Dynamic Planning Under Uncertainty for Theater Airlift Operations

by

Kiel M. Martin

B.S. Operations Research
United States Air Force Academy, 2005

SUBMITTED TO THE SLOAN SCHOOL OF MANAGEMENT IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE IN OPERATIONS RESEARCH
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

Signature of Author: ______________________________
Sloan School of Management
Interdepartmental Program in Operations Research
May 17th, 2007

Approved by: ______________________________
Stephan E. Kolitz
The Charles Stark Draper Laboratory, Inc.
Technical Supervisor

Certified by: ______________________________
Cynthia Barnhart
Professor, Civil and Environmental Engineering
Co-Director, Operations Research Center
Thesis Advisor

Certified by: ______________________________
Dimitris Bertsimas
Boeing Professor of Operations Research
Co-Director, Operations Research Center

[This Page Intentionally Left Blank]

Dynamic Planning Under Uncertainty for Theater Airlift Operations

by

Kiel M. Martin

Submitted to the Sloan School of Management on May 17th, 2007 in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Operations Research

Abstract

In this thesis, we analyze intratheater airlift operations, and propose methods to improve the planning process. The United States Air Mobility Command is responsible for the air component of the world wide U.S. military logistics network. Due to the current conflict in Iraq, a small cell within Air Mobility Command, known as Theater Direct Delivery, is responsible for supporting ongoing operations by assisting with intratheater airlift.

We develop a mathematical programming approach to schedule airlift missions that pick up and deliver prioritized cargo within time windows. In our approach, we employ *composite variables* to represent entire missions and associated decisions, with each decision variable including information pertaining to the mission routing and scheduling, and assigned aircraft and cargo. We compare our optimization-based approach to one using a greedy heuristic that is representative of the current planning process. Using measures of efficiency and effectiveness, we evaluate and compare the performance of these different approaches. Finally, we adjust selected parameters of our model and measure the resulting changes in operating performance of our solutions, and the required computational effort to generate the solutions.

Technical Advisor: Andrew Armacost
Lt Col, United States Air Force
Associate Professor and Director of Operations Research

Technical Supervisor: Stephan Kolitz
Title: Principal Member of the Technical Staff
The Charles Stark Draper Laboratory, Inc.

Thesis Advisor: Cynthia Barnhart
Title: Professor, Civil and Environmental Engineering

[This Page Intentionally Left Blank]

ACKNOWLEDGMENTS

There are many who have given me support, encouragement, and guidance over these last two years.

First, I would like to thank everyone both at the Air Force Academy and here at MIT that had a part in giving me the opportunity to pursue my degree at MIT.

Stephan Kolitz, my advisor at Draper Laboratory, thank you for your continual advice and guidance. Cindy Barnhart, thanks for taking the time and having the patience to meet with me regularly, even with an ever hectic schedule. Lt Col Andy Armacost, thanks for your help to make this thesis a reality. Your devotion to students, current and past, is extraordinary. I appreciate your time and efforts helping me narrow down on a thesis topic during our trip to AMC and your comments during the editing process. Francis Carr, I never would have been able to learn Python on my own. I appreciate the many hours you spent helping me learn and working with me to develop my program.

I also have to thank John, Joe, Kevin, Stefan, Keith, and Chris. Having friends to share the laughs, frustrations, headaches, and homework made these two years much more enjoyable. To the current and future Draper Fellows, work hard, play hard, and make sure to get started early. To my family, thanks for your encouragement and help during the never ending editing process.

Finally to Leilea, my finance, thank you for your sustaining love and support over this trying time. You alone know my fears, struggles, desires, and dreams.

Kiel M. Martin, 2nd Lt, USAF
17 May 2007

[This Page Intentionally Left Blank]

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1 Introduction

"My logisticians are a humorless lot. They know if my campaign fails, they are the first ones I will slay." -*Alexander the Great*

Logistics has always been and will continue to be an essential part of military operations. The United States military logistics network is a complex system comprised of several major commands and spanning countries across the globe. From large-scale strategic initiatives, to models of individual processes, analysis has been performed on almost every aspect of this network, trying to improve the overall system.

1.1 Research Scope – Theater Direct Delivery

This thesis takes a look at one specific piece of the United States military logistic network, known as Transportation Command (TRANSCOM), by trying to improve intratheater airlift. TRANSCOM is broken into air, land, and sea components. The air component, Air Mobility Command (AMC), has the responsibility to "deliver maximum war-fighting and humanitarian effects for America through rapid and precise global air mobility" [13]. AMC is responsible for airlift operations that span the globe. One key delineation in airlift is intertheater and intratheater airlift.

Intertheater airlift refers to airlift linking theaters of operations, and is under the purview of AMC. The US definition of theater of operations is a term that refers to the administrative command responsible for a given area. Intratheater airlift refers to airlift conducted within a theater and is usually conducted by a geographic combatant commander using assigned assets, instead of AMC. In the current conflict in Iraq, these lines have been blurred as AMC aircraft and personnel have been tasked to support Central Command (CENTCOM) with intratheater airlift. This thesis looks at how a unit within AMC, Theater Direct Delivery (TDD), supports current operations in the Middle East by assisting with intratheater airlift.

