_s Minimum, maximum speed of shuttle ship s [knots]

$fuel_s(speed)$	Fuel consumption as a function of shuttle speed. Standard navy fuel consumption tables are in gallons per hour or barrels per day versus speed in knots. The function here maps knots to [fuel units]
$inptTAT_s$	Inport turn-around time to reload shuttle ship s [days]
p_{lat_p}, p_{lon_p}	Coordinates of port p [degrees]
$x_{lat_{s,t,a}}, x_{lon_{s,t,a}}$	Coordinates fixed by the scheduler for shuttle s to occupy at the start of time period t to perform action a [degrees]
$g_{lat_{g,t}}, g_{lon_{g,t}}$	Coordinates of strike group g at the start of time period t [degrees]
$fuel_cost$	Cost of shuttle own-fuel, shuttle diesel fuel marine (DFM) [\$/fuel unit]
$port_cost_{s,p}$	Cost of a visit by shuttle s to port p [\$]
$g_uses_{g,t,c}$	Consumption by g during time period t of commodity c [c-units]
$g_mxload_{g,c}$	Maximum capacity of g to carry commodity c [c-units]
$g_starting_c_{g,c}$	Inventory at start of planning horizon of commodity c [fraction of $mxload_{g,c}$]
$g_limit_c_{g,c,l}$	Commodity limit triggering a shortage violation (i.e., safety stock) [c-units]
$g_penalty_c_{g,c,l}$	Positive multiplicative penalty for a shortage violation [\$/c-unit violation]
$g_priority_g$	Weight assigned to RAS volume delivered to strike group g [scalar]
npv_t	Net present value discount term [fraction]; this term is often referred to as the “fog of future planning” discount.
$s_capacity_{s,c}$	Shuttle ship s capacity for commodity c [c-units]
$s_init_load_{s,c}$	Shuttle ship s initial inventory of commodity c [fraction of $s_capacity_{s,c}$]
$pier_s$	Pier capacity used by shuttle ship s [pier capacity]

$pier_cap_p$	Port capacity [pier capacity]
$c_cost_{p,c}$	Commodity c cost at port p [\$/c-unit]
$c_priority_c$	Priority of commodity c [scalar]
$reward_c$	Reward for delivery of commodity c [\$/c-unit]

Derived Data

v_cost_v	Voyage costs (excluding loaded commodity costs) [\$]
$fuel_burned_{v,t}$	Own-fuel burned by shuttle s on voyage $v \in V_s$ during time period t [fuel units]

Provided and derived data refers to specific values regarding fuel consumption, as a function of ship class and speed (Brown et al. 2017). Transit times, between ship locations and ports of replenishment, as well as time required to load stores and fuel in-port, are calculated to determine schedule feasibility. Fuel and port costs are applied to replenishment events to appropriately affect the constraints of the objective function. Supply and customer ship storage capacities, by commodity, are fed into RASP to determine replenishment event scheduling requirements; the time needed to load and ability to physically store quantities required are inputs RASP also considers determining feasibility of events occurring. Lastly, RASP accounts for safety stock limits, which are specified by combatant commander, to calculate penalties associated with event scheduling. Going below a designated threshold for a given commodity triggers a penalty in the objective function. “RASP maintains a catalog of customer ships with their commodity capacities and consumption rates for a variety of employment activities” which are compiled to determine quantities required, per ship, and schedule event cycles accordingly (Brown et al. 2017).

Decision Variables

$VOYAGE_v$	Binary indicator that shuttle voyage v is selected
$VISIT_{g,d}$	Binary indicator that at least one shuttle visits g on day d
$LOAD_{s,t,p,c}$	Amount of commodity c loaded by shuttle s at start of time period t at port p [c-units]
$HOLD_{s,t,c}$	Shuttle s commodity c contents at start of time period t [c-units]
$RAS_{s,g,t,c}$	Amount of shuttle s delivery to g during time period t of commodity c [c-units]
$VIOLATION_{g,t,c,l}$	Amount of inventory deficiency of c for g , at start of planning period t below level l [c-units]

Decisions made include specific paths (voyages) ships could take to get from current location to the location of future tasking. Initial commodity load levels are specified, providing starting conditions for the RASP to determine replenishment timelines based on usage rates. Decisions also include quantities, by commodity, delivered to customer ships as well as quantities that are not delivered; RASP will minimize the amount not delivered, which is the difference between quantity requested and available for delivery.

Formulation

$$\begin{aligned}
& \text{MIN} \\
& \text{VOYAGE, VISIT,} \\
& \text{LOAD, HOLD,} \\
& \text{RAS, VIOLATION} \\
& \sum_{\substack{s \in S \\ v \in V_s}} v_cost_v VOYAGE_v + \sum_{\substack{s \in S, t \in T, \\ p \in P, c \in C}} c_cost_{p,c} LOAD_{s,t,p,c} \\
& + \sum_{\substack{g \in G, t \in T, \\ c \in C, l \in L}} npv_t g_penalty_c_{g,c,c_level} VIOLATION_{g,t,c,l} \\
& - \sum_{\substack{s \in S, G \in G, \\ t \in T, c \in C}} g_priority_g * c_priority_c reward_c RAS_{s,g,t,c} \\
& + \text{elastic penalties}
\end{aligned} \tag{1}$$

$$\begin{aligned}
s.t. \quad & HOLD_{s,t,c} - \sum_{v \in V_S | c = \text{'DFM'}} fuel_burned_{v,t} VOYAGE_v \\
& - \sum_{g \in G} RAS_{s,g,t,c} + \sum_{p \in P} LOAD_{s,t,p,c} \stackrel{\circ}{=} HOLD_{s,t+1 | t < \|T\|} \quad \forall s \in S, \\
& \quad \quad \quad t \in T, c \in C \quad (2)
\end{aligned}$$

$$\sum_{\substack{s \in S, \\ \tau < t}} RAS_{s,g,\tau,c} \stackrel{\circ}{\leq} \sum_{\tau < t} g_uses_{g,\tau,c} + [g_mxload_{g,c}(1 - g_starting_cg, t, c)]_{t=1} \quad \forall g \in G, \quad t \in T, c \in C \quad (3)$$

$$\begin{aligned}
& \sum_{\substack{s \in S, \\ \tau < t}} RAS_{s,g,t,c} + \sum_{l \in L} VIOLATION_{g,t,c,l} \\
& + g_starting_c_{g,c} mxload_{g,c} \\
& \geq \sum_{\tau < t} g_uses_{g,\tau,c} - g_mxload_{g,c} (1 - g_limit_{g,c}, 'SAFETY') \\
& \forall g \in G, t \in T, c \in C
\end{aligned} \tag{4}$$

$$RAS_{s,g,t,c} \leq \min\{g_mxload_{g,c}, s_capacity_{s,c}\} \sum_{v \in V_S | \{g,t\} \in V_S} VOYAGE_v \quad \forall s \in S, \forall g \in G, t \in T, c \in C \quad (5)$$

$$\sum_{v \in V_S} VOYAGE_v \leq 1 \quad \forall s \in S \quad (6)$$

$$\sum_{v \in V_S | \{g,t\} \in GT_v \wedge t \in T_d} VOYAGE_v \leq VISIT_{g,d} \quad \forall g \in G, d \in D \quad (7)$$

$$\sum_{d \text{--} window_g \leq \delta \leq d} VISIT_{g,\delta} \leq 1 \quad \forall g \in G, d \in D \quad (8)$$

$$LOAD_{s,t,p,c} \leq \sum_{v \in V_S | \{t,p\} \in TP_v} s_capacity_{s,c} VOYAGE_v \quad \forall s \in S, t \in T, p \in P, c \in C \quad (9)$$

$$\sum_{\substack{s \in S, \\ v \in V_S | \{t,p\} \in TP_v}} pier_s VOYAGE_v \leq pier_cap_p \quad \forall t \in T, p \in P \quad (10)$$

$$VOYAGE_v \in \{0, 1\} \quad \forall s \in S, v \in V_S$$

$$VISIT_{g,d} \in \{0, 1\} \quad \forall g \in G, d \in D$$

$$0 \leq LOAD_{s,t,p,c} \leq s_capacity_{s,c} \quad \forall s \in S, t \in T, p \in P, c \in C$$

$$0 \leq HOLD_{s,t,c} \leq s_capacity_{s,c} \quad \forall s \in S, t \in T, c \in C$$

$$HOLD_{s,1,c} = s_init_load_{s,c} s_capacity_{s,c} \quad \forall s \in S, c \in C$$

$$0 \leq RAS_{s,g,t,c} \leq \min\{g_mxload_{g,c}, s_capacity_{s,c}\}$$

$$\forall s \in S, g \in G, t \in T, c \in C$$

$$0 \leq VIOLATION_{g,t,c,l} \leq g_mxload_{g,c}$$

$$\times (g_limit_c_{g,c,l} [-g_limit_c_{g,c,l-1}]_{l>1})$$

$$\forall g \in G, t \in T, c \in C, l \in L \quad (11)$$

The objective function expresses fuel cost associated with execution of replenishment events scheduled. All feasible events are scheduled, and penalties will be assessed for missed events, quantities not delivered, and percentages of each commodity that are below safety stock.

The quantities, per commodity, cannot exceed the customer ship's storage capacity nor the

usage quantity for the period of the individual replenishment event; exceptions can be made if a customer ship's initial conditions begin a scenario with storage capacities are not at 100%. Quantities delivered, by commodity, must cover safety stock plus amounts required for the customer ship to operate until the next replenishment. Customer ships must replenish each cycle, via RAS or INREP and cannot replenish outside of a designated event window. This allows for deliberate tracking of replenishment events to determine fleet effectiveness in meeting demands (Brown et al. 2017).

CHAPTER 3:

Statistical Sensitivity Analysis

3.1 Introduction

We conduct statistical sensitivity analysis to quantify the effects varying RASP inputs have on schedules produced from RASP. Results of the analysis performed will be used to develop intuition regarding schedule feasibility and establish statistical models that will enable operational logistics planners to provide timely recommendations to decision-makers. Daily, planners are posed with many questions like the following:

- What would occur if the number of days before requiring replenishment is changed?
- If there is a change to supply ship availability, or ports available for replenishment, what effect would that have?
- If customer ships are no longer allowed to pull into port to load fuel and stores, would the supply ship experience savings in fuel consumed?

