Proceedings of the 2004 Winter Simulation Conference
R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

SIMULATING TRANSPORTATION PRACTICES IN MULTI-INDENTURE MULTI-ECHELON (MIME) SYSTEMS

Joshua B. McGee
Manuel D. Rossetti
Scott J. Mason

Department of Industrial Engineering
University of Arkansas
Fayetteville, AR 72701, U.S.A.

ABSTRACT

Military supply chains encompass a complicated network of customers and suppliers, and deal with a wide variety of items. Demand inside the network is generated at the unit level at a specific base. The demand from the bases is aggregated to military service depots, which comprise the wholesale level in the network. The many layers of the supply chain often result in unnecessary cost and delay times, as well as low network reliability. Better integration between the multiple levels of the supply chain may be achieved through the effective utilization of transportation modes and criterion. In this paper, we present a simulation for quantifying the effect of transportation options (i.e. truckload shipping, less-than-truckload shipping, transshipments, and express air shipping) on shipping costs and operational availability.

1 INTRODUCTION

Since the end of the Cold War military budgets have been declining drastically, and the Department of Defense's logistical system has been asked to be more flexible and responsive with less money. In the past, they have met their needs by relying on massive inventories. But, the Department of Defense now seeks to implement quicker, more agile logistics systems which will reduce the inventory dollars on hand (Condon 1999). To this end, the Armed Forces 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.

This research presents a simulation model based on the Multi-Indenture Multi-Echelon (MIME) repairable inventory system used by the United States Air Force (USAF). The MIME system discussed in this paper is similar to the systems analyzed by Sherbrooke (1968 and 1986), Muckstadt (1973), Nahmias and Rivera (1979), and Graves (1985). Using simulation, this research assesses the effect of applying commercial practices to military supply chains, and then evaluates the results by using metrics currently used by the Air Force.

It is important to the accuracy of our results that our model be a close representation of the current repairable parts supply chain system. Throughout the modeling process we communicated with our contacts at the Air Force. We received process data from them as well as provided them with validation statistics and model data. This open line of communication allowed us to gain a full understanding of the system we are modeling. The model we describe in this paper is an approximation of the USAF system, and is not meant to be an exact representation of their repairable parts system.

This paper is arranged as follows. Section 2 gives an overall description of the simulation model. The shipping practices are described in section 3. Sections 4 and 5 discuss the experiment presented in this paper along with the results of that experiment. In section 5 a brief summary of conclusions is provided.

2 MODEL DESCRIPTION

We use Arena © 7.01 to develop a simulation model of the USAF's current MIME supply chain. This model is used to compare various logistic practices that potentially could be adopted by the Air Force to improve supply chain efficiencies.

2.1 Supply Chain Structure

In the model, 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 split into two regions, with three bases in each region. Figure 1 details the structure for the model.

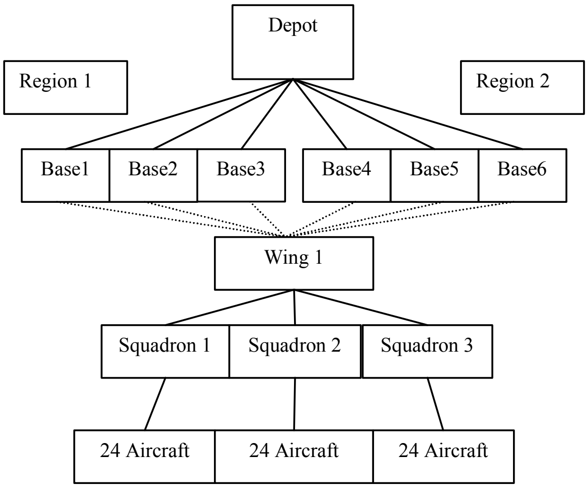

Figure 1: Supply Chain Structure

2.2 Weapon Systems and Bases

The model represents weapon systems (for the purposes of this paper, the weapon systems are aircraft) as objects with two levels of indenture. Initially, each aircraft is assigned a base number, index number, and tail number. The tail number is a model-wide unique number assigned to each aircraft. This number allows the user to compare aircraft individually across bases. The index number is unique to each aircraft at a given base. The base number indicates the base at which the aircraft is stationed. The model can accommodate a variable number of bases, and each base can have a variable number of aircraft (both values are set by the user). Table 1 displays the relationship between these three identification numbers.