1.2 Overview of Thesis

In this thesis we provide an overview of the current Theater Direct Delivery planning and execution process; we present an optimization-based model to improve the theater airlift planning process; and we conduct an analysis of our results. The remainder of this thesis is organized as follows:

Chapter 2: Theater Direct Delivery and Current Mission Planning Process

The purpose of this chapter is to introduce theater airlift in the context of Theater Direct Delivery. We summarize the organizational structure and hierarchy of organizations responsible for intratheater airlift. We then describe the planning and execution process for TDD missions supporting intratheater airlift within CENTCOM. Furthermore, we highlight the shortcomings of the current process, concluding with important metrics and motivation for further study.

Chapter 3: Functional Analysis & Modeling Scope

In this chapter we present a functional description of the TDD mission planning process and introduce our model. We discuss the scope of our model, including assumptions and data generation. Last, we classify our problem and present our modeling approach in terms of our variable definitions.

Chapter 4: Modeling

In this chapter we discuss our model in detail. We break the model down into functions and explain each in detail. We introduce our mathematical formulation to optimize scheduled missions and briefly describe the greedy heuristic we insert into our model. We briefly describe the greedy algorithm we use later to compare to our optimized approach. We also discuss possible improvements to our model as well as alternate approaches.

Chapter 5: Results and Analysis

In this chapter we discuss our results, comparing our model to a greedy heuristic, representative of the current mission planning process. We describe the metrics, objective functions, and parameters used in our modeling. We also analyze different parameter settings and their effects on both operational metrics and computation time.

Chapter 6: Summary and Future Work

This chapter summarizes the work presented in this thesis and discusses areas of future research.

[This Page Intentionally Left Blank]

Chapter 2 Theater Direct Delivery and the Current Mission Planning Process

Chapter 1 provided an introduction to the overall problem and an overview of how intratheater airlift fits into military logistics system worldwide. This chapter has three objectives:

1. Summarize the organizational hierarchy for those organizations involved in intratheater airlift;
2. Describe the process for planning and executing the missions of the C-17s that support intratheater airlift operations; and
3. Discuss important operational metrics and motivation for further analysis and modeling of the system.

In any major conflict, logistics pipelines must surge material to meet the needs of those fighting on the front lines. The current (as of 2007) conflict in Iraq highlights the complexity and difficulty of supplying hundreds of thousands of troops over half a world away. Theater commanders rely on the concept of *reachback* to sustain operations. *Reachback* is defined as the "process of obtaining products, services, and applications, or forces, equipment, or materiel from Air Force Organizations that are not forward deployed" [1]. Theater Direct Delivery (TDD) is an example of Central Command (CENTCOM) using the *reachback* concept to provide additional airlift capacity. In this thesis, we will examine the work of TDD as one small element of this logistics pipeline, intratheater airlift.

2.1 Organizational Hierarchy

The job of the Theater Direct Delivery (TDD) cell within the Tanker Airlift Command Center (TACC) under Air Mobility Command (AMC) at Scott AFB, MO is to "provide responsive airlift to support the war fighter's intratheater airlift requirements" [1]. TDD supports airlift requirements for CENTCOM by moving intratheater cargo (PAL) and passengers (PAX). We define intratheater airlift requirements as any movement of cargo by aircraft, such that the origin and destination are within a single area of operations.

TDD was initially set up soon after the recent conflict in Iraq to handle oversized and outsized cargo within CENTCOM using three aircraft. Oversized and outsized cargo is defined as cargo that only fits on large body aircraft such as the C-17 or C-5. Later, the unit was expanded to contain six aircraft in the hopes of reducing the number of ground convoys. As the conflict in Iraq progressed, the existence of TDD was a factor in the decision to send reserve units of C-130s assigned to CENTCOM back home to the states as the larger C-17s moved in to pickup the slack. Currently, TDD contains twelve C-17s assigned through a 'verbal' agreement with AMC to support CENTCOM requirements. However, TDD has a unique relationship with CENTCOM in that the TACC retains both operational and tactical control of TDD in supporting CENTCOM's needs.

2.1.1 Structure

The AMC organizational structure can be broken down into the force providers of the 18th Air Force and the major departments that support operations. The 18th AF serves as the war fighting component of AMC. All AMC wings and groups, as well as both Expeditionary

Mobility Task Forces (EMTF), report to 18th AF. This includes more than 54,000 personnel, both military and civilian. The TACC also falls under 18th AF and "serves as the organization's air operations hub, planning and directing tanker and transport aircraft operations around the world" [1]. The major departments are responsible for strategic planning, training, and organization. These departments were recently reorganized into the standard Air Force staff structure. This change added the A9 function: Analysis, Assessments and Lessons Learned [7] as a full directorate and created a staff consistent with the rest of the military structure.

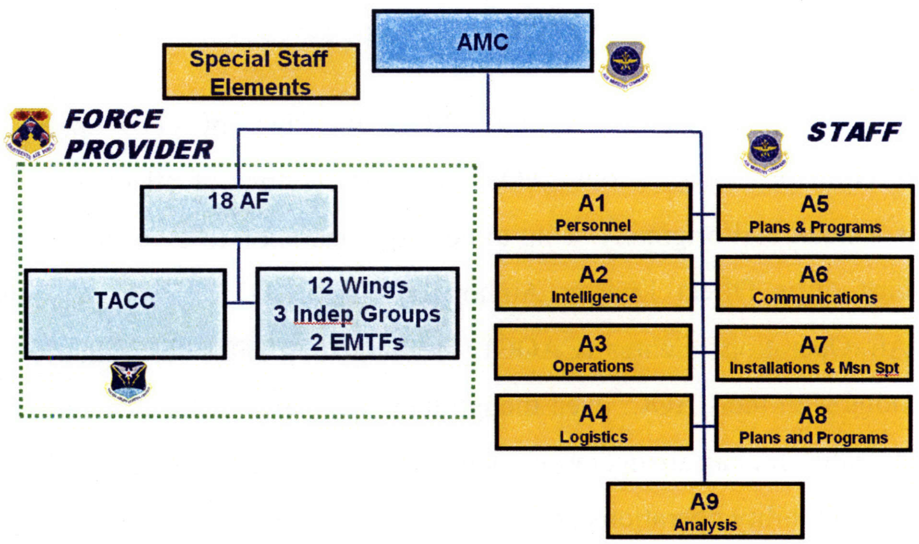

Figure 2-1: AMC Organizational Structure [1]

