QUANTIFYING THE EFFECT OF TRANSPORTATION PRACTICES IN MILITARY SUPPLY CHAINS

Joshua B. McGee, B.S.I.E. / Research Assistant, University of Arkansas
Manuel D. Rossetti, Ph.D., P.E. / Associate Professor, University of Arkansas
Scott J. Mason, Ph.D., P.E. / Associate Professor, University of Arkansas

ABSTRACT

Traditional multi-echelon inventory and readiness-based models have not fully examined the ability of effective transportation utilization to reduce cost, delay times, and improve readiness in the overall military logistics network. In this paper, we develop a simulation-based methodology for quantifying the effect of transportation options (i.e. truckload shipping, less-than-truckload shipping, transshipments, and express air shipping) on shipping costs, customer wait times, abort rates, and operational availability. Simulation was used to develop a multi-echelon (depots, bases) model of regional supply chain support for aircraft spare part maintenance activities. The resulting model was used for experimentation and to develop response surface equations for the behavior of the system. The logistics implications of the results are discussed as well as managerial insights into the behavior of such systems. Our analysis indicates that focusing more on local inventory and local repair can have a significant impact on the operational availability of the system. This shift should be looked at in terms of the cost of local repair resources compared to less transportation costs.

1 INTRODUCTION

Muckstadt (2005) presents a comprehensive overview of multi-indentured multi-echelon (MIME) spare part inventory systems and their analytical treatment. The MIME system is one that has been studied in some depth over the years. Through the course of these studies many mathematical-based models have been developed to explore the MIME system. Sherbrooke (1968 and 1986) developed one of the first mathematical models known as the METRIC model, and then extended his model in the 1986 Vari-METRIC model. Similar models were analyzed in papers by Muckstadt (1973), Nahmias (1979), Slay (1984 and 1996), and Graves (1985). These models provide valuable insight into the MIME system; however, they must make many limiting assumptions which hamper their ability to provide detailed analysis. Attempts to relax these limiting assumptions inside the mathematical models lead to intractability. Through the development of simulation-based models, these limiting assumptions can be relaxed, providing a model that captures more of the subtlety and variation in the system. For the purposes of this paper a simulation model was developed based of a simplified representation of an Air Force MIME repairable parts system. Our model has the following characteristics:

- Supply chain that includes bases, depots, and suppliers.
- Multiple part types with different failure characteristics.
- Multi-Indenture part structure with a parent component consisting of many child components. The failure of the parent component is dependent on the failure its (child) components. Multiple child components can be failed after an operation.

- Less-than-truckload and truck-load shipping options are included and transit times can be stochastic
- Stochastic repair can occur both at the base and the depot.
- End-item availability is the key performance measure.

We know of no analytical model which encompasses all of these aspects, especially with regards to operating under stochastic conditions. Thus, we conclude that standard queuing and inventory models would not be applicable for this situation. It is also doubtful that solutions provided by analytical models under these circumstances would be tractable, and certainly they will not be as flexible as simulation.

In the past, military planners have achieved readiness rates by relying on large inventories; however planner now seek to implement quicker, more agile logistics systems which will reduce inventory on hand while improving readiness with the same or fewer dollars (Condon 1999). To this end, military logisticians have undertaken a variety of initiatives, such as Lean Logistics and Velocity Management, to improve responsiveness and reduce the total cost of inventory by decreasing logistics pipeline times. Within private, multi-echelon inventory systems, similar to that of the military's, commercial practices have significantly reduced the need for inventory stockpiles by reducing pipeline times. Using simulation, this paper assesses the effect of applying such commercial practices to military supply chains, and then evaluates the results.

2 DISCRETE-EVENT SIMULATION DESCRIPTION

A stochastic discrete-event simulation model was developed in Arena © 7.01 to simulate a simplified representation of a military repairable parts supply chain. This model was used to compare various commercial logistic practices. The simulation model is both generic (it can handle any number of bases, aircraft, and parts) and data driven. In the instance of the model discussed in this paper, there are six independent bases supported by a single depot. There are twenty-four aircraft assigned to each squadron, three squadrons assigned to each wing, and one wing assigned to each base. In this structure, there are a total of 72 aircraft assigned to each base. This results in a total of 432 aircraft within the system. The six bases are divided into two regions, with three bases in each region. Figure 1 details the structure for the model.

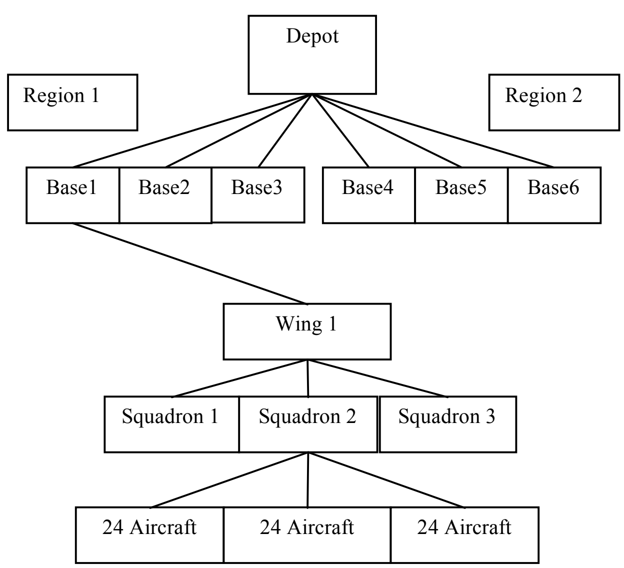

Figure 1: Supply Chain Structure

The model represents weapon systems (aircraft) as objects with two levels of indenture within the component part structure. The first level of indenture entails aircraft which are made up of multiple Line Replaceable Units (LRUs). These LRUs are in turn comprised of multiple Shop Replaceable Units (SRUs) constituting the second level of indenture. Each of the 432 aircraft in the system is comprised of six LRUs, one of each type. Each of the six LRU types is comprised of four SRUs yielding a total of 24 distinct SRUs per aircraft. LRUs and SRUs of the same type are identical and interchangeable. Figure 2 illustrates the two levels of indenture used in the model. The subscript i denotes the LRU type, while j denotes the SRU type.