RASP often provides the answer to all these questions but can require more time than available to do so. Therefore, it is our effort to provide statistical models to planners, giving them the means to perform “back of the envelope” calculations should questions arise and running RASP again is not feasible.

3.2 Statistical Sensitivity Analysis

Statistical sensitivity analysis, in our application, is a means to measure how robust RASP solutions are to changes to inputs and how changes affect results. The purpose of statistical sensitivity analysis is to find connections between model inputs and outputs that develop intuition and explain which inputs are most influential to outputs (Glen 2022). This type of analysis is also useful because explaining as much uncertainty in the data as possible helps improve prediction capability (Pichery 2022).

3.3 Datasets: Scenario CSG-1 and Scenario CSG-2

This thesis analyzes two datasets, one supports a one-CSG scenario, the other a two-CSG scenario; they are referred to as Scenario CSG-1 and Scenario CSG-2 respectively. Both scenarios are concerned with CSGs, leaving from either Bremerton, WA or San Diego, CA, and their operational employment in the Pacific Theater. Typically a deployed CSG is comprised of large and small combatants, submarines, and the embarked air wing aircraft. However, RASP is primarily concerned with fuel consumption, i.e., ship fuel (DFM) and navy aviation fuel (JP5), thus submarines are not modeled since their fuel usage is comparatively minimal. Additionally, supply ships may accompany the CSGs to provide logistics support; the supply ships become the primary focus. The following is the composition of ships in a CSG, by class and quantity, that are the focus of each scenario:

Carrier Strike Group (Scenario CSG-1 and CSG-2)

- One Nimitz Class Aircraft Carrier (CVN) (Figure 3.1)
- One Ticonderoga Class Guided Missile Cruiser (CG) (Figure 3.2)
- Two Arleigh Burke Class Guided Missile Destroyers (DDG) (Figure 3.3)
 - Nominally, a DDG consumes five to 10% of its fuel capacity daily during normal underway transits depending on speed.



Figure 3.1. Nimitz Class Aircraft Carriers: USS John C. Stennis (CVN 74) and USS Carl Vinson (CVN 70). Source: Mizokami (2021).



Figure 3.2. Ticonderoga-Class Cruiser: USS Chancellorsville (CG 62). Source: Zaffar (2021).



Figure 3.3. Arleigh Burke-Class Destroyer: USS Gravely (DDG 107). Source: Sea Forces (2022).

Also in the scenario, there are four or three supply ships modeled based on Scenario CSG-1 or Scenario CSG-2 respectively.

Supply Ships

- Henry J. Kaiser Class Replenishment Oiler (Figure 3.4)
- Royal Australian Navy Fleet Replenishment Oiler (Figure 3.5)



Figure 3.4. Henry J. Kaiser Class Fleet Replenishment Oiler. Source: Naval Technology (2022).



Figure 3.5. Royal Australian Navy Replenishment Oiler: HMAS Sirius. Source: Royal Australian Navy (2022).

Statistical Model Inputs Included

- CSG CYCLE: The number of sustainment cycle days between required customer ship replenishment (based upon fuel consumption rates). Sustainment cycle refers to the number of days a ship can self-sustain before needing replenishment.
- CLF AVAIL: The specific supply ships available to support RAS events in the schedule.
- CLF REGION: The constraints placed on supply ships supporting replenishment schedules. We refer to “regional constraints” as restrictions placed on a supply ship

that constrain its operating region. “TETHER” indicates supply ships are required to remain within boundaries of their region assigned. “FLOAT” indicates supply ships are authorized to transit outside boundaries of their region by a specified number of nautical miles.

- **PORT DELAY:** The number of days a supply ship is delayed from returning to sea after INREP.
- **PORT AVAIL:** The specific ports available to customer and supply ships for replenishment.
- **REP TYPE:** Specifies whether INREP is available to customer ships, or if RAS is the only option.

CLF AVAIL and PORT AVAIL are converted to categorical variables to properly account for the appropriate relationship between values.

Statistical Model Outputs Included

- **CLF UW Days (%):** The percent of time supply ships are scheduled to be underway over the specified time horizon.
- **Estimated Supply Ship Fuel (Gallons):** The amount of fuel, in gallons, the supply ships are estimated to burn executing the schedule produced; this value includes fuel provided to customer ships.
- **Maximum Percent Below:** The maximum percentage of stores level below the specified safety stock experienced by any one customer ship considered by RASP.

3.4 Statistical Models


Multi-Variable Linear Regression (MVLR) and classification lead to predictive statistical models for supply ship fuel consumption, supply ship underway percentage, and maximum percent below the safety stock in both scenarios. The goal is to construct a statistical model to predict RASP outputs; planners benefit from the ability to get close estimates of model outputs of interest, without requiring a new optimized solution. Least squares regression allows comparison of statistical model predicting power, using different model input interactions; we seek as simplified a statistical model as possible.

Only linear interactions between model inputs provide baselines from which to reduce Root

Mean Square Error (RMSE) and are the basis for statistical models offered for calculations performed by hand. Cross-input interactions are included in this thesis as well; however, certain interactions reduce interpretability (e.g., multiplying CSG CYCLE with itself is not intuitive but aids in explaining more variation in a dataset and reducing model prediction error).

3.5 One-Hot Encoding

One-hot encoding is a method of handling categorical variables used for linear regression; this encoding converts categorical variable values so machine learning algorithms can attempt to improve predictions. Categorical variables represent data that is not numeric in nature; they are labels (i.e., values for the variable are words such as “blue” or “dog”) but can sometimes have a natural ordering. For example, clothing has a natural order to sizing (e.g., small, medium, and large) but neither value indicates a “better” one for predictability. Since words inherently do not have worth, categorical variables must have their values translated into integers for use by an optimization model. One-hot encoding transforms categorical variable values into the binary ones: “1” and “0,” creating separate columns for each of the possible values of the original categorical variable. An example where this thesis uses one-hot encoding is when PORT AVAIL is converted into separate columns, each representing a possible situation, such as Pearl Harbor, Hawaii or Yokosuka, Japan not being available. Figure 3.6 provides an example of conversions.



Type	AA_Onehot	AB_Onehot	CD_Onehot
AA	1	0	0
AB	0	1	0
CD	0	0	1

Figure 3.6. One-Hot Encoding Representation: This is a pictorial representation of the output of one-hot encoding. Categorical variables are transformed into a subset of columns, one for each possible value within the original categorical variable. Adapted from Fawcett (2022).

One hot encoding is useful for categorical data where there is no implied significance between possible values; typically machine learning models will naturally assign importance

to higher numerical values, and this isn't always true (Fawcett 2022).

3.5.1 Multi-Variable Linear Regression

MVLR is a statistical technique used to predict dependent variable values based on inputs from two or more independent variables e.g., (James et al. 2013). MVLR generally takes the following form:

$$y_i = \beta_0 + \sum_{j=1}^n \beta_{0j} x_{ij} + \sum_{j=1}^n \sum_{k=1}^n \beta_{jk} x_{ij} x_{ik} + \epsilon_i \quad | \quad \forall j \leq k \quad (3.1)$$

where:

- y_i is the output value
- β_0 is the intercept, i.e., the value of y when x_{ij} and x_{ik} are all 0
- x_{ij} and x_{ik} is the j -th or k -th input of the i -th observation respectively
- β_n are the regression coefficients representing the change in y relative to a one-unit change in respective x_{ij}
- $x_{ij}x_{ik}$ are 2-way interactions between model inputs
- ϵ_i is the statistical model's random error (residual) term

3.5.2 Statistical Classification

Classification is a statistical method, within supervised learning, that determines if an observation “belongs” to a specific group or category. Data with known class labels trains the statistical model, developing a classification rule used to assign new data to one of the classes e.g., (Pizer et al. 2022). Statistical classification is among the most structured of supervised learning methods, where rules are essentially dictated by the user, prior to starting the model, instead of being left open for the method to develop its own classifications, e.g., (James et al. 2013).

A multitude of classification methods exist supporting supervised learning. All consider as little as two groups of observations; the maximum number of groups is limited by computational resources available, and the dataset being analyzed. Support Vector Machine (SVM), a method within supervised learning, is a machine learning algorithm that classifies observations into one of two categories. The algorithm seeks a hyperplane that distinctly classifies data, maximizing the margin between all possible data points and the hyperplane

itself. This thesis used SVM for classifications. Figure 3.7 offers a pictorial representation where of a classification model input.




Figure 3.7. Support Vector Machine: SVM seeks a hyperplane that separates data points into two classes of data, while maximizing the margin around each available data point. The margin around each point is based upon a distance calculation determined by how the SVM algorithm is implemented. Source: Ghandi (2018).

In general, an SVM formulation is described by Royset and Wets (2021) as the following:

$$\min_{a,b} \frac{1}{N} \sum_{i=1}^N \max\{0, 1 - y^i(a^\top x^i + b)\} + \rho \|a\|_1,$$

where x^i is the vector of model inputs associated with observation i and y^i is the label assigned to observation i . Here the last term involves a non-negative tuning parameter ρ that encourages sparse solutions (i.e., the vector a has many zero components). The problem can be reformulated as:

$$\begin{aligned}
& \min_{a,b,z,u} \frac{1}{N} \sum_{i=1}^N z_i + \rho \sum_{j=1}^n u_j \\
& \text{s.t. } 1 - y^i(a^\top x^i + b) \leq z_i \text{ for } i = 1, \dots, N \\
& \quad 0 \leq z_i \text{ for } i = 1, \dots, N \\
& \quad -a_j \leq u_j, a_j \leq u_j \text{ for } j = 1, \dots, n.
\end{aligned}$$

This linear problem can be solved efficiently and reliably using simplex or interior point methods for problem instances considered in this thesis. For optimal a and b , an input x will be predicted as having label “1” if $a^\top x + b > 0$ and “−1” if $a^\top x + b < 0$.

