

Modeling & Simulation in Support of MALSP II
Report of Findings

Primary Investigator:

Chad W. Seagren
Major, USMC
Assistant Professor
Operations Research Dept
Naval Postgraduate School
Monterey, CA

8 March 2013

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, 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 the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.						
1. REPORT DATE (DD-MM-YYYY) 15-05-2013		2. REPORT TYPE FINAL REPORT		3. DATES COVERED (From - To) 23 SEP 2011 - 15 APR 2013		
4. TITLE AND SUBTITLE THE OPTIMIZING AVIATION SUPPLY DISTRIBUTION (MALSP II) STUDY				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Major Chad Seagren Jessica Hancock				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Marine Corps Combat Development Command Operations Analysis Division, Analysis Branch (C 191) 3300 Russell Road Quantico, VA 22134-5130				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Headquarters, United States Marine Corps Deputy Commandant for Aviation, Aviation Logistics Support Branch 3000 Marine Corps Pentagon, Washington, DC 20350-3000				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution Statement A. Approved for public release, distribution is unlimited.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT This project employs modeling and simulation analysis in an effort to assess various aspects of the Marine Aviation Logistics Support Program II (MALSP II) doctrine currently in development. MALSP II leverages recent innovations in information technology, logistics chain management, and continuous process improvement in order to improve logistical support to deployed Marine air assets. The model develop in this project is a Discrete Event Simulation, implemented in Java using the SimKit programming libraries. It is employed to specifically analyze aspects of MALSP II doctrine; however, it is sufficiently flexible to analyze various aspects of Legacy MALSP concepts as well. The model is relatively straightforward and has a small number of inputs. In most cases, sufficient data exist to confirm the suitability of modeling a given empirical distribution in a particular manner. In some instances where data are more difficult to obtain, compelling theoretical reasons exist as justification for selecting a particular distribution or algorithm.						
15. SUBJECT TERMS Marine Aviation Logistics Support Program (MALSP); SimKit; Discrete Event Simulation; Main Operating Base (MOB)						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Jessica Hancock	
Unclass	Unclass	Unclass	SAR	109	19b. TELEPHONE NUMBER (Include area code) 703-432-8383	

Table of Contents

Executive Summary	4
1. Introduction to Marine Aviation Logistics	7
2. The SimKit Model	19
3. Input Validation	27
4. Assessment of MALSP II Packages	34
5. Analysis of Inclusion Decisions.....	55
6. Analysis of TRR Uncertainty.....	62
7. Future Work and Recommendations	71
8. Glossary of Terms.....	73
9. List of External Files Referenced	80
10. Mathematical Appendix.....	81
11. References.....	110

Executive Summary

This project employs modeling and simulation analysis in an effort to assess various aspects of the Marine Aviation Logistics Support Program II (MALSP II) doctrine currently in development. MALSP II leverages recent innovations in information technology, logistics chain management, and continuous process improvement in order to improve logistical support to deployed Marine air assets.

The model developed in this project is a Discrete Event Simulation, implemented in Java using the SimKit programming libraries. It is employed to specifically analyze aspects of MALSP II doctrine; however, it is sufficiently flexible to analyze various aspects of Legacy MALSP concepts as well.

The model described in this report is relatively straightforward and has a small number of inputs. In most cases, sufficient data exist to confirm the suitability of modeling a given empirical distribution in a particular manner. In some instances where data are more difficult to obtain, compelling theoretical reasons exist as justification for selecting a particular distribution or algorithm. However, it is still true that this report relies on data on high-priority requisitions from Joint Task Force – Horn of Africa (HOA) from 2006 to 2011, as well as two years of data on repairables for CH-53Es from Marine Aviation Logistics Squadron 16 (MALSP-16). While it is likely that demand patterns for other platforms behave similarly under other circumstances, it would be necessary to re-validate the model for other Type/Model/Series (T/M/S). That said, the process to validate the use of particular input parameters for a different situation would be identical.

The following research questions in this document were addressed (note: numbering is preserved from the Statement of Work):

1. How can allowancing at the Parent Marine Aviation Logistics Squadron (PMALS) improve or facilitate the building and managing of MALSP II packages?

The relative abilities of the different allowance structures (Legacy and MALSP II) were assessed for repairable items to support a notional MALSP II nodal laydown. The nodal laydowns that were examined include different quantities (0, 4, 8, and 16) of deployed aircraft examined with and without the implementation of an En-Route Support Base (ESB). Employing the proposed MALSP II allowancing structure to support a MALSP II logistics network significantly improves certain system performance measures (both practically and statistically) relative to supporting the same system using Legacy MALSP allowances.

-MALSP II allowances unambiguously improve performance for MALS with no deployed squadrons.

-MALSP II allowance packages achieve superior Main Operating Base (MOB) *Response Times* relative to Legacy allowance packages in all cases considered.

-Implementing an En-Route Support Base (ESB) increases *Response Time* at both the MOB and PMALS, for Legacy and MALSP II allowance packages. Recommend against stocking repairables at the ESB, except for certain low-density items under a narrow range of circumstances.

2. *What is the optimal/robust criteria for including an item in a packup?*

- If high utilization rates (relative to training hours) for the deployed aircraft are expected, ensure National Item Identification Numbered (NIINs) items with greater than 8 demands in the previous 24 months are included in the packup.

-Ensure NIINs with exceptionally high demand (e.g., greater than 70 demands in previous 24 months) are stocked using medium (95th percentile) or lower risk (that is, higher percentile) demand filtering.

- Using 80th percentile demand filtering for buffer sizing is sufficient for all types of NIINs for cases in which the aircraft at the Forward Operating Base (FOB) are not expected to experience substantially greater utilization rates than while in peace-time.

-Using 95th percentile demand filtering for buffer sizing is sufficient for all types of NIINs for cases in which the aircraft at the FOB are expected to experience up to three times the utilization rate relative to training hours.

4. *How does uncertainty regarding Actual Time to Reliably Replenish (TRR) effect Response Time at deployed nodes?*¹

-When *Actual TRR* exceeds the *Design TRR*, the node typically experiences high levels of *Response Time*. This effect is greater as *Design TRR* decreases and/or as number of supported aircraft (a/c) increases.

-NIINs with *Demand Frequencies* greater than 1 demand per month are most affected by differences between *Actual* and *Design TRR*.

-No risk demand filtering (100th percentile) provides robust protection in nearly all cases.

Time constraints preclude addressing the following questions and are therefore left for future work:

3. *What is the optimal/robust criteria for removing an item from a packup?*

5. *What is the optimal/robust criteria for positioning low density items?*

[Note: The detailed analysis in Chapter 3E certainly informs this issue.]

¹ On the Statement of Work, this question appears as “How frequently should buffers be re-sized? Are there useful leading indicators (e.g., say, if Actual Time to Reliably Replenish (TRR) exceeds Design TRR by a certain amount)”.

6. *Where should repair capability (i.e. Intermediate Maintenance Activity (IMA), or T-AVB, etc.) be placed in the nodal lay-down, if at all?*

This project demonstrates the utility of the modeling and simulation in assessing the performance of allowance packages and various business rules. As such, the following recommendations are proposed for future action:

- Address remaining research questions (questions 3, 5, and 6, above).
- Develop this model (or something like it) into a tool accessible to aviation logisticians in the operating forces to use for planning logistical networks.
- Consider employing this model (or something like it) together with the cost optimization algorithms in order to provide additional insight on operational impacts of allowancing decisions.
- Develop a canonical MALSP II scenario that will, among other things, enable analysts to assess more objectively whether an allowance package provides a sufficient level of support.
- Pursue an additional assessment of the utility of *Response Time* as a measure of effectiveness.

1. Introduction to Marine Aviation Logistics

This project employs modeling and simulation analysis in an effort to assess various aspects of the Marine Aviation Logistics Support Program II (MALSP II) doctrine currently in development. MALSP II leverages recent innovations in information technology (IT), logistics chain management, and continuous process improvement in order to improve logistical support to deployed Marine air assets.

A. The Legacy MALSP Doctrine

The Marine Corps developed the Marine Aviation Logistics Support Program (MALSP) in the late 1980s to facilitate the effective logistical support of aviation units deployed to confront Cold War opponents on the plains of Europe and elsewhere. The MALSP construct consists of a system of packages of parts, personnel, support equipment, and mobile facilities that are modular and flexible enough to tailor to nearly any imagined contingency. Marine Corps Warfighting Publication (MCWP) 3-21.2 (Aviation Logistics) is the publication that made MALSP into part of Marine Corps doctrine.

Figure 1-1 is a graphical depiction of MALSP's modular nature. The pie in the top-center of the diagram indicates that the logistics support consists of spare parts, people, support equipment, and mobile facilities. The wedges comprising this pie are typically drawn from the Marine Aviation Logistics Squadrons (MALS) that support the flying squadrons involved in operational deployment.




Figure 1-1. The modular MALSP construct (source: MCWP 3-21.2)

The other pies in Figure 1-1 represent how parts packages integrate to support an Aviation Combat Element (ACE). A MALS possesses specific support packages for each Type/Model/Series (T/M/S) aircraft, but it also may possess “common” packages that support multiple T/M/S. The primary idea is that the packages build upon one another and integrate with each other to support the ACE.

An important manner in which capability integrates under the MALSP construct is at various levels of maintenance. In naval aviation, the Organizational Level (O-Level) refers to the maintenance conducted by Marines assigned to flying squadrons; the squadron may also be referred to as an Organizational-level Maintenance Activity (OMA). This level of maintenance primarily consists of removal and replacement of defective parts; inspections; servicing, and incorporation of Technical Directives. Alternatively, the Intermediate-level Maintenance Activity (IMA) refers to maintenance conducted at the MALS. This primarily refers to performance of maintenance on components and related support equipment; calibration of designated equipment; providing technical assistance, etc. It should also be noted that flying squadrons assigned to the Marine Aircraft Group (MAG) also possess a certain amount of personnel structure that augments the IMA. These augmentees deploy whenever their respective squadron deploys and they help to form the MALS detachment that supports the deployment.

While there are many differences between MALSP and MALSP II, the most important differences as they pertain to this study are their Scheme of Maneuver, Allowancing, and Level-setting. We examine each from the perspective of the MALS.

1. Scheme of Maneuver

Typically, the MALSP operational Scheme of Maneuver envisions the deployment of squadron-level (or higher) aviation assets. Thus, in a MALSP contingency, a MAG would receive the order to deploy a flying squadron to a particular part of the world (say, a Humanitarian Assistance Disaster Relief mission to Haiti, etc). At this time, the MALS would prepare the Fly-In Support Package (FISP) associated with the T/M/S for deployment and begin assembling the Peculiar Contingency Support Package (PCSP). The flying squadron would self-deploy to their destination in theater and be followed closely by arrival of the FISP via strategic airlift. The aviation logistics personnel at the deployment site would support the squadron from the FISP until arrival of the PCSP, usually within 30 days. Upon arrival of the PCSP, the FISP may or may not be redeployed. (Doctrinally, the FISP is supposed to be redeployed back to the MALS in order to be ready for the next contingency. In practice, however, the FISP typically remains in theater.) Finally, any additional requirements to sustain the deployment would arrive in the Follow-On Support Package (FOSP). Timing of the arrival of the FOSP is dependent on availability of lift resources.

While the above is a simple sketch of a regional contingency, there are a number of key characteristics that define it as a MALSP operation. First, the operational units constituting the deployment are squadrons or higher. Second, logistical support for the deployed aircraft envisions accumulating a large collection of personnel, parts, and equipment and placing them in the same location as the deployed aircraft.

2. Allowancing

In terms of repair parts packages, MALSP envisions a building block concept. Each MALS has the following packages of parts, usually for each type of aircraft the MALS supports: FISP; PCSP; Common Contingency Support Package (CCSP); FOSP; and possibly a Training Squadron Allowance (TSA) for those MALS who support training squadrons.

The FISP comprises sufficient O-Level parts (parts that the OMA may remove and replace) to support a squadron flying wartime hours for 30 days. When not deployed, the FISP is part of the Aviation Supply Officer's "protected stock" and as such may not be used to fill local demand. The PCSP typically supports a squadron of aircraft, or multiple squadrons of aircraft, and contains both O- and I-Level parts to support aircraft flying wartime hours for 30 days. Finally, the CCSP contains items that support multiple T/M/S. Rotary Wing MALSs typically each possess a Rotary Wing Common package, while Fixed Wing MALSs typically possess a Fixed Wing Common package. The FOSP contains approved allowances in excess of those contained in the other packages. Thus in garrison, a MALS's assigned allowance packages typically comprise a FISP for each squadron, a PCSP for each squadron, one CCSP, and a FOSP for each T/M/S.

To fully understand allowancing, it is important to recognize the distinction between consumable and repairable material. Aircraft components that have been deemed economical to repair (e.g., transmissions, avionics, etc.) are known as repairables. Components that do not possess this quality (e.g., bolts, washers, etc.) are classified as consumables and can be thrown away after they are removed from the aircraft.

Due to their high economic value, special rules apply to the management and processing of repairable parts. One such rule is that in order for a MALS to order a repairable item from the Wholesale System, the MALS must have an available quantity. In other words, the MALS's on-hand quantity must be less than the MALS's total allowance quantity, which is the sum of all the allowances for that item in all the MALS's packages. This prevents the MALS from ordering more repairables than they are allowed to possess. The allowance levels are managed at the Naval Inventory Control Point – Philadelphia (NAVICP), which ultimately sets such quantities for all MALSP packages throughout the fleet. Thus, while the MALSP packages contain both types of items, the MALS Aviation Supply Officer has much wider authority to set allowance levels for consumable items.

The value and relative scarcity of repairable items, coupled with the fact that allowance quantities are set by an activity external to the MALS, tends to make repairable asset availability a significant constraint to aviation logistics system performance. This is the primary reason why, for the most part, the subsequent analysis only considers repairable items. The primary take-away from the allowancing process is the fact that the algorithm employed by NAVICP to calculate allowances assumes that the parts packages and the deployed aircraft the packages are supporting are co-located.

3. Level-setting

The term “level-setting” is not a MALSP term per se, rather it is a Legacy methodology for determining the appropriate number of consumable items a MALS ought to retain in its inventory. The algorithm is employed by the R-Supply database to determine the Reorder Objective (RO) for each consumable item, which is the quantity to which stock is ordered.

While the algorithm may take a number of different factors into consideration, it generally calculates Average Monthly Demand (AMD) for the given item on the basis of the past 24 months of demand history and then sets the RO to approximately three times that amount. For items that experience especially high demand frequency, this methodology can result in larger than necessary inventory levels. For items that are ordered infrequently, it can often result in RO levels that are insufficient to support all usage.

In summary, the MALSP construct envisions sending large collections of parts, personnel, and support equipment to support large groups of deployed aircraft. When packages are prepared to support deployed aircraft, algorithms to determine RO levels leverage AMD making them susceptible to expected value propagation problems.

B. MALSP II

MALSP II is the next generation of Marine Corps aviation logistics doctrine. It leverages Continuous Process Improvement (CPI) techniques such as Theory of Constraints, Lean, and Six Sigma to achieve desired or greater readiness with fewer resources. Embracing MALSP II requires a nearly complete transformation of the aviation logistics culture, argues Garant in his 2004 article, “The Transformation of Marine Aviation Logistics.” (*Marine Corps Gazette* pg. 88). Garant highlights that one of the innovations of the MALSP II concept is the manner in which support packages are created. Rather than simply taking x times *Average Monthly Demand* for an allowance level for a particular part, creating a *Buffer* requires consideration of two concepts. The first concept is the *Pattern of Demand*, which consists of the ordering history for the item in question. The second concept is the *Time to Reliably Replenish* (TRR), which takes into account the “worst case” for the time required to replenish (or repair) an item. Properly leveraging these concepts enables logisticians to dramatically improve the effectiveness and responsiveness of the support they provide.

The goal of the new program is to “provide logistics support to deployed and non-deployed core capable units at higher levels of performance while also decreasing the infrastructure and resource inventory” (Steward, 2008 pg. 40). The new concept is “horizontally and vertically integrated from end-to-end” and “focused on material management, maintenance, transportation, information systems, and planning” (Steward, 2008 pg. 41).