Table 1: Identification of Aircraft

Tail Number	Index Number	Base Number
1	1	1
2	2	1
3	3	1
4	1	2
5	2	2
6	3	2
7	1	3
8	2	3
9	3	3

In Table 1, the user has set the model to simulate a MIME system comprised of three bases, with aircraft stationed at each base. Note that each aircraft has a unique tail number, but aircraft from different bases may be assigned the same index number. For example, there are three index number 2 aircraft, one at each base.

Each weapon system has two levels of indentures. The first level of indenture entails aircraft which are made up of multiple Line Replaceable Units (LRU). These LRUs are in turn comprised of multiple Shop Replaceable Units (SRU) constituting the second level of indenture. The number of SRUs per LRU type can vary as set by the user; however, the number of LRUs per aircraft remains constant system wide. Each of the 432 aircraft in the system are comprised of six LRUs, one of each type. LRUs of the same type are identical and interchangeable. Figure 2 illustrates the two levels of indenture in the model. In Figure 2, the subscript i denotes the LRU type, while j denotes the SRU type.

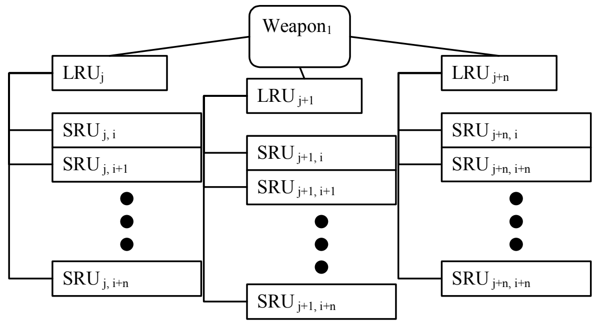

Figure 2: Hierarchy of Weapon System

In the model each of the six LRU types are comprised of four SRUs yielding a total of 24 distinct SRUs per aircraft. The SRU types for each LRU type are unique and cannot be shared between LRU types, but within the same LRU type, the component SRUs are identical. For example, the four SRUs comprising LRU type 1 of aircraft index number 1 are identical to the four SRUs comprising LRU type 1 of any other aircraft in the system.

2.3 Weapon Status

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

- 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.
- Non-Mission Capable (NMC). An aircraft is designated NMC when one or more of its critical 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.
- Phase Inspection (PI). 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. Operational Availability is defined as the percentage of time an aircraft spends in the Mission Capable state.

2.4 Failures

The failure of an SRU results in the failure of an LRU and therefore the weapon system. While operating on the aircraft, each SRU's Time to Failure (TTF) is modeled as an entry in a 2-dimensional array (Table 2).

Table 2: Time to Failure Matrix

	Base	1					
	Tail	1					
SRU Type ↓	LRU Type →	LRU 1	LRU 2	LRU 3	LRU 4	LRU 5	LRU 6
(SRU, 1)							
(SRU, 2)							
(SRU, 3)							
(SRU, 4)							

Table 2, is expanded for the first aircraft at base 1 and will be referred to as the TTF matrix. Each cell of the TTF matrix contains the TTF for the SRU corresponding to that cell. This value is generated by a distribution held in the expression array "mean-time-to-failure" (MTTF). Each cell of the MTTF expression array contains the distribution used to generate the TTF for the SRU corresponding to that cell. The model contains three levels of MTTF (in hours), each of which is modeled as an exponential distribution with some mean value: high-exponential (500), medium-exponential (400), low-exponential (300). There are eight SRUs assigned to each of these three levels.