3.6 Assessing Statistical Model Accuracy

All predictions have a certain degree of uncertainty associated with them. Typically, fidelity of model outputs is measured in terms of R^2 and RMSE. An R^2 value represents the proportion of variation explained by the inputs in the statistical model, and their interactions with each other. R^2 is a value between 0 and 1; see e.g., (James et al. 2013). RMSE measures overall error present. This thesis uses both R^2 and RMSE to compare sub-cases of the main scenarios to each other.

3.7 Scenarios Constructed

Most recent versions of RASP, beginning with version 2.10, include a batch processing functionality, greatly reducing hands-on requirements for multiple scenarios. RASP batch processing allows multiple permutations to be performed and is only limited by the maximum number of scenario values that can be managed through the “Scenario Manager” functionality within Microsoft Excel (Microsoft Technical Support 2022).

Establishing a framework for developing decision rules was initialized through Scenario CSG-1. Varying input values results in approximately 50,000 exhaustive permutations which are input into RASP to produce corresponding outputs. Solve times are between four and 52 seconds, and constructing the dataset took an aggregate of approximately 72 hours. Table

3.1 describes all inputs and the levels for each used to create the dataset.

Table 3.1. Overview of Inputs Used in Scenario CSG-1

Input (Type of Variable)	Levels and Description
CSG CYCLE (Continuous)	2, 3, 4, 5, 6, 7, 8; CSG CYCLE represents the number of sustainment cycle days between required customer ship replenishment (based on fuel consumption rates).
CLF AVAIL (Categorical)	ALL, NO GDL, NO PEC, NO JEC, NO SIRI, GDL = USNS GUADALUPE, PEC = USNS PECOS, JEC = USNS JOHN ERICSSON, SIRI = HMAS SIRIUS CLF AVAIL represents the specific supply ships available to support RAS events in the schedule.
FBE DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; FBE DELAY represents the number of days a supply ship is delayed from returning to sea after INREP.
GUM DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; GUM DELAY represents the number of days a supply ship is delayed from returning to sea after INREP.
PLH DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; PLH DELAY represents the number of days a supply ship is delayed from returning to sea after INREP.
SDO DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; SDO DELAY represents the number of days a supply ship is delayed from returning to sea after INREP.
PORT AVAIL (Continuous)	ALL, NO YOK, NO SUB, NO SIN, NO SAS, NO GUM, NO PLH, NO FBE, NO TAO FBE, NO TAO GUM, NO TAO PLH, NO TAO SIN, NO TAO SUB, NO TAO YOK, NO TAO SAS, NO TAO SDO; PORT AVAIL represents the specific ports available to customer and supply ships for replenishment.
REP TYPE (Continuous)	1, 2; 1 = BOTH, 2 = RAS ONLY; REP TYPE specifies whether INREP is available to customer ships, or if RAS is the only option.

Varying input values for Scenario CSG-2 results in approximately 100,000 exhaustive permutations and produces a dataset that represents a wide range of situations. This dataset supports future research and “what if...?” cases without needing to gather additional data elements. Solve times are between five and 512 seconds, with 86% of solves occurring within 10 seconds. Growing this dataset took approximately 6 weeks to produce a solution for all permutations. Table 3.2 describes all inputs and the levels for each used to create the dataset.

Table 3.2. Overview of Inputs Used in Scenario CSG-2

Input (Type of Variable)	Levels and Description
CSG CYCLE (Continuous)	2, 3, 4, 5, 6, 7, 8; CSG CYCLE represents the number of sustainment cycle days between required customer ship replenishment (based on fuel consumption rates).
CLF AVAIL (Categorical)	X1 = ALL, X2 = NO GDL, X3 = ONLY GDL, X4 = NO PEC, X5 = NO JEC, X6 = ONLY PEC, X7 ONLY JEC; Each input has values of 0 and 1; CLF AVAIL represents the specific supply ships available to support RAS events in the schedule.
CLF REGION (Categorical)	1, 2; 1 = TETHER: Supply ships are required to remain within boundaries of their region assigned; 2 = FLOAT: Supply ships are authorized to transit outside boundaries of their region by a specified number of nautical miles; CLF AVAIL represents the specific supply ships available to support RAS events in the schedule.
SDO DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; SDO DELAY represents the number of days a supply ship is delayed from returning to sea after INREP in San Diego, California.
PLH DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; PLH DELAY represents the number of days a supply ship is delayed from returning to sea after INREP in Pearl Harbor, Hawaii.
GUM DELAY (Continuous)	0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10; GUM DELAY represents the number of days a supply ship is delayed from returning to sea after INREP in Guam.
PORT AVAIL (Continuous)	Y1 = ALL, Y2 = NO SDO, Y3 = NO PLH, Y4 = NO GUM, Y5 = NO YOK, Y6 = NO SAS, Y7 = NO SIN, Y8 = NO FBE NO TASO FBE, NO TAO GUM, NO TAO PLH, NO TAO SIN, NO TAO SUB, NO TAO YOK NO TAO SAS, NO TAO SDO; PORT AVAIL represents the specific ports available to customer and supply ships for replenishment.
REP TYPE (Continuous)	1, 2; 1 = BOTH, 2 = RAS ONLY; REP TYPE specifies whether INREP is available to customer ships, or if RAS is the only option.

3.8 Customary Discussion of Optimization Model

For Scenario CSG-1, the RASP formulation has 3,294 continuous variables, 2,012 discrete variables, and 6,084 constraints. For Scenario CSG-2, the RASP formulation has 3,222 continuous variables, 7,475 discrete variables, and 5,818 constraints. A typical model run on a Intel Xeon 6230R (2.10 GHz) machine with 128 GB of memory has a Passmark (2022) rating of 33,733.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:

Statistical Model Predictive Capability

This chapter focuses on results of statistical sensitivity analysis. Statistical models presented suggest predictability of outputs is possible; however, future analysis may improve prediction capability of the statistical models. Results of predicting fuel consumed by supply ships are discussed first, followed by supply ship underway percentage. The discussion then shifts from addressing Scenario CSG-1 (Cases A through A.2) to Scenario CSG-2 (Cases B through B.5.2) which again focus on predicting fuel consumption. Finally, a SVM model classifying maximum percent below the safety stock threshold (Case C) is included, providing an additional statistical model for quick-analysis calculations.

For reference, Table 4.1 provides an overview of the cases and sub-cases used in the statistical sensitivity analysis contained within this thesis. All cases try to predict observations related to a five-day logistics sustainment cycle over a 60-day time horizon.

Table 4.1. Overview of Cases Explored.

Scenario	Case	Method	Output	Statistical Model Input Limitations
CSG-1	A.1	Regression	CLF Fuel Consumption	NO OMISSIONS
CSG-1	A.2	Regression	CLF UW Days %	NO OMISSIONS
CSG-2	B	Regression	CLF Fuel Consumption	NO OMISSIONS
CSG-2	B.1	Regression	CLF Fuel Consumption	CLF AVAIL: ONLY GDL, ONLY PEC, ONLY JEC
CSG-2	B.2	Regression	CLF Fuel Consumption	PORT DELAY: SDO Delay (< 3 Days)
CSG-2	B.3	Regression	CLF Fuel Consumption	PORT DELAY: PLH Delay (1,2, > 4 Days)
CSG-2	B.4	Regression	CLF Fuel Consumption	PORT AVAIL: NO GUM
CSG-2	B.5	Regression	CLF Fuel Consumption	CLF AVAIL: ONLY GDL, ONLY PEC, ONLY JEC, RAS ONLY
CSG-2	B.5.1	Regression	CLF Fuel Consumption	CLF AVAIL: ONLY GDL, ONLY PEC, ONLY JEC, RAS ONLY PORT AVAIL: NO GUM, NO YOK
CSG-2	B.5.2	Regression	CLF Fuel Consumption	CLF AVAIL: ONLY GDL, ONLY PEC ONLY JEC, RAS ONLY; PORT AVAIL: NO GUM, NO YOK, NO PLH, NO FBE
CSG-2	C	Classification	Maximum % Below Threshold	CLF AVAIL: ONLY GDL, ONLY PEC ONLY JEC, RAS ONLY; PORT AVAIL: NO GUM, NO YOK, NO PLH, NO FBE

4.1 Scenario CSG-1 Results

4.1.1 Case A.1: Predicting Supply Ship Fuel Consumption

To begin efforts predicting supply ship fuel consumption, a baseline for error must be established; thus, a statistical model is constructed using the raw dataset supporting Scenario CSG-1 – this becomes Case A.1. Categorical variables are transformed, via one-hot encoding as described in Section 3.5, but no exclusions or further changes are made. Without any effort to finesse interactions, the resulting Least Squares Linear Regression explains approximately 86% of variation within the data. Figure 4.1 shows estimated fuel consumption

versus predicted values; generally, performance is good and most of the variability in the dataset is explained. If a second-degree factorial design (i.e., all interactions between inputs are two-way interactions) was used, R^2 improves dramatically from 0.86 to 0.96 and RMSE improves from 62,041 to 39,145 gallons. This change in regression modeling improves average error by roughly 45% and is therefore a significant improvement. However, this improvement comes at a cost to the model's interpretability.

Figure 4.1 notes RMSE as approximately 62,000 gallons. With an average estimated fuel consumption of 751,840 gallons, this error is approximately 8% of total fuel consumed. This error spans the entire time horizon (60 days) and is technically divided among all ships but is summarized as supply fuel consumption prediction error. Thus the average error, per ship, per RAS event is 277 gallons, a negligible amount. Generally, based on the author's logistics experience an error this low for a fluid commodity value, such as fuel, is acceptable as flow rates are difficult to regulate. Each meter used to measure flow is not calibrated at the same time, nor by the same individual or method; this introduces variability in measurements before even being installed in the refueling station.