1. Scheme of Maneuver

In the most general terms, the innovation of the MALSP II concept is to place a series of physical buffers between the Parent Marine Aviation Logistics Squadron (PMALS) and the

deployed aircraft at the distant end of the supply chain. The central purpose of each of the physical buffers is to provide protection in the form of a physical inventory of material against uncertainties in obtaining replenishment parts from the next node in the network. Figure 1-2 is a graphical depiction of the canonical MALSP II supply network.




Figure 1-2. Notional nodal network

The colored boxes in Figure 1-2 represent the physical buffers that reside at each of the locations. The solid arrows depict material flowing from left to right, while the dotted lines depict demand signals flowing from right to left. The nodes are separated geographically, but more importantly, they are separated by the amount of time it takes to reliably ship material from one node to the next. TRR is the 90th percentile of the distribution of shipping times between two nodes. For example, the four-day TRR between the En-Route Support Base (ESB) and the Main Operating Base (MOB) indicates 90 percent of all material shipped from the ESB to the MOB arrives in four days or fewer. Note, in this example we have a single Forward Operating Base (FOB) and a single MOB, it is entirely possible to have multiples of both types of nodes. While the O-Level and I-Level maintenance construct is maintained under MALSP II, it is important to note that I-Level maintenance occurs exclusively at the PMALS, while O-Level maintenance can occur wherever aircraft are present (e.g., MOB, FOB, etc.).

Operationally, a major difference between the sorts of deployments envisioned under MALSP and those envisioned under MALSP II, is that the latter specifically recognizes the potential for operations at the lower intensity end of the Range of Military Operations. Thus, MALSP II explicitly considers smaller than squadron-sized deployments of aircraft. Further, MALSP II could even support multiple Forward Operating Bases (FOBs) within a theater from which 2 to 4 aircraft might operate.

The network leverages the concept of successive parent-child relationships between the nodes. For example, the FOB is typically the terminal node in the system. A FOB may consist of a detachment with as few as two aircraft and a single supply Marine equipped with a man portable supply package deployed to some remote desert location. Alternatively, a FOB may be as large as two or three squadrons of aircraft supported by a substantial package that could include

dynamic components. The FOB's parent node is typically the MOB. The MOB is usually only one day away from the FOB and is ideally located out of harm's way. Thus, the FOB only requires a single day's worth of material in its buffers, because the MOB can reliably replenish every day. The MOB's parent node is the ESB, which is an optional node whose purpose is to reduce the time-distance between the MOB and the PMALS. If there is no ESB present, then the MOB's parent node is the PMALS. Finally, while it is not a single node, conceptually, the wholesale supply system is the PMALS's parent and ultimately replenishes the entire system.

When the child node receives a demand signal, either in the form of demand from its own flight line or demand from *its* child, it either fulfills the demand with material on hand in its own buffer, or it passes the demand signal along to its parent. If the child is able to issue a part from its buffer to satisfy the demand, then it also originates a request to replenish its buffer that it passes to its parent. Thus, material flows from parent to child, while demand signals flow from child to parent.

2. Allowancing

MALSP II envisions a revised system of allowances for repair parts. They are the Fly-in Support Allowance (FSA), MAG Support Allowance (MSA), I-Level Contingency Allowance (ICA), and Strategic Support Allowance (SSA). The MALSP II Allowance Handbook contains the details and characteristics of these packages.

The FSA is designed to support a specific number (e.g., 4, 8, etc.) and type of aircraft for a period of 10 to 30 days flying at wartime hours. In other words, the bulk of the support (in terms of repair parts) to deployed aircraft is expected to come from the FSA. It is envisioned that the FSAs will contain primarily O-Level remove and replace type parts, but it is also recognized that it may be necessary to draw some parts from the MSA, ICA, or SSA in order to fully support a deployment.

MSAs contain both I and O-Level parts intended to support a specific T/M/S assigned to a given MALS. The allowances will be based on 90 days of usage at peacetime flying rates. ICAs are intended to supplement the MALS's stockpile of I-Level parts in their inventory.

It is important to note that while there exists some conceptual differences between MALSP packages and MALSP II packages, NAVICP still uses the "ARROWS/SPO" model to predict usage and thereby allocate allowances to these packages. NAVICP essentially models the FSA and ICAs as modified PCSPs, while they model the MSA as a modified FOSP.

A major part of this project is to assess the performance of the new MALSP II allowance packages. Thus, implicit in this charge is to assess the validity of the assumption that the new MALSP II allowance packages can be effectively developed using Legacy modeling methodology.

3. Buffer-sizing

The supply chain managers at the PMALS use software that implements the following algorithm to determine the buffer sizes at each node. For a given demand history for a particular item over a period of time of length T days, let the demand experienced for that item on day t be given by d_t . Suppose, for the node in question, that the Design TRR is r days. Let subtotal X_s satisfy the following equation:

$$X_s = \sum_{t=s}^{s+r-1} d_t, s = 1, 2, \dots, T - r + 1$$

There are a total of $T-r+1$ such subtotals for a given length of time over which the demand history is recorded. Let S denote the set of all such subtotals. The maximum buffer level B is given by:

$$B = \max(X_s) \quad \forall s \in S$$

Thus, the maximum buffer level for a particular item is simply given by the maximum historical demand experienced at the node through any duration of time equal to the Design TRR. The equation for B is the theoretical buffer level with the least risk in that it acts upon the maximum observed demand. The buffer sizing software enables the user to choose various levels of risk by filtering out some of the more extreme values in the observed demand pattern. (Note, also, that buffer levels in this context correspond to Reorder Objective under Legacy terminology and may be used interchangeably in this document.)

Consider the case of an item with the Demand Pattern shown in Figure 1-3. The figure shows the quantity demanded each day for one month. Suppose this month is representative of other months in the history so that the Average Monthly Demand for this item is approximately 75. Thus, under Legacy level-setting procedures, the RO would be set at approximately 225 (that is, 3×75). However, buffer sizing avoids the problems associated with expected value propagation, and takes into account the tails of the demand distribution.




Figure 1-3. Daily demand for one month for a notional item.

Consider the problem of setting the appropriate buffer level for a node that has a 4-day TRR from its parent. Thus, examine the demand pattern based on the TRR. Using the notation from above, the first subtotal for a 4 day TRR is given by $X_1 = 12$. This calculation is graphically depicted in Figure 1-4. (According to Figure 1-3, $d_1 = 9$; $d_2 = 0$; $d_3 = 0$; and $d_4 = 3$. So, $X_1 = \sum_{t=1}^4 d_t = 9 + 0 + 0 + 3 = 12$)




Figure 1-4. Calculating the first subtotal, X_1 , using $TRR=4$.

The next subtotal is given by $X_2 = 3$ and is displayed graphically in Figure 1-5. ($X_2 = \sum_{t=2}^5 d_t = 0 + 0 + 3 + 0 = 3$)




Figure 1-5. Calculating the second subtotal, X_2 , using $TRR=4$.

Notice that there are 27 such subtotals for this thirty-day period of demand data. Considering the entire two-year demand history for this item would generate 727 elements in the set S . Suppose the set of all X_S exhibits the empirical cumulative distribution depicted in Figure 1-6.




Figure 1-6. Cumulative Distribution of X_s for notional item and TRR=4.

The maximum quantity ever demanded during any four-day period in the previous two years is 32. Thus, setting the RO to 32 bears little risk of running out of material provided this demand pattern continues and the replenishments arrive consistent with the TRR. The supported unit would have to order more material in four days than at any other time in the previous two years for an NIS to occur. However, there may be managerial reasons to exclude some of the more extreme values in this distribution. Such reasons include if it is known that the largest demand spike(s) occurred because of a one-time order to comply with a technical directive; or if managers may be willing to take on some risk in exchange for reducing the material requirement at the node.

Systematically neglecting the extreme values of this distribution is known as demand filtering. The tool that managers use to set the buffer sizes (RO) for each item is formerly known as Enterprise Logistics Analysis Tool (ELAT). ELAT also has the capability to apply the levels of demand filtering outlined in Table 1-1.

Table 1-1. Demand Filtering Levels

Percentile	Risk Level	Buffer Size (RO)
100	No Risk	32
97	Low Risk	28
94.7	Med Risk	25
87	High Risk	22

In practice, logistics managers tend to opt for “High Risk” demand filtering, because such a buffer level tends to require less material in order to stock the buffers. In the example above, the difference between No Risk and High Risk filtering is 10 items. This tendency to select High Risk filtering is partly an attempt to reduce the Iron Mountain, and partly an effort to reduce the burden placed on the inventory at the PMALS of supporting the deployed detachment. Note, in certain scenarios examined in this project, we may employ non-standard levels of demand filtering, such as 80th or 90th percentile.

In summary, MALSP II envisions a more dynamic and responsive logistics system intended to support aircraft deployments over a wider range of military operations. In addition, the MALSP II concept improves the necessary inventory management process of level setting by accounting for the more extreme values of the demand pattern for each item.

2. The SimKit Model

The model is implemented in Simkit, an open source Discrete Event Simulation package for the Java programming language.² The model consists of an *Execution* class that contains the main method and controls the entire simulation to include the collection and aggregation of various statistics. The model provides a separate class for each of the different nodes in the network. In addition, the class *WholesaleSystem* models the behavior of the wholesale supply system, namely satisfying requests for stock replenishments from the PMALS, as well as Direct Turn-Over (DTO) requirements. Finally, a class entitled *RSupply* monitors outstanding requirements at each node and tracks the status of incoming documents.

In essence, the model applies an algorithm to generate demand from respective flight-lines, and then manages the response of each of the nodes to that demand. When a squadron or detachment orders a part, a requirement for that part resides at the flight-line. The supply node at the flight-line responds by issuing the part and satisfying the requirement or, if none are on-hand, requesting the part from its parent. When a parent fulfills a demand request from its child, it also generates a new requirement in order to replenish its buffer. Each node possesses instance variables that simply count its buffer level, the number of requirements owed to its children, and the number of parts due from its parent. When the part reaches the requirement, the requirement is satisfied.

For simplicity, the model only considers a single item per run as none of the prospective measures of effectiveness requires joint consideration of all items in the system or package.

A. Processes and Events

In addition to the customary SimKit events such as *Run*, each node in the model may possess the capability to execute as many as five other events, specifically related to the MALSP II process. These events are as follows:

DemandArrival

This event signifies the generation of local Organizational (squadron) Level demand. When a node receives a local demand signal, the first step in processing is to determine if sufficient quantities are on hand to satisfy the demand. If so, the part is issued locally and the node schedules an *OrderReplenishmentFromTo* event for its parent. If the item is Not-In-Stock (NIS), the node generates a *PassDTO* event for its parent node.

OrderReplenishmentFromTo(from, to)

A child node schedules an *OrderReplenishmentFromTo* for its parent node in order to communicate a request for replenishment. After the parent node confirms that they are the

²For more information, see: <http://diana.nps.edu/Simkit/>

intended recipient of the replenishment request signal, the node determines if a sufficient quantity is on hand in the buffer to satisfy the request. If so, the parent node schedules a *ReceiveReplenishmentFromTo* event for their child, which effectively issues the child a part from its buffer and it schedules an *OrderReplenishmentFromTo* for its parent in order to replenish the issue. In the event that the current node cannot fulfill the original replenishment request, the node schedules an *OrderReplenishmentFromTo* for its parent, in effect, simply passing the signal along.

PassDTO(from, to)

Child nodes request DTOs from their parents by scheduling a *PassDTO*. The parent first confirms they are the intended recipient of the request. Next, the node determines if sufficient on hand quantities exist to fill the requirement. If so, the current node schedules a *ReceiveDTOFromTo* for their child node and issues the part, which requires the current node to also schedule an *OrderReplenishmentFromTo* for their parent. If insufficient quantity is on hand, the current node passes the DTO signal to its parent with a *PassDTO*.

Note that while the request for DTO signal is passed successively from child to parent until it reaches the PMALS, it is possible for material to be sent from the PMALS and skip the ESB.

ReceiveReplenishmentFromTo(from, to)

Upon receipt of a replenishment document, the receiving node first determines whether any of its children have any outstanding DTOs (if rescreen issues are allowed) or any outstanding replenishment requests. The node also appropriately considers whether any documents are in transit towards those nodes. The following is the hierarchy of outstanding requirements.

- Outstanding DTO at FOB
- Outstanding DTO at MOB
- Outstanding DTO at PMALS
- Outstanding Buffer at FOB
- Outstanding Buffer at MOB
- Outstanding Buffer at ESB
- Outstanding Buffer at PMALS

If a higher order outstanding requirement exists, the current node schedules a *ReceiveReplenishmentFromTo* for its child node, thus forwarding the material towards the requirement. If no higher priority requirements exist, then the node stocks the part on its shelf.

ReceiveDTO(from, to)

Upon receiving an item marked as a DTO, the node checks whether any of its children have outstanding DTO requests. If so, the current node schedules a *RecieveDTO* for the appropriate child node. If no higher order requirements are currently outstanding, the node checks whether it has any outstanding DTOs. If so, it satisfies the requirement. If not, the node stocks the part on its shelf.

B. Nodes

The four types of MALSP II nodes (PMALS, ESB, MOB, and FOB) each possess a common set of instance variables. First, each node possesses a buffer to manage, which includes an RO and current inventory level.

```
private int ro;  
private int inventory;
```

Each node accounts for all DTOs and Replenishment requests it owes to its children.

```
private int childDTOOwed;  
private int childReplenOwed;
```

Each node accounts for all DTOs and Replenishment due to it from its parent.

```
private int ownDTOsDue;  
private int totalDTOsDue;  
private int ownBufferDue;  
private int totalReplenDue;
```

If the node has a flight-line to support, it accounts for outstanding requirements owed to supported squadrons.

```
private int ownFlightLineOwed;
```

Finally, each node possesses the following tally variables for tracking performance statistics.

```
private int issues;  
private int localDemands;  
private int totalDemands;  
private int rescreenIssues;  
private int niss;  
private int ncs;  
private int dtos;
```

Each node manages its own buffer, accounts for parts owed to its children, and accounts for parts due from its parent. However, given characteristics of the nodal network (i.e., FOBs do not have children) nodes differ in the events they are allowed to perform. Table 2-1 outlines the events each node can perform.

Table 2-1. Event List by Node

Event	PMALS	ESB	MOB	FOB	FlightLine	WholesaleSystem
<i>demandArrival</i>	X		X	X	X	
<i>OrderReplenishmentFromTo(from, to)</i>	X	X	X			X
<i>PassDTOFromTo(from, to)</i>	X	X	X			X
<i>ReceiveReplenishmentFromTo(from, to)</i>	X	X	X	X		
<i>ReceiveDTO(from, to)</i>	X	X	X	X		

Finally, the *SimEventListener* construct enables different Simkit objects to “hear” events other objects schedule for them on the event list. Each child-parent pair are *SimEventListeners* of each other. In addition, every MOB and FOB are also *SimEventListeners* of the PMALS node so that they may respond directly to *ReceiveDTO* events from the PMALS.

C. Simplifying Assumptions / Why

1. *Communication between nodes is complete and instantaneous.*

In reality, communication between dispersed nodes is sometimes interrupted or incomplete. However, such deficiencies are typically not consistent enough to substantially affect long-term system performance. In addition, it is unlikely that the system configuration and business rules analyzed in this project will mitigate this problem in any way, if it does exist. Finally, newly developed information technology systems, such as the Expeditionary Pack Up Kit (EPUK), are intended to further improve communication within and between nodes.

2. *Only Hi-Priority Demands are considered.*

Only high-priority demands have the potential to directly impact readiness, thus, they are the focus of this research.

3. *Inventory Management is perfect and complete.*

While in reality, an item might erroneously be declared “Not-In-Stock” due to an inventory discrepancy, this problem is abstracted because the target of this project (i.e., system level business rules and allowancing decisions) is unlikely to mitigate or otherwise affect this problem in any way.

4. *All nodes - especially FOB/PMALS - can “see” the requisition’s scheduled delivery date.*

In reality, aviation logisticians make maintenance and supply decisions on the basis of the best known status of incoming requisitions and expected delivery dates.

5. *Wholesale System sends all material to PMALS.*

In practice, it is sometimes possible to arrange for the Wholesale System to send both DTO and replenishment requests directly to the node that ordered the requirement. This arrangement is

only possible for a certain portion of consumables and is expressly prohibited in the case of repairable items. For the sake of conformity we do not allow this capability in the model.