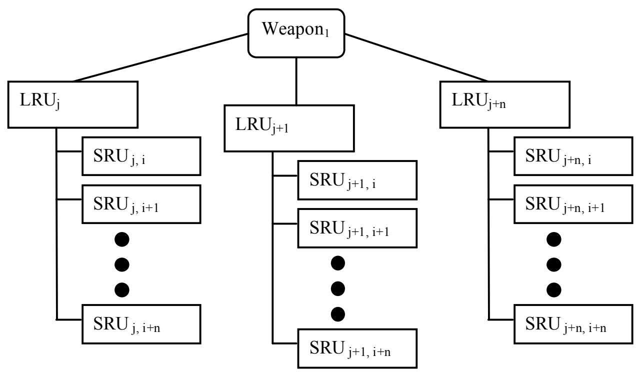

Figure 2: Hierarchy of Weapon System

For the purposes of this model, aircraft are always categorized as being in one of three states:

(i) Mission Capable (MC). An aircraft is designated MC when it is capable of flying a sortie. This status can correspond to an aircraft that is currently flying a sortie or is waiting to be assigned to a sortie.

(ii) Non-Mission Capable (NMC). An aircraft is designated NMC when one or more of its SRUs fails. This status corresponds to an aircraft that is down either awaiting a spare part or currently in the process of spare part installation. NMC aircraft cannot fly sorties.

(iii) Phase Inspection (PI). An aircraft enters phase inspection after it has accumulated a certain number of operating hours. During phase inspection the aircraft is inspected, repaired, and refurbished according to maintenance specifications. An aircraft is designated PI when it enters the phase inspection module. While in phase inspection the aircraft is not available to fly sorties; therefore, an aircraft listed as PI is also considered NMC.

The percentage of time each aircraft is in each state is tracked and reported as a key performance metric of the simulation model. There are other states which an aircraft in the real world system can occupy, but for this research the number of aircraft states was reduced to simplify the model. In further studies the number of weapon system states could be expanded to include states such as Partially Mission Capable, Cannibalization, etc.

The failure of an SRU results in the failure of an LRU and therefore the weapon system. We modeled SRU failure as a stochastic failure process in which SRU Time to Failure (TTF) values are generated from a user specified Mean Time to Failure (MTTF) distribution according to Monte-Carlo methods. The baseline model contains three levels of MTTF (in hours), each of which is modeled as an exponential distribution with some mean value. The three values which were used are as follows: high-exponential (500), medium-exponential (400), low-exponential (300). Eight SRU TTF values were generated from each of these three levels. The TTF is tracked for each of the SRUs operating on each of the aircraft in the simulation. While an aircraft is operational it accrues operating hours, and the TTF for each SRU on that aircraft is decremented equivalently. Aircraft failure occurs when any of the component SRU's TTF values equals or drops below zero. When an SRU fails, the part is sent to the repair process and a spare is requested from the base's inventory. If there is not a spare available within the base's inventory, an order is created and given a backorder status. Repair and replenishment is discussed further in the next few paragraphs.

When a SRU is deemed defective, it enters the repair process. The model first decides whether or not the part can be repaired at the base level. Due to repair resource constraints it is highly unlikely that a failed SRU can be repaired at the base level (Miller 1992). In the majority of cases, the SRU must be sent to the depot for repair. If the SRU can be repaired at the base the SRU enters the queue for the base repair process. If the SRU must be repaired at the depot, the SRU is delayed a shipping time generated from a user specified random distribution, and then enters the queue for the depot repair resource.

In the case that the part must be sent to the depot for repair, an order for the part is generated and sent to the depot. This order waits in the order queue at the depot. Priority is given to backorders in this queue. For every part that is sent to the depot, an order is generated for a part to be sent back to the base. This practice holds with a one-for-one inventory policy.

Upon completion of the repair process, the SRU becomes functional and the part is sent to inventory. If the SRU was repaired at the base, the base's inventory is incremented. If the part was repaired at the depot, the depot inventory is incremented. It is from this depot inventory that the orders are filled. The first order in queue, of the same type as the repaired SRU, is filled. After an order is filled, the order is shipped to the base where the order originated. Once the base receives the shipment, the repair part is entered into the base's inventory. When a base's inventory is incremented, each of the aircraft in the queue holding NMC aircraft is checked, and the first aircraft in queue needing a part of the same type moves to the installation process. If there are no aircraft in need of the SRU, the part remains in the base's inventory.

3 SHIPPING

Parts can be shipped between echelons in two ways: ground shipping and express air shipping. Most parts are shipped via trucks that pick up and drop off parts at the bases and depot; however, when an aircraft is listed as Mission Impaired Capability Awaiting Parts (MICAP) its part shipments are expedited. These parts are express shipped, usually arriving at the final destination in one or two days but at a higher cost. The next few paragraphs discuss the shipping options explored in this paper and when each is used.

When an SRU failure occurs, its parent aircraft is listed as NMC due to the fact that all SRUs in our simulation model are required for mission capability. A search is then made through the aircraft's assigned base's inventory for the needed SRU. If the SRU is found within the base's inventory, the SRU is installed on the aircraft immediately. In this case, standard shipping is used to move the failed part to the depot and the ordered part back to the base, if required. However, if the SRU is not available at the base, the SRU is given a backorder status and the aircraft is listed as MICAP. In this case, the express shipping is used to expedite the shipments. The effect is to expedite the shipping of the part from the base to the depot for repair. When the depot receives an order that has a backorder status, it fills the order by shipping the first available part of that type as MICAP back to the base.