There is a fundamental flaw in only considering averages in analysis (Stanford University 2022). Thus, the upper and lower bounds on absolute error are found, using percent difference calculations. In entirety, the maximum and minimum absolute errors become 61.5% and 0.0003% respectively. Omitting outliers representing the extremely infeasible situation where all resupply ports are unavailable through the entirety of the scenario and no supply ships are available to support operations the maximum error reduces to less than 20%. For the sake of completeness, all predictions are included in Figure 4.2, which shows the distribution of error associated.




Figure 4.1. Scenario CSG-1 Fuel Prediction: Predicted Supply Ship Fuel Consumed versus RASP Estimated Values (gallons). The least squares linear regression representing the raw data contains minimal error. An R^2 value of 0.86 is significantly higher than expected without performing typical data massaging or input interaction to improve prediction error. Average absolute percent error is 9.6%.

4.1.2 Fuel Consumption Prediction Error

Figure 4.2 shows most prediction errors fall within $\pm 20\%$ of estimated values; in fact, values are most often under-predicted. Knowing the statistical model under-predicts most often gives planners insight into prediction fidelity. Thus, when budgeting resources, fuel and money, planners understand when predictions need to be supplemented by either adding or subtracting amounts of fuel for a better estimate. Situations producing larger amounts of error are attributed to compounding issues of unavailability (i.e., when ports in the AOR are unavailable, supply ships are unavailable, and cycle days are lower all at the same time, there is a reduction in ability to predict fuel consumed). These scenarios are very rare; thus, we recommend a re-solve of RASP in these situations.




Figure 4.2. Scenario CSG-1 Fuel Consumed Prediction Error: Predicted values come from the statistical model and estimated values are from RASP. Average percent error is -4.6% .

Statistical Model: Predicting Supply Ship Fuel Consumption (Scenario CSG-1)

When predicting supply ship fuel consumption, we offer the following statistical model from our 7TH Fleet AOR example Scenario CSG-1.

$$\begin{aligned}
\text{Fuel Consumed (K Gal)} = & 110.501 - 67.9619 \text{ CSG CYCLE} + 55.1034 \text{ ALL CLF} \\
& + 26.6207 \text{ NO GDL} + 45.393 \text{ NO PEC} - 117.916 \text{ NO JEC} \\
& - 0.3105 \text{ FBE DELAY} - 0.52 \text{ GUM DELAY} \\
& - 3.9244 \text{ PLH DELAY} - 2.5421 \text{ SDO DELAY} \\
& - 1.4988 \text{ ALL PORTS} + 9.396 \text{ NO YOK} \\
& - 2.4208 \text{ NO SUB} - 0.7474 \text{ NO SIN} \\
& + 21.9374 \text{ NO SAS} - 3.567 \text{ NO GUM} \\
& - 109.778 \text{ NO PLH} + 95.9277 \text{ NO FBE} \\
& - 1.3052 \text{ NO TAO FBE} - 1.3264 \text{ NO TAO GUM} \\
& - 1.5864 \text{ NO TAO PLH} - 0.9858 \text{ NO TAO SIN} \\
& - 0.5858 \text{ NO TAO SUB} - 0.3305 \text{ NO TAO YOK} \\
& - 1.4867 \text{ NO TAO SAS} + 5.3019 \text{ REP TYPE},
\end{aligned}$$

where Fuel Consumed is expressed in thousands of gallons; NO TAO__ refers to supply ships with United States flag not being allowed into the respective port; GDL = USNS GUADALUPE, PEC = USNS PECOS, JEC = USNS JOHN ERICSSON, SIRI = HMAS SIRIUS, FBE = Fleet Base East, Australia; GUM = Guam; PLH = Pearl Harbor, Hawaii; SAS = Sasebo, Japan; SDO = San Diego, California; SUB = Subic Bay, Philippines; YOK = Yokosuka, Japan; and SIN = Singapore.

In this specific scenario, we identify that as the number of sustainment cycle days increases the fuel supply ships consume decreases (68,000 gallons for each day in a sustainment cycle). Additionally, when USNS John Ericsson and Pearl Harbor are unavailable the statistical model suggests supply ships generally consume approximately 118,000 and 110,000 gallons less fuel executing the schedule respectively. Whether or not this decrease in fuel consumption is realized is situation dependent. This statistical model is easily implemented in an Excel spreadsheet for convenient use by the planner.

4.1.3 Case A.2: Predicting Supply Ship Underway Percentage

Like CLF fuel consumption, predicting CLF underway percentage is possible in Scenario CSG-1. The statistical model fitting the raw dataset seems to perform better than the one

predicting fuel consumption. Errors observed relate to the same infeasible cases described in Section 4.1.2. Again, no observations were omitted, to portray the full range of predictions.

Figure 4.3 shows an RMSE of 2.8%, which equates to approximately 2 days. The average and median absolute percent difference of all predictions is approximately 21% and 11.7% respectively. This indicates absolute error will typically be roughly 12% (7 days) or less, over the 60-day time horizon. If a second-degree factorial design is used, R^2 improves from 0.87 to 0.96, and RMSE decreases from 2.8% to 1.8%. A full-factorial design sees no additional improvement in R^2 , but an improved RMSE of 1.5%. Although second-degree and full factorial models offer better predictive capability, they reduce interpretability of the model like Case A.1. An alternate view removes calculating absolute value to determine whether predictions tend to be greater or less than estimated values RASP produces on average. Figure 4.4 shows the distribution of absolute percent difference of predictions, while Figure 4.5 shows the distribution of percent difference between estimated values from RASP and predicted values from the statistical model.




Figure 4.3. Scenario CSG-1 Supply Ship Underway Percentage Prediction: Predicted supply ship underway percentage versus RASP estimated values (percent of days underway during a 60-Day Time Horizon). Most error, between 30% and 50% (predicted values), is generated from situations that are highly unlikely in the scenario (i.e., replenishment ports are unavailable and supply ships are constrained in number available).

Figures 4.4 and 4.5 show distribution of prediction error for supply ship underway percentage. Most often, RASP values are under-predicted; this is useful insight. It is now understood that the statistical model in Section 4.1.4 will often under-predict underway percentages so an additional amount could be added to the estimate for more accurate predictions. Without adding the additional time supply ships will be underway, planners may create situations where customer ship replenishment requirements overlap, which can obviously be infeasible if the locations are large distances apart. The percentage added to predictions is dependent upon the proximity of ports when closed in relation to the CSG, coupled with the sustainment cycle of the CSG (e.g., if the closest port available, say Guam, becomes unavailable, and cycle days is four). Percentages added will be a maximum of 10%, decrease to zero as cycle days increase to six, and remain zero as cycle days increase further.




Figure 4.4. Scenario CSG-1 Supply Ship Underway Percentage Absolute Error: The majority of error is less than 20% from estimated values. An upper bound of 20% on absolute error seems good; however, 20% of a 60-day time horizon is 12 days. An error of 12 days is often not acceptable. Average absolute percent error is 21.4%.




Figure 4.5. Scenario CSG-1 Supply Ship Underway Percentage Error: The statistical model under-predicts most often; planners plan for this and supplement predicted values on a case-by-case basis. Average percent error is -13.7%.

4.1.4 Statistical Model: Predicting Supply Ship Underway Percentage

When predicting supply ship underway percentage, we offer the following statistical model from our 7TH Fleet AOR example Scenario CSG-1.

$$\begin{aligned}
 UWDays(\%) = 100 \times & (0.3512 - 0.0203 \textit{CSG}CYCLE + 0.0363 \textit{ALL} \textit{CLF} \\
 & + 0.0113 \textit{NO} \textit{GDL} + 0.0292 \textit{NO} \textit{PEC} - 0.0838 \textit{NO} \textit{JEC} \\
 & - 0.000002 \textit{FBE}DELAY + 0.001 \textit{GUM}DELAY - 0.0034 \textit{PLH}DELAY \\
 & - 0.0014 \textit{SDO}DELAY - 0.0159 \textit{ALL} \textit{PORTS} - 0.0127 \textit{NO} \textit{YOK} \\
 & - 0.0161 \textit{NO} \textit{SUB} - 0.0155 \textit{NO} \textit{SIN} - 0.0105 \textit{NO} \textit{SAS} \\
 & + 0.0229 \textit{NO} \textit{GUM} + 0.0242 \textit{NO} \textit{PLH} + 0.1498 \textit{NO} \textit{FBE} \\
 & - 0.0158 \textit{NO} \textit{TAO} \textit{FBE} - 0.0155 \textit{NO} \textit{TAO} \textit{GUM} \\
 & - 0.0159 \textit{NO} \textit{TAO} \textit{PLH} - 0.0156 \textit{NO} \textit{TAO} \textit{SIN} - 0.0158 \textit{NO} \textit{TAO} \textit{SUB} \\
 & - 0.0157 \textit{NO} \textit{TAO} \textit{YOK} - 0.0158 \textit{NO} \textit{TAO} \textit{SAS} \\
 & - 0.0065 \textit{REPTYPE}),
 \end{aligned}$$

where NO TAO__ refers to supply ships with United States flag are not allowed into the respective port; GDL = USNS GUADALUPE; PEC = USNS PECOS; JEC = USNS JOHN ERICSSON; SIRI = HMAS SIRIUS; FBE = Fleet Base East, Australia; GUM = Guam; PLH = Pearl Harbor, Hawaii; SAS = Sasebo, Japan; SDO = San Diego, California; SUB = Subic Bay, Philippines; YOK = Yokosuka, Japan; and SIN = Singapore.

In this specific scenario, we identify that when USNS John Ericsson is unavailable and when Australia is closed as a resupply port supply ships will be underway 8% (5 days) and 15% (9 days) more respectively. Additional days underway as specified are based on a 60-day time horizon. This statistical model is easily implemented in an Excel spreadsheet for convenient use by the planner.