D. Additional Modeling Rules

In addition to the general description of the model above, we create three additional modeling rules that incorporate elements of reality into the model as necessary. The first rule enables the inclusion of both consumable and repairable items in the model; the second rule enables child nodes to provide material to parent nodes; and the final rule models the behavior of the Intermediate Maintenance Activity.

1. Quantity (Q)

In its baseline form, the model assumes a quantity of one for each document. However, in order to consider consumable items, the model must accommodate document quantities greater than one. Thus, we activate rule **Q** whenever the set of NIINs under consideration include consumable NIINs.

The following assumptions apply when rule **Q** is activated:

Q1. Allow partial issues. Partial issues occur when the on-hand quantity at a given node is greater than zero, but less than the quantity demanded. The on-hand portion of the requirement is immediately fulfilled and a document for the remainder is referred to the node's parent. However, if rule **L** is active, the current node solicits lateral support from each of its child nodes prior to referring the document.

Note: Given **Q**'s application to consumable items only, it is incompatible with rule **M** below. However, it is compatible with rule **L**.

2. Intra-Network Lateral Support (L)

If physical buffers for a NIIN exist at every node in the MALSP II network, then it is typically the case that if the Wholesale Supply System fails to keep up with demand, the buffer at the PMALS will "dry up" first, followed by the ESB (if present), then the MOB, etc. In other words, parents tend to run out of scarce items before their children do. This is also true in the event that scarce low density items are placed at forward / junior nodes.

When rule **L** is activated, any node that determines its current inventory is insufficient to meet an immediate flight-line demand (in other words, the node determines it is about to go NIS on that item) it will solicit a lateral support request from all of the nodes junior to it in the network.

The following assumptions apply when **L** is activated:

L1. Lateral Support requests solicited from nodes in order of proximity to requesting node. This is the simplest way to implement this rule, though it is entirely possible that other implementations are superior in some way.

Note: Rule **L** is compatible with analysis of both consumable and repairable items.

3. Maintenance Actions (**M**)

In reality, in order for the squadron to obtain a Ready-For-Issue (RFI) repairable item from the Supply Officer they must turn in a Non-RFI (NRFI) carcass. The system also recognizes a small proportion of repairable items as Remain-In-Place (RIP), a squadron that orders such an item is granted an additional 24 hours from the time they were issued the RFI item to provide a turn-in to the Supply Officer.

When the PMALS receives a repairable demand from a squadron on its flight line, if the item is in stock it is issued and the NRFI carcass is obtained in exchange. The carcass is inducted into the IMA for repair as a Supply Officer Asset (SOA). If the repair is successful and results in an RFI, the item is placed in the PMALS inventory. If the repair effort is unsuccessful and results in a Beyond Capability of Maintenance (BCM) action, the carcass is returned to the wholesale system (most likely sent to Depot or commercial vendor for repair) and an *OrderReplenishmentFromTo* event is scheduled to replenish the PMALS's stock.

If the item is NIS, the NRFI asset is inducted into the IMA as an Expeditious Repair (EXREP). If the IMA repairs it, the item is returned to the squadron to fulfill their requirement. If the item is BCMd, the NRFI asset is turned into the system and a Direct-Turn-Over (DTO) document is ordered on behalf of the squadron. This process is implemented in the model, subject to the caveats described below concerning RIPs, etc.

In practice, as in the model, when a child node receives a repairable demand and the item is in stock, the RFI item is issued and the carcass is sent back to the PMALS for induction into the IMA. The disposition of that item is identical to that of the SOA from the PMALS.

Finally, if the child node, e.g., the MOB, experiences a demand for a repairable item but is NIS for that item, the NRFI turn-in becomes an EXREP. In practice, it may be sent back to the PMALS for repair (this can be time consuming), however, in this model it is assumed such EXREPs are immediately BCMd and an internal DTO document request is sent from the child to the parent.

It is important to note that the IMA does not have permission to attempt repair on all repairable components. These items, typically identified with an "X1" Individual Component Repair List (ICRL) Code, in practice are technically inducted into the IMA, but are processed for BCM nearly immediately.

The following assumptions apply when **M** is activated:

M1. No RIPs are allowed. This is necessary for simplification of the process. For the most part, this assumption primarily impacts the time required to replenish the Supply Officer's shelves at various nodes, rather than directly affecting the time that an aircraft is down for a part.

M2. All EXREPs forward of PMALS are BCMd. This assumption avoids the computationally complex task of tracking an EXREP carcass back to the IMA and then fulfilling the requirement upon successful repair of the item. In many cases, the shipment-repair-shipment of this item would take longer than it would take for another node in the network to fulfill the requirement or for an external activity to provide support (e.g., lateral support from a Navy activity). In practice, few logistics managers would wait for the shipment-repair-shipment of a carcass from a forward node anyway, so this assumption should have minimal influence on system performance.

M3. All NRFI items that are Supply Officer Assets are inducted. This should be uncontroversial, as it closely matches reality.

M4. Items for which the IMA lacks repair capability are inducted and immediately BCMd. This should also be uncontroversial. The implication is that repair times for all Non-X1 items are drawn from the same distribution, regardless of the ultimate outcome. In other words, the times to repair and the times to BCM an item are drawn from the same distribution.

E. Analytical Plan

The model developed in this project is a Discrete Event Simulation, implemented in Java using the SimKit programming libraries. It is employed to analyze aspects of MALSP II doctrine, however, it is sufficiently flexible to analyze aspects of the Legacy MALSP concepts as well.

The same general plan applies for analyzing each of the research questions with which we are tasked. First, determine the optional special business rule(s) that are necessary to include in the model to address the given research question. Second, determine the nature of the inputs that best address the question. For example, actual demand data is used for the NIINs included in the allowance packages to address question one, while generic NIINs with notional demand data is used for all other questions. Third, develop and execute the experiment. Given that most of the variables are categorical, as well as the relatively fast simulation time, full factorial experimental designs are implemented.

Research Questions

The following research questions are addressed in this document (Note: question numbering is preserved from the Statement of Work):

1. *How can allowancing at the PMALS improve or facilitate the building and managing of MALSP II packages?*

The relative abilities of the Legacy MALSP allowance packages, and proposed MALSP II packages for CH-53Es at MALS-16, are assessed to support a notional MALSP II nodal laydown. Employing the proposed MALSP II allowancing to support a MALSP II logistics network significantly improves certain system performance measures (both practically and statistically) relative to supporting the same system using Legacy MALSP allowances under a wide array of circumstances.

2. What is the optimal/robust criteria for including an item in a packup?

An experiment was designed and implemented with the intent to glean information regarding the expected performance of various types of items stocked in a logistical network. The rules deduced in this section are robust against numerous sources of variance and are applicable to many sorts of MALSP II deployments.

4. How does uncertainty regarding Actual TRR effect Response Time at deployed nodes?³

In this section, an experiment was designed and implemented with the intent to flex the model with respect to the effect of differences between the Design and Actual TRRs. The extent to which such uncertainty may adversely impact system performance in terms of Response Time was identified, in addition to the identification of business rules to mitigate these effects.

Time constraints preclude addressing the following questions and are therefore left for future work:

3. What is the optimal/robust criteria for removing an item from a packup?

5. What is the optimal/robust criteria for positioning low density items?

[Note: The detailed analysis in Chapter 3E informs this issue.]

6. Where should repair capability (i.e. IMA, or T-AVB, etc.) be placed in the nodal lay-down, if at all?

³ On the Statement of Work, this question appears as “How frequently should buffers be re-sized? Are there useful leading indicators (i.e., say, if Actual Time to Reliably Replenish (TRR) exceeds Design TRR by a certain amount).”

3. Input Validation

In this section, the rationale for selecting the input parameters for the model is outlined. In this sense, the model is rather straight-forward in that inputs involve modeling Demand Frequency, Demand Quantity, Inter-Nodal Shipment times, and Maintenance Actions.

A. Demand Frequency

The frequency of demand events (i.e., the number of documents ordered regardless of total quantity) for a particular item for a given period of time, say day or month, is commonly modeled as a Poisson Process (Law and Kelton, 2003, pg. 325). Thus, the number of requisitions ordered per month can be modeled as a Poisson distribution with rate parameter λ , where λ is the mean number of requisitions/month.

The ability to effectively model demand frequency as a Poisson process enables the leveraging of the following two characteristics to greatly simplify the model: (a) Inter-arrival times are distributed exponentially ($1/\lambda$); and (b) the sum of two independently distributed Poisson random variables is also a Poisson distribution.

Suppose the count of demands for a given item against a particular aircraft is a Poisson distribution (λ). Since the focus is on CH-53Es in this project, and CH-53E squadrons are assigned 16 aircraft, then it is the case that the count of demand events for a given item for a particular squadron is a Poisson distribution ($16\cdot\lambda$). In this manner, the distribution of inter-arrival times can be scaled to account for any number of aircraft at a particular node. For example, the inter-arrival times for a detachment of 4 aircraft would be Exponentially distributed $\left(\frac{1}{4\cdot\lambda}\right)$.

Figure 3-1 (next page) displays histograms of demand inter-arrival times of four different NIINs managed at MALS-16.




Figure 3-1. Inter-arrival times for four MALS-16 NIINs.

While this is hardly an exhaustive list, it is indicative in that those NIINs with the most demands tend to fit an exponential fairly well. Typically, if there is a divergence, it is because the tails of the exponential are lighter than the tails of the empirical distribution.

Thus, strong theoretical justification exists for assuming exponential inter-arrival times. Comparing empirical distributions for NIINs with sufficient data confirms that such an assumption is at least plausible. And while a large portion of NIINs have low demand and therefore insufficient data to draw strong conclusions regarding an empirical distribution, the fact is that these items have lower demand anyway; thus, the actual choice of distribution (e.g., exponential, uniform) is less likely to dramatically affect the results.

B. Demand Quantity

Relatively few models of logistics systems model the demand signal along two margins, both frequency and quantity, as is done in this analysis. Modeling demand in this manner makes analysis of consumable items (which may have document quantities greater than 1) and

repairable items (which all have document quantities equal to 1), significantly easier. But it also brings a higher level of fidelity overall at only a small cost of increased complexity.

Pre-Expended Bin items are not considered in this analysis. Such items are typically ordered in quantities of 100 or greater. They are typically easy to obtain in the wholesale system and are of such high demand and high volume that Supply Officers essentially allow each squadron to manage their own inventories of these items in Pre-Expended Bins. It is for these reasons that extremely high demand quantities can be neglected from this analysis.

In the following analysis, all repairable demand is modeled with a quantity per document of 1. Documents for consumable items may also have a quantity per document of 1. (This is indicated by the DemandQuantType = “single”.) All other document quantities use a Triangular (a, b, c) distribution. Triangular distributions are intuitively sensible because the parameters are minimum, maximum, and most likely.

Figure 3-2 is a histogram of all document quantities for High-Priority (Hi-Pri) documents between 2006 and 2011 ordered by the detachment at HOA. A document is deemed Hi-Pri if the aircraft cannot fly or otherwise perform one of its missions without the part. Notice that well over half of all documents during this timeframe have a quantity equal to 1, and at least 90% have a document quantity less than 4.




Figure 3-2. Distribution of Quantity per Document for HOA: 2006-2011

One important benefit of the two-dimensional modeling of demand frequency and quantity is that it facilitates independence between these two aspects of demand. So, as frequency distributions are scaled up and down according to number of aircraft or wartime intensity, the distribution governing quantity per document remains unchanged.

C. Shipping Times

An important innovation for this project relative to other research projects that address this and similar topics is modeling the inter-nodal shipping times as a distribution. Previous efforts have simply set the ship time between two nodes deterministically at the given TRR. This simultaneously overestimates the average time it takes to ship items between those nodes (i.e., in reality, 90% of items arrive *sooner* than the TRR), while not allowing any shipment to “break” TRR. Modeling the inter-nodal shipping times as a distribution avoids both of these unnecessary departures from reality.

In this project, all inter-nodal shipping times are modeled as Lognormal (μ , σ^2). According to Law and Kelton, the theoretical reasons for selecting a Lognormal distribution are for “Time to perform some task ... [and for] quantities that are the product of a large number of other quantities (by virtue of the Central Limit Theorem),” (Law and Kelton, 2003, pg 307). Both provisions apply in this case.

In addition, the empirical reasons for selecting the Lognormal distribution are also quite strong. Figure 3-3 contains a histogram of all shipments between the ESB and FOB during a six month period in 2011. The red line is a continuous fit of a Lognormal distribution.




Figure 3-3. ESB to FOB Ship-times

Each inter-nodal shipping time is characterized by three parameters: the μ and σ^2 of the internal Normal distribution, as well as a shift parameter. The shift parameter is necessary because the minimum value of the Lognormal is 0, however, this would allow nearly instantaneous shipments between nodes. Thus, either the minimum empirical value, or reasonable minimum, is assigned

to the shift parameter. The parameters μ and σ^2 are then selected to obtain appropriate values for the mode and 90th percentile.

D. BCM/RFI Rate

The BCM rate is the proportion of carcasses for a given NIIN that are BCMd after induction into the maintenance cycle at the IMA, over a particular interval of time. It is obtained as a population proportion where $P_{BCM} = \text{Number of BCMs} / \text{Number of Completed Maintenance Actions}$. If the BCM rate for an item is P_{BCM} , then the RFI rate is given by $P_{RFI} = 1 - P_{BCM}$.

A histogram of the empirical BCM rates for the 432 repairable NIINs contained in the proposed MALS-16 packages is provided in Figure 3-4. The take-away here is that given the spread of BCM rates, we assign each NIIN a unique BCM rate equal to its empirical rate.




Figure 3-4. Distribution of BCM rates

The modal value of estimated BCM rates is 1.0, due to the relatively large proportion of X1 items. X1 items are those repairables that neither the O- nor I-Levels are allowed to attempt repair and are therefore sent to Depot. Notice that the rest of the rates are rather uniformly distributed between 0 and 1. Thus, each NIIN is assigned its own unique BCM rate equal to its empirical rate.

E. Maintenance Time: TRR_M

Given the wide variance in the quantity of data for maintenance times, the time that a carcass spends in maintenance as a Triangular (0, c, 0) distribution was modeled. While there probably exists sufficient theoretical justification for modeling maintenance time as a Lognormal

distribution (or perhaps others), since many NIINs under consideration lack sufficient observations, it is more reasonable to employ a more general distribution.

The Triangular (0, c, 0) distribution is a generally accepted distribution to use in the absence of sufficient empirical data. In addition, it has the intuitively appealing property of user assigned minimum, maximum, and modal values. Just as in reality, the most frequently observed time is typically very small. For the maximum value parameter, c, $t+b$ is selected such that t matches the NIIN's empirical TRR_M or some other reasonable value.

Solving for b in Figure 3-5 requires trigonometry. Suppose that the area of Triangle OHV = 1.0 and that the area of Triangle NMV = 0.1.




Figure 3-5. Construction of Triangular Distribution of Maintenance Times

The following quadratic equation describes the relationship between t and b : $9b^2 - 2tb - t^2 = 0$. In this case, t is set equal to either the empirical TRR_M , or a reasonable estimate thereof, and solve for b . This ensures that 90% of the observations drawn from the random distribution are smaller than the empirical TRR_M .

F. Buffer Sizes

While the process of modeling the intricacies of the arrival of demand signals is described above, a necessary, though indirect input, into the model are buffer sizes, or more accurately, the amount of demand filtering that corresponds to the determined ideal buffer level. In practice, supply Marines use a tool formerly known as ELAT to determine the appropriate physical buffer levels at each node. The inputs to that algorithm include the empirical demand data (namely quantities and inter-arrival times), a desired amount of filtering of outliers, and the TRR between the node in question and its parent. The study team did not have access to ELAT, or its latest

incarnation; although, it is not clear whether access would have helped, due to the copious use of notional data.

Thus, the ELAT algorithm was implemented in Java, using demand generated exactly the way in which it is generated in the model. In addition, demand filters may be selected from the 80th percentile to the 99.5th percentile in the same way as implemented in ELAT. So, while the output from our version of ELAT has not been compared to the actual version of ELAT, the actual ELAT's algorithms are rather straightforward and easily replicated. (See Chapter 1, Section B for an in-depth description of buffer sizing and demand filtering.)