While an aircraft is operational it accrues operating hours, and each cell corresponding to that aircraft in the TTF array (i.e. every cell representing a component SRU for that aircraft) is decremented equivalently. Aircraft failure occurs when any of the component SRU cells equals or drops below zero.

Before an aircraft can fly a sortie, a pre-flight check is performed to see if all of its component SRUs, and hence all of its component LRUs, are functional. In the construction of the model, pre-flight inspection equates to checking if all of the aircraft's cells in the TTF matrix are greater than zero. If this is not the case, the aircraft's status is set to NMC and the weapon system enters the repair process.

2.5 Sortie Assignments

Sorties are generated at the beginning of every day and are assigned to specific bases. The number of sorties assigned to each base is generated from a discrete uniform distribution over the range of 56 and 66 per day. The sorties for each base are divided into two groups (or "goes"), the first scheduled at 8:00 am and the second scheduled at 12:00 pm. Fifty-five percent of the generated sorties for each day are scheduled for the first run, while the rest are scheduled for the second run. These numbers are intended to simulate approximately 12 planes flying in the first group and 10 planes in the second group, "12 turn 10", for each squadron of 24 aircraft. At the scheduled run time, the sorties search for available aircraft. If no aircraft are available at the base when the sortie is scheduled to be flown, the sortie is delayed for a time sampled from a triangular distribution with parameters (5,10,15) minutes. Following this delay, the sortie again searches for an aircraft. If no available aircraft are found the process is repeated once more. If an available aircraft is not available on the third attempt, the sortie is cancelled; however, if an aircraft becomes available at any point, the sortie is assigned to that aircraft and the aircraft moves on to pre-flight operations. This process of searching for a sortie is a simplified representation of the complex sortie assignment process. The multiple searches for available aircraft along with the delay between searches simulates the time window given to a flight crew to initiate their assigned sortie.

2.6 Pre-Flight Operations

Once a sortie has been assigned an aircraft, the aircraft begins flight preparation. The first operations to be performed are refueling and weapons loading. The times required to complete these operations are sampled from triangular distributions with parameters (8,10,12) and (25,30,35) minutes respectively. The aircraft then moves to final preparation, which includes engine start, final systems check, and taxiing. The final preparation time is generated from a triangular distribution (7,10,12) minutes. Since the aircraft's engines are started during final preparation time, the aircraft's operating hours continue to accrue; therefore, after final preparation is completed, the total elapsed time since engine start is decremented from the TTF values associated with that aircraft.

A pre-flight check of all aircraft component SRUs is then performed. If any of the aircraft's component SRUs have failed, the failed part(s) is/are removed from the aircraft and then sent to the repair process. The aircraft is then forced to wait for spare parts. This is called a ground abort. If no failures are found, the aircraft flies its assigned sortie.

2.7 Sortie Flights

Once the aircraft has passed the pre-flight inspection, it is ready for takeoff. The time it takes for each aircraft to takeoff is generated from the triangular distribution with parameters (2,3,4) minutes; however, before the aircraft can takeoff it must first wait until a base runway is available. After seizing a runway and taking off, the aircraft undergoes an in-flight check. If any SRUs are found to be failed, the sortie is aborted, a runway is seized, the aircraft lands, and then moves to the repair process as previously described. This is termed an air abort. If the aircraft passes the in-flight check, it continues to fly the sortie.

The sortie duration is generated from a triangular distribution with parameters (.5,1.35,2) hours. After completing the sortie, the aircraft identifies a runway and lands. Landing time in minutes is generated from a triangular distribution with parameters (14,15,16). Once the aircraft has landed, it undergoes its post flight check. The duration of the sortie is decremented from all of the corresponding cells in the aircraft's TTF matrix. If any SRUs are found to be failed, the failed parts are removed from the aircraft and proceed to the repair process, and then the aircraft moves to wait for spare parts. If the aircraft passes the post-flight check, it will continue on to wait for another sortie. An aircraft will continue to fly sorties until a system fails in the air or a failure is found on the ground during maintenance checks.