4.1.5 Scenario CSG-1 Model Input Correlation

A final portion of analysis for Scenario CSG-1 discusses correlation between inputs. We find that variance in predicting fuel consumption is predicated on small negative correlations between port delays in three specific ports. One hypothetical example of port delay modeled is a situation where weather precludes a supply ship from getting underway after INREP to support the replenishment schedule. Another example could be that scheduled maintenance availabilities preclude a customer or supply ship from returning to sea and meeting planned RAS events. As delay in one of the ports increases, predicting power remains relatively unchanged if delays in the other ports are minimal and/or do not overlap other port delays.

Scenario CSG-1 provides a proof of concept, and a basis from which to pursue a better predictive tool using Scenario CSG-2. The statistical models produced for Scenario CSG-1 are reliable, providing fleet commanders useful and timely information.

4.2 Scenario CSG-2 Results

4.2.1 Case B: Predicting Supply Ship Fuel Consumption

We shift to Scenario CSG-2, predicting supply ship fuel consumed on a five-day logistics sustainment cycle over a 60-day time horizon. We trained the statistical model using observations regarding two-through-eight cycle days (omitting 5-day cycles) and tested

the statistical model using observations of five cycle days. Figures 4.6 and 4.7 establish a baseline to explore reduction in estimation error.

Case B Model Inputs

Aside from inputs omitted as noted above, the base case varies all other available inputs: CSG CYCLE, CLF AVAIL, PORT AVAIL, and RAS TYPE. Regression of the base case explains roughly 70% of variation through linear input interactions. The majority of observations exhibit error when ports and/or supply ships assigned to the area are not available. Changes to number of days in the sustainment cycle seem to have negligible effects. Since supply ships can aid other regions (in certain circumstances) the fuel consumption is based on the proximity customer ships are to INREP ports and supply ships when replenishment of a customer ship is needed. As time passes, customer and supply ship positioning changes, and introduces variability in how RASP will accommodate changes to related input variables.




Figure 4.6. Scenario CSG-2 (Case B: Base Case) Supply Ship Fuel Consumed Prediction: Situations with higher amounts of fuel consumed represent scenarios where resupply ports and supply ships are required to transit further between ports and customer ships.

Figure 4.7 shows an overwhelming majority of values being under-predicted. Although they have large amounts of error, these predictions still offer value; planners can assume predictions will be reliably inaccurate by roughly 75% on average when predicted values are at extremes of possible values. This base case shows predicting supply ship fuel consumption is possible; planners can more reliably predict fuel usage when error is reduced. Figure 4.8 shows the range of absolute error. From this basis, we will try and reduce error by establishing sub-cases, limiting inputs used to train the statistical model, in hopes of producing better predictions.




Figure 4.7. Scenario CSG-2 (Case B: Base Case) Supply Ship Fuel Consumed Prediction: Situations with higher amounts of fuel consumed represent scenarios where resupply ports are not available requiring supply ships to transit further between replenishment ports and customer ships.




Figure 4.8. Scenario CSG-2 Case B Absolute Error: There is a large variance in percent error; most of which tends toward 50% or more.

4.3 Improving Error Through Subsets of Case B

Since the base case has reduced prediction power, Cases B.1 through B.5 are created seeking to reduce error. Each sub-case reduces the scope of the problem to a smaller subset of observations used to train the statistical model. In each sub-case, specific observations associated with CLF AVAIL and PORT DELAY (e.g., USNS Pecos not available) are omitted from the training set focusing on a limited set of outcomes.

4.3.1 Case B.1: Predicting Supply Ship Fuel Consumption – Supply Ship Availability

Case B.1 Model Inputs

Starting with Case B, Case B.1 uses observations where either USNS Guadalupe, USNS Pecos, or USNS John Ericsson is included and available to meet replenishment requirements (i.e., no combination of the supply ships is used in the same scenario). Choosing only these supply ship options results from Principal Component Analysis (PCA) as shown in Figure 4.9; ranges of possible values decrease when these inputs (supply ships) are included. Figure 4.9 suggests that better predictions may result when only one supply ship is available to support the replenishment schedule instead of when multiple supply ships are available at the same time.




Figure 4.9. Scenario CSG-2 (Case B.1) Principal Component (CLF AVAIL): Observations associated with X3 (ONLY GDL: USNS Guadalupe), X6 (ONLY PEC: USNS Pecos), and X7 (ONLY JEC: USNS John Ericsson) produce predictions that are more tightly grouped than others. These inputs suggest areas where improvements to error could occur. Every combination of three supply ship availabilities is analyzed.

The lower left portion of Figure 4.10 shows observations that align with one estimated value, but have different predicted values. The observations associated with 110K Gal (estimated value) appear when only changes to delay of USNS PECOS leaving San Diego occur. When this happens, the statistical model does not seem to know how to handle the delay after customer ships are outside of the San Diego region and seems to try and continue including its influence in calculations. In actuality, a delay from getting underway out of San Diego in excess of 3 days will have no effect on customer and supply ships in Scenario CSG-1 and CSG-2 because they have traveled outside the San Diego region and are supported by Pearl Harbor. This is a clear example where simple linear regression may not be as reliable for predicting fuel consumption for Scenario CSG-2. Additionally, the observations associated with 180,000 Gal (estimated value) appear when the sustainment cycle is either 6, 7, or 8 days; customer ships are only allowed to RAS; and USNS Guadalupe (Pearl Harbor Region) is the only supply ship that is available. Again we see the statistical model is unable to see

when the CSG is no longer in the Pearl Harbor region and it continues including Pearl Harbor delay in the calculus. When situations like these two arise in practice, the planner must intervene to prevent schedule infeasibility.




Figure 4.10. Scenario CSG-2 (Case B.1) Supply Ship Fuel Consumed Prediction: Observations with larger amounts of fuel (estimated value) have large errors. The statistical model often under-predicts in these situations.

Regression (least squares) results in the top three, most influential, inputs as CLF REGION, CSG CYCLE, and X6 (only USNS Pecos is available). CLF REGION being the most important input is intuitive, requiring a supply ship to go outside its assigned region to support the replenishment schedule could yield larger amounts of fuel consumed since transit distances between INREP and RAS location could be much greater depending on supply and customer ship locations. Cycle days is also intuitive: customer ships with a shorter sustainment (a smaller number of days in cycle) would need replenishment more frequently, which would require supply ships to transit back to port for fuel on-load more often. The most interesting input is the availability of USNS Pecos, which is assigned to the San Diego Region, because after the first five days of the scenario, the CSGs are no longer in USNS Pecos' region. The conclusion here is that when USNS Pecos is the only

supply ship available, the feasibility of supporting the RAS schedule essentially becomes zero, especially when sustainment cycles are short. Unless customer ships are allowed to INREP or CLF REGION is set to “FLOAT,” allowing USNS Pecos to come outside of her assigned region, the scenario becomes infeasible quickly.

Case B.1 Prediction Error

To begin developing intuition regarding how well we can predict fuel consumption, we use error associated with Case B.1 as a starting point. As seen in Figures 4.11 and 4.12, the statistical model under-predicts most often, with the majority of predictions being at most 50% below estimated values; approximately 70% of observations have error of 50% or less. 50% error is a large amount of error, so reducing error is desirable. This can be achieved through interactions between inputs that are other than strictly linear; however, loss in simplicity of the statistical model could result.




Figure 4.11. Scenario CSG-2 (Case B.1) Supply Ship Fuel Consumed Prediction Error: Most error is attributed to when under-predicting fuel consumption occurs, with most error being within 50% of estimated values.




Figure 4.12. Scenario CSG-2 (Case B.1) Supply Ship Fuel Consumed Prediction Absolute Error: The majority of error, roughly 82%, has error within 50% of estimated value.

4.3.2 Case B.2: Predicting Supply Ship Fuel Consumption – San Diego Delay

Case B.2 Model Inputs

Limiting port availability is a natural next step in analysis. Case B.2 only uses data associated with port delay; we start with delays of one and two days by the San Diego assigned supply ship USNS PECOS. We wish to see effects SDO DELAY has on predictions of supply ship fuel consumption. Of course, if delays persist longer than cycle days, then customer ships will have to return to port to replenish fuel if INREP events are authorized. Figure 4.13 indicates delays of one or two days have the most variability, thus offer a broader range of data points to train the statistical model and provide a better basis for predictions when additional observations are presented.




Figure 4.13. Scenario CSG-2 (Case B.2) Principal Component (SDO DELAY): Predicting Supply Ship Fuel Consumed seems to have least error when delays of one or two days are experienced.

Figure 4.14 shows that explaining the source of error improves, but RMSE does not improve, compared to Case B and B.1. This indicates that delay of a supply ship leaving port has impacts on the amount of fuel consumed, but we are unable to quantify this change in delay, day for day, when the scenario progresses past the five initial days. Port delay is modeled by assessing delays from the beginning of the timeline, instead of at a point designated by the planner. Being able to describe effects delay has at any point in the timeline that the planner wishes would be powerful.




Figure 4.14. Scenario CSG-2 (Case B.2) Supply Ship Fuel Consumed Prediction: The statistical model shows increasing error as fuel consumption values get larger. This sub-case focuses on delay of the San Diego-assigned supply ship leaving port. Using San Diego port delay accounts for more sources of prediction error than Case B.1, but RMSE increases. The same is true when comparing this sub-case to Case B. This model exhibits signs of over-training (see e.g., James et al. (2013)) and does not adequately reflect reality.

Case B.2 Prediction Error

Most of the error shown in Figure 4.15 is within 50% of estimated values; in fact, on average, predictions have an absolute error of roughly 51%. Further, 68% of observations have error less than 50% and 81% of observations have error less than 60%. Supply ship fuel consumption prediction values range from approximately 300,000 to 500,000 gallons. This small range invites poor model performance when compared to RASP outputs because the range of RASP outputs is between 350,000 to 1,150,000 gallons. In this sub-scenario, all predictions were above estimated values (i.e., the statistical model always over-predicts). Following Case B.1, B.2 sees the range of error increase from 98.5% to 120%; of course, this is opposite of the goal. Thus, we move to Case B.3, to see if delays from Pearl Harbor offer better insight.




Figure 4.15. Scenario CSG-2 (Case B.2) Supply Ship Fuel Consumed Prediction Error: Approximately 80% of all observations have error less than 60%. 60% error could range anywhere between 69,000 to 350,000 gallons of fuel. This statistical model always over-predicts fuel consumption.