G. Summary of Validation

The model described in this report is relatively straightforward and has a small number of inputs. In most cases, sufficient data exist to confirm the suitability of modeling a particular empirical distribution in a particular manner. In cases where data is more difficult to obtain, by NIIN, compelling theoretical reasons exist as justification for selecting a particular distribution or algorithm. However, it is still true that this report relies on High-Priority data ordered from HOA from 2006 to 2011, as well as two years' of data on repairables for CH-53Es from MAL-S-16. While it is likely other items behave similarly under other circumstances, it would be necessary to re-validate the model for other T/M/S. The process to validate the use of particular input parameters for a different situation would be identical to the one performed in this study.

4. Assessment of MALSP 2 Packages

In the following section, the primary research question answered is: *1. How can allowancing at the PMALS improve or facilitate the building and managing of MALSP II packages?*

The relative abilities of the different allowance structures (Legacy and MALSP II) are assessed to support a notional MALSP II nodal laydown. Employing the proposed MALSP II allowancing structure to support a MALSP II logistics network significantly improves certain system performance measures (both practically and statistically) relative to supporting the same system using Legacy MALSP allowances.

- MALSP II allowances unambiguously improve performance for MALS with no deployed squadrons.

- MALSP II allowance packages achieve superior MOB Response Times relative to Legacy allowance packages in all cases considered.

- Implementing an ESB increases Response Time at the node under consideration for both Legacy and MALSP II Allowance packages.

- Recommend not stocking repairables at the ESB, except for certain low density NIINs under a narrow range of circumstances.

A. Scope

The dataset “Candidates_rac.xlsx” contains recommended MALSP II allowances for support of CH-53Es at MALS-16. Its range is 565 repairable NIINs. It is intended to replace the Legacy MALSP allowances currently in place. This section addresses the question of whether the new MALSP II allowances sufficiently support the MAG’s ability to deploy a detachment of aircraft and support them via a MALSP II nodal laydown.

The Legacy allowances consist of the MALSP packages that reside at MALS-16. For the purposes of the current analysis, these are:

- 2 Peculiar Contingency Support Packages (PCSP)
- 1 Fly-In Support Package (FISP)
- 1 Intermediate Operations Support Package (IOSP)
- 1 Rotary Wing Common Support Package (CCSP-RW)

The MALSP II allowances consist of the following proposed packages:

- 1 MAG Support Allowance (MSA)
- 1 I-Level Contingency Allowance (ICA)

- 1 Strategic Support Allowance (SSA)
- 2 Forward Support Allowances (FSA)

The range and depth of each package is drawn from the Microsoft Access database “ch53e_MALSP II_Mar12”. The quantity of packages was confirmed by a Subject Matter Expert (SME) or was an assumption based on Study Performer's knowledge.

Rules employed: **L, M**
NIINs: actual

In summary, the relative performance of Legacy allowances versus MALSP II allowances in support of a MALSP II logistical support scheme must be assessed. The logistical system in the model is capable of intra-network lateral support, meaning that if a requirement for a part of an aircraft at the PMALS is NIS, the PMALS can solicit lateral support requests to both the ESB (if applicable) or the FOB. The Intermediate Maintenance Activity is modeled as well. Thus, NRFI Supply Officer Assets originating from anywhere in the network, as well as EXREPs from aircraft at the PMALS, are inducted into the IMA and delayed in a manner consistent with empirical maintenance times.

B. Data and Inputs

The Aviation Financial Analyst Tool (AFAST) is a database that aggregates financial information, to include demand data, from all naval aviation activities. The AFAST dataset contains all requisite demand history data; while the dataset contains 565 NIINs, 432 were successfully matched to the “AFAST_Repairables_M16.xlsx” dataset. The 432 NIINs with demand history successfully joined is found in “M16_cand_afast_join_29_july.xlsx”.

Finally, of the 432 repairable NIINs, the MALSP II Program Office successfully matched those to 416 NIINs with maintenance history (i.e., total BCMs, total RFIs, repair times, etc.).

1. Demand Frequency

As mentioned above, the NIINs are matched with recommended allowances and with demand data from AFAST. The intent is to identify only the high-priority demand NIINs from squadrons supported by MALS-16 over a period of 24 months. MALS-16 supported three squadrons during this time period (calendar year [CY] 2010-2011). Figure 4-1 contains a histogram of the Average Monthly High-Priority Frequency (AMF) for each NIIN during this time period.




Figure 4-1. Mean (G-Series Document per Month) for each NIIN for CH-53E Squadrons at MALS-16

Thus, for each NIIN, the AMF is divided by 3 to obtain the expected AMF for a single squadron. The reciprocal of that value provides the μ parameter in that NIIN's Exponential distribution (μ) from which are drawn the demand inter-arrival times. This distribution is not varied in the experimental design described below, except to scale them as appropriate to account for the number of aircraft at a particular node.

2. Maintenance Parameters

The dataset "niin_sum_w_TRR_sub_M.xls" contains the maintenance history as provided by the MALSP 2 Program Office.

The BCM rate employed in the model for each NIIN is a simple proportion of number of BCMs / number of Maintenance Actions. This parameter does not change in the experimental design.

It is important to identify a special category of repairable item (ICRL code: X1) for which the PMALS does not have permission and/or the capability to attempt to repair. In reality, such items are automatically BCM'd and immediately returned to the Depot (or equivalent) for repair (thereby enabling the PMALS to order a stock document to replace the BCM'd item). While the datasets provided lack the ICRL or SM&R codes necessary to unequivocally determine whether an item is X1, we assume that all items with 100% BCM rates and extremely low maintenance times are X1 items. In the model, all X1 items are immediately BCM'd and are subject to the TRR for wholesale stock of 25 days. The TRR for stock documents is denoted TRR_{Stock} .

The final element of maintenance activity is the time taken to either successfully repair or BCM the item. For NIINs with a sufficiently large number of observations (e.g. $N > 30$), TRR_M is set to the 90th percentile of the distribution of those repair times. The observations of all other low demand, non-X1 items are aggregated into a single distribution. Table 4-1 shows a break-down of the 416 NIINs along these two margins.

Table 4-1. Distribution of NIINs across maintenance data categories.

		Frequency		
		High	Low	Totals
Repair Capability	Yes	67	187	254
	No	10	152	162
Totals		77	339	416

Thus, repair capability exists on 254 of the NIINs, and just over one quarter of those have sufficient observations to assign a unique TRR_M .

Consistent with the assumptions contained in “MALSP II Allowance Handbook”, the TRR_{Stock} for each item under consideration is set to 25 days. Setting the physical buffer levels at the PMALS requires the calculation of a composite TRR. This is due to the fact that the length of time it takes for a carcass to be replenished, from NRFI carcass turn-in to RFI on the Supply Officer’s shelf, depends greatly on whether the item is successfully repaired at the IMA. The formula for the composite TRR is given below.⁴

$$TRR_{Comp} = P_{BCM} \cdot (TRR_{Stock} + TRR_M) + (1 - P_{BCM}) \cdot TRR_M$$

Notice that TRR_M appears in both terms because the maintenance time is always drawn from the same distribution, regardless of outcomes. In other words, it is assumed that the time the item spends in maintenance is independent of the outcome of the maintenance action. Finally, notice TRR_{Stock} only appears in the event that the item is BCM’d.

C. Experimental Design

A full factorial design is employed using the factors and levels described in figure 6, below. It is slightly unbalanced due to the fact that the factors Deployed a/c, ESB, and TRR do not apply to the scenario with zero deployed aircraft.

⁴ This is from the author’s recollection of a conversation with Mr. Lauren Eck in the summer of 2009.

$$\begin{array}{ccccccccccc}
 & \text{NIINs} & & \text{deployed a/c} & & \text{ESB} & & \text{TRR} & & \text{Homeguard PMALS Scenario} & & \text{Allowance Structure} \\
 (& 432 & \times & 3 & \times & 2 & \times & 2 & + & 432 &) \cdot 2 & = 11,232 \\
 & & & 16, 8, 4 & & \text{Yes, No} & & \text{High, Low} & & & & \text{MALSP, MALSP 2}
 \end{array}$$

NIIN – Each of the 432 NIINs are considered. Note that each has unique demand frequency, BCM rate, and TRR_M .

Deployed a/c – The number of aircraft deployed to the FOB. The experiment envisions a total of three squadrons of 16 aircraft each. Any aircraft not deployed remains with the PMALS.

ESB – An ESB is present in certain configurations considered.

TRR – TRR between the PMALS and the MOB is varied from 6 days (Low) and 15 days (High).

Allowance Structure – These are the primary factors of the experiment.

The parameters associated with the 432 NIINs are described in the previous section. Demand and maintenance parameters remain unchanged and are unique to each NIIN. Figure 4-2 outlines the details and nomenclature related to the network configurations considered in the experiment.




Figure 4-2. Network Configurations to include TRR levels.

Notice that Configuration 1_ employs a single node, denoted the MOB, in support of the deployed aircraft. Configuration 2_ employs an ESB. Finally, the *a*, *b*, and *c* designators refer to the number of aircraft deployed to the MOB.

Each design point is run for a duration of 720 days and is replicated 100 times. Common Random Numbers is employed in an effort to reduce the variance between replication and improve the power of subsequent statistical analysis.

Filling Physical Buffers

In order to fill the physical buffers at each node, an ELAT like algorithm is employed. High risk demand filtering (87th percentile) is applied in order to maintain consistency with the assumption of the "MALSP II Allowance Handbook". (Note: Subsequent analysis could easily consider different levels of demand filtering.) The following steps apply:

1. Determine Buffer Levels (Ideal Reorder Objective) at each node using ELAT methodology.
2. Sum the Ideal ROs across all nodes in the network. This is the total material requirement for the NIIN, given the particular configuration of the logistics network.
3. Assuming all material in either allowance structure (i.e., Legacy or MALSP II) is on-hand and unrestricted (i.e., material in the FISP is accessible), calculate Fill Proportions and Deficient NIINs.

Allocate allowances using the following algorithm for Configuration 0:

4. Fill PMALS

Allocate allowances using the following algorithm for Configuration 1_:

4. Fill MOB\
5. Fill PMALS
6. Fair share any excess

Allocate allowances using the following algorithm for Configuration 2_:

4. Fill ESB (nearly all ESB allowances = 1)
5. Fill MOB
6. Fill PMALS
7. Fair share any excess between PMALS and MOB

The fill algorithm places a high priority on filling the buffers at forward nodes. Admittedly, this is a one-size-fits-all rule that is easily implemented in a spreadsheet to efficiently process numerous NIINs for thousands of unique scenarios. However, in practice the MALS Supply Officer will likely place a similarly high priority on support to deployed aircraft, thus ensuring that buffers at deployed nodes are filled first in many or most cases.

Assessment of Capacity of Packages to Fill Physical Buffers

Before executing the experimental design, each allowance structure's relative capacity to fill the ideal physical buffer levels at the required nodes must first be assessed. Let a NIIN's Fill Proportion be: total allowance available / total required at each node. It follows that the Average Fill Proportion for an Allowance Structure is the same measure averaged across all NIINs. Panel A in Figure 4-3 shows the Average Fill Proportion using Legacy allowances and Panel B shows the same for MALSP II allowances.




Figure 4-3. Average Fill Proportion for Legacy (Left) and MALSP II (Right) Allowances

First consider Configuration_0 displayed in the left third of each of the charts. Notice that it only contains one blue bar because only the PMALS has any allowances in this configuration. In the Legacy chart (left), the Fill Proportion for the PMALS is approximately 0.9, while the Fill Proportion for the PMALS is nearly 1.0 in the MALSP II chart (right). Thus, MALSP II allowances enable greater Average Fill Proportion relative to Legacy allowances for these NIINs.

Consider Configuration_1a (32 a/c at PMALS; 16 a/c at MOB) displayed in middle third of each chart. Notice they contain both blue and green bars because the PMALS and MOB both have physical buffers in this configuration. In this case, the Fill Proportion at the PMALS is higher with the Legacy allowances (left) than MALSP II allowance (right) (0.8 vs. 0.7); while the Fill Proportion at the MOB is higher with MALSP II allowances (right) (0.99 vs. 0.9).

Finally, consider Configuration_2a (32 a/c at PMALS; 16 a/c at MOB; w/ ESB) displayed in right third of each chart. In this case, the MALSP II packages achieve a higher Fill Proportion at the ESB (1.0 vs. 0.9) relative to Legacy and at the PMALS (0.64 vs. 0.58); but the Legacy packages achieve higher Fill Proportions at the MOB.

A Deficient NIIN is one for whom their total ideal requirement (the sum of Ideal ROs across all nodes) exceeds the total allowance for a given package. The shortfall is the amount by which the allowance is deficient. Panel A (left) of Figure 4-4 compares the number of deficient NIINs by allowance structure and configuration. Panel B (right) of the same figure compares the magnitude of their respective shortfalls.




Figure 4-4. Deficient NIINs (Left) and Total Shortfall (Right) Comparisons

For Configuration_0, MALSP II allowances yield fewer deficient NIINs relative to Legacy allowances. Under all Configuration_1 structures, MALSP II allowances yield a greater number of deficient NIINs relative to Legacy allowances. In fact, 30% of the NIINs in the package are deficient under MALSP II allowances. Finally, under all Configuration 2 structures, MALSP II allowances yield fewer deficient NIINs relative to Legacy allowances. However, 40% of the NIINs in the package are deficient under MALSP II allowances.

A tradeoff appears to exist between the Fill Proportions at the MOB, ESB, and PMALS. To determine whether the appropriate trade-off has been made, the performance of the system with respect to Response Time and Supply Effectiveness at each node must be examined.

D. Results

In this section, the output of the experiment is examined to determine any conclusions that might be drawn.

1. Measures of Effectiveness

Response Time

The amount of time required to fulfill a particular flight-line requirement is known as Document Time. The duration of each simulation replication is 720 days. So, for a given replication and a given NIIN, the NIIN Level *Response Time* at a particular node is simply the sum of all the Document Times for each document experienced at that node during that run. Since each design point in the experiment is replicated 100 times, the NIIN Level *Response Times* is aggregated over all replications. So, for a given design point, there is one Mean NIIN Level *Response Time* for each node in the configuration for each of the 432 NIINs. The Package Level *Response Time*, or simply *Response Time*, is the sum of these 432 Mean NIIN Level *Response Times* at the

respective node. Thus, a node's *Response Time* (MOB, PMALS, etc.) is the sum of the 432 Mean NIIN Level *Response Times* experienced at that node.

Net Supply Effectiveness

In practice, a MALS Net High-Priority Effectiveness is the proportion of all carried demand for high-priority items that are filled from on-hand stock. Because in this project only high-priority items were considered, the “high-priority” aspect of the nomenclature was occasionally suppressed. Note also that Net Supply Effectiveness was treated as a characteristic of a particular node. However, it would not be improper to aggregate this measure across all nodes to measure the performance of the entire system in the aggregate.

For a particular 720-day replication, the Net Supply Effectiveness for a given NIIN at a given node is the number of documents ordered at that node and fulfilled immediately from stock on-hand at that node divided by the total number of documents for that NIIN ordered at that node. Because each design point is replicated 100 times, the numerator is technically the sum of all documents fulfilled immediately (in all 100 replications) divided by the sum of all documents ordered (again, in all 100 replications). In order to obtain the Net Supply Effectiveness for a given node, take the average of all NIIN-level Net Supply Effectiveness measures weighted appropriately by volume of demand.

2. Configuration_0

The first configuration is the easiest to assess. All three squadrons are at “home” with the PMALS. Figure 4-5 compares the Expected PMALS Response Time exhibited for each allowance structure.




Figure 4-5. PMALS Response Time for Configuration 0.

For the 432 NIINs under consideration, the MALSP II allowances yield a ~65% reduction in PMALS Response Time.

3. MOB Response Time

Figure 4-6 compares MOB Response Time for the chosen configurations and two allowance structures. (Note TRR = High. For the corresponding graphs pertaining to TRR = Low, refer to the Mathematical Appendix.)




Figure 4-6. Aggregate MOB Response Time (TRR = High)

Notice, that Configuration_1a and Configuration_2a (both to the left of the first dashed vertical line) are the scenarios that involve 1 full squadron deploying to a MOB. Configuration_1a is without an ESB, while Configuration_2a includes an ESB. It makes sense to compare these configurations under Legacy and MALSP II allowances simultaneously. Essentially, this enables the decision of whether or not to include an ESB contingent on the relative performance of the different allowance structures.

MALSP II allowances provide superior performance relative to Legacy allowances in each of these configurations. In addition, this advantage appears to increase as the number of deployed aircraft increases. Finally, and perhaps unsurprisingly, MOB *Response Times* decrease as number of deployed a/c decrease.