2.8 Phase Inspections

The total operating hours for each aircraft is tracked. Once an aircraft accrues 280-320 operating hours, it must undergo a phase inspection. While an aircraft is in the phase inspection process its weapon system status is set to PI. A phase inspection is a complete check of the aircraft from top to bottom. In our model, when planes leave phase inspection all aircraft components are assigned a new MTTF, simulating this top to bottom schedule maintenance activity. Only two planes at each base can be in phase inspection simultaneously. Aircraft that have accumulated operating hours that are within this range check the phase inspection process for their base each time after passing the post-flight inspection. If there are already two aircraft in process at phase inspection for that base, the aircraft will continue to fly sorties. Once an aircraft exceeds 320 flight hours, it cannot fly another sortie until completing phase inspection. The phase inspection delay is generated from a triangular distribution with parameters (7,10,11) days.

2.9 The Replacement Process

The model checks each SRU associated with the aircraft sequentially in each of the three system checks. The first time a cell in an aircraft's TTF matrix is less than or equal to zero, the aircraft is considered in failure. When a failed SRU is detected, the aircraft is marked as being failed, and the model removes the SRU from the aircraft. This SRU is sent to the repair process, which is described later. The model then continues to check for other SRU failures on that aircraft. Once a failed aircraft completes the systems check, the model performs an inventory check for the failed parts associated with that aircraft. Inventory levels at the bases and at the depot are modeled using two separate matrices, similar to those shown in Table 3 and Table 4. These matrices will be referred to as the base inventory matrix and the depot inventory matrix, respectively.

The number in each cell of Table 3 and Table 4 represents the number of spare parts of a given SRU type that are available at the corresponding location. The initial value of the cells in these matrices is defined by the user. When a failed aircraft initiates an inventory check, the model begins by checking if a spare SRU of the same type is available at the aircraft's base (determined by the aircraft's base number). If a spare SRU of the same type is not available at the aircraft's base, an order is placed to the depot. This order is assigned a backorder status. If a spare part is not available at the depot, the order is held in a queue at the depot with priority given to backorders waiting on a part of that SRU type to be repaired. The aircraft waits in a First-In-First-Out (FIFO) queue for the next available spare of correct SRU type to arrive at the corresponding base. When the order is filled at the depot, the part is shipped to the base. Once the part arrives at the base, the inventory level for that SRU type at the base is incremented and a signal is sent to the aircraft queue. This signal indicates that new parts have arrived, initiating an inventory check. When an aircraft finds a needed spare in inventory, installation of that part begins. If there are multiple failures on a given aircraft, the aircraft will wait in queue until all corresponding SRUs are available, but each of the spare parts are installed as they arrive.

Once a needed spare part is available at the base, the installation process begins. At the beginning of the installation process, the aircraft must wait for the part to be issued from supply. This simulates the delay between the time the part arrives at the base and the time the part is ready to be installed on the aircraft, and is generated from a triangular distribution with parameters (35,130,165) minutes. After the part is issued from supply it is ready to be installed on the aircraft. Installation times are generated from a triangular distribution with parameters (60,84,120) minutes. Upon completion of the installation process, the object representing the spare part is disposed, and the corresponding cell in the TTF matrix is re-initialized to a

Table 3: Inventory at the Base Level

	Base	1			2		
SRU Type ↓	**LRU Type →**	LRU 1	LRU 2	LRU 3	LRU 1	LRU 2	LRU 3
(SRU 1)		3	3	3	3	3	3
(SRU 2)		3	3	3	3	3	3
(SRU3)		3	3	3	3	3	3
(SRU 4)		3	3	3	3	3	3

Table 4: Inventory at the Depot Level

		Depot		
SRU Type ↓	**LRU Type →**	LRU 1	LRU 2	LRU 3
(SRU 1)		5	5	5
(SRU 2)		5	5	5
(SRU3)		5	5	5
(SRU 4)		5	5	5

number generated by the MTTF expression array. Once all failed parts have been replaced on the aircraft, the aircraft is ready to fly sorties.