4.3.3 Case B.3: Predicting Supply Ship Fuel Consumption – Pearl Harbor Delay

Case B.3 Model Inputs

Case B.3 only uses data associated with port delays of three and four days from Pearl Harbor, Hawaii. Focusing on these values of PLH DELAY is in attempt to capture whether delays from Pearl Harbor influence schedule supportability alone and/or in combination with delays from San Diego. It takes roughly seven days transit time from San Diego to Pearl Harbor, so delays could compound with one another. If port delays persist longer than sustainment cycle days, then customer ships must return to port to replenish fuel. Figure 4.16 indicates a delay of three days offers the most variability, with four-day delays reducing the variability.




Figure 4.16. Scenario CSG-2 (Case B.3) Principal Component (PLH DELAY): Predicting Supply Ship Fuel Consumed seems to have least error when delays of three or four days are experienced. This is the basis for analysis in Case B.3.

From Figure 4.17 we see that using Pearl Harbor port delays in this manner decreases prediction RMSE compared to Case B.2. Compared to Case B, Case B.3 accounts for more sources of variation within the dataset; however the associated RMSE increases which indicates the statistical model developed becomes over-trained (the statistical model fits the observations too well and will not predict reality well).




Figure 4.17. Scenario CSG-2 (Case B.3) Supply Ship Fuel Consumed Prediction: Case B.3 focuses on predicting fuel consumed, using data related to three or four-day delays from Pearl Harbor. Prediction RMSE increases compared to Case B.2, but more variation in the dataset is explained. This model exhibits signs of over-training and does not adequately reflect reality.

Case B.3 Prediction Error

Prediction error in Case B.3 is almost identical to Case B.2; Figure 4.18 represents error for Case B.3, which is relatively unchanged from Case B.2. However, the minimum and maximum percent error values are 13% and 133% respectively, and the average error seen is 49%. An error of 49%, in this context, amounts to at least 253,000 gallons, which is not negligible. Just like Case B.2, using delays from Pearl Harbor, Hawaii does not look to be a key factor. More variance in the data is explained (71% versus 70%), but error in prediction is worse.




Figure 4.18. Scenario CSG-2 (Case B.3) Supply Ship Fuel Consumed Prediction Error: Case B.3 over-predicts values RASP. Approximately 68% of observations have error less than 50% and zero have error less than 10%.

Moving to Case B.4, it seems that PORT DELAY, as indicated in Cases B.2 and B.3, does not have as strong of an effect on prediction power as originally thought. Ultimately, in an effort to not over-fit, we decide that modeling port delay in this fashion does not offer enough insight; a different set of inputs may reduce error. Therefore, we close port availability completely for customer and supply ships instead of delaying ships from returning to sea.

4.3.4 Case B.4: Predicting Supply Ship Fuel Consumption – Closing Guam as a Resupply Port

Case B.4 Model Inputs

Case B.4 considers Guam being unavailable for INREP by any ship, regardless of type. The choice to restrict Guam specifically comes from PCA captured in Figure 4.19. The graphic shows three inputs that influence supply ship fuel consumption when not available: Guam,

Pearl Harbor, and San Diego. Being the furthest west of the three options, Guam seems it would have most effect on scenarios since transit across the Pacific in “normal” cycle-days conditions would not see limitations attributed to Pearl Harbor or San Diego availability and customer ships could get to Guam within normal sustainability limits.




Figure 4.19. Scenario CSG-2 (Case B.4) Principal Component (PORT AVAIL): Predicting Supply Ship Fuel Consumed seems to have least error when either Guam, Pearl Harbor, or San Diego are not available for replenishment. Case B.4 chooses to consider effect on predictability when Guam is not available.

Looking at Figure 4.20, closing Guam for replenishment explains the same amount of the variability as Cases B.2 and B.3. However, RMSE increases compared to Case B.3, but is decreased from Case B.2. Additionally, when compared to Case B, this sub-case accounts for more sources of variation within the dataset, but RMSE increases. Therefore, use of this case to predict supply ship fuel consumption leads to mixed results and the use of this model is less reliable. Port availability clearly affects fuel consumption, but to quantify its effects directly requires including input interactions that will complicate the statistical model. Using Guam’s availability alone does not provide enough reliable insight to fuel consumption predictability.




Figure 4.20. Scenario CSG-2 (Case B.4) Supply Ship Fuel Consumed Prediction: This sub-case focuses on whether Guam is available as a replenishment port. Varying whether Guam is available does not provide improvement to explaining sources for error, and it increases RMSE relative to previous cases. Comparing these results to Case B, shows that limiting use of Guam as a port does not significantly effect prediction power.

Case B.4 Prediction Error

Case B.4 also does not show improvement in prediction error comparatively (Figure 4.21 supports this assertion). Like Cases B.2, and B.3, explaining variation in the dataset improves to 71%, however RMSE worsens to 75,000 gallons, the worst of all scenarios created. Like Cases B.1 through B.3, approximately 61% of observations have error less than 50%. Closing a single port for INREP (here it is Guam) does not help improve error bounds alone. Therefore, closing Guam as a resupply port for supply and customer ships will not have a large impact on schedule feasibility if it is the only port closed.




Figure 4.21. Scenario CSG-2 (Case B.4) Supply Ship Fuel Consumed Prediction Error.

4.3.5 Case B.5: Prediction Supply Ship Fuel Consumption – Supply Ship Availability and RAS Only

Case B.5 Model Inputs

Case B.5 limits inputs used in Case B.1 and Figure 4.22 shows the regression results. This regression only uses observations associated with individual supply ships solely supporting the schedule (i.e., GDL, PEC, and JEC individually support the schedule), coupled with a requirement that customer ships can only replenish at-sea. The idea is to create an environment where the CLF fleet is in extremis and on-station requirements for customer ships cannot afford time required transiting between the operating area and replenishment ports.




Figure 4.22. Scenario CSG-2 (Case B.5) Supply Ship Fuel Consumed Prediction: Predictions begin separating when fuel consumption values are between 320,000 gallons and 340,000 gallons; this suggests a relationship among input values occurs (here the value of CLF REGION is the relationship).

Case B.5 Prediction Error

Finally, we see some improvement to error from previous cases. As a reminder, Case B.1 sees error range from a minimum of 0.03% and a maximum of 98%. The explanation of variation in this regression worsens from Case B.1, but only slightly (58% versus 59%), and RMSE increases from 67,000 gallons to 69,000 gallons. This is an example of where the average value for a data set is misleading; there needs to be a relative point shared by both regressions to truly assess which one is better. Despite these changes to R^2 and RMSE, we see bounds on error tighten; the lower and upper bounds to percent error become 7% and 44% respectively. Even more insightful is these bounds to error are actually based on the value of one particular input. When CLF REGION is “TETHER” the statistical model predictions are always under RASP values with lower and upper bounds to error become -28% and -44% respectively. Alternatively, when CLF REGION is “FLOAT”

predictions are always over RASP values with lower and upper bounds to error of 7% and 17% respectively. This is a significant finding that yields a tightened range of error based on whether values are under or over-predicted. Figure 4.23 provides a pictorial representation of error produced by the statistical model.




Figure 4.23. Scenario CSG-2 (Case B.5) Supply Ship Fuel Consumed Prediction Error: Error distribution is bi-modal, groups of error values are attributed to the value of CLF REGION. Supply ships are either strictly assigned to their region, or are able to execute replenishment events up to a distance outside of an assigned region.

Using Case B.5 as a benchmark, we further pursue a case where lower and upper error bounds come closer to each other. Case B.5.2 builds upon Case B.5.1, producing a more restricted operational environment. Even though the number of observations in the training set gradually become less, error values are based on the same test set originally constructed with Case B.5.

4.3.6 Case B.5.1: Predicting Supply Ship Fuel Consumption – Supply Ship Availability, RAS Only, and Guam and Yokosuka Resupply Ports Closed

Case B.5.1 Model Inputs

Case B.5.1 further builds upon Case B.5, making replenishment ports in Guam and Yokosuka, Japan unavailable. Our idea is to develop a situation where weather precludes entering port, or perhaps diplomatic clearances have expired and United States ships are no longer allowed to enter these ports, to see impacts on fuel predictions. Figure 4.24 shows results of performing this regression.




Figure 4.24. Scenario CSG-2 (Case B.5.1) Supply Ship Fuel Consumed Prediction: Fuel consumption predictions begin converging to bi-modal means, one centered around approximately 260,000 gallons and the other around 410,000 gallons. This separation occurs as a result of the value CLF REGION selected when running RASP.

Case B.5.1 Prediction Error

Figure 4.25 shows error distribution to again be bi-modal. Analysis shows over- or under-predictions also stem from the value of “CLF REGION,” just like in Section 4.3.5. Again, when supply ships are confined to their assigned region, the statistical model always under-predicts fuel consumption. The error ranges between -26.9% and -13.4%, an improvement on bounds to error from Case B.5. Conversely, when allowed to transit outside their region fuel consumption is over-predicted; the error here ranges between 17.8% and 26.6%, an improvement from Case B.5.




Figure 4.25. Scenario CSG-2 (Case B.5.1) Supply Ship Fuel Consumed Prediction Error: Error distribution remains bi-modal, from Case B.5; however, bounds to the ranges of error, per mode, decrease with respect to Case B.5 values. Closing Guam and Yokosuka as resupply ports gives better predictions to supply ship fuel consumption.

4.3.7 Case B.5.2: Predicting Supply Ship Fuel Consumption – Supply Ship Availability; RAS Only; and Guam, Yokosuka, Pearl Harbor, Australia Resupply Ports Closed

Case B.5.2 Model Inputs

Case B.5.2 builds upon Case B.5.1, adding conditions where the ports of Pearl Harbor, Hawaii and Australia are also unavailable. We wanted to create an extreme situation where essentially the entire region's ports are unavailable for replenishment and supply ships are required to transit long distances to support the fleet's operations. Here we test the strength of the logistics chains, trying to “break” their ability to provide fuel and stores as required. Figure 4.26 shows the result of this regression, which is almost identical to Figure 4.24.