Most surprising of all, adding an ESB dramatically *increases* response times. Most other analyses of the ESB assume unlimited material availability, or at least sufficient allowances to fill all buffer requirements. Section E is a detailed analysis into the causes and implications of this behavior.

4. PMALS Response Time

Figure 4-7 below illustrates Expected PMALS *Response Time* for a number of configurations and allowance structures.




Figure 4-7. Aggregate PMALS Response Time (TRR = High).

As above, the left third of the graph displays scenarios in which 16 aircraft are deployed, the 8 aircraft deployed scenario is in the middle panel, and the 4 aircraft deployed scenario is in the right panel of the graph. Legacy allowances without an ESB achieve the best performance, followed closely by both MALSP II implementations; however, the difference is slight and not practically significant.

PMALS *Response Time* increases as the number of deployed aircraft *decreases*. This is due to the fact that a particular amount of material must be removed from the PMALS and placed at additional nodes (i.e., MOB, ESB) in order to support any deployment. Increasing the number of deployed aircraft results in a directly proportional decrease in local demand at the PMALS, but because of the interaction between the TRR to the distant nodes and demand pattern, the amount of additional material required is substantially less than proportional in most cases. Thus, the PMALS experiences considerably less local demand with only slightly less material in its local inventory, which tends to improve PMALS *Response Time*.

As with MOB *Response Time*, adding an ESB *increases* PMALS *Response Time*. The effect is much greater for Legacy Allowances. Notice that the magnitude of the MOB *Response Times* is substantially less than PMALS *Response Times*. This is due to the buffer filling algorithm placing a higher priority on filling MOB buffers.

Table 4-2 is a summary of the differences in *Response Times* at each node for those MALSP II scenarios with an ESB and those without.

Table 4-2. Summary of MALSP II Differences in Response Times with and without ESB

		TRR = High			TRR = Low		
Response		config_a	config_b	config_c	config_a	config_b	config_c
MOB	difference	621.6	316.2	133.2	223.8	112.8	46.8
	p-value	0.028	0.026	0.195	0.072	0.063	0.420
PMALS	difference	795.0	1018.5	1014.2	169.6	176.6	153.0
	p-value	0.162	0.080	0.095	0.870	0.890	0.934

Each of the differences are positive, as also indicated in Figures 4-6 and 4-7, which illustrates that the inclusion of an ESB increases *Response Time* at each node. Further, those differences that are statistically significant at the 0.10 level of significance are shown in bold. Differences for MOB Response Time are statistically significant for all scenarios except the smallest deployments of aircraft (4 aircraft). In contrast, the only statistically differences for PMALS Response Time are the deployments of 4 and 8 aircraft with High TRR.

5. Net Supply Effectiveness

Net High-Pri Supply Effectiveness is still a useful metric to evaluate the performance of a package. Figure 4-8 compares Net Supply Effectiveness at the MOB.




Figure 4-8. MOB Net High-Pri Effectiveness (TRR = High)

MALSP II Allowances yield slightly higher MOB Supply Effectiveness in all scenarios; however, the differences are not practically significant. Consistent with findings regarding *Response Times*, adding an ESB slightly degrades MOB Supply Effectiveness.

Figure 4-9 compares Net High-Pri Effectiveness at the PMALS in the various scenarios.




Figure 4-9. PMALS Net High-Pri Effectiveness (TRR = High).

Legacy Allowances achieve slightly higher PMALS Supply Effectiveness in configurations without ESB, though the differences are not of practical significance. Similarly, adding ESB slightly degrades PMALS Effectiveness.

E. Detailed Analysis on Effect of ESB

The conclusions regarding the ESB are certainly the most surprising and counterintuitive. Most previous research has found the ESB improves the expected performance of the logistical system. However, previous projects have not examined actual allowances and have assumed sufficient material availability to (at least attempt) to fill all buffers as necessary.

First consider the effect of adding an ESB on the total material requirement for a particular NIIN in a given deployment scenario.

Assume a particular NIIN is ordered one per day, each day. If the TRR from the PMALS to MOB is 10 days, then the ideal buffer at the MOB (without an ESB) is 10. Insert an ESB between the PMALS and MOB, with the PMALS-to-ESB TRR = 7 and the ESB-to-MOB TRR = 3. In this case, adding the ESB requires no additional material (ESB buffer = 7; MOB buffer = 3; $7+3 = 10$). However, it is theoretically possible that the addition of an ESB might enable efficiencies that result in decreases in the intermediate TRRs, say PMALS-to-ESB TRR = 6 and ESB-to-MOB TRR = 2 ($6+2=8$). Thus, implementing an ESB could reduce the material requirement for an item if it results in intermediate TRRs whose sum is less than the TRR from PMALS to MOB.

Alternatively, consider a NIIN that is ordered one per month, each month. If the TRR from PMALS to MOB is 10 days, then the ideal buffer at the FOB (without an ESB) is 1. In fact, the

ideal buffer at any node whose TRR < 30 is 1. Thus, even if adding an ESB results in inter-nodal shipping efficiencies, the material requirement from adding an ESB is always greater than without (e.g., 1 in buffer at FOB + 1 in buffer at ESB > 1 in buffer at FOB).

Adding an ESB may reduce the overall material requirement when (a) the sum of the intermediate TRRs is less than the TRR direct from PMALS to the node; and (b) the NIIN's average inter-arrival time is less than the TRR direct from PMALS to the node. Lower demand NIINs typically have lower allowances, so fully implementing the MALSP II nodal laydown with an ESB tends to require a substantial relative increase in material (i.e., a quantity of 1 at each node). A large proportion of NIINs in the MALSP II allowance package are relatively low demand NIINs.

Next, each of the 432 NIINs in the package are examined to determine the relationship between total allowance available and the additional material requirement incurred as a result of adding an ESB. Each NIIN is assigned to a category that describes this relationship. A Category "A" NIIN is one for whom the MALSP II total allowance is sufficient for configurations that both do and do not employ an ESB. Category "D" NIINs are deficient in both cases. That is, the total MALSP II allowance is insufficient to fill buffers at the PMALS and FOB, let alone buffers at the PMALS, FOB, and ESB.

Category "B" NIINs are those with sufficient allowance levels without an ESB, but not enough to fill the buffer requirement once an ESB is required. Finally, Category "C" describes those NIINs that are deficient in configurations that do not employ an ESB, but because of efficiencies from adding an ESB, are not deficient in those configurations. None of the 432 NIINs fall into Category "C." Note that buffer levels in this section are constructed using high risk (87th percentile) demand filtering.

Table 4-3 shows the distribution of NIINs among the categories.

Table 4-3. Distribution of NIINs by Deficiency Category

Category	High TRR		Low TRR	
	Config_a	Config_c	Config_a	Config_c
A	255	256	256	256
B	46	45	45	45
C	0	0	0	0
D	131	131	131	131
Total	432	432	432	432

As described above, adding an ESB tends to increase aggregate response time at both the MOB as well as the PMALS. Table 4-4 outlines how that increase is distributed among the categories. (Note: Table 4.4 applies MALSP II allowances, 16 deployed a/c to the MOB, and a High TRR. For the corresponding graphs pertaining to TRR=Low, refer to the Mathematical Appendix. All results presented in this detailed analysis section are for MALSP II allowances.)

Table 4-4. Configuration_a MOB Response Time (total response times in document days)

Category	NIINs	without ESB	with ESB
A	255	88.4	263.5
B	46	292.5	239.5
D	131	819.7	1319.2
Total	432	1200.6	1822.2

As Table 4-4 confirms, the vast majority of the increase of Response Time at the MOB is from Category “D” NIINs. Recall, in all Configuration_1 scenarios, the very first buffer filled is the MOB. So, for NIINs with total allowance quantity of 1, that quantity is sent to the MOB. For all Configuration_2 scenarios, that quantity is sent to the ESB, because for those configurations with an ESB, the ESB is filled first. Thus, for low density items, adding an ESB results in removing the part from the MOB and placing it at the ESB. This must necessarily increase MOB Response Times.

Table 4-5 contains the distribution of Response Times by category for the PMALS, for the exact same scenario. (Note: MALSP II allowances, 16 a/c deployed to MOB, High TRR).

Table 4-5. Configuration_a PMALS Response Time

Category	NIINs	without ESB	with ESB
A	255	808.2	953.3
B	46	213.4	1284.7
D	131	4535.9	4114.6
Total	432	5557.6	6352.6

First, notice that the sums at the bottom of each column in Tables 4-4 and 4-5 correspond to values found in Figures 4-6 and 4-7. Next, notice that Category “D” NIINs make up the vast majority of the extra response time that is created at the MOB upon implementing an ESB (from 819.7 to 1319.2, see Table 4-3). However, the contribution of Category “D” NIINs actually decreases at the PMALS (from 4535.9 to 4114.6).

Again, the vast majority of Category “D” NIINs are those with just one allowance quantity. Thus, in scenarios without an ESB, that item resides at the MOB where it responds quickly to demand among the deployed aircraft, but a lateral support request from the PMALS must travel the entire distance from the MOB to PMALS. In scenarios with an ESB, that single item resides at the ESB. This MUST increase response time for all MOB requirements; however, this reduces the distance between the item (now at the ESB) and the PMALS, thus tending to reduce the PMALS response time.

A very similar narrative applies to items with allowance quantity of 2. Without the ESB, the MOB and PMALS each have quantity = 1. Implementing the ESB moves the item from the PMALS to the ESB. So, such items experience little change in MOB *Response Time* and substantial increase in PMALS *Response Time*.

Finally, one might expect that Category “A” NIINs, i.e. those with sufficient allowances to fully implement all nodes in the network, would experience improved performance with the ESB. However, as an aggregate, the Category “A” NIINs do not fare better under the ESB.

These relationships hold as the number of deployed aircraft decreases. Tables 4-6 and 4-7 display the MOB and PMALS *Response Times* for Configuration_c, in which 4 aircraft are deployed to the MOB. (Note: Applies MALSP II allowances, 4 a/c deployed to MOB, and High TRR.)

Table 4-6. Configuration_c MOB Response Time

Category	NIINs	without ESB	with ESB
A	256	4.8	21.9
B	45	26.2	25.6
D	131	182.6	299.2
Total	432	213.6	346.8

Table 4-7. Configuration_c PMALS Response Time

Category	NIINs	without ESB	with ESB
A	256	927.1	1036.5
B	45	257.4	1651.6
D	131	6092.5	5603.1
Total	432	7277.0	8291.2

The relationships are qualitatively similar to scenarios with more deployed aircraft. For example, the aggregate *Response Time* increases for both the MOB and the PMALS when an ESB is added to the network. However, for Category “D” NIINs, the increase in MOB *Response Time* from adding an ESB (from 182.6 to 299.2) is less than the decrease in PMALS *Response Time* (6092.5 to 5603.1). This implies that circumstances exist in which it is favorable to stock certain items at the ESB based narrowly on material availability.

A “Low Density” NIIN is defined as an item with an allowance quantity less than or equal to two. All other NIINs are defined as “Regular Density” NIINs. Table 4-8 provides a breakdown of Low and Regular Density NIINs by category. There are no Low Density NIINs in Category “A,” while all but five Category “D” NIINs are Low Density.

Table 4-8. NIIN Density Type by Category

Category	Density	
	Low	Regular
A	0	255
B	17	29
D	126	5
Total	143	289

For each of the 432 NIINs in the package, consider the case of Configuration_1a (with High TRR). Recall, Configuration_1a is the scenario in which 16 aircraft are deployed to a FOB and the logistics network does not make use of an ESB. Composite *Response Time* (CRT) is defined as the sum of the PMALS and MOB *Response Times* for each NIIN. Next consider Configuration_2a and calculate CRT. (That is, add the ESB). Compare the CRT for each NIIN with and without an ESB. Table 4-9 illustrates the outcome from this process.

Table 4-9. Effect of Adding ESB on Composite Response Time by Category (Configuration _a).

Adding ESB	low		low Total	reg			reg Total	Grand Total
	B	D		A	B	D		
better		123	123	94	1		95	218
same				34			34	34
worse	17	3	20	127	28	5	160	180
Grand Total	255	29	5	289	17	126	143	432

Notice that 218 NIINs have better CRTs with an ESB than without. However, the ESB actually makes 180 NIINs worse off according to this measure.

The CRT measure in this example is in need of a refinement. The main defect is that it weights *Response Time* for documents in support of deployed aircraft presumably flying combat missions equal to the *Response Time* for documents in support of aircraft flying training missions at the PMALS. Thus, the proposed concept of Weighted Composite *Response Time* (WCRT) is contained in the following formula:

$$WCRT_i = PMALS_Response_Time_i + \beta \cdot MOB_Response_Time_i$$

In this case, β is a user-defined weight intended to convey the priority of support to aircraft at the FOB over those at the PMALS, and i is an index that indicates the NIIN. (Note: In the case of Table 4-9, $\beta = 1.0$.)

It is important to consider the appropriate interpretation of β . From a certain perspective, β is a penalty assessed to *Response Time* accumulated at the MOB. If the objective is to minimize WCRT, then $\beta > 1$ encourages logisticians to prioritize the reduction of *Response Time* at the MOB. Logisticians often make choices that result in a continental United States (CONUS) aircraft waiting a number of days for a part in order to save a deployed aircraft from waiting (on the ground) for that same part. Such choices are a manifestation of this priority.

The actual β for a given deployment/network in reality is just not accessible to researchers or logisticians. It is based on many things, to include the nature of the deployment, the nature of the training mission, the logisticians' understanding of commander's intent with respect to logistics support, etc. However, it is possible to reflect on the revealed behavior of logisticians in the past. It is certainly common for a logistician to make a decision that saves a deployed aircraft one day of downtime but results in as many as 10-15 days of downtime for an aircraft at the PMALS. Such decisions imply that β 's as high as 10 or 15 are not altogether implausible.

Table 4-10 is an update of Table 4-9, this time using WCRT where $\beta = 2$.

Table 4-10. Effect of Adding ESB on Weighted Composite Response Time ($\beta = 2.0$; Configuration_a)

Adding ESB	low		low Total	reg			reg Total	Grand Total
	B	D		A	B	D		
better		1	1	68	4		72	73
same				33			33	33
worse	17	125	142	154	25	5	184	326
Grand Total	17	126	143	255	29	5	289	432

Notice that even with this relatively small level of β , all but one Category D NIINs is worse off after adding the ESB, as are an additional 27 (154-127) Category “A” NIINs. Figure 4-10 displays the relationship between β and the number of NIINs that experience an increase in WCRT of an ESB for Configuration_a.




Figure 4-10. Configuration_a NIINs with Higher WCRT from Adding ESB

As illustrated in Figure 4-10, over 350 of the 432 NIINs in the package experience higher WCRT at even the very conservative $\beta = 2.8$. Given that so many NIINs exhibit higher WCRT with an ESB, and a large proportion of the remaining NIINs exhibit no change, the following is recommended: For squadron level deployments, recommend not stocking repairables at ESB.

As the number of deployed aircraft decreases (simultaneously increasing the number of aircraft at the PMALS node) this conclusion is less clear. Table 4.11 outlines the NIINs by category with better or worse WCRT when $\beta = 1.0$.

Table 4-11. Effect of Adding ESB on Weighted Composite Response Time ($\beta = 1.0$; Configuration_c)

Adding ESB	low		low Total	reg			reg Total	Grand Total
	B	D		A	B	D		
better		122	122	138			138	260
same				33			33	33
worse	17	4	21	85	28	5	118	139
Grand Total	17	126	143	256	28	5	289	432

Notice in this case, nearly all Category “D” items are made better off if stocked at the ESB. All Category B NIINs are worse off and Category A NIINs are split.

Figure 4-11 illustrates the relationship β and the number of NIINs that are made worse off by the inclusion of an ESB for Configuration_c as measured by WCRT. It is not until $\beta = 13$ that all Category D NIINs are rendered worse off.




Figure 4-11. Configuration_c NIINs with Higher WCRT from Adding ESB

Notice for deployments of 4 a/c, the number of NIINs for which an ESB is counterproductive increases far less quickly with increases in β , than in the previous case with 16 deployed a/c.