In the model, we simulate the use of both Less-than-Truckload (LTL) and the Truckload (TL) commercial carriers. These two options constitute standard shipping as previously described. The standard shipping feature is controlled through two variables, truck capacity and minimum batch size. The truck capacity dictates the number of SRUs each truck can hold. Minimum batch size is a percentage, which is multiplied by the truck capacity. The resulting value is the smallest number of SRUs that warrant a truck trip. For example, in the model the truck capacity is set to 20 SRUs. To turn on the LTL option, the minimum batch size is set to 20%; therefore, a shipping point must have at least 4 SRUs waiting to be shipped to warrant a truck trip to that location. If that location has less than 4 SRUs waiting to be shipped, a pickup is not ordered from the LTL carrier; however, if that location has 4 or more SRUs waiting, a pickup is ordered and all parts waiting to be shipped from that location are picked up by the carrier. To simulate the TL scenario in the model, the minimum batch size percentage is set to 100%. This means that 100% of the truck capacity must be waiting at a shipping point before a pickup is ordered. Currently, a single check of the items awaiting shipment at each location is made each day at 8:00 a.m. This is true for both the LTL and TL case.

When parts receive a backorder status they are shipped with the MICAP designation. Parts receiving a MICAP designation wait in a separate queue for express shipment. At 8:00 am each day, a commercial air shipping service picks up all the parts needing air shipping and ships them to their respective locations both at the bases and the depot. MICAP shipping times are generated from a triangular distribution with parameters (22,24,26) hours. These parameters were supplied by Air Force personnel as representative of next day delivery. The model assumes that the express shippers have unlimited capacity. This allows the model to rely on MICAP if the standard shipping option cannot keep up with the shipping volume, just as the Air Force uses MICAP to expedite shipping.

A lateral transshipment (LTS) is defined as a shipment between locations on the same echelon of the model structure, in this case a shipment between bases. In the model there are two regions of three bases each. If the LTS feature is turned on, when a failure occurs, the model will first check the base inventory for a spare, then the bases within the region, and finally the depot. When a search is made of the bases within a region the model selects the base with the most inventory available for that specific SRU. Once such a selection is made, a shipment is initiated from the selected base. If none of the bases in the region have inventory available, the order is sent directly to the depot. The transshipment scenario assumes that the bases within a region are closer to each other than to the depot, and therefore can fill the need in a time effective manner.

In Chapter 5 of their book, Law and Kelton (2000) discuss "six classes of techniques for increasing the validity and credibility of a simulation model." We performed the following five of the six techniques:

1. Collect high-quality information and data on the system – We completed a thorough review of Air Force maintenance and supply policies as well as conducted detailed conversations with Air Force logistics managers.

2. Interact with managers on a regular basis – Air Force logistics managers supported our effort by providing us with initial system descriptions and data as well as providing feedback throughout the model development phase.
3. Maintain an assumptions document and perform a structured walk-through – Throughout the model development process, we documented all simplifying assumptions and performed regular model walk-throughs with our Air Force contacts.
4. Validate the output from the overall simulation model - A series of simulations were run using realistic data inputs allowing us to compare our model outputs to target values set by the Air Force. In particular, we compared our operational availability rates and aircraft sortie abort rates to Air Force operational specifications and concluded that we were within tolerance.
5. Animation – Animation was used in testing and debugging the simulation to verify and validate the detailed workings of the model.

After performing these verification and validation techniques, we then developed a set of experiments to explore the effect of using commercial shipping practices along with other factors on the Air Force supply chain. These experiments provided an understanding of the factors that contribute to the availability of aircraft within the Air Force's supply chain. This understanding will allow logistics planners to better grasp the effect that their decisions may have on the operational performance of the system.

4 EXPERIMENTAL DESIGN

A factorial experimental design was used in our experimental studies. In a full factorial design, design points are investigated at all possible factor combinations. The experiments identify the main effects and the interactions between the factors. In this paper, there are 11 factors under study, and each factor has two levels. This is represented as a 2^k factorial design. If a full factorial design were run for the 11 factor experiment described above it would require 2^{11} = 2048 runs. Therefore, in our experiments, a fractional factorial design was utilized. In a fractional factorial design, a reduced number of runs are used to analyze the main affects and interactions between the factors, albeit with less granularity. A 1/16 fractional design was chosen, requiring 128 runs of the experiment, rather than the full 2048 runs. This experimental design is termed a Resolution V Design. In a Resolution V Design, no main effect or two factor interaction is confounded with any other main effect or two factor interaction (Montgomery and Runger 1999). Table 1 lists the 11 factors examined in this paper with a brief description of each. Table 2 outlines the factor values used during experimentation. The values listed in Table 2 are important later in understanding the results of our experiments. The factors, Shipping Option (A), MICAP (D) and Transshipment (K) are either present or not present, with (-1) indicating not present and (+1) indicating present. The factors, Sortie Duration, Sortie Frequency, Repair Time, Time to Failure, Pre/post Flight Operations, and Unscheduled Maintenance have two levels where the value assigned to the low level (-1) is the probability distribution specified in the table and the value assigned to the high level (+1) represents a 20% increase over the low level. The Inventory Position factor controls the allocation of inventory to local or depot level. There are 38 spares total per SRU in the system under each scenario. For the Low (-1) level there were 3 spares for each SRU held at each of the

6 bases and 20 spares for each SRU held at the depot. For the High (+1) level there were 6 spares for each SRU held at each of the 6 bases and 2 spares for each SRU held at the depot. For the Local Repair factor low (-1) indicates that only 1% of the parts can be repaired locally while the high (+1) indicates that 25% of the parts can be repaired locally.