Within the TACC, several directorates exist to plan, manage, and execute ongoing AMC operations. For the purposes of this paper, it is relevant to know that TDD is officially referred to as XOCR under the Command and Control Directorate (XOC). More in-depth information concerning the overall organizational structure of the TACC can be found in Air Force Doctrine Document 2-6 on Air Mobility Operations [1]. As shown in *Figure 2-3,* it is important to note that both planning and execution for TDD assigned aircraft fall within the TDD cell. TDD maintains necessary coordination with the Air Mobility Division (AMD) within the Air Operations Center in theater.

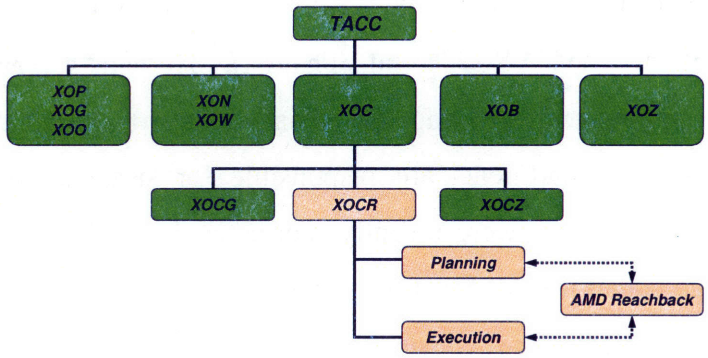

Figure 2-2: TACC Organizational Chart [1]

While TDD is organizationally under XOC, the chief of TDD operationally reports directly to the Director of Operations (XOZ), who has control of all operations currently in execution. As a small cell originally setup to be surge operation to assist airlift requirements in the Middle East, it is Direct Liaison Authorized (DIRLAUTH) with personnel under the CENTCOM Air Operations Center (CAOC), meaning that TDD has authority to consult or coordinate actions with the CAOC, even though they operate as difference commands [4]. The CAOC is responsible for all air operations within CENTCOM. The Air Mobility Division is responsible for all mobility operations and has an Airlift Control Team (ALCT) that plans, tasks, and coordinates intratheater airlift operations.

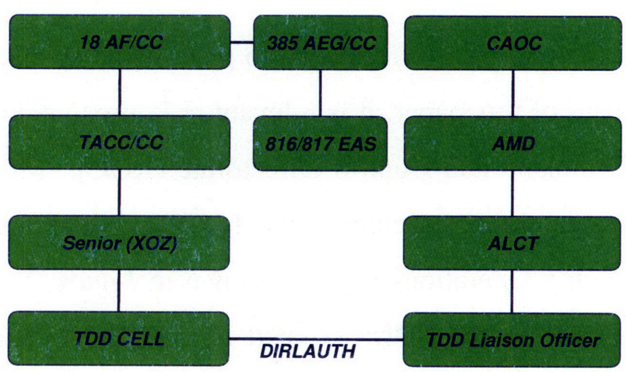

Figure 2-3: 18th Air Force & CAOC Relationship [1]

2.1.2 Assets

As mentioned previously, TDD directly controls twelve C-17s positioned in the theater to carry on its mission: eleven aircraft at Al Udeid, Qatar and one aircraft at Manas, Kyrgyzstan.

Additionally, one aircraft in Turkey flies missions supporting both TDD and channel routes. This aircraft can fly into Turkey on a channel mission and be temporarily used to fly TDD legs into Iraq before returning to Turkey, then resume flying its channel route.