Figure 4.26. Scenario CSG-2 (Case B.5.2) Supply Ship Fuel Consumption Prediction: Fuel consumed values are expressed in thousands of gallons. Errors in prediction seem to be converging toward respective bi-modal means. Observations are more densely grouped, indicating error is improving.

Case B.5.2 Prediction Error

Figure 4.27 shows that ranges of error improve, but only slightly compared to Case B.5.1; in fact, the change is almost negligible, making the additional ports unavailable had almost no effect on feasibility of meeting requirements. The reason for this effect could be they are so far away that no consideration was made as being viable resupply options in the first place. Or their capacities were such that when not available the remaining port available (Sasebo, Japan) has enough capacity to absorb demand from Pearl Harbor and Australia not being available. Of course, there are other situations dictating port availability (e.g., port loading restrictions) that are not considered; the statistical model assumes a port scheduled will have availability when the requirement occurs.

Bounding prediction error can continue in this fashion; however, additional steps may begin over-fitting the statistical model. Seeing the small change in prediction error lower and upper bounds between Cases B.5.1 and B.5.2 indicates a plateau may have been reached, and further attempts to constrain the scenario will result in relaxation of error bounds. Building a series of sub-cases to the dataset shows that predicting fuel consumption is possible and we have reached a point where predictions are reliable, while keeping the resulting statistical model simplified enough for quick hand-calculations when necessary.

The big take-aways for a scheduler from analysis are:

1. Predicting supply ship fuel consumption and supply ship underway percentage are possible.
2. Supply ship fuel consumption predictions are most influenced by the number of days in the sustainment cycle, regional boundary restrictions enforced upon supply ships, and the number of supply ships available in order of importance.
3. Closing ports has an effect on supply ship fuel consumption when compounded with limitations focused around sustainment cycle days and the number of supply ships available.




Figure 4.27. Scenario CSG-2 (Case B.5.2) Supply Ship Fuel Consumed Prediction Error: Bounds to prediction error ranges, per mode, decrease from Case B.5.1. Bi-modal distribution of error is associated to the corresponding CLF REGION value considered.

4.3.8 Case B Overall Error Summary

Figure 4.28 summarizes error from each of the cases explored in Case B category. As indicated in previous sections, each sub-case tries to reduce the range of error associated with predictions. Clearly, Cases B.5.1 and B.5.2 limit error the most, and are the basis for statistical models provided regarding Scenario CSG-2.




Figure 4.28. Scenario CSG-2 Error Summary: The range of prediction error values varies between cases. Case B.1 is the basis for creating Cases B.5, B.5.1 and B.5.2. Case B.5.2 has the best range of error, and the values are separated based on the CLF REGION parameter. Refer to Section 4.3.5 for description of CLF REGION's values.

4.3.9 Statistical Model: Predicting Fuel Consumed Scenario CSG-2

When predicting supply ship fuel consumption, we offer the following statistical model from our 7TH Fleet AOR example Scenario CSG-2:

$$\begin{aligned}
 \text{Fuel Consumed (K Gal)} = & 114.2760 - 4.3801 \text{ CSG CYCLE} + 15.1575 \text{ ONLY GDL} \\
 & + 7.9637 \text{ NO PEC} + 152.9025 \text{ CLFREGION} \\
 & - 2.1069 \text{ SDO DELAY} - 2.3154 \text{ PLH DELAY} \\
 & - 1.0438 \text{ GUM DELAY} - 0.0102 \text{ ALL CLF} \\
 & + 21.6812 \text{ NO SDO} + 0.1173 \text{ NO SAS},
 \end{aligned}$$

where GDL = USNS GUADALUPE; PEC = USNS PECOS; GUM = Guam; PLH = Pearl Harbor, Hawaii; SAS = Sasebo, Japan; SDO = San Diego, California; SAS = Sasebo, Japan.

In this specific scenario, we identify that having only USNS Guadalupe available and not having USNS Pecos available increases supply ship fuel consumption. We also identify that

when supply ships are not allowed to transit outside their assigned region supply ship fuel consumption increases by approximately 152,000 gallons. This statistical model is easily implemented in an Excel spreadsheet for convenient use by the planner.

4.4 Case C: Classification of Maximum Percent Below Observations

4.4.1 Case C

This section focuses on classifying whether a particular RASP input produces an output (for “Maximum Percent Below”) above or below a 30% safety stock threshold. Such a threshold may be specified by a commander to provide a minimum amount of safety stock required to execute missions. For this purpose, we develop a Support Vector Machine (SVM) model, using the L1-regularization technique, in the Python-based, open-source optimization modeling (PYOMO) language. For supporting information regarding L1-regularization, see Section 3.5.2. Using the raw data from Case B, the SVM model trains using 24,192 data points. In the training set 37% and 63% of the data points are labeled “1” and “-1” respectively. In the test set (4,032 data points) 43% and 57% are labeled “1” and “-1” respectively. A label of “1” indicates a data point is “above” the threshold and a label of “-1” indicates a data point is “below” the threshold. Outputs of the classification model predict whether a customer ship drops below the safety stock threshold.

Case C Model Inputs

Figures 4.29, 4.30, and 4.31 show results of classification. Classification fidelity is indicated by the number of correct classifications and the optimal value produced by SVM. In Case C, we see that the best percent of correct classifications and optimal value reached are 82.14% and 0.215 respectively. An optimal value as close to zero as possible coupled with the point where percent correctness is highest and false classifications are lowest is best. Larger values of ρ have a tendency in resulting in an over-trained classification model and less accurate predictions. With this best choice of ρ the SVM model correctly classifies “-1” labels 94% of the time and “1” labels 66% of the time.




Figure 4.29. Correct Classification Percentage: as the SVM penalty parameter ρ increases, we expect the percent of correct classifications to decrease because the margin around each point increases. We see that maximum correct classification percentage is 82.14%. Threshold for classification is -30% .




Figure 4.30. Classification Optimal Value: As the SVM penalty parameter ρ increases, we expect optimal value to increase. An optimal value of 1, in absolute value, indicates performance is no better than a 50-50 game of chance; in these cases, a horizontal line is fit to the data. The lowest optimal value reached is 0.215.




Figure 4.31. False Positive and False Negative Classifications: as ρ increases, the number of false classifications vary. When ρ is roughly 0.0115, percent of correct classifications is highest and false classifications are at their “best” respective values.

SVM Classification Model

Analysis of the data suggests that the optimal value for ρ is 0.011 since at this value not only does the number of false positives become among the lowest across all ρ values (144), the percentage of correct classifications is at its maximum value as well. Using this value for ρ , the following SVM model applies:

$$\begin{aligned} &\text{If } -11 + \text{CSG CYCLE} + 3 \text{ ALL CLF} + \text{NO PEC} - \text{ONLY PEC} \\ &\quad + 2 \text{ CLF REGION} + \text{RAS TYPE} > 0, \text{ then predict below the threshold.} \end{aligned}$$

Otherwise, predict above threshold.

Outputs to the above classification model equal to "1" interpret as situations where at least one customer ship in the scenario drops below the commander's threshold (here the threshold is 30%). CSG CYCLE values are integers between "2" and "8", CLF REGION and RAS TYPE are either "1" or "2", and the remaining inputs are binary. An example of an observation classified as dropping below the threshold has the following values: CSG CYCLE = 7, ALL CLF = 0, NO PEC = 0, ONLY PEC = 1, CLF REGION = 2, and RAS TYPE = 2. These values represent a scenario when customer ships are on a seven-day sustainment cycle, only USNS PECOS supports the replenishment schedule, and customer

ships are only able to RAS. As CSG CYCLE increases, the chance customer ships will fall below the threshold decreases. This statistical model is easily implemented in an Excel spreadsheet for convenient use by the planner.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5:

Conclusions and Future Research

5.1 Scenario Results

5.1.1 Scenario CSG-1

Predicting fuel consumption is possible with three inputs having the biggest contribution to statistical model outputs. In order of influence, those inputs are logistics sustainment cycle days, supply ship availability, and port availability. Among all observations, approximately 25% have error less than 20% and roughly 50% of observations have error less than 10% between values predicted by the statistical model and RASP.

Enhancing the statistical model by including more complex interactions between inputs, such as two-way interactions, significantly improves prediction power. A second-degree factorial design predicting CLF fuel consumption improves R^2 and RMSE dramatically. If a full factorial design is implemented, R^2 gets worse and RMSE increases, indicating we are trying too hard to explain all sources of data variation. After all, the R^2 seen in the second-degree factorial statistical model far exceeds standard practice, indicating any further efforts to explain variation could be wasteful. Using second-degree and full factorial statistical models will reduce prediction error but also decrease interpretability. Explaining a resulting output from the statistical model may become untenable.

Looking at predicting supply ship underway percentage, we see improvements by implementing both a second-degree factorial and full-factorial design. What is different from supply ship fuel consumption, predicting supply ship underway percentage using a full-factorial design improves RMSE over a second-degree design.

However, seeking a simple statistical model for planner use, a full factorial design is unreasonable. A second-degree factorial design is easier to interpret than a full-factorial design, as the number of calculation elements is restricted significantly, and is much less time-consuming to perform associated calculations. The return on investment for the full-factorial design, related to simplicity, does not favor the fleet scheduler. A second-degree

factorial statistical model renders a much more simplified statistical model, and is more functional; but the calculations are still extensive. The focus of this thesis is to make available a usable calculation tool to inform decision-makers without having to re-run RASP. The statistical models resulting from a full-factorial and second-degree factorial designs are too cumbersome, containing over 5,000 and 200 interactions respectively. Both are more involved than a quick calculation during a meeting, thus the statistical model in Section 4.1.4 provides a proper balance of complexity and time required.