2.10 Repair Process

When a SRU is deemed defective, the model creates an entity representing the defective SRU. It is highly unlikely that this failed SRU can be repaired at the base (Miller 1992). In the model, the probability that a part can be repaired at the base is set to 0.01. 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 resource. If the SRU must be repaired at the depot, the SRU is delayed some shipping time, and then enters the queue for the depot repair resource. Table 5 details the matrix used to generate the shipping times between different locations in the model. The three different distributions used in the shipping time matrix are outlined in Table 6

Table 5: Shipping Time Distributions

Num	Distribution (Hours)	Description
1	TRIA(12,81.6,184.8)	This distribution is used to generate shipping times between bases in the same region.
2	TRIA(31.2,170.4,348)	This distribution is used to generate shipping times between the depot and the bases in region 1.
3	TRIA(76.8, 266.4, 453.6)	This distribution is used to generate shipping times between the depot and the bases in region 2. This distribution is also used to generate the shipping time between bases which are not in the same region.

Table 6: Shipping Time Matrix

	Base1	Base2	Base3	Base4	Base5	Base6
Base1	0	1	1	3	3	3
Base2	1	0	1	3	3	3
Base3	1	1	0	3	3	3
Base4	3	3	3	0	1	1
Base5	3	3	3	1	0	1
Base6	3	3	3	1	1	0
Depot	2	2	2	3	3	3

The repair stations at all bases and the depot give priority to backorders for repair jobs. Repair times at each base and the depot are random distributions set by the user. In the model, the repair times are generated from an exponential distribution with a mean of eight hours.

If 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 as mentioned earlier. This practice holds with a one for one inventory policy. In other words, for every part that is sent to the depot an order is generated for a part to be sent back to the base, a one-for-one replenishment policy. Again in this queue, back-orders are given priority.

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 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. When the depot inventory is incremented, a signal is sent to the queue holding unfilled orders. When this signal is received, all orders are checked. The first order in queue, of the same type as the repaired SRU, is filled. After an order is filled, the order is shipped back to the base where the order originated. Once the shipment as been received it is entered into the base's inventory. When a bases inventory is incremented, a signal is sent to the queue holding NMC aircraft. Each of the aircraft are checked, and the first aircraft in queue needing a part of the same type which was entered into the base's inventory moves to the installation process. If there are no aircraft in need of the SRU, the SRU remains in the base's inventory.

3 SHIPPING

Bases only ship out failed SRUs and receive only functional SRUs. Conversely, the depot only receives failed SRUs and only ships out functional SRUs. Spare and failed parts can be shipped between echelons in two ways: ground shipping and express air shipping. Most parts are shipped on trucks that pick up and drop off parts at the bases and depot; however, MICAP parts are air shipped, usually arriving at the final destination in one or two days.

3.1 LTL/TL Shipments

In the model, we can simulate the use of both Less-than-Truckload (LTL) and the Truckload (TL) commercial carriers. These features are controlled through two variables, truck capacity and minimum batch size. The truck capacity dictates the number of SRUs that each truck can hold. Minimum batch size is a percentage, which is multiplied by the truck capacity. The resulting number 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 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 waiting to be shipped at each location is made each day at 8:00 a.m. This is true for both the LTL and TL case.

3.2 Direct Shipments (MICAP)

When parts receive a backorder status they are shipped with the MICAP designation. MICAP shipments are express air shipments. A part can be designated MICAP when the base needs to ship the failed part to the depot or when the depot needs to ship a spare part to a base. The base designates a failed part as MICAP when the base does not have a spare part in its inventory to replace the failed part. 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. Parts which receive the MICAP designation wait in a separate queue for air shipping. 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 shipments are express air shipments with shipping times generated from a triangular distribution with parameters (22,24,26) hours. The model assumes that the air shippers have unlimited capacity. This allows the model to rely on MICAP if the regular shipping is not able to keep up with the shipping volume, just as the Air Force uses MICAP to expedite shipping.