This is due to the fact that implementing an ESB tends to result in material being shifted from the MOB to a position closer to the PMALS. Because fewer deployed aircraft implies relatively more demand experienced locally at the PMALS, implementing an ESB improves WCRT in more cases. (Alternatively, it implies β must be higher in order to conclude that moving material from the MOB to the ESB was not an improvement.) Another insight from this graph is that stocking Category “D” NIINs at the ESB is performance enhancing under a range of β values.

The Mathematical Appendix contains the corresponding tables and figures for the Low TRR scenario. In general, the qualitative relationships are similar. These revelations imply the following recommendations for deployments of between 4 and 16 aircraft:

- Recommend against stocking Category “A” repairable NIINs at the ESB. As Tables 4-9 through 4-11 indicate, stocking Category “A” NIINs at the ESB reduces WCRT for some while increasing it for others. However, closer investigation fails to reveal the additional factors that determine this outcome. For example, the magnitude of the excess of allowances over requirement does not affect this outcome. Therefore, reliance is on the information in Tables 4-5 through 4-7 (pp. 48, 49, 50), which confirm in the aggregate that stocking Category “A” NIINs at the ESB unambiguously increases *Response Time* at both the MOB and PMALS.

- Recommend against stocking Category “B” NIINs at the ESB. As with Category “A” NIINs, stocking Category “B” NIINs at the ESB increases WCRT for some while decreasing it for others. Thus, reliance is on Tables 4-5 through 4-7, which confirm in the aggregate that stocking Category “B” NIINs at the ESB increases *Response Time* at both the MOB and PMALS.

- Stocking Category “D” NIINs at the ESB can result in improved WCRT for sufficiently low β levels. This effect is stronger as both the number of deployed aircraft and TRR is reduced. In other words, recommend stocking Category “D” NIINs at the ESB if the mission to which the detachment is assigned is not substantially more important than the local training mission that the PMALS supports directly.

F. Summary

In this section, the behavior is simulated of actual NIINs contained in the proposed MALSP II allowances for MALS-16 in support of CH-53Es. The performance of the MALSP II allowances is compared with the performance of the Legacy MALSP allowances under a variety of circumstances. Overall, the MALSP II allowances significantly outperform Legacy allowances both statistically and practically.

Most surprisingly, we find that implementing an ESB tends to reduce system performance both in terms of *Response Time* and Supply Effectiveness. Closer examination confirms this phenomenon and suggests that the main driver behind it is the behavior of low density items. In short, moving scarce resources to an ESB and away from deployed aircraft tends to only be worthwhile in very limited circumstances. (Note: Conclusions regarding the ESB only apply to repairable items.)

Perhaps most importantly, as a “proof of concept” this is a successful project. The process works and there is every reason to believe that the process is generalizable to other T/M/S and other MALS, etc. The specific results, however, do not necessarily apply to other platforms, or even necessarily to other MALS that support CH-53Es.

Finally, this section demonstrates the utility of *Response Time* as a Measure of Effectiveness. In addition to the fact that very little was gained by looking at Supply Effectiveness in this section, *Response Time* has substantial intuitive appeal in that it incorporates aspects of the time domain, in addition to aspects of material availability.

5. Analysis of Inclusion Decisions

In this section, the important determinants of whether to include an item in a deployed buffer are determined.

2. *What is the optimal/robust criteria for including an item in a packup?*

- If high utilization rates (relative to training hours) for the deployed aircraft are expected, ensure NIINs with greater than 8 demands in the previous 24 months are included in the packup.

- Ensure NIINs with exceptionally high demand (i.e., greater than 70 demands in previous 24 months) are stocked using Medium (95th percentile) or lower risk (i.e., higher percentile) demand filtering.

- Using 80th percentile demand filtering for buffer sizing is sufficient for all types of NIINs for cases in which the aircraft at the FOB are not expected to experience substantially greater utilization rates relative to training hours.

- Using 95th percentile demand filtering for buffer sizing is sufficient for all types of NIINs for cases in which the aircraft at the FOB are expected to experience up to three times the utilization rate relative to training hours.

A. Scope

The scope for this question is much more general than the scope of the previous section. The examined NIINs are notional and possess a broad range of different qualities and their performance is scrutinized under a wide array of logistical systems.

Rules employed: **Q**

NIINs: Notional

The NIINs are notional, in that they are generically constructed using a wide range of demand frequencies and quantities. No distinction is made between consumable or repairable, thus the behavior of the IMA is neglected. An overarching simplifying assumption in this section is essentially the availability of sufficient allowances to fill buffers throughout the network.

The Supply Officer's decision-making process (upon being tasked with supporting a deployment of aircraft detached from one of its home squadrons) was mirrored. Typically, the Supply Officer examines the demand history for the particular squadron over the previous 24 months and makes range decisions on the basis of this information.

B. Experimental Design

A full factorial experimental design, with the following factors and levels, was implemented:

Demand Frequency		Demand Quantity Type		PMALS TRR _{Stock}		PMALS TRR _{DTO}		ESB TRR		MOB TRR		FOB TRR		PMALS Aircraft		Buffer Filter		Wartime Intensity
6	×	3	×	2	×	2	×	2	×	2	×	2	×	3	×	4	×	3
Very Hi, High, High-Mid, Mid, Low-Mid, Low		High, Mid, Single		30, 50		10, 20		6, 10		3, 8		1, 3		16, 8, 4		100, 95, 80, 0		1, 2, 4

The experiment is comprised of 11,232 design points. Each design point is replicated 50 times. Common Random Numbers are employed in an effort to reduce the variance between replications thereby increasing the power of subsequent statistical analysis.

Demand Frequency

Six different levels of demand frequency were considered, as outlined in Table 5-1.

Table 5-1. Demand Frequency Types

Demand Freq	Mean[hits / mo]	Mean[24*hits / mo]
Very-High	30	720
High	3	72
Med-High	0.99	23.76
Med	0.33	7.92
Med-Low	0.165	3.96
Low	0.081	1.944

The levels correspond to demand frequencies exhibited by a single squadron of 16 aircraft. Again, the idea is to mirror the process through which an Aviation Supply Officer might build a package for a detachment. Upon being tasked with supporting a detachment of x aircraft from a given squadron, the Supply Officer would likely examine the previous two years' of demand history for that squadron. So, Table 5-1 translates the demand frequencies considered in the experiment to notional empirical data.

Demand Quantity

Three different levels of quantity per document were considered.

Table 5-2. Demand Quantity Types

	Minimum	Mode	Maximum
Single	1	1	1
Medium	1	2	4
High	1	20	50

The parameters in the table correspond to the parameters of the Triangular distribution from which the quantities for each document are drawn.

TRR

The TRRs are implemented by drawing shipping times from a Lognormal Distribution that has a reasonable mode and a theoretical 90th percentile that corresponds to the TRR required in the experimental design. In the design nomenclature, the label assigned to the TRR refers to the shipping times between the given node's parent. Thus, TRR for the ESB refers to the characteristics of the shipping times between the PMALS and the ESB.

Configurations

The following three network configurations were considered:




Figure 5-1. Design Configurations

ESBs are included in Configurations II and III because their efficiency is not tested or examined in this section.

Buffer Levels / Demand Filtering

The buffer levels in this experiment correspond to various levels of demand filtering. The following risk levels were considered: no risk (100th percentile), medium risk (95th percentile), high risk (80th percentile), and not carried (0th percentile). The difference between system performance for “not carried” and “high risk” is the benefit of including the item in the package.

Wartime Intensity

A major challenge for any Supply Officer attempting to build a package to support a deployed detachment of aircraft under the MALSP II construct is the lack of an effective means with which to predict future demand. The common practice is to examine MALS-wide demand for the platform in question over a period of 24 months. All NIINs with demand over a particular threshold, such as four hits per year, are considered as candidates for inclusion. However, little consideration is given to the fact that the detachment might only be four aircraft, while the empirical data upon which range decisions are made may include as many as 48 or more aircraft.

Demand for a NIIN at a node is, among other factors, a function of the number of aircraft. This fact provides theoretical justification for scaling the frequency of demand to account for the number of aircraft deployed to a particular node. However, experience shows that deployed

aircraft are often flown more and exposed to harsher environments. The Wartime Intensity Factor addresses this difficulty. It is a scaling factor that increases the demand frequency at the MOB and/or FOB. In reality, this factor is neither chosen nor accessible to measurement; however, implementing it in the experimental design and exposing the model to various levels helps to determine if this factor substantially affects system performance.

C. Results

The general Scheme of Maneuver for this section involves running the experiment and obtaining the Mean *Response Time* at each node for each design point. (Each design point corresponds to the behavior of some notional NIIN under a particular set of conditions regarding the logistics network.) The intent of the experiment is to generate a response surface that will provide evidence of which factors, if any, substantially affects the behavior of the response variable(s). Classification trees were built using FOB and MOB *Response Times* as response variables in order to gain initial insight into potentially important factors (see Mathematical Appendix) and create an Ordinary Least Squares multivariate regression model of main effects (and use Stepwise techniques to select among them) to confirm the statistical significance of these factors (see Mathematical Appendix). The relationships between these factors are then graphically examined.

FOB Response Time

Preliminary classification trees and regression analysis confirms that Demand Frequency, Demand Filtering, Wartime Intensity, and PMALS TRR_{DTO}, have the greatest effect on FOB *Response Time* (see Mathematical Appendix). This leaves the number of deployed aircraft, network configuration, demand quantity type, and most levels of TRR as factors that do not substantially affect FOB *Response Time*. Thus, all conclusions are robust with respect to these factors as well.

For example, consider the graph in Figure 5-2. This graph corresponds to a Wartime Intensity = 3 and TRR_{DTO} = 10. The values of FOB *Response Time* for all four Demand Filtering levels at each of the six demand frequency levels are displayed along the x-axis. Note that most of these values are zero. However, for *mid*, *mid-high*, *high*, and *very-high* demand frequency types, there are non-zero values for the 0th percentile Demand Filtering, as indicated by the red bars. Thus, under the given circumstances, failing to place *mid*, *mid-high*, *high*, and *very-high* type NIINs in the physical buffer at the FOB results in (possibly large) positive values for expected FOB *Response Time*. The existence of positive values of response time for Demand Filtering = 80 for *high* and *very high* demand frequency types (as indicated by the green bars) indicate that greater higher percentiles of Demand Filtering are required to reduce or eliminate expected response time for those types of NIINs.




Figure 5-2. FOB Response for Warptime Intensity = 3; TRR_{DTO} = 10.

Figure 5-3 shows the relationship between the primary factors. The figure contains six graphs. (The graph from Figure 5-2 is in the upper right corner.) Each graph relates to a particular combination of Warptime Intensity and TRR_{DTO}. Each individual graph shows the FOB *Response Time* for each of the six levels of demand frequency at the four possible buffer levels.




Figure 5-3. FOB Main Comparison

Looking at all six graphs provides insight on the effect of different levels of Wartime Intensity and TRR_{DTO} . Thus, *FOB Response Time* tends to increase monotonically as Wartime Intensity (utilization) increases. Likewise, *FOB Response Time* increases as TRR_{DTO} increases.

These relationships imply the following robust rules for package development:

- If high utilization rates (relative to training hours) for the deployed aircraft are expected, ensure NIINs with greater than eight demands in the previous 24 months are included in the pickup.

- Ensure NIINs with exceptionally high demand (i.e., greater than 70 demands in previous 24 months) are stocked using Medium (95th percentile) or lower risk (i.e., higher percentile) demand filtering.

- Using 80th percentile demand filtering for buffer sizing is sufficient for all types of NIINs for cases in which the aircraft at the FOB are not expected to experience substantially greater utilization rates relative to training hours.

- Using 95th percentile demand filtering for buffer sizing is sufficient for all types of NIINs for cases in which the aircraft at the FOB are expected to experience up to three times the utilization rate relative to training hours.

MOB Response Time

The results for *MOB Response Time* are nearly identical in terms of magnitude and significance to that for *FOB Response Time*. As above, preliminary classification trees and regression analysis suggests that Demand Frequency, Demand Filter, Wartime Intensity, and PMALS TRR_{DTO} , have the greatest effect on *MOB Response Time*. The charts analogous to Figure 5-3 for *MOB Response Time* is also nearly identical. Thus, the conclusions and recommendations for stocking items at the MOB are identical to those conclusions regarding packages built for the FOB.

D. Summary

In this section, an experiment was designed and implemented with the intent to glean information regarding the expected performance of various types of items stocked in a logistical network. The rules deduced in this section are robust against numerous sources of variance and are universally applicable to many sorts of MALSP II deployments.

6. Analysis of TRR Uncertainty

In this section, we examine aspects of designing an MALSP II logistics network under uncertainty, and then managing the network as information is revealed.

4. How does uncertainty regarding Actual TRR affect Response Time at deployed nodes?⁵

-When an Actual TRR exceeds the Design TRR, the node typically experiences high levels of *Response Time*. This effect is greater as Design TRR decreases and as number of supported a/c increases.

-NIINs with Demand Frequencies greater than one demand per month are most affected by differences between Actual and Design TRR.

-No risk demand filtering (100th percentile) provides robust protection in nearly all cases.

A. Scope

The intent of this experiment is to examine the process through which support for deployed aircraft is initiated and maintained. Design and implementation of the MALSP 2 nodal laydown requires certain assumptions about the state of the world that will remain unknown until the network is in place. One of the most important of these factors is the distribution of shipping times between nodes. In order to address this uncertainty, the Aviation Supply Officer generally declares a Design TRR, a number based on experience and intuition. However, because the TRR directly determines (though, in a non-linear fashion) material requirements in the buffer, it is critical that (a) the Design TRR is as accurate as possible; and (b) if it is not accurate, action is taken as soon as possible to correct it.

Rules employed: **Q, L**

NIINs: Notional

B. Experimental Design

Demand Frequency		Demand Quantity Type		Design TRR		Delta		Aircraft		Buffer Filter		Wartime Intensity
3	×	3	×	8	×	8	×	4	×	4	×	2
Very Hi, High-Mid, Low		High, Mid, Single		1,3,5,8,11,17,25,32		-10,-6,-2,2,4,6,8,10		4,8,12,16		100, 95, 90, 80		1,3

⁵ On the Statement of Work, this question appears as “How frequently should buffers be re-sized? Are there useful leading indicators (i.e. say, if Actual Time to Reliably Replenish (TRR) exceeds Design TRR by a certain amount).

Demand Frequency

Three different levels of demand frequency were considered; they are *Very-High*, *Mid-High*, and *Low*. These types exhibit mean hits per month of 30, 1, and 0.081. (See Table 4-1 for more detail.)

The levels correspond to demand frequencies exhibited by a single squadron of 16 aircraft. Again, the idea is to mirror the process through which an Aviation Supply Officer might build a package for a detachment. Upon being tasked with supporting a detachment of x aircraft from a given squadron, the Supply Officer would likely examine the previous two years' of demand history for that squadron. So, table x translates the demand frequencies considered in the experiment to notional empirical data.

Demand Quantity

Three different levels of demand quantity, or quantity per document, were considered.

Table 6-1. Demand Quantity Types

	Minimum	Mode	Maximum
Single	1	1	1
Medium	1	2	4
High	1	20	50

The parameters in the table correspond to the parameters of the Triangular distribution from which the quantities for each document are drawn.

Design TRR

The Design TRR takes on values of 1, 3, 5, 8, 11, 17, 23, and 30 days. This is the TRR that the node is “expected” to operate under. Thus, the buffer level at the node is calculated using the given Design TRR.

Delta

Delta is the difference between the Actual (as modeled) TRR implemented in the experiment and the Design TRR. Negative values of delta indicate that the Actual (as modeled) is less than the design.

Aircraft

Aircraft is the number of aircraft supported at the node may be 4, 8, 12, or 16 aircraft.

Buffer Levels / Demand Filtering

The buffer levels in this experiment correspond to various levels of demand filtering. The following risk levels were considered: no risk (100th percentile), medium risk (95th percentile), high risk (90th percentile) and high-high risk (80th percentile). (See Chapter 1 for an explanation of the relationship between buffer levels and demand filtering.)

Wartime Intensity

The Wartime Intensity Factor is simply a scaling factor that increases the demand frequency at the MOB and/or FOB. In reality, this factor is neither chosen nor accessible to measurement. However, implementing it in the experimental design and exposing the model to various levels will help determine if this factor substantially affects system performance.