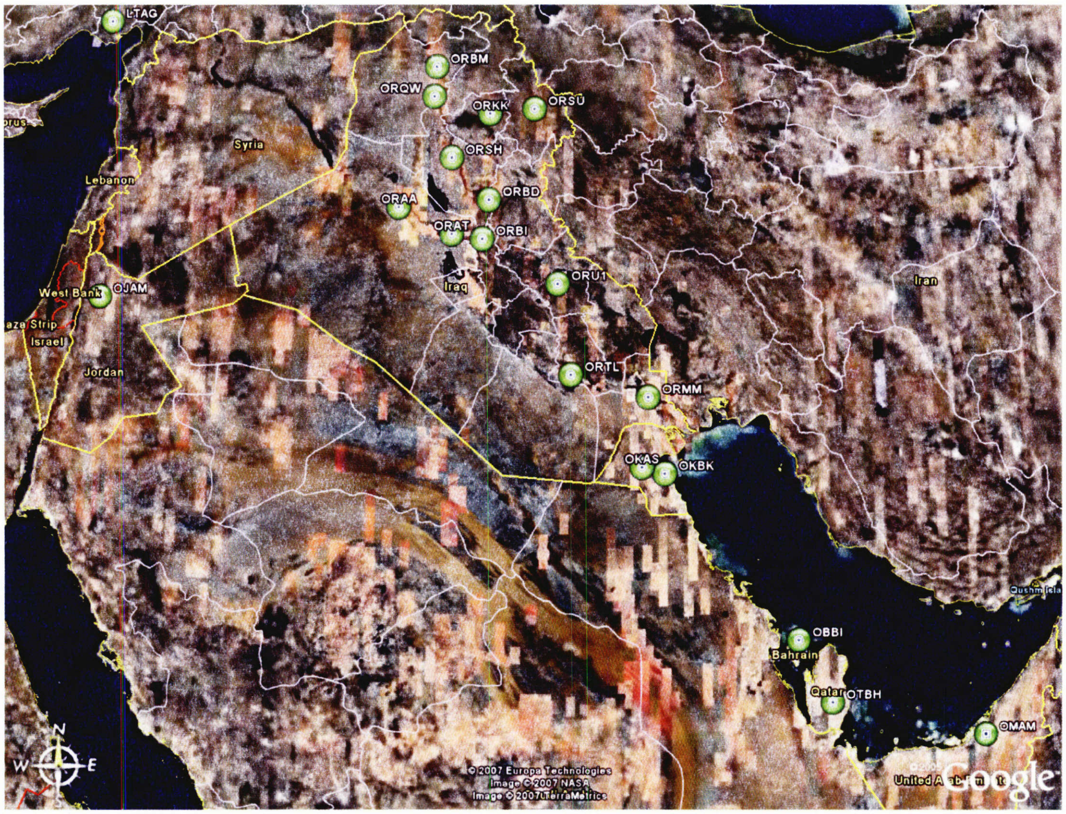

Figure 2-4: Relevant Aerial Ports within CENTCOM served by TDD [10]

Figure 2-4 illustrates a snapshot of the Middle East highlighted with the major aerial ports served by TDD in green. Other ports exist that are not shown on the map. TDD flies missions into Afghanistan and parts of Africa. Most TDD missions start and end at Al Udeid (OTBH). Aircraft usually fly a positioning leg to either Kuwait or a port in Iraq and then fly a few cargo legs before returning to Al Udeid.

The C-17 is the newest cargo aircraft in the mobility fleet. It can be used for strategic airlift, carrying large payloads between continents or for theater airlift, landing at austere airfields close to the fight. One of its most important capabilities is the ability to deliver all types of cargo directly to forward bases in a deployment area. The C-17 also has a high reliability rate of over 80% [3].

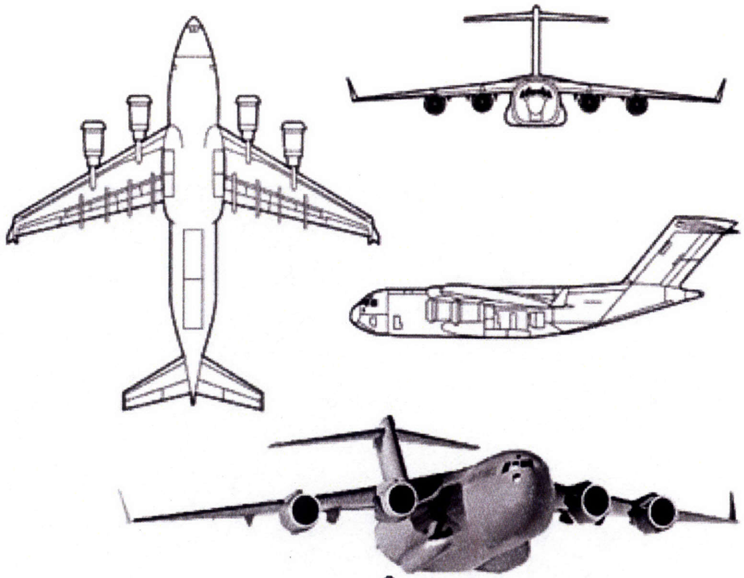

Figure 2-5: C-17 Globemaster III [3]

The C-17 has several cargo configurations and can also be used for several other purposes such as carrying airborne troops, aero-medical evacuation, and airdrops. It can carry up to 189 troops/paratroops, 48 litter and ambulatory patients and attendants, or 170,900 pounds of palletized cargo in up to 18 pallet positions. In terms of cargo, the C-17 can reach capacity either in weight or volume. Most palletized loads carried for intratheater airlift do not max out the weight of the aircraft, but rather max the volume by using all 18 pallet positions.

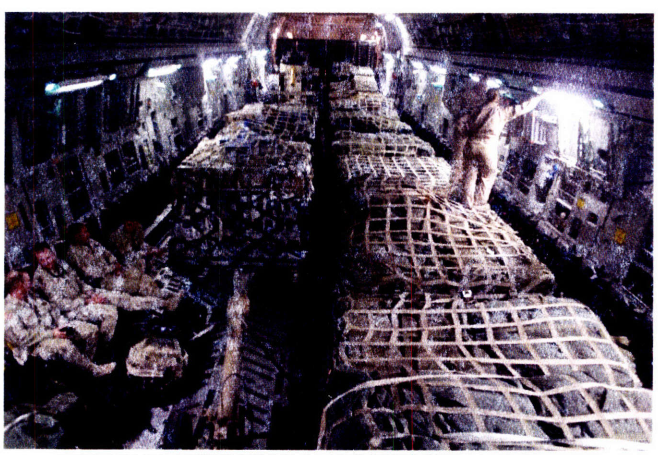

Figure 2-6: C-17 Double Row Pallet Configuration

Figure 2-7: C-17 Passenger Configuration with Middle Row of Seats

TDD also has access to other assets, both military and commercial aircraft. Additional military (organic) aircraft scheduled by other TACC directorates may become available after completing their given missions. The use of other assets is referred to as an *in-system select,* in which an asset can be temporarily borrowed by TDD. The use of commercial aircraft, such as the IL-76 and AN-124, are also scheduled through TDD. The IL-76 is often employed for daily missions. The AN-124 is only used when necessary to handle excessive cargo loads in the system. These commercial contracts require additional coordination within AMC to request and accept bids from the commercial carriers. However, the carriers submit their schedule and the TDD planners are only responsible for inputting these missions into the overall schedule prior to execution [10].