Table 1: Factors Examined and Descriptions

Factors	Description
Standard Shipping	This determines whether Truck Load or Less Than Truck Load shipping will be used.
Sortie Duration	This factor refers to the actual length of a sortie.
Sortie Frequency	The number of sorties assigned to a base each day.
MICAP	This determines whether express shipments will be used.
Repair Time	Repair time is the delay time for the repair process.
Inventory Position	In the model, Inventory is set up to be either centralized or decentralized. Centralized indicates that more of the system wide inventory is held at the depot while, decentralized means that more of the inventory is held at the bases.
Time To Failure	This factor refers to the time to failure of individual SRUs.
Pre/Post Flight Maintenance	This factor refers to all the maintenance operations that are required to prepare an aircraft for flight and maintenance operations, which are performed after the flight has taken place. The operations currently included in this factor are: Refuel/Weapons Load, Engine Start, Final Systems Check, and Taxiing, Pre-Flight Check, Parking and Recovery, and Service/Debrief.
Unscheduled Maintenance	This encompasses all operations associated with the failure of a part. The operations included are: Troubleshooting, Remove Part, Wait for Part to Issue From Supply, Delay for Paperwork, Installation, Operational Check, Operational Check, Signoff Discrepancy, Document Corrective Action
Local Repair	This dictates the percentage of parts that can be repaired at the base level.
Lateral Transshipment	This factor indicates whether or not transshipments can be used as a source of supply.

Table 2: Factor Values Used in Experiments

	Factor	**Low**	**High**
A	Shipping Option	LTL	TL
B	Sortie Duration	Triangular (.333,1.747,2)	Triangular (.333,1.747,2)*1.2
C	Sortie Frequency	ANINT(Uniform(56,67))	ANINT(Uniform(56,67))*1.2
D	MICAP	On	Off
E	Repair Time	Exponential (8)	Exponential (8)*1.2
F	Inventory Position	Depot	Local
G	Time to Failure	Exponential (300) Exponential (400) Exponential (500)	Exponential (300)*1.2 Exponential (400)*1.2 Exponential (500)*1.2
H	Pre/Post Flight Operations	Normal Levels	Normal Levels +20%
I	Unscheduled Maintenance	Normal Levels	Normal Levels +20%
J	Local Repair	1% of parts repaired locally	25% of parts repaired locally
K	Transshipment	On	Off

Each simulation was given a warm-up period, a run length, and a specified number of replications. Initial tests were used to establish the warm-up period and run length for the experiments. In these tests, time persistent response data was analyzed both statistically and graphically. These initial tests showed that after six months the model appeared to reach a steady state, and that a year of data collection would allow us to make statistical inferences based on our results. The simulation was set to run 128 instances, each of which represents a different combination of factors or a single design point within the experiment. At the beginning of each of these instances, the level of each of the factors is read into the model. The simulation is warmed up at the beginning of each instance, and the system statistics were cleared after each of the runs. Therefore, the simulation model collects data for 128 independent design points and does not use the method of common random numbers. By not using common random numbers, we are able to directly apply standard statistical analysis techniques and do not need to rely on the more complicated techniques required when the design points are not independent. Each of these 128 design points was replicated five times using a different stream of random numbers for each of the five replications, yielding a total of 640 independent observations.

5 DATA AND DATA ANALYSIS

For this experiment, eight different responses were used to measure the effect that the factors had on the system. The Table 3 lists these responses along with a brief description.

Table 3: Responses Used in Experiment and Descriptions

Responses	Description
Operational Availability	This is the ratio of time a plane is either available to fly or flying to the time a plane is unavailable due to scheduled or unscheduled maintenance.
Abort Rate	This is the ratio of sorties aborted to sorties scheduled. Sorties may be aborted due to lack of planes or a failed part in pre-flight or in-flight inspection.
Customer Wait Time	Customer Wait Time refers to the time in hours from when a plane fails and enters unscheduled maintenance until the plane is available to fly again.
Total Transportation Cost	In our model only factors connected to shipping contribute to the total cost. These factors are MICAP, ground shipping, and transshipment. Each of these factors was assigned a cost per shipment. Data was collected for the number of each type of shipment, and that number was multiplied by the derived cost per shipment to yield the cost o each factor. The Total Transportation Cost is the sum of these three factor costs.
Sorties Flown	This is the cumulative number of sorties flown by an individual aircraft over the course of an experimental run.
Flight Hours	This is the total number of flying hours accrued by an individual aircraft over the course of a replication.
Times Failed	This is the total number of failures incurred by an individual plane over the course of a replication.
Total Backorders	This refers to the total number of backorders that occurred within a replication. A backorder occurs when a part fails and a replacement is not available in the bases inventory.

The first four responses listed in Table 3, Operational Availability, Abort Rate, Customer Wait Time, and Total Transportation Cost, were considered the most important metrics for scenario performance. The remaining responses were taken into account while reviewing scenarios to identify outliers and to evaluate the operational validity of the scenarios. Summary statistics were calculated for each of the eight responses to provide insight as to how the data behaves across all scenarios. Table 4 lists the summary statistics for the data collected on each response. The summary statistics included in Table 4 are: $\bar{x}$ - Sample Mean, s- Standard Deviation, and s.e. - Standard Error. The standard error is the standard deviation divided by the square root of the sample size. In this case there were 640 observations. The data values given in Tables 4 and 5 are for a 365 day year.

Table 4: Response Summary Statistics

Response	$\bar{x}$	s	s.e.
Operational Availability	75.28	5.75	0.2270
Abort Rate	0.14	0.08	0.0033
Customer Wait Time (Hours)	70.78	19.80	0.7830
Total Transportation Cost	72,244	44,770	1,770
Sorties Flown	294.26	20.21	0.8000
Flight Hours	413.66	27.33	1.0800
Times Failed	27.05	2.42	0.0960
Mean Total Backorders	3713.20	1466.9	58.0

Table 5 gives statistics that describe the distribution of each of the responses. Each data set is divided into four quartiles containing a quarter of the data points each. The four quartile values are defined as follows: min is the minimum value, Q_1 is the data point one-quarter through the data set, $\tilde{x}$ is the median, Q_3 is the data point three-quarters through the data set, and max is the maximum value.