Finally, considering all conclusions regarding the Scenario CSG-1 statistical model, solve times are extremely short. A maximum solve time for any of the permutations was 54 seconds, and 94% of observations saw solve times within 10 seconds from start to finish; typically, it would take longer than 10 seconds for someone to use a calculator. The statistical model presented in Section 4.1.4 provides a good prediction of fuel consumption; the prediction is accurate enough to determine whether RASP should be re-run. Of course, running RASP again making the appropriate changes will be much more accurate; however, the effort of this thesis is to develop intuition and insight about important inputs to the statistical model.

We conclude that when scenarios are relatively simple, those containing one CSG and requisite supply ship support can predict fuel burned with acceptable amounts of error and are reliable to decision-makers without needing an optimized solution from RASP.

5.1.2 Scenario CSG-2

Predicting Fuel Consumption

There are numerous inputs that when considered will provide predictions of higher quality using a statistical model. However, as the number of inputs increases complexity increases as well. The statistical model offered in Section 4.3.9 is an improvement from the one offered in Section 4.1.2 only in the sense of time required to perform calculations by hand or calculator. Each one offers its value based on the scenario being considered, one CSG versus two. The statistical model for Scenario CSG-2 offers better predictability because its range of error is bounded not only by an upper or lower limit, but also based on supply ship regional constraints. Being able to reduce the range of error based on one value alone is powerful and knowing when to either add or subtract a certain percentage of fuel to the

prediction helps understand requirements better. Further, when fuel prediction errors are 50,000 gallons, for example, we can confidently say that the error is negligible, not only for one specific event (of course depending on the number of ships in the scenario) but across the time horizon. 50,000 gallons of fuel, spread across eight ships, is roughly 6,000 gallons per ship, or 277 gallons for one event; these values are negligible and are accepted by United States Navy leadership as “costs of operations.”

Predicting Maximum Percent Below

Classifying maximum percent below is possible with a margin of error that can be improved upon. Classifications are based upon whether a ship (supply or customer) drops below a 30% safety stock threshold. The best value for ρ is 0.011 because the maximum percent of correct classifications overall (82%) is reached as well as false positives and false negatives are relatively low. The SVM model classifies ships that drop below the 30% safety stock threshold with 94% accuracy, and those that do not with 66% accuracy.

In certain scenarios, we are able to identify specific vessels and resupply ports that are more important than others when predicting supply ship fuel consumption, underway percentage and maximum percent below. Each scenario sees different supply ships and resupply ports as more influential than others to predicting outputs.

5.2 Key Take Aways

In the Scenario CSG-1 setting, we predict that CLF fuel consumption decreases by 67,962 gallons per customer ship sustainment cycle day increase. It decreases by 117,916 gallons when the supply ship from Guam (USNS John Ericsson in this case) is unavailable. It decreases by 109,778 gallons when Pearl Harbor is unavailable for supply and customer ship INREP. If the cost of fuel is \$2.70 per gallon, the reductions save \$187,497, \$318,373, and \$296,401 respectively. These savings will have a significant impact on how to support combatant force employment.

Also in the Scenario CSG-1 setting, we predict that supply ship underway percentage decreases by 2 per customer ship sustainment cycle day increase. With a 10-day sustainment cycle, supply ships spend 12 fewer days underway which significantly reduces operational costs associated to supply ships being underway. Underway percentage increases when

Guam, Pearl Harbor, and Australia are unavailable as resupply ports by 2 for each port individually. If all three ports are unavailable concurrently, an increase of 6% underway equates to 3.6 days.

In the Scenario CSG-2 setting, we predict CLF fuel consumption decreases by 152,903 gallons when supply ships are restricted to their assigned region for a savings of \$528,938 spent executing the schedule. Thus, planners should aim to keep supply ships in their own regions as best as possible.

Of course CLF fuel consumption savings typically come at the expense of other metrics planners consider (e.g., fuel inventory levels onboard combatants). Trade-offs occur where fuels savings are not always the leading metrics determining supply and customer ship schedules. The respective statistical model for each scenario is easily implemented in an Excel spreadsheet for convenient use by the planner to predict some outputs of RASP from various inputs.

5.3 Future Research

Future research can and will improve upon our efforts. The following list is not all inclusive.

- Model delays from replenishment ports more appropriately. Delay experienced at a specified time in the scenario, instead of from the beginning of a planning horizon, will be more realistic. Answering questions like: “What would happen if USNS John Ericsson was delayed from getting underway from Guam by two days?” would be useful. Forecasting resupply port closure can be based on weather trends, trend of current diplomatic relations, or as part of exploratory what -if analysis.
- Re-define PORT AVAIL implementation, creating environments where port availability is limited in certain periods of the time horizon, instead of for the entire duration. Being able to model port closure for a specific number of days through the entire time horizon (starting at any point) can lead to forecasting supply ship fuel consumption differently.
- Compound model input interactions in an effort to explain more of the sources of variability and further bound prediction error. Linear input interactions are easiest to interpret; however, more variation can be explained, and prediction error can be

improved with more complex interactions between inputs.

- Vary classification thresholds to determine schedule feasibility, vice only focusing on one value, to develop “go and/or no-go” criteria. Schedules could be deemed infeasible at a quick glance, then RASP can be re-run to determine outputs at optimality.

THIS PAGE INTENTIONALLY LEFT BLANK

List of References

- Brown GG, Carlyle WM (2008) Optimizing the US Navy's Combat Logistics Force. *Naval Research Logistics*, 55, 800-810.
- Brown GG, DeGrange WC, Price WL, Rowe AA (2017) Scheduling combat logistics force replenishments at sea for the US Navy. *Naval Research Logistics*, 64, 677-693, <https://doi.org/10.1002/nav.21780>.
- Brown GG, Dell RF, Wood RK (1997) Optimization and Persistence. *INTERFACES* 27(5):15-37.
- Cardillo JA (2004) *Optimizing Global Operations Plans for the Combat Logistics Force*. Master's thesis, Naval Postgraduate School, Monterey, CA, september.
- Demand (2022) *Merriam-Webster* (Merriam-Webster, Incorporated), accessed April 8, 2022, <https://www.merriam-webster.com/dictionary/demand>.
- Density (2022) *Merriam-Webster* (Merriam-Webster, Incorporated), accessed April 8, 2022, <https://www.merriam-webster.com/dictionary/density>.
- Diaz A (2010) *Heuristics in Global Combat Logistics Force Operational Planning*. Master's thesis, Naval Postgraduate School, Monterey, CA, <https://calhoun.nps.edu/handle/10945/5456>.
- Fawcett A (2022) *Data Science in 5 Minutes: What is One Hot Encoding?* (Educative, Inc.), accessed July 12, 2022, <https://www.educative.io/blog/one-hot-encoding>.
- Ghandi R (2018) *Support Vector Machine - Introduction to Machine Learning Algorithms* (Towards Data Science), accessed August 22, 2022, <https://towardsdatascience.com/support-vector-machine-introduction-to-machine-learning-algorithms-934a444fca47>.
- Glen S (2022) *Sensitivity Analysis ("What -if"): Definition* (Statistics How To), accessed August 25, 2022, Cited from StatisticsHowTo.com: Elementary Statistics for the rest of us! <https://www.statisticshowto.com/sensitivity-analysis>.
- James G, Witten D, Hastie T, Tibshirani R (2013) *An Introduction to Statistical Learning: with Applications in R* (New York, NY: Springer New York), ISBN 978-1-4614-7138-7, URL http://dx.doi.org/10.1007/978-1-4614-7138-7_3.
- Microsoft Technical Support (2022) *Switch between various sets of values by using scenarios* (Microsoft), accessed July 13, 2022, <https://www.educative.io/blog/one-hot-encoding>.

- Mizokami K (2021) Why 11 Navy Aircraft Carriers Simply Aren't Enough. Accessed August 4, 2022, <https://www.popularmechanics.com/military/navy-ships/a35928241/navy-aircraft-carriers/>.
- Naval Technology (2022) Henry j. kaiser fleet replenishment oilers. Accessed August 4, 2022, <https://www.navy.gov.au/hmas-sirius>.
- Passmark (2022) CPU Benchmarks. Accessed September 8, 2022, https://www.cpubenchmark.net/cpu_list.php.
- Pichery C (2022) *Sensitivity Analysis* (Elsevier B.V., ScienceDirect), accessed April 17, 2022, Cited from Encyclopedia of Toxicology (Third Edition), 2014 <https://www.sciencedirect.com/science/article/pii/B9780123864543004310>.
- Pike J (1999) Underway replenishment (UNREP). Federation of American Scientists (FAS). Accessed March 16, 2022, Updated Saturday, March 06, 1999 12:25:38 PM, <http://www.fas.org/man/dod-101/sys/ship/unrep.htm>.
- Pizer SM, Hong J, Vicory J, Liu Z, Marron J, Choi Hy, Daman J (2022) *Object shape representation via skeletal models (r-reps) and statistical analysis* (Elsevier B.V., ScienceDirect), accessed April 17, 2022, <https://www.sciencedirect.com/science/article/pii/B9780128147252000145>.
- Royal Australian Navy (2022) HMAS Sirius. Accessed August 4, 2022, <https://www.navy.gov.au/hmas-sirius>.
- Royset JO, Wets RJB (2021) *An Optimization Primer* (Springer), ISBN 978-3-030-76275-9, URL http://dx.doi.org/10.1007/978-3-030-76275-9_2.
- Sea Forces (2022) Ddg 107 - USS Gravely. Accessed August 4, 2022, <https://www.seaforces.org/usnships/ddg/DDG-107-USS-Gravely.htm>.
- Stanford University (2022) *Sam Savage* (Stanford University), accessed July 13, 2022, <https://online.stanford.edu/instructors/sam-savage>.
- Supply (2022) *Merriam-Webster* (Merriam-Webster, Incorporated), accessed April 8, 2022, <https://www.merriam-webster.com/dictionary/supply>.
- Zaffar H (2021) US Navy to Decommission Five Ticonderoga-Class Cruisers. Accessed August 4, 2022, <https://www.thedefensepost.com/2022/03/29/us-navy-decommission-ticonderoga-cruisers/>.

Initial Distribution List

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California