2.2 Current Planning & Execution Process

The planning process for TDD begins in the CENTCOM Deployment Distribution Operations Center (CDDOC). The DDOC concept is a fairly new concept within the US military and is an effort to better organize and coordinate the logistics efforts within a command. The DDOC mission is to "support the geographic combatant commander's operational objectives by synchronizing and optimizing the intertheater and intratheater distribution aspects…of material, and other forms of sustainment in support of the geographic combatant commander missions"

[4]. The DDOC concept was initiated in 2003 to meet the challenges of distribution within CENTCOM and then integrated within all other commands.

The TDD planning process is still new and has yet to be codified in doctrine. The first step within the CDDOC is to determine the PAX/PAL requirements for airlift in theater using AMC assets, meaning large bodied aircraft such as the C-17 or C-5. Personnel input the requirements into the Intratheater Airlift Request system (ITARS) and must validate these requirements before giving them to the AMD [10].

2.2.1 Air Mobility Division

Within each major command there exists an air operations center. The Air Mobility Division is responsible for planning, coordinating, tasking, and executing air mobility missions for the air operations center. Within the AMD, the Director of Mobility, Chief of the AMD, and ALCT determine what AMC assets are necessary to facilitate the movement of intratheater cargo. Personnel within the AMD work with the CDDOC, USTRANSCOM, and the TACC as necessary to coordinate efforts. Excessive backlog in the system or the nature of the cargo can require additional assistance from AMC, in the form of large body aircraft, such as the C-17, to carry oversized or outsized cargo.

To assign cargo on TDD aircraft, AMD creates a spreadsheet document called the *cargo tracker* that identifies the loads that are to be moved on AMC assets. Each requirement entered into the *cargo tracker* has several attributes:

- Priority level from 1 to 10;
- Estimated number of pallets (PAL);
- Estimated tonnage;
- Number of passengers (PAX);
- Aerial port of embarkation (APOE);
- Aerial port of debarkation (APOD);
- Available Load Day (ALD); and
- Required Delivery Day (RDD).

It is important to note that if a requirement has already been scheduled, it also contains the day scheduled and the days left until the RDD. These attributes help the planners sort the requirements when figuring out the next requirement(s) to plan. The cargo tracker is updated daily and the updates can include new entries or changes to existing requirements, such as a change in the size or delivery time window.

AMD is also responsible for de-conflicting CENTCOM C-130 movements so that multiple aircraft are not scheduled to pick up the same load causing wasted capacity. Any lack of coordination causes frustration both in CENTCOM and in the TACC when an aircraft arrives at a port to pick up a load that has already been moved or is not there. Many times cargo will be loaded on an opportune aircraft that happens to have extra capacity and is going to the desired port. The problem arises when the cargo might not be visible in the Global Air Transportation Execution System (GATES) as moved, which will be explained later in this chapter. When this occurs, an aircraft can still be scheduled to pick it up from its previous location, greatly reducing efficiency and effectiveness.

2.2.2 Theater Direct Delivery: TACC/XOCR

After requirements from AMD are received by TDD in the form of the *cargo tracker*, actual mission planning begins. TDD uses several web-based tools to assist them in the mission planning process: Global Decision Support System II (GDSS2), Consolidated Air Mobility Planning System (CAMPS), maximum on ground (MOG) tools, and the AMC Policy Matrix. TDD relies on a graphical representation, the *TOMCAT,* of all current and future planned missions as a tool to see mission details in an organized format.

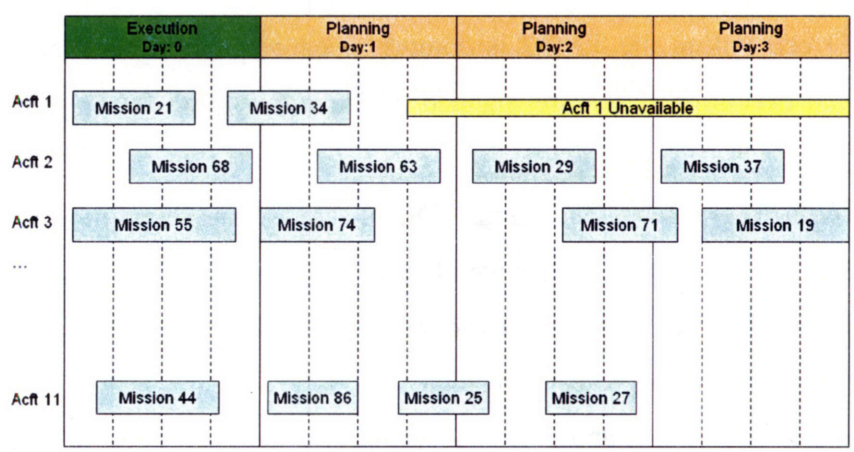

Figure 2-8: Representation of the TOMCAT

As shown in the previous figure, the TOMCAT is broken down by day and shows missions as blocks of time in which certain aircraft are unavailable. Each mission also contains information describing the cargo, start times, permission to land (PPR), diplomatic clearances (DIP), and any other necessary details, illustrated in *Figure 2-9.* The TOMCAT provides an overview of all missions on the schedule. This helps TDD to organize currently scheduled missions and help make decisions to add or adjust missions when necessary.