Table 5: Response Distribution Statistics

Response	min	Q_1	$\tilde{x}$	Q_3	max
Operational Availability	57.253	72.06	76.82	79.42	84.52
Abort Rate	0.017	0.07	0.120	0.23	0.36
Customer Wait Time (Hours)	45.628	56.65	64.01	84.83	129.25
Total Transportation Cost	11,580	37,247	63,528	101,929	198,805
Sorties Flown	238.240	283.59	289.96	306.04	346.68
Flight Hours	336.470	391.76	418.49	439.38	454.36
Times Failed	22.660	25.38	26.45	29.11	31.48
Mean Total Backorders	578	2579	3631.5	4717.5	7478

The values in Table 5 show that the data collected on each response covered a wide range of values, and a significant portion of the data falls near realistic response targets. In our experimental runs 25% of the data for: Operational Availability was above 79.42%, Abort Rate was below 7%, and Customer Wait Time was below 56 hours. The wide range of response values was due to the large number of factors varied in our experiment. The large number of the factor combinations explored in this experiment would never logically be used in the real system (e.g. shifting inventory to the depot level and repair resources to the base level); however, these factor combinations were important in our statistical study of how the factors affected the response values. Due to the fact that this experiment is a relative comparison, some deviation in the mean of the response values from an actual Air Force repairable parts system do not affect the main results presented in this paper.

Table 6 provides more information about the distribution of the response data. The information provided in Table 6 details the probabilities that the associated response fulfills the logical statement listed. For example, in the case of Operational Availability, $P(X >= 80) = 0.186$ indicates that *across all* design points there is an estimated probability of 0.186 that operational availability is greater than or equal to 80%. Also, 65% of the data for Customer Wait Time is below 72 hours or 3 days. This data shows that for each response value our experimental results ranged very close to the targets set for an actual Air Force repairable parts system. The data establishes the reasonableness of our experimental design by matching our experimental results with an actual system. While this type of data was primarily used for validation, it is also of interest because it indicates that we have captured a wide range of interesting and realistic design scenarios in our experiments.

Table 6: Response Probabilities

Operational Availability	**P(x >= 80)**	0.1859
	P(x >= 75)	0.6172
Abort Rate	**P(x <= .05)**	0.1750
	P(x <= .15)	0.5453
	P(x <= .20)	0.7016
Customer Wait Time (Hours)	**P(x <= 24)**	0.0469
	P(x <= 72)	0.6547
	P(x <= 96)	0.8734
Total Transportation Cost	**P(x <= 25,000)**	0.1547
	P(x <= 50,000)	0.3516
	P(x <= 75,000)	0.6125

The relationships between some of the responses were analyzed graphically. Of specific interest is the relationship between Total Transportation Cost and Operational Availability. To reduce the total number of plotted points from 640 to 128, the five replications of each design point were averaged providing an estimate of the response for each design point. This reduced the variability in the graph to provide a clearer picture of the data trends. Figure 3 plots Operational Availability vs. Total Transportation Cost and shows the diminishing return between Operational Availability and Total Cost. This is a common trend when comparing performance metrics with total cost. There is usually a point at which spending increases faster than the improvement provided by the increased expenditure. Shipping plays a large role in this trend of diminishing returns. There are many ways to reduce customer wait time and increase operational availability through expedited or express shipments i.e. MICAP, but the cost of such practices grows at a rate that soon diminishes or even overtakes the value returned.

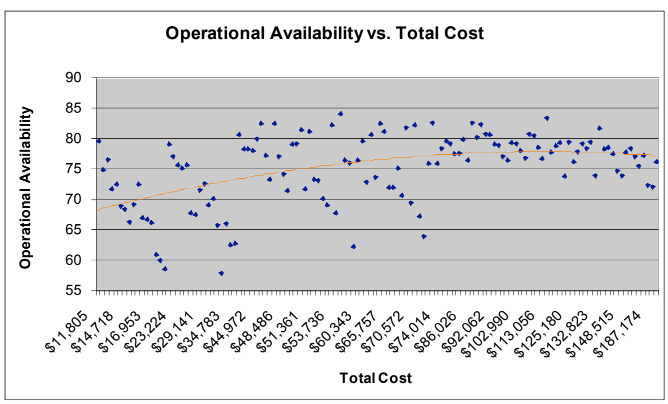

Figure 3: Operational Availability vs. Total Transportation Cost

Given the summary statistics on the behavior of each response across all scenarios, we know that our experimental design is representative of realistic operating conditions and that our response values cover the design space well. The next step was to take a detailed look at how each of our factors affected the response values discussed. Linear regression was used to quantify the effect that each of the factors had on each of the responses. This was completed using the software package MINITAB™. A response surface model was developed for each of the responses using the fractional factorial design to fit a linear regression equation between the factors of interest and the responses. Included in this analysis were all the main effects along with all first order interactions. The high-resolution or our design allows the inclusion of all first order interactions with complete confidence that there will be no confounding coefficients. Each of the eight regression formulas contain all the main effects as well as their interactions coming to a total of 66 terms. The regression formulas take the general form

$$Y = \beta_0 + \sum_{h=1}^{p} \beta_h x_h + \sum_{h=1}^{p-1} \sum_{j=h+1}^{p} \beta_{hj} x_h x_j + e \qquad \textbf{(1)}$$

where Y is the response level, β_0 is the intercept term, x_h is the factor level, and β_h is the first-order factor coefficient describing the effect the factor has on the response, and $\beta_{h,j}$ is the second order factor coefficient. This general form contains 66 terms and is very cumbersome. To reduce the number of terms included in the regression formulas, we evaluated the contribution that each of the included factors has within the regression models. We accomplished this by using the p-value. The p-value is the smallest level of α (alpha) at which the term can be deemed significant, where α is defined as the acceptable probability of error (Montgomery 1999). A p-value is generated for each of the main effects as well as the first order interactions. For this experiment, we chose an α of 0.01, and any factor with a p-value less than or equal to 0.01 was deemed to have a statistically significant effect on the response for this sample size.