3.3 Lateral Transshipments

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. This is a common commercial practice used to expedite shipping. For this scenario the bases are split into regions based on geographical location. 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. The first thing to be done when checking the bases within the same region is to create a list of bases that have inventory available. This list is stored in an array. A selection is made from this array based on a user-defined criterion. Currently, this criterion is set to choose the base with the most inventory on-hand for that particular part. Once 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 on 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.

The above sections give a complete description of the model. From this model we develope a set of experiments to explore the effect using commercial shipping practices along with other factors have on the Air Force supply chain.

4 EXPERIMENTAL DESIGN

A factorial experimental design is used in our experimental studies. A full factorial design allows for design points to be investigated for all possible factor combinations. The experiments identify the main effects and the interactions between the factors. In this research there are 11 factors and each of the factors have two levels. Because of the large number of factors, a fractional factorial design was chosen. In a fractional factorial design, a reduced number of runs can be used to analyze the main affects and interactions between the factors, albeit with less granularity. A 1/16 fractional design is used, allowing for 128 runs of the experiment to be made, rather than $2^{11} = 2048$ runs needed for a full factorial design. The experiment is 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. For the purposes of this paper we will focus on 3 of the 11 factors: Shipping Option, MICAP, and Transshipment. Table 7 lists these three factors with a brief description.

Table 7: Factors and Descriptions

Factors	Description
Shipping Option	This determines whether truck load or Less Than Truck Load shipping will be used.
MICAP	This determines whether express air deliveries will be used to expedite backorders.
Transshipment	This factor indicates whether or not transshipments at the base level can be used as a source of supply.

Table 8 outlines the factor values used in experimentation. Each factor has two factor levels. The values listed in Table 8 will be important later in understanding the results of our experiments.

Table 8: Factors Values

Factor	Low	High
Shipping Option	LTL	TL
MICAP	On	Off
Transshipment	On	Off

Each simulation is set up to have a warm-up period, a run length, and a specified number of replications or runs. We use a warm up period of six months, and then collect data for one year. Our simulation is 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 data that is collected for each instance is written to an Excel worksheet after the run has completed. The simulation is warmed up at the beginning of each instance, and the system is cleared after each of the runs; therefore, the simulation model collects data for 128 independent design points. Each of these 128 design points is replicated five times using a different stream of random numbers, yielding a total of 640 independent observations.

5 DATA AND DATA ANALYSIS

This section will outline the data from the experiment along with some analysis of that data. For this experiment, eight different responses are set up to measure the effect the factors had on the model. In this paper we will discuss two of these responses. Table 9 lists these two responses along with a brief description.

Table 9: Responses 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.
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 of each factor. The Total Transportation Cost is the sum of these three factor costs.

Summary statistics are calculated for each of the responses to provide insight as to how the data behaves across all scenarios. Table 10 lists the two responses being focused on in this paper followed by the summary statistics for the data collected on each response. The summary statistics included in Table 10 are: $\bar{x}$ - Sample Mean, s-Standard Deviation, and s.e. - Standard Error.

Table 10: Response Summary Statistics

Response	$\bar{x}$	s	s.e.
Operational Availability	75.28	5.75	0.2270
Total Transportation Cost	72,244	44,770	1,770

The relationship between some of the responses is also 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 are averaged providing an estimate of the response for each design point. This will reduce the noise in the graph and 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 other 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 roll 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.

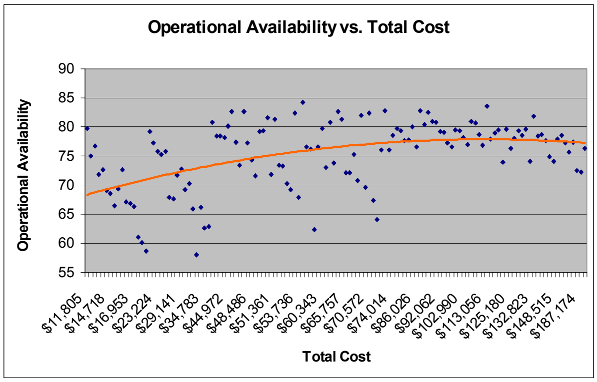

Figure 3: Operational Availability vs. Total Transportation Cost

6 CONCLUSIONS

The MICAP factor is one of the most influential factors in our experimentation as determined by regression analysis. This factor simulated the use of express carriers to expedite shipping times. Over the past decade the logistical defense related budgets have been reduced. This in turn has had an effect on the way the military supply chain operates. Inventory levels in the supply chain have been falling along with the budgets. The pressure to reduce both inventory and spending has induced a lot of 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 very high. The cost of MICAP shipments was the largest cost component in our simulation model.