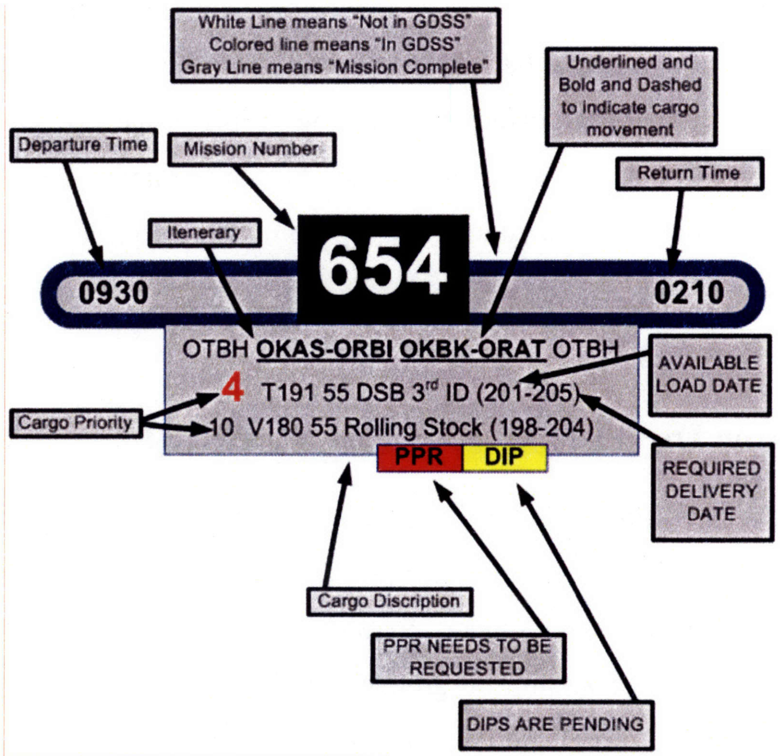

Figure 2-9: Example of Mission Information in TOMCAT

The TDD cell within the TACC is compromised of planning and execution functions. Both planning and execution desks are manned 24 hrs a day, with fewer personnel during the night shift.

As shown in *Figure 2-10,* the current TDD mission (msn) scheduling process consists of planning and execution, with input from several external sources. External inputs include:

1. Regularly Scheduled missions to deliver backlog cargo;
2. List of priority shipments to be scheduled;
3. Theater Data/Constraints; and
4. Higher Priority Conflicts.