The estimated effects and the regression coefficients were used to determine the practical significance of the factor in relation to the response. The estimated effect is defined as the change in the response produced by a change in the factor (Montgomery and Runger 1999). The regression coefficient is the actual coefficient in the regression formula associated with the factor. An examination of these statistics for a factor indicates the estimated change in the response experienced if the level of the factor is changed. In other words, these statistics describe the sensitivity of the response to changes in the associated factors. It is important to note that even if a factor is deemed statistically significant an examination of the estimated effect and regression coefficient may reveal that the factor is practically insignificant.

With the knowledge gained through these statistics, the regression models for each of the responses were simplified. This was done by removing the factors from the model that were either statistically or practically insignificant. The R^2 (coefficient of multiple determination) value was used as a metric for the change induced in the model's fit. Reducing the number of factors in the model will reduce the amount of variation explained by the model, but we can afford to be slightly less explanatory for the sake of simplicity. Table 2 associates a letter with each of the factors considered in the experiments. Table 7 lists the simplified equations that were developed for the four most important responses along with the resulting R^2 value. We will focus on these four responses for the rest of our discussion. The factor interactions are signified by an asterisk between two factors (i.e., A*B). The Appendix lists the ANOVA results for the operational availability response surface equation. Similar analysis was done for the other equations, but for the sake of brevity not included in this paper. A full detailed report provided to the Air Force Research Laboratory contains these additional details and is available upon request from the authors.

Table 7: Simplified Regression Equations

Simplified Regression Equation	R^2
Operational Availability (OA)	
OA = 75.279 + (-1.432 A) + (-1.104 B) + (-3.593 D) + (0.549 F) + (2.488 G) + (0.566 H) + (-1.05 I) + (2.034 J) + (-0.865 (A * D)) + (0.684 (D * G)) + (1.272 (D * J))	93.99%
Abort Rate (AR)	
AR = 0.13783 + (0.03938 B) + (0.06234 C) + (0.01912 D) + (-0.01295 G) + (0.01821 H) + (-0.01146 J)	91.86%
Customer Wait Time (CWT)	
CWT = 72.783 + (5.8 A) + (13.858 D) + (-3.23 G) + (3.749 I) + (-8.116 J) + (3.931 (A * D)) + (-5.5 (D * J))	92.42%
Total Transportation Cost (TTC)	
TTC = 72,244 + (-10,224 A) + (-33,774 D) + (-9,014 F) + (-9,265 G) + (-14,563 J) + (-5,317 K) + (-6,525 (A * D)) + (8,793 (D * F)) + (7,165 (D * G)) + (10,978 (D * J))	97.16%

The initial regression models contained 66 terms while the average number of terms included in the reduced models for these four responses is only 8.5. The factors within the equations listed in Table 7 have a large impact both statistically and practically on the value of the associated response. The coefficient assigned to each factor in the regression equations indicates the magnitude of the effect that a specific factor will have on the value of the response. For example, in the equation for Operational Availability factor D has a negative impact with magnitude 3.593 while factor G has a positive impact of 2.488. For every case where factor D takes on its high value, Operational Availability is reduced by 3.593. In this manner these equations can be used as rough predictors for the response based on factor values. Table 8 gives a break down of the most influential factors across all eight responses. The Number column lists the number of simplified regression models that include the associated factor, and the percent column lists the percent of the simplified regression models that include the associated factor. Table 8 accounts for all eight of the responses for which data was collected.

Table 8: Factor Influence All Responses

Factors	Number	Percent
A	6	75%
B	6	75%
C	4	50%
D	8	100%
E	0	0%
F	2	25%
G	8	100%
H	4	50%
I	2	25%
J	8	100%
K	1	12.5%
A*D	4	50%
B*C	2	25%
D*F	1	12.5%
D*G	4	50%
D*J	6	75%

The most influential factors are MICAP, TTF, and Local Repair. These factors are included in all of the simplified regression models. Table 8 indicates which factors, when changed, influence the greatest number of responses. It is notable that one factor, repair time, was not significant in any of the simplified models. This is due to its relative length in time as compared to other delays that have more effect on the system i.e. shipping time. This result does not mean that repair time is unimportant. In fact this result indicates that future modeling should focus on developing the relationship between the repair process and operational availability.

The regression equations developed from the simulation model allow direct what-if analysis without rerunning the simulation model. The regression equations also provide managerial insight into the behavior of this class of system. For example, within the range of the experimental design, in order to minimize transportation cost without any performance constraints, the regression equations recommend Scenario 1 in Table 9. In Table 9, Scenario 2 represents the case of maximizing operational availability with no constraints, and Scenario 3 represents minimizing transportation cost while ensuring an operational availability of at least 75%. For Scenario 2 high operational availability is achieved at a significant increase in the cost. This scenario also achieves low abort rates and low customer wait times, and it is achieved essentially by using less-than-truck load shipping. Scenario 3 indicates that we can still get low transportation cost and achieve reasonable operational availability by using full truck load shipping, improving reliability, and using MICAP expediting. All three scenarios used more local inventory and more local repair capability.