C. Results

The general Scheme of Maneuver for this section is to run the experiment and obtain Mean *Response Time* for each design point. (Each design point corresponds to the behavior of some notional NIIN under a particular set of conditions regarding the logistics network.) The intent of the experiment is to generate a response surface that provides evidence of which factors, if any, substantially affects the behavior of the response variable(s). Classification trees were built using *Response Time* as the response variable in order to gain initial insight into potentially important factors. (See Mathematical Appendix.) An Ordinary Least Squares multivariate regression model that includes only main effects (and uses Stepwise techniques to choose among them) was created to confirm the statistical significance of these factors. (See Mathematical Appendix.) The relationships between these factors are then graphically examined. Figures 6-1 through 6-4 illustrate these relationships.

The classification trees and regression models confirm that Demand Frequency Type, Buffer, number of aircraft, Design TRR, and Delta are statistically significant. This leaves Demand Quantity and Wartime Intensity as factors that do not significantly influence the results.

First, consider Figure 6-1, which shows the *Response Time* for a node with a Design TRR = 1 and a *Very-High* Demand Frequency Type. It also employs 80th percentile demand filtering and supports 4 aircraft.




Figure 6-1. Response Time for 4 a/c; 80th Percentile Demand Filtering; Demand Frequency Type = Very High; and Design TRR = 1.

According to the graph, under these conditions, when Delta, that is the difference between the Actual TRR and the Design TRR, is 2 days or higher, the node experiences significant positive expected *Response Times*. Furthermore, as the magnitude of the delta between Actual and Design TRR increases, so too does the magnitude of the *Response Time* experienced at the node.

Figure 6-2 shows how the *Response Time* is affected as the Delta between Actual and Design TRR changes and the differing levels of Demand Frequency Types are applied. It is comprised of 16 individual graphs, corresponding to two Demand Frequency Types (*Med-High* and *Very-High*) and the eight factor levels for Design TRR from 1 to 30 along the x-axis. The colored bars in each grid correspond to the *Response Times* that result from the given level of Delta.




Figure 6-2. Response Times for 4 Aircraft and 80th Percentile Demand Filtering.

Consider the sub-graph in the lower-left corner of Figure 6-2, which corresponds to Design TRR = 1 and Demand Frequency Type = *Very-High*. The colored bars in the sub-graph arranged from left to right correspond to delta levels of -10 to 10. So, according to the graph, for nodes that are set up with a Design TRR = 1, even small deviations, such as the empirical TRR = 3 results in high levels of expected response time. The response increases as the size of the error increases (i.e., as delta increases).

The following relationships are evident from all 16 graphs. Demand Frequency Type = *low* results in zeros across the board, thus differences between Design and empirical TRRs tend not to affect low frequency NIINs. For *Very-High* frequency type NIINs, the effect of the Delta on Response Time diminishes as Design TRR increases. Finally for *Med-High* demand type items, increasingly higher Design TRR values result in modest increases of Response Time for nearly all levels of Delta. Recall that *Med-High* demand frequency type NIINs exhibit an average of 1 demand per month. Variance surrounding the arrival of demands is likely reason for the gentle increase in *Response Time*. In other words, this phenomenon is likely the result of the demand filter level. (That it disappears at lower risk levels is consistent with this hypothesis – see Figure 6-4.)

The same configuration (i.e., 4 a/c), with 100th percentile demand filtering results in zero expected *Response Time* for all combinations of Demand Frequency Type, Design TRR, and Delta. Thus, for deployments with small numbers of aircraft, implementing buffers based on no risk demand filtering is robust against any uncertainty associated with the Actual TRR.

Figure 6-3 shows 80th percentile demand filtering, but with a full squadron of 16 aircraft. The relationships are qualitatively similar to the scenario with four aircraft. The only apparent difference is that the magnitudes of the relationships are larger.




Figure 6-3. Response Time for 16 Aircraft and 80th Percentile Demand Filtering.




Figure 6-4. Response Time for 16 Aircraft and 100th Percentile Demand Filtering.

Finally, Figure 6-5 shows no demand filtering with a full squadron of 16 aircraft. With 16 aircraft, using no demand filtering (*no risk*) is not quite sufficient to eliminate *Response Time* under all circumstances. For relatively low levels of Design TRR, namely three days and below, positive values of *Response Time* still exist. That said, no demand filtering (i.e., 100th percentile) is sufficient to eliminate *Response Time* that may result in all other combinations of Demand Frequency Type and Design TRR.

The analysis implies the following recommendations and conclusions:

- When an Actual TRR exceeds the Design TRR the node typically experiences high levels of *Response Time*. This effect is greater as Design TRR decreases or as number of supported a/c increases.

- NIINs with Demand Frequencies greater than 1 demand per month are most affected by differences between Actual and Design TRR.

- No risk demand filtering (100th percentile) provides robust protection in nearly all cases.

D. Summary

In this section, an experiment was designed and implemented with the intent to flex the model with respect to the effect of differences between the Design and Actual TRRs. The extent to which such uncertainty may adversely impact system performance was identified. The next step was to leverage statistical techniques to determine how a logistician may (a) reliably identify a discrepancy between Actual and Design TRRs in real life and, based on the reliability of that test and relative costs of resizing the buffer of the deployed pickups, (b) determine when it is appropriate to change the Design TRR for a given node.

In addition, it is important to note that this section primarily examines uncertainty with regard to trans-shipment times. Another component of uncertainty is related to Demand Frequency. While the Wartime Intensity term does address the possibility that an item may be ordered more frequently than originally planned, more could be done along these lines. Namely, more could be done with regard to determining the amount of evidence required to conclude that a buffer should be resized due to changes in demand patterns.

7. Future Work and Recommendations

This project demonstrates the utility of modeling and simulation in assessing the performance of allowance packages and various business rules. Perhaps the most obvious recommendation is to continue down the path initiated by this project. Three research questions were left unaddressed due to time constraints. In addition, while the MALS-16 allowances for CH-53Es were thoroughly analyzed, such analysis should be conducted for allowances at other MALS and other T/M/S aircraft. Such an effort would go a long way towards determining whether the conclusions from that aspect of the study (i.e., that stocking repairables at ESB is counter-productive) are applicable to the rest of naval aviation.

The model developed in this project could be used to quantify the expected benefit of transitioning from MALSP to MALSP II. While Chapter 3 considers the relative performance of Legacy allowances and MALSP II allowances in supporting a MALSP II type deployment, it would certainly be possible to measure the overall MALSP II concept (i.e., MALSP II allowances and MALSP II network) relative to the Legacy MALSP concept (i.e., MALSP allowances and MALSP type support of deployed aircraft).

A significantly more ambitious recommendation is to couple the model with an allowance costing algorithm. Allowancing models base range and depth decisions on the basis of many factors to include costs and expected effect on readiness. Using this model (or one of similar fidelity) could improve the costing models' assessment of the operational impact of marginal changes in a NIIN's allowance level.

Part of the model's utility lies in its ability to estimate the expected results of various logistics management decisions on each NIIN. For example, on the basis of existing empirical data, the model can predict expected performance associated with decisions of whether to include a NIIN at a particular node in the network. Such a capability would be invaluable to aviation logisticians at the MALS level who are charged with developing and managing MALSP II-style logistics networks in support of aircraft deployments. Developing the model so that any aviation logistician in the fleet can use it is highly recommended. This would involve developing a graphical user interface (GUI), as well as methods with which to easily import empirical demand and allowance data.

One challenge faced throughout the project was the lack of a canonical MALSP II style deployment. The MALSP construct envisions a squadron self-deploying on short notice to Germany to repel a Soviet invasion of Europe. The aircraft are followed closely by the FISP, which is intended to support wartime hours for 30 days until the PCSP arrives in theater. There exists no corresponding canonical description of such a scenario under MALSP II. Thus all configurations depicted in this project were arbitrarily selected. This greatly inhibits any assessment of a given allowance package's performance, because other analysts may arrive at completely different conclusions solely on the basis of the nature of the operations they envision. Creating a canonical, or baseline, scenario would solve this problem.

One innovation associated with MALSP II is the increased awareness of the time domain. Towards that end, node *Response Time* was found to be a useful measure of effectiveness in nearly all stages of this project. Developing this concept, both theoretically and operationally, is recommended.

8. Glossary of Terms

a/c - aircraft

AFAST - Aviation Financial Analyst Tool

AFAST is a database that, among other things, aggregates the demand history of all retail supply activities in the naval aviation enterprise.

BCM – Beyond the Capability of Maintenance

An NRFI item that is inducted into the repair process but cannot be successfully repaired is categorized as a BCM. The term BCM may also be used as a verb such as when the IMA determines or declares an item to be BCM.

CCSP – Common Contingency Support Package

The Common Contingency Support Package contains items that support multiple T/M/S. Rotary Wing MALs typically each possess a Rotary Wing Common package, while Fixed Wing MALs typically possess a Fixed Wing Common package.

Deficient NIIN

An item for which Actual RO < Ideal RO.

Document Time

Document time is the time that transpires between the instantiation of a document and the time that it is completed out.

DTO – Direct Turn-Over

DTO is a type of document that is created when a MALs receives a document from a squadron for a part that it cannot fulfill (i.e. the part is either NIS or Not Carried). The document is then referred to the wholesale system but is considered a DTO, because when the MALs receives it, the part is turned over directly to the intended squadron.

ESB – En-route Support Base

The node between the PMALS and the MOB or FOB intended to reduce the time transporting items between the Continental United States and the deployed aircraft in theater.

EXREP – Expeditious Repair

When the squadron orders a repairable component that the MALSP is either NIS for or does not carry, then the NFI carcass is inducted into the IMA as an EXREP. They are given the highest priority in the repair cycle because these components are holding down an aircraft.

Fill Proportion

This is a characteristic of a particular item at a given node. It is given by: Actual RO / Ideal RO, for the item at that node.

Fill Proportion, Average

This is a characteristic of a particular node. It is simply the average of all the Fill Proportions for all the items that reside at that node.

FISP – Fly-In Support Package

The FISP comprises sufficient O-level parts (that is, parts that the OMA may remove and replace) to support a squadron flying wartime hours for thirty days. When not deployed, the FISP is part of the Aviation Supply Officer's "protected stock" and as such may not be used to fill local demand.

FOB – Forward Operating Base

The FOB is the distant end of the MALSP II logistical support network. The FOB supports typically between 2 and perhaps 32 aircraft.

FOSP – Follow-On Support Package

The Follow-On Support Package contains approved allowances in excess of those contained in the other packages.

FSA – Fly-in Support Allowance

The FSA is a MALSP II allowance package that is intended to provide PMALS with the means (in terms of repair parts) to support deployed aircraft. The FSA will primarily consist of O-Level parts.

Hi-Pri – High Priority

A document is deemed high-priority if the aircraft either cannot fly or cannot perform one of its missions without the part.

ICA – Intermediate Contingency Allowance

The ICA is a MALSP II allowance package that provides the PMALS with I-Level parts. The parts are primarily intended to support the PMALS and the IMA.

ICRL - Individual Component Repair Listing

The ICRL is a list that contains all the items a MALS or IMA has permission to repair, to include the types of repair actions it may perform.

IMA – Intermediate Maintenance Activity

The IMA resides at the MALS, and primarily refers to performance of maintenance on components and related support equipment; calibration of designated equipment; providing technical assistance; etc.

MAG – Marine Aircraft Group

The MAG is an operational organization (analogous to a Regiment) that generally consists of four to eight flying squadrons and one MALS.

MALS – Marine Aviation Logistics Squadron

Each Marine Aircraft Group is comprised of a certain number of flying squadrons and one MALS. Both the Intermediate Maintenance Activity and the group's entire inventory of repair parts reside at the MALS.

MALSP – Marine Aviation Logistics Support Program

The Marine Corps developed the Marine Aviation Logistics Support Program (MALSP) in the late 1980s in order to facilitate the effective logistical support of aviation units deployed to confront Cold War opponents on the plains of Europe and elsewhere. The MALSP construct consists of a system of packages of parts, personnel, support equipment, and mobile facilities that are modular and flexible enough to tailor to nearly any imagined contingency. MCWP 3-21.2 (Aviation Logistics) is the publication that enshrines MALSP into Marine Corps doctrine.

MALSP II – Marine Aviation Logistics Support Program II

MALSP II is the next generation of Marine Corps aviation logistics doctrine. It leverages Continuous Process Improvement (CPI) techniques such as Theory of Constraints, Lean, and Six Sigma to achieve desired or greater readiness with fewer resources.

MOB – Main Operating Base

The MOB's primary function is to be parent node to the FOB.

MSA - MAG Support Allowance

The MSA is a MALSP II allowance construct that is primarily intended to provide support to aircraft performing training missions (flying peacetime hours) directly supported by the PMALS.

NAVICP – Naval Inventory Control Point

A Wholesale Supply Activity that manages all Depot Level Repairables. NAVICP sets all allowances for repairable items for each MALSP.

NIIN – National Item Identification Number

Aviation Logisticians may use the terms NIIN and item interchangeably.

NIS – Not-In-Stock

NIS is a local status code that indicates that a node received a demand for an item, but the node had insufficient quantity on-hand in order to fulfill it.

Nodal Laydown

A term that describes the placement and structure of the nodes in the MALSP II network.

NRFI – Not Ready For Issue

Repairable parts that are not in working condition are classified as NRFI.

OMA – Organizational Maintenance Activity

This level of maintenance primarily consists of removal and replacement of defective parts; inspections; servicing, and incorporation of Technical Directives. In the Marine Corps, the OMA resides at the squadron level, therefore, it is common to refer to squadrons as Organizational Maintenance Activities.

PCSP – Peculiar Contingency Support Package

The Peculiar Contingency Support Package can support a squadron, or multiple squadrons, of aircraft. It contains both O and I level parts to support aircraft flying wartime hours for thirty days.

PEB – Pre-Expended Bin

Each squadron has a PEB which contains items with relatively high demand frequencies. The squadrons use so many that it is deemed worthwhile to allow them to maintain large quantities of these items rather than constantly ordering them from the MALSP.

PMALS – Parent MALS

The PMALS is the MALS that is ultimately supporting the entire MALSP II logistical network.

RFI – Ready For Issue

Repairable parts that are in working condition are classified as Ready For Issue. RFI may also be used as a verb, in the sense that the IMA may RFI a previously defective part.

RIP – Remain In Place

When a squadron orders a repairable, it usually needs to turn the NRFI carcass into the MALS before the MALS will issue a RFI component. There is a class of components, however, for which this exchange may be delayed (i.e. landing gears). These are known as RIP items.

RO – Reorder Objective

RO is essentially the maximum number of items a MALS should have on their shelves for a particular item. In all cases, on-hand quantity plus quantity on order should equal RO.

RO, Actual

Given constraints such as material availability, this quantity is the RO assigned to the node in practice. For example, suppose that ELAT calls for an RO of 2. However, if the MALS only has an allowance for 1 item, then the Actual RO for the node will be 1.

RO, Ideal

Under a given set of circumstances (i.e. TRR, number aircraft, etc), this quantity is the RO that ELAT determines is most appropriate for the item in question at the particular node.

R-Supply Database

Each MALS possesses their own stand-alone R-Supply Database they use to manage their inventories and finances (specifically it is used to set the RO for consumables). The database typically contains approximately 24 months of demand history for that MALS.

SOA – Supply Officer's Asset

SOA is a local classification that indicates that the part is “destined” for the Supply Officer's Inventory. For example, if a NRFI component in the repair process at the IMA is a SOA, it means that once the part is successfully RFI'd, it is to be returned to the Supply Department and will be added to its on-hand inventory.

Shortfall

Shortfall is the amount by which Ideal RO exceeds Actual RO for a given item at a particular node.

T-AVB

The T-AVB is the logistics support ship. The mission of the T-AVB is to provide rapid and dedicated sealift for employment of a tailored aviation Intermediate Maintenance Activity (IMA) to support deployment of US Marine Corps fixed and rotary wing aircraft.

T/M/S – Type / Model / Series

T/M/S is a way to reference a particular type of aircraft. E.g., one T/M/S that MALSP-29 in New River supports is the CH-53E. Most MAGs support multiple T/M/S.

TRR – Time to Reliably Replenish

TRR is a MALSP II concept that is intended to capture the “worst case” time required before a part is returned to the inventory. Note, this could either mean the time required for a stock document to arrive, a pickup replenishment, or the time required to successfully repair an item. The TRR is set at the 90th percentile of the distribution of all such times.