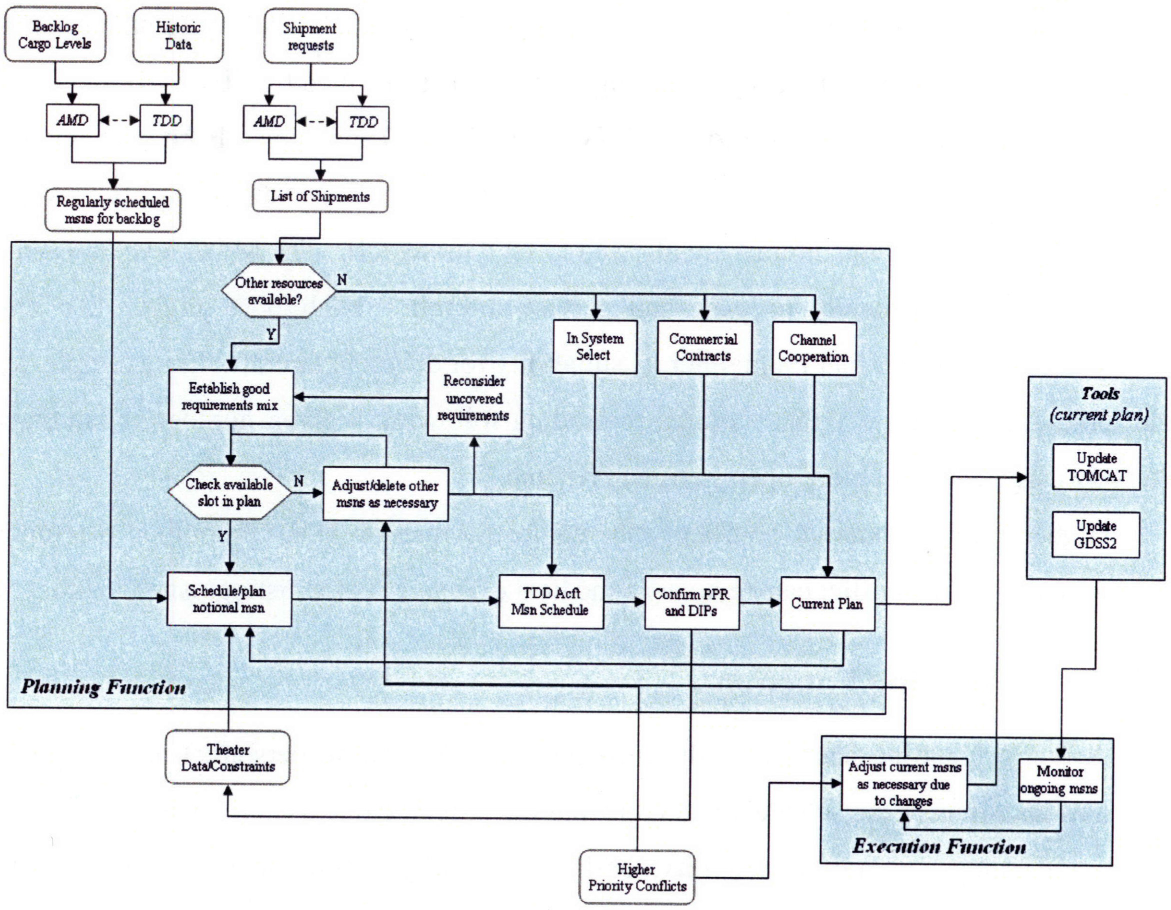

Figure 2-10: Current TDD Mission Scheduling Process

2.2.2.1 External Inputs

Backlog cargo levels are constantly monitored and can be accessed through an online database system called GATES. GATES shows the cargo levels by origin/destination for each port within the theater. However, GATES is not always accurate due to users in theater not updating the system. For example, a shipment may sit at a port for multiple days and only be input into the system shortly before it is scheduled to be loaded onto an arriving cargo aircraft. Backlog levels are important to TDD as a metric of overall performance. TDD leadership becomes concerned when backlog levels get excessively high.

TDD coordinates with the AMD to plan regularly scheduled missions to haul backlog cargo, as shown in *Figure 2-10*. These missions are created based on backlog cargo levels and historic TDD mission data. TDD schedules these missions from 1-4 weeks in advance, as opposed to 1-7 days for missions carrying specific priority cargo. Currently, approximately three

STARS missions run daily to handle backlog cargo. TDD and AMD also coordinate to create a list of shipments. This list contains priority shipments to be scheduled by TDD.

Several constraints, specific to CENTCOM, limit TDD planning. Each port within the theater has unique characteristics: operating hours, MOG, and aircraft suitability. A port's operating hours define when aircraft are allowed to land or takeoff. The MOG level at each port describes the port's capacity to park and service aircraft. Each port might have unique circumstances to consider that might alter its suitability for different aircraft. For example, a port might have a gravel runway that requires re-grading following a given number of landings by wide body cargo aircraft before operations can resume.

Higher priority external conflicts can require the use of TDD aircraft on a temporary basis. For example, a conflict can be an aero-medical evacuation or a distinguished visitor.

2.2.2.2 Planning Function

Actual scheduling starts with the mission planners. These personnel are in charge of adding new missions to the schedule and adjusting the schedule due to unforeseen changes, such as execution delays, aircraft availability, or higher priority requirements. The bulk of the planning typically occurs between two weeks and 72 hours before execution.

Similar to domestic airlines' operations, delays in execution ripple through the schedule causing later missions assigned to the delayed aircraft to be pushed back. In some instances, this simply changes the start time of the next mission, but sometimes later missions must be cancelled due to current delays. Aircraft can become unavailable due to maintenance, aircrew needs, diplomatic clearance restrictions, weather, MOG restrictions, cargo handling maintenance, or unpredicted cargo fluctuations. Aircraft can be temporarily assigned to other higher priority missions such as Aerial Evacuation (AE) of wounded personnel.

For example, an aircraft might be reassigned to take injured American troops from a port in Iraq all the way to Germany. In this case, that aircraft would be unavailable for several days. The arrival of higher priority requirements on the cargo tracker might also force a mission planner to reschedule currently scheduled missions in order to accommodate the new cargo.

Mission planners use several rules to try to improve the scheduling process.

- Only schedule 9 of 11 aircraft;
- Leave holes in schedule until 48 hours out; and

- Plan extra slack time between missions.

Based on the historic utilization rates of C-17 aircraft, which are approximately 80% [3], planners try to schedule a similar percentage of their aircraft. The extra aircraft are only used when necessary, but often must be used due to unexpected changes. Also, planners try to leave gaps in the schedule until the day before execution. This allows planners the ability to add a mission on the day before execution in the required time slot without adjusting any other missions. Planners try to adapt as best they can when the schedule gets overloaded with more than the expected number of aircraft unavailable, or if the aircraft scheduled are full and additional high priority requirements appear in the system.

Planning a new mission to cover existing requirements is accomplished with the greedy algorithm in which the most important requirements are considered first. Planners generally sort the cargo tracker by priority and then by time window. The planner references the *TOMCAT*, which contains all schedule information for each aircraft as a graphic time line. A mission is tentatively considered to fill an open time on the schedule, comprised of the required cargo from the top of the list, along with any other cargo that can fit on the aircraft or be carried by scheduling additional legs.

The planner runs through several checks to confirm that the mission is feasible. First, considered are individual limitations for each airfield, such as the airfield's operating hours, suitability, and any other special considerations. Aircraft can only use an airfield during its operating hours and each airfield has unique operating hours. Planner's consider MOG at each airfield and ensure that the mission uses each port when there is sufficient capacity at the airfield to handle the aircraft. The planned mission time is calculated by estimating the flight times and turn times at each port.