Table 9: Managerial Insight Scenarios

	Factor	Scenario 1	Scenario 2	Scenario 3
A	Shipping Option	Use FTL	Use LTL	Use FTL
B	Sortie Duration	Shorter sorties	Shorter sorties	Shorter sorties
C	Sortie Frequency	Less sorties	Less sorties	Less sorties
D	MICAP	Use MICAP	No MICAP	Use MICAP
E	Repair Time	Speed up repair	Speed up repair	Speed up repair
F	Inventory Position	More local	More local	More local
G	Time to Failure	Lower failure rate	Higher failure rate	Lower failure rate
H	Pre/Post Flight Operations	Normal	Normal + 20%	Normal + 20%
I	Unscheduled Maintenance	Normal	Normal	Normal
J	Local Repair	Repair more locally	Repair more locally	Repair more locally
K	Transshipment	No transshipment	No transshipment	No transshipment
	Operational Availability	78.004%	85.274%	79.136%
	Abort Rate	0.0126	0.01079	0.04903
	Customer Wait Time	75.777	47.46	75.77
	Transportation Cost	$10678	$44842	$10678

6 CONCLUSIONS

The focus of this research was to determine the effect of different shipping policies have on a military MIME supply chain. Also of interest were the relative effects when compared with other influential factors present in the system. In order to quantify the factor effects, we developed a set of experiments which allowed us to create regression meta-models for key performance measures. The regression meta-models described the effects of our experimental factors and their combinations on the performance measures and provided us with insights into the influence of the factors on a MIME system.

The largest contributing factors were MICAP, TTF, Local Repair, Shipping Option, Sortie Duration, and Inventory Position. Time to Failure and Sortie Duration were also influential in our experiments; however, the fact that they are significant may be less interesting from a practical standpoint. It is easy to see that more reliable parts within a MIME system will allow better performance. It does point out that investing in more component reliability should not be neglected. Along the same lines, if sortie duration is increased, more flight hours will accrue per aircraft resulting in an increase in failures; however, it is difficult to reduce the duration of sorties by any practical means without compromising mission objectives. The other four factors have greater implications.

Over the past decade (absent active combat) logistical defense budgets have been reduced. This in turn has impacted the way the military supply chain operates. Inventory levels in the supply chain have been falling along with the budgets. Today the military supply chain is being asked to be "more flexible and responsive" with less inventory and "at a lower total cost."(Condon and Cunningham 1999) The pressure to reduce both inventory and spending has induced stress on the military supply chain. As the inventory levels fell through the 1990s and into the present, it became harder to maintain a reliable flow of material. The Air Force has compensated for the low inventory levels by using express carriers, and they have been successful; however the cost of relying on these express carriers is high. The MICAP factor simulated the use of express carriers to expedite shipping times, and was one of the most influential factors in our experimentation. The cost of MICAP shipments was the largest cost component in our simulation model. Figure 3 illustrates a diminishing returns relationship between transportation cost and operational availability. The cost of MICAP shipments was the largest driver in the shape of this curve.

The use of express carriers is beneficial up to a point, but relying on the carriers soon becomes counter productive, becoming a money sink without providing proportional benefit. Reducing the reliance on MICAP in an Air Force type supply chain would both reduce transportation cost and force the exploration of new opportunities for improvement. MICAP does provide some benefit, but exploring other means for increasing operational availability could provide more benefit for the cost involved. Masciulli and Cunningham (2001) indicate that a redefinition of the policies governing MICAP shipments along with selection rules when it comes to choosing a commercial carrier would be beneficial from a cost standpoint. Their research indicated only a small percentage change in cost, while our research indicates significant cost reductions through greatly reducing if not eliminating MICAP without drastic effects on the other responses. Our experiments indicate that the emphasis should not be internal to the MICAP policies, but should instead be external, focusing on the inventory policies, inventory structure, and other lower cost shipping options.

There are many alternatives which can be explored to reduce the Air Forces reliance on MICAP shipments. Carter and London (2002) explore the SRU inventory levels. They indicate that current inventory levels for specific SRUs are not meeting demand while others are overstocked. In addition, they argue that the probability of LRU failure should drive inventory levels. They also make the point that, when setting inventory levels for repairable parts, a cost balance should be reached between the one time cost of purchasing the item and the cost of a backorder for that item. Larvick (2000) discusses the logistics system as a whole and describes “reach-back capability”. Reach-back capability refers to the greater ability of the upper echelon levels to respond to variation at the lower levels. In Larvick’s model, increased sortie duration and frequency result in more failures at the base level. He refers to reach-back as the ability of the higher echelons to respond to this change. This concept ties directly into the idea of increased supply chain visibility. Murphy (1999) explores the possibility of “Collocating Air Force weapon systems inventory with the Defense Logistics Agency premium service facility.” These research efforts touch on just a few areas where there are opportunities for system improvements which could result in a military supply chain that is robust and reliable, but does not rely so heavily on costly MICAP shipments.

The two other shipping factors that were investigated in our experiments were Shipping Option (LTL/TL) and Transshipment. The Shipping Option factor explored the difference in using LTL vs. TL shipping. In our experiments, a cost benefit was seen when using the TL shipping option; the lowest costs were realized in scenarios using TL shipping. These cost differences, however, were overshadowed by the cost of MICAP shipping. In the same light, the Transshipment factor did not have a large effect in our experiments. This, again, was due to the fact that the MICAP option had a dominating effect. In future models, these shipping options as well as direct shipments and scheduled deliveries should be explored in a more detailed fashion outside the shadow of MICAP.

The Inventory Position factor also played an important role in the response values. Shifting inventory to the base level was beneficial. This worked in conjunction with increasing the Local Repair capabilities. In other words the best performing scenarios had more repair resources at the base level along with more of the spares inventory being pushed to the bases. The cost of these capabilities was not fully explored in our experiments. Extending more of the repair resources to the base level would be a costly operation, but the benefits could be far-reaching. The Air Force’s new MICAP prevention program through proactive demand leveling (PDL), see Blazer et al. (2004) is a step in this direction. In this approach, demand is analyzed at both the base and the system level to gain a better understanding of bench stock requirements. Bench stock is material on hand at the base for repair at the base. This approach recognizes (at a fundamental level) how bench stock can improve maintenance and repair activities. By shifting the focus to having materials closer to initial demand and point of use, operational availability can be directly improved. More repair capability at the base implies a return to three levels of maintenance. Our experiments in this regard showed benefit over two levels of maintenance in operational availability and transportation cost. Our study showed that using three levels of maintenance is a viable option to reduce transportation cost and improve fleet availability, and should be investigated in future models along with other inventory and repair options.