TRR, Actual

In contrast to the Design TRR, the Actual TRR is the empirical TRR experienced between two nodes in a MALSP II system.

TRR, Design

When a MALSP II system is in its planning stages, the TRR between two nodes is not necessarily known with certainty and must be estimated. The Design TRR is the initial estimate for the TRR between the nodes and is the TRR used to initially set all buffers.

TRR_{DTO}

The subscript DTO indicates that this TRR describes the distribution of shipping times from the wholesale system (possibly to include lateral support from external activities) to the PMALS. Such documents are typically of high-priority status, so TRR_{DTO} should be relatively low in practice.

TRR_M

The subscript M indicates that this TRR describes the distribution of maintenance times for a particular (repairable) NIIN. Maintenance time is counted from the time a part is turned into the IMA from the time the maintenance action is completed.

TRR_{Stock}

The subscript Stock indicates that this TRR describes the distribution of shipping times from the wholesale system for stock replenishment documents to the PMALS. Such documents have low priority, so TRR_{Stock} can be relatively high in practice.

WCRT - Weighted Composite Response Time

WCRT is a method to aggregate the values of Response Time at each node in the network in order to obtain an overall measure of the performance of the entire network. The parameter β is a notional weighting of the importance of the mission of the deployed aircraft relative to the training mission of the PMALS local demand.

$$WCRT_i = PMALS_Response_Time_i + \beta \cdot MOB_Response_Time_i$$

The subscript i is the index of the NIIN.

X1

X1 is a code from the ICRL List that indicates that a IMA does not have permission to attempt a repair on that item.

9. List of External Files Referenced

ch53e_MALSPII_Mar12.accdb

Access database that contains the prospective MALSPII allowances for CH-53Es. This database was the source for all allowance and AFAST demand data for those NIINs.

AFAST_Repairables_M16.xlsx

Excel spreadsheet summarizes the AFAST demand data for repairables at MALS-16. (Queried from ch53e_MALSPII_Mar12.accdb)

Candidates_rac.xlsx

Excel spreadsheet that contains NIINs recommended for inclusion into MALS-16 MALSP II packages. (Queried from ch53e_MALSPII_Mar12.accdb)

M16_cand_afast_join_29 July.xlsx

This spreadsheet joins the data contained in Candidate_rac.xlsx with the data contained in AFAST_Repairables_M16.xlsx. This spreadsheet was used to develop the unique NIIN-level demand parameters for Chapter 4.

Niin_sum_w_TRR_sub_M.xlsx

Spreadsheet that contains all maintenance actions and maintenance times experienced at MALS-16 over two year time on NIINs contained in the MALSP II packages. Provided by the MALSP II Program Office.

10. Mathematical Appendix

A. Items from Chapter 4

The items included in this section of the Mathematical Appendix are primarily the Low TRR versions of the Hi TRR graphs presented above, as well as the ANOVA tables for comparisons of mean Response Time for every scenario.




Figure A-1. MOB Response Time. (TRR = Low)

Figure A-1 is the TRR=Low equivalent to Figure 4-6.





Figure A-2. PMALS Response Time (TRR = Low)

Figure A-2 is the TRR=Low equivalent to Figure 4-7.

1. Configuration 0 ANOVA.



Means for Oneway Anova

Level	Number	Mean	Std Error	Lower 95%	Upper 95%
Legacy_no_ESB	100	2833.13	92.944	2649.8	3016.4
MALSPIL_no_ESB	100	964.82	92.944	781.5	1148.1

Std Error uses a pooled estimate of error variance

Comparisons for all pairs using Tukey-Kramer HSD


	q*	Alpha
	1.97202	0.05
Abs(Dif)-HSD		
	Legacy_no_ESBMALSPIL_no_ESB	
Legacy_no_ESB	-259.2	1609.1
MALSPIL_no_ESB	1609.1	-259.2

Positive values show pairs of means that are significantly different.

Figure A-3. ANOVA for Configuration 0.

Figure A-3 shows the ANOVA for Configuration 0 and corresponds to Figure 4-5. The Tukey HSD test confirms that the difference between the Legacy and MALSP II package performance is statistically significant.

2. MOB Response Time: Configuration _a; High TRR



Means for Oneway Anova

Level	Number	Mean	Std Error	Lower 95%	Upper 95%
Legacy_no_ESB	100	2808.07	158.49	2496.5	3119.7
Legacy_with_ESB	100	3044.74	158.49	2733.2	3356.3
MALSPII_no_ESB	100	1200.58	158.49	889.0	1512.2
MALSPII_with_ESB	100	1822.22	158.49	1510.6	2133.8

Std Error uses a pooled estimate of error variance


Level	Mean
Legacy_with_ESB A	3044.7363
Legacy_no_ESB A	2808.0745
MALSPII_with_ESB B	1822.2244
MALSPII_no_ESB C	1200.5754

Levels not connected by same letter are significantly different.

Figure A-4. ANOVA for Configuration _a; High TRR

Figure A-4 is the ANOVA for Configuration _a and High TRR and corresponds to the left-hand panel of Figure 4-6.

3. MOB Response Time: Configuration _b; High TRR



Means for Oneway Anova

Level	Number	Mean	Std Error	Lower 95%	Upper 95%
Legacy_no_ESB	100	1087.47	79.418	931.3	1243.6
Legacy_with_ESB	100	1199.00	79.418	1042.9	1355.1
MALSPIL_no_ESB	100	414.37	79.418	258.2	570.5
MALSPIL_with_ESB	100	730.57	79.418	574.4	886.7

Std Error uses a pooled estimate of error variance

Level	Mean
Legacy_with_ESB A	1199.0029
Legacy_no_ESB A	1087.4710
MALSPIL_with_ESB B	730.5667
MALSPIL_no_ESB C	414.3744

Levels not connected by same letter are significantly different.

Figure A-5. ANOVA for Configuration _b; High TRR

Figure A-5 is the ANOVA for Configuration _b and High TRR and corresponds to the middle panel of Figure 4-6.

4. MOB Response Time: Configuration _c; High TRR




Figure A-6. ANOVA for Configuration _c; High TRR

Figure A-6 is the ANOVA for Configuration _c and High TRR and corresponds to the right-hand panel of Figure 4-6.

5. MOB Response Time: Configuration _a; Low TRR




Figure A-7. ANOVA for Configuration _a; Low TRR

Figure A-7 is the ANOVA for Configuration _a and Low TRR and corresponds to the left-hand panel of Figure A-1.

6. MOB Response Time: Configuration_b; Low TRR




Figure A-8. ANOVA for Configuration_b; Low TRR

Figure A-8 is the ANOVA for Configuration_b and Low TRR and corresponds to the middle panel of Figure A-1.

7. MOB Response Time: Configuration _c; Low TRR




Figure A-9. ANOVA for Configuration _c; Low TRR

Figure A-9 is the ANOVA for Configuration _c and Low TRR and corresponds to the right-hand panel of Figure A-1.

8. PMALS Response Time: Configuration _a; High TRR




Figure A-10. ANOVA for Configuration _a; High TRR

Figure A-10 is the ANOVA for Configuration _a and High TRR and corresponds to the left-hand panel of Figure 4-7.

9. PMALS Response Time: Configuration _b; High TRR




Figure A-11. ANOVA for Configuration _b; High TRR

Figure A-11 is the ANOVA for Configuration _b and High TRR and corresponds to the middle panel of Figure 4-7.

10. PMALS Response Time: Configuration _c; High TRR




Figure A-12. ANOVA for Configuration _c; High TRR

Figure A-12 is the ANOVA for Configuration _c and High TRR and corresponds to the right-hand panel of Figure 4-7.

11. PMALS Response Time: Configuration _a; Low TRR




Figure A-13. ANOVA for Configuration _a; Low TRR

Figure A-13 is the ANOVA for Configuration _a and Low TRR and corresponds to left-hand panel of Figure A-2.

12. PMALS Response Time: Configuration _b; Low TRR





Figure A-14. ANOVA for Configuration _b; Low TRR

Figure A-14 is the ANOVA for Configuration _b and Low TRR and corresponds to middle panel of Figure A-2.

13. PMALS Response Time: Configuration _c; Low TRR



Means for Oneway Anova

Level	Number	Mean	Std Error	Lower 95%	Upper 95%
Legacy_no_ESB	100	4465.21	181.67	4108.0	4822.4
Legacy_with_ESB	100	7669.47	181.67	7312.3	8026.6
MALSPIL_no_ESB	100	3906.43	181.67	3549.3	4263.6
MALSPIL_with_ESB	100	4059.38	181.67	3702.2	4416.5

Std Error uses a pooled estimate of error variance

Level	Mean
Legacy_with_ESB A	7669.4731
Legacy_no_ESB B	4465.2070
MALSPIL_with_ESB B	4059.3807
MALSPIL_no_ESB B	3906.4251

Levels not connected by same letter are significantly different.

Figure A-15. ANOVA for Configuration _c; Low TRR

Figure A-15 is the ANOVA for Configuration _c and Low TRR and corresponds to right-hand panel of Figure A-2.

14. Effect of Adding ESB on Response Time for Low TRR design points.

Table A-1. MOB Response Time for Configuration _a by Category (TRR=Low)

Category	NIINs	without ESB	with ESB
A	256	22.5	56.6
B	45	67.5	47.7
D	131	384.7	594.3
Total	432	474.7	698.5

Table A-2. PMALS Response Time for Configuration _a by Category (TRR=Low)

Category	NIINs	without ESB	with ESB
A	256	335.7	368.0
B	45	110.9	550.9
D	131	2461.7	2159.0
Total	432	2908.3	3077.9

Table A-3. MOB Response Time for Configuration _c by Category (TRR=Low)

Category	NIINs	without ESB	with ESB
A	256	1.1	3.8
B	45	10.0	6.0
D	131	92.8	140.9
Total	432	103.9	150.7

Table A-4. PMALS Response Time for Configuration _c by Category (TRR=Low)

Category	NIINs	without ESB	with ESB
A	256	409.6	430.8
B	45	127.9	708.3
D	131	3368.9	2920.3
Total	432	3906.4	4059.4

These tables are the TRR=Low equivalents to Tables 4-3 through 4-6.




Figure A-16. Configuration_a NIINs with Higher WCRT from Adding ESB.

Figure A-16 is the TRR=Low equivalent to Figure 4-10.




Figure A-17. Configuration_c NIINs with Higher WCRT from Adding ESB.

Figure A-16 is the TRR=Low equivalent to Figure 4-11.

B. Items from Chapter 5

The experiment described in this chapter generates response surfaces for FOB *Response Time* and MOB *Response Time* as functions of Demand Frequency Type, Demand Quantity Type, PMALS TRR_{Stock}, PMALS TRR_{DTO}, ESB TRR, MOB TRR, FOB TRR, PMALS aircraft, Demand Filter Level, and Wartime Intensity. A partition tree is run first on each response surface to identify the most important factors.

	RSquare	RMSE	N	Number of Splits	AICc
Training	0.824	144.11154	6912	14	88360.9
Validation	-0.02	1768.2695	58332		




Figure A-17. Left-Hand portion of partition tree: Response = FOB Response Time.




Figure A-18. Right-Hand portion of partition tree: Response = FOB Response Time.

Scaled Estimates

Nominal factors expanded to all levels

Continuous factors centered by mean, scaled by range/2

Term	Scaled Estimate		Std Error	t Ratio	Prob> t
Intercept	1.5407091		0.011538	133.53	<.0001*
Dem_Freq[Low]	-1.210506		0.025801	-46.92	<.0001*
Dem_Freq[Mid-Low]	-0.999818		0.025801	-38.75	<.0001*
Dem_Freq[Mid]	-0.505738		0.025801	-19.60	<.0001*
Dem_Freq[Mid-Hi]	0.2772495		0.025801	10.75	<.0001*
Dem_Freq[Hi]	0.7354398		0.025801	28.50	<.0001*
Dem_Freq[Super Hi]	1.703373		0.025801	66.02	<.0001*
TRRpdto	0.0328655		0.011538	2.85	0.0044*
buffer_categorical[0]	2.2353038		0.019985	111.85	<.0001*
buffer_categorical[80]	-0.107786		0.019985	-5.39	<.0001*
buffer_categorical[95]	-0.586808		0.019985	-29.36	<.0001*
buffer_categorical[100]	-1.540709		0.019985	-77.09	<.0001*
Wartime Intensity	0.3120133		0.014132	22.08	<.0001*




Figure A-19. Multivariate Regression of Main Effects for FOB Response Time.

	RSquare	RMSE	N	Number of Splits	AICc
Training	0.942	78.40316	9792	14	113243
Validation	-0.03	1815.2914	55452		




Figure A-20. Left-Hand portion of partition tree: Response = MOB Response Time.




Figure A-21. Right-Hand portion of partition tree: Response = MOB Response Time.

Nominal factors expanded to all levels

Continuous factors centered by mean, scaled by range/2

Term	Scaled Estimate		Std Error	t Ratio	Prob> t
Intercept	1.5292096		0.010258	149.08	<.0001*
Dem_Freq[Low]	-1.28118		0.026763	-47.87	<.0001*
Dem_Freq[Mid-Low]	-0.907222		0.022236	-40.80	<.0001*
Dem_Freq[Mid]	-0.43161		0.022236	-19.41	<.0001*
Dem_Freq[Mid-Hi]	0.3620243		0.022236	16.28	<.0001*
Dem_Freq[Hi]	0.9683342		0.022236	43.55	<.0001*
Dem_Freq[Super Hi]	1.2896537		0.022236	58.00	<.0001*
TRRpdto	0.0422856		0.010136	4.17	<.0001*
buffer_categorical[0]	2.1464354		0.017557	122.26	<.0001*
buffer_categorical[80]	0.0400079		0.017557	2.28	0.0227*
buffer_categorical[95]	-0.582518		0.017557	-33.18	<.0001*
buffer_categorical[100]	-1.603925		0.017557	-91.36	<.0001*
Wartime Intensity	0.4620149		0.012898	35.82	<.0001*




Figure A-22. Multivariate Regression of Main Effects for MOB Response Time.

C. Items from Chapter 6

The experiment described in this chapter generates a response surface for *Response Time* as a function of Demand Frequency Type, Demand Quantity Type, Design TRR, Delta, Aircraft, Demand Filtering Level, and Wartime Intensity. A partition tree is run first on each response surface to identify the most important factors.

	RSquare	RMSE	N	Number of Splits	AICc
Training	0.408	280.30881	16128	13	227591
Validation	-0.13	2793.4202	16128		




Figure A-22. Left-Hand portion of partition tree: Response = *Response Time*.




Figure A-23. Right-Hand portion of partition tree: Response = *Response Time*.

Scaled Estimates

Nominal factors expanded to all levels

Continuous factors centered by mean, scaled by range/2

Term	Scaled Estimate		Std Error	t Ratio	Prob> t
Intercept	1.5950352		0.010821	147.41	<.0001*
Demand Freq Type[Low]	-1.081211		0.015303	-70.65	<.0001*
Demand Freq Type[Mid-High]	0.3750541		0.015303	24.51	<.0001*
Demand Freq Type[Very High]	0.7061572		0.015303	46.15	<.0001*
Demand Qaunt Type[hi]	0.3338717		0.015303	21.82	<.0001*
Demand Qaunt Type[med]	0.1046212		0.015303	6.84	<.0001*
Demand Qaunt Type[single]	-0.438493		0.015303	-28.65	<.0001*
designTRR	-0.371871		0.016812	-22.12	<.0001*
delta	0.8647829		0.018897	45.76	<.0001*
pmals_ac	0.9647441		0.014517	66.45	<.0001*
Buffer Sizes[80]	0.7969659		0.018742	42.52	<.0001*
Buffer Sizes[90]	0.4657224		0.018742	24.85	<.0001*
Buffer Sizes[95]	0.1760118		0.018742	9.39	<.0001*
Buffer Sizes[100]	-1.4387		0.018742	-76.76	<.0001*

Figure A-24. Multivariate Regression of Main Effects for *Response Time*.

11. References

Aviation Logistics, *Marine Corps Warfighting Publication* 3-21.2

Law, A. and W. Kelton (2000). *Simulation Modeling and Analysis*. 3d Ed. Boston: McGraw Hill.

Garant, P. (2004). “The Transformation of Marine Aviation Logistics.” *Marine Corps Gazette* 88 (5).

Steward, D. (2008). “Pushing a Pull System.” *Defense AT&L*, July-August, 40 – 43.