Figure 3 illustrates the diminishing returns relationship between transportation cost and operational availability. The cost of MICAP shipments is a large driver in the shape of this curve. If the MICAP cost component were removed this curve would take on more of a linear shape. The diminishing returns relationship would not disappear, but it would be reduced. Reducing the role MICAP plays in the Air Force supply chain will both reduce cost greatly and force new opportunities for improvement to be explored.

The two other shipping factors that are investigated in our experiments are 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 is seen when using the TL shipping option. In fact the lowest costs are realized in scenarios using TL shipping. These cost differences, however, are overshadowed by the cost of MICAP shipping. In the same light, the Transshipment factor does not have a large effect in our experiments, 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.

There are many opportunities for expansion of the simulation model presented in 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. Future work has already been funded and is in the beginning stages for expanding the model presented in this paper to explore the Sortie Generation process. The goal of this new project is to extend the current simulation and mathematical modeling methodologies to assist unit-level logistics managers in analyzing the effects of different resource allocation policies and identify risks in logistical plans. The model will encompass sortie generation, maintenance activities, and the effect of limited equipment and inventory.

ACKNOWLEDGMENTS

This material is based upon work supported by the United States Air Force under Contract No. F33615-99-D-6001. 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. This work has been supported through the efforts of: Mohsen Manesh, Ashlea Bennett, Ryan Houx, Amanda Mills B.S.I.E. University of Arkansas.

REFERENCES

Condon, T., et al.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.

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

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

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.

AUTHOR BIOGRAPHIES

JOSHUA B. MCGEE, B.S.I.E. is a research assistant and graduate student in the Industrial Engineering Department at the University of Arkansas. He has held several positions with local manufacturers working on projects ranging from automated system design to lean manufacturing activities. His research interests include design and analysis of manufacturing and transportation systems through computer simulation and mathematical models. He is a member of IIE and Alpha Pi Mu. His email address is `<jbm04@uark.edu>`.

MANUEL D. ROSSETTI, Ph. D., P. E. is an Associate Professor in the Industrial Engineering Department at the University of Arkansas. He received his Ph.D. in Industrial and Systems Engineering from The Ohio State University. Dr. Rossetti has published over thirty journal and conference articles in the areas of transportation, manufacturing, health care and simulation and he has obtained over $1.5 million dollars in extra-mural research funding. His research interests include the design, analysis, and optimization of manufacturing, health care, and transportation systems using stochastic modeling, computer simulation, and artificial intelligence techniques. He serves as an Associate Editor for the International Journal of Modeling and Simulation and is active in IIE, INFORMS, and ASEE. He is serving as co-editor for the WSC 2004 conference. His email address is `<rossetti@uark.edu>` and his web address is `<www.uark.edu/~rossetti>`.

SCOTT J. MASON, Ph. D., P. E. is an Assistant Professor in the Department of Industrial Engineering at the University of Arkansas and a registered Professional Engineer in the state of Arkansas. Prior to his current position, he spent eight years working on factory modeling, simulation, and capacity analysis projects at SEMATECH, Advanced Micro Devices, Intel, Wright Williams & Kelly, and National Semiconductor. Dr. Mason received his B.S.M.E. and M.S.E. from The University of Texas at Austin, and his Ph.D. from Arizona State University. His research interests include production systems scheduling, modeling, and analysis and manufacturing and transportation applications of operations research. He is a member of INFORMS and a Senior Member of IIE.