There are many opportunities for expansion of the simulation model developed for this paper. The following are areas where model expansion would be of benefit to future studies: repair process, cannibalization, queue prioritization, sortie generation and assignment, inventory policies and costing, shipping alternatives, policies, and interaction. Investigating these areas would allow us to further explore the findings presented in this paper, specifically, the influence of inventory and repair resource allocation policies and their tradeoff with reliance on express shipping. Future work has already been funded and is in the beginning stages for expanding the model presented in this report to explore the Sortie Generation process. The goal is to extend the current simulation and mathematical modeling methodologies to assist unit-level maintenance managers in analyzing the effects of different sortie scheduling policies and identifying risk in different scheduling plans. The extended model will encompass sortie generation, maintenance activities, and the effects of limited equipment and inventory.

ACKNOWLEDGMENT

This material is based upon work supported by the Air Force Research Laboratory Logistics Readiness Branch (AFRL/HEAL). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

APPENDIX

Estimated Effects and Coefficients for Operational Availability

Term	Effect	Coef	SE Coef	T	P
Constant		75.279	0.05626	1338.02	0.000
A	-2.864	-1.432	0.05626	-25.46	0.000
B	-2.209	-1.104	0.05626	-19.63	0.000
D	-7.187	-3.593	0.05626	-63.87	0.000
F	1.098	0.549	0.05626	9.76	0.000
G	4.975	2.488	0.05626	44.21	0.000
H	1.131	0.566	0.05626	10.05	0.000
I	-2.101	-1.050	0.05626	-18.67	0.000
J	4.068	2.034	0.05626	36.15	0.000
A*D	-1.730	-0.865	0.05626	-15.37	0.000
D*G	1.369	0.684	0.05626	12.16	0.000
D*J	2.544	1.272	0.05626	22.61	0.000

Analysis of Variance for Operational Availability

Source	DF	Seq SS	Adj SS	Adj MS	F	P
Main Effects	8	18069.5	18069.5	2258.69	1E+03	0.000
2-Way Interactions	3	1814.3	1814.3	604.76	298.52	0.000
Residual Error	628	1272.2	1272.2	2.03		
Lack of Fit	116	1210.3	1210.3	10.43	86.22	0.000
Pure Error	512	62.0	62.0	0.12		
Total	639	21156.0				

R^2 – 93.99%

REFERENCES

Blazer, D., McCormick, B., Naylor, S., Nicholson, R. A., Oates, R., Parrish, T. “Air Force Regional Stockage Policy Opportunities – Part I: Regional Stock Levels for Non-Communication Electronic Low Demand Items”, AFLMA Report LS200126302, Air Force Logistics Management Agency, Maxwell AFB, Gunter Annex AL 36114-3236

Carter, M. W. and London, R. 2002. “Why so many AWP LRus?”, *Air Force Journal of Logistics*, Winter, 26 (4) 31-33.

Condon, T., and Cunningham, T. 1999. “A Comparison of Air Force Organic Airlift and Commercial Air Express Distribution Performance”, *Air Force Journal of Logistics* 23(1) 8-12.

Graves, S.C. 1985. “A Multi-Echelon Inventory Model for a Repairable Item with One for-One Replenishment”, *Management Science* 31(10) 1247-1256.

Larvick, J.A. 2000. “Improving Support to AMC”, *Air Force Journal of Logistics* 24(4) 2-41.

Law, A.M. and Kelton, W. D.. 2000. *Simulation Modeling and Analysis*, 3rd Ed. New York: McGraw-Hill.

Masciulli, J.L., W.A. Cunningham III. 2001. "Air Force MICAP Shipping Policies: Are they optimal from a cost standpoint?", *Air Force Journal of Logistics* 25(3) 1-41.

Miller, L.W. 1992. "DRIVE (Distribution and Repair in Variable Environments): Design and Operation of the Ogden prototype", RAND Report, (R-4158-AF).

Montgomery, D.C., Runger, G.C. 1999. *Applied Statistics and Probability for Engineers*. 2nd Edition. John Wiley & Sons, Inc. New York.

Muckstadt, J.A. 1973. A model for a Multi-Item, Multi-Echelon Multi-Indenture Inventory System. *Management Science* 20(4) 472-481.

Muckstadt, J. A. 2005. *Analysis and Algorithms for Service Parts Supply Chains*, Springer Science+Business Media, Inc. New York, New York, USA.

Murphy, M.J. 1999. "Collocating Air Force Weapon Systems Inventory with the Defense Logistics Agency Premium Service Facility", *Air Force Journal of Logistics* 23(3) 29-41.

Nahmias, S., H. Rivera. 1979. "A deterministic model for a repairable item inventory system with a finite repair rate", *International Journal of Production Research* 17(3) 215-221.

Sherbrooke, C.C. 1968. "METRIC: Multi-echelon technique for recoverable item control", *Operations Research*. 16, 122-141.

Sherbrooke, C.C. 1986. "VARI-METRIC: Improved Approximations for Multi-Indentured, Multi-Echelon Availability Models", *Operations Research* 34(2) 311-319.

Slay, M.F. 1984. "VARI-METRIC: An Approach to Modeling Multi-Echelon Resupply when the Demand Process is Poisson with a Gamma Prior", *Logistics Management Institute*, LMI Working Note AF301-3. Slay, M.F.,

Bachman, T. C., Kline, R. C., O'Malley, T. J., Eichorn, F. L., King, R. M. 1996. "Optimizing Spares Support: The Aircraft Sustainability Model. Logistics Management Institute", LMI Report AF501MR1.