Planners reference a flight time calculator from CAMPS that uses historic data to estimate flight times. Estimated port turn times are determined based on what operations are necessary at each stop and the location of the port. For example, 2 hr 15 min is generally used as the turn time for a C-17 at a port, but 1 hr 45 min is used for all ports within Iraq. Based on the mission time, the planner must confirm the mission will not break the Crew Duty Day (CDD), which is 18 hours for a standard crew and 24 hours for an augmented crew. Permission to land at each airfield must be requested before the mission starts. Finally, any required diplomatic clearance must be granted if necessary and planners must send requests to a specific cell within the TACC that handles this issue. Some clearances require a few days notification and others take longer to process depending on the specific characteristics of the cargo and the country to be flown over.

After these steps have been taken, the planner updates the *TOMCAT* to reflect the newly added mission and then enters the mission information into the GDSS2, as shown in *Figure 2-10.* Planners continue to add missions to the schedule until all current requirements are covered or aircraft are exhausted.

2.2.2.3 Execution Function

From the moment a mission is within 24 hours of takeoff, to the time the aircraft ends the mission, the mission is monitored at the execution desk. As mentioned in *Section 2.1.1*, this differs from most other aircraft scheduled by the TACC in that most planning and execution functions are operated by different directorates. Unforeseen changes or delays in the current missions require adjustments to the schedule. The execution desk coordinates with AMD personnel and the TDD planner to minimize disruptions to the schedule, while delivering as much of the required cargo as possible. For example, if an aircraft scheduled to takeoff with high priority cargo goes down for maintenance, other aircraft scheduled to carry lower priority cargo can be reassigned if no other options exist.

As a mission on the schedule comes closer to execution, changes require more work. Diplomatic clearances and permission to land at an airfield must be redone. Air crews must be notified of changes within 16 hours and flight managers responsible for monitoring individual aircraft must be notified if changes occur within 6 hours of takeoff. Depending on the circumstances, a planned mission might not be able to be rescheduled if changes occur. Because TDD is a small unit, it can be flexible and adjust to the inherent uncertainty in the mission scheduling process if necessary.

2.3 Motivation for Modeling & Analysis

The objective of TDD is to support CENTCOM and respond to needs for intratheater airlift. From the perspective of troops in theater, it is critical to have effective airlift. This means getting the right cargo to the right place at the right time. For TDD, planners are also concerned with efficiency in terms of utilization rates and number of aircraft used. There is a constant struggle to balance effectiveness and efficiency. TDD is effective when it can satisfy all of the requirements given from the AMD in CENTCOM.

2.3.1 Important Metrics

Several metrics are reported weekly as a summary for mission planners within TDD and the leadership of the TACC. This summary includes the scheduled requirements, scheduled capacity, delivered requirements, and actual capacity for both passengers and pallets broken down by day. These metrics show planners the aircraft capacity utilization rate and allow them to compare how close their planned schedules match up with what actually gets delivered. The number of aircraft used per day and a forecast of requirements for the next week are also reported. Additionally, backlog cargo levels for each port are reported and leadership becomes concerned when these levels are high.

The comparison between planned and actual capacities indicates another unspoken metric: robustness. Another way to improve effectiveness and efficiency is to build robust plans that are more flexible to unforeseen change in the system. Planners want to build plans that minimize the deviation between actual and expected plans. By minimizing disruptions to plans, some inefficiencies, such as cancelled missions, can be eliminated.

According to TDD leadership, current operations are generally effective at delivering the given requirements, but there is plenty of room for improvements in efficiency. These improvements would most likely be reflected as an increase in the backlog cargo delivered or additional available capacity to haul passenger that would otherwise be assigned to C-130 aircraft.

2.3.2 Relevant Modeling Efforts

In the past year, efforts have been made to improve the planning process for intratheater airlift. The Naval Postgraduate School built a model that has been used by personnel in the AMD to aid planners scheduling airlift. The Air Tasking and Efficiency Model (ATEM) provides routing and cargo decisions so that planners can better maximize utilization of intratheater cargo aircraft. ATEM has been used to solve weeklong scenarios with up to 30 aircraft, 20 ports, two commodities, five aircraft types, and multiple aircraft configurations [15].

To further improve the model, a heuristic has been developed that has generated similar results as the IP, but with shorter runtimes. This heuristic does not require the use of any commercial optimization software and interfaces easily with MS Excel. ATEM still contains some limitations. Transshipment and throughput over multiple ports is not considered. Also, cargo is not given with any delivery time windows or priority. However, it has shown improvements over the previous process of manually building routes.

Chapter 3 Functional Analysis & Modeling Scope

The previous chapter presented the current TDD mission scheduling process from an operational perspective. This chapter has four main objectives:

1. Develop a functional analysis of the current TDD mission scheduling process;
2. Describe the input and output we will use for our model;
3. Define the scope of modeling efforts presented in the remainder of this thesis; and
4. Discuss the problem classification and possible modeling approaches.

3.1 Overall Functional Analysis

Currently, the actual mission planning is a labor intensive process. Several aspects of the mission planners' task can be captured using a mathematical model. For instance, the schedule for TDD aircraft can be generated using this model that requires fewer flight hours to deliver required shipments and/or takes less time for planners to create. Shorter solve times allows planners the time to analyze further the solution and compare alternatives. The improved solution will better meet operational objectives of efficiency and effectiveness – introduced in *Section 2.3*.

The following figures illustrate how a model can simplify the TDD planning process. *Figure 3-1* illustrates an overlay of the functions within the mission scheduling process that can be captured within a model. We then collapse these functions into a single model, as illustrated in *Figure 3-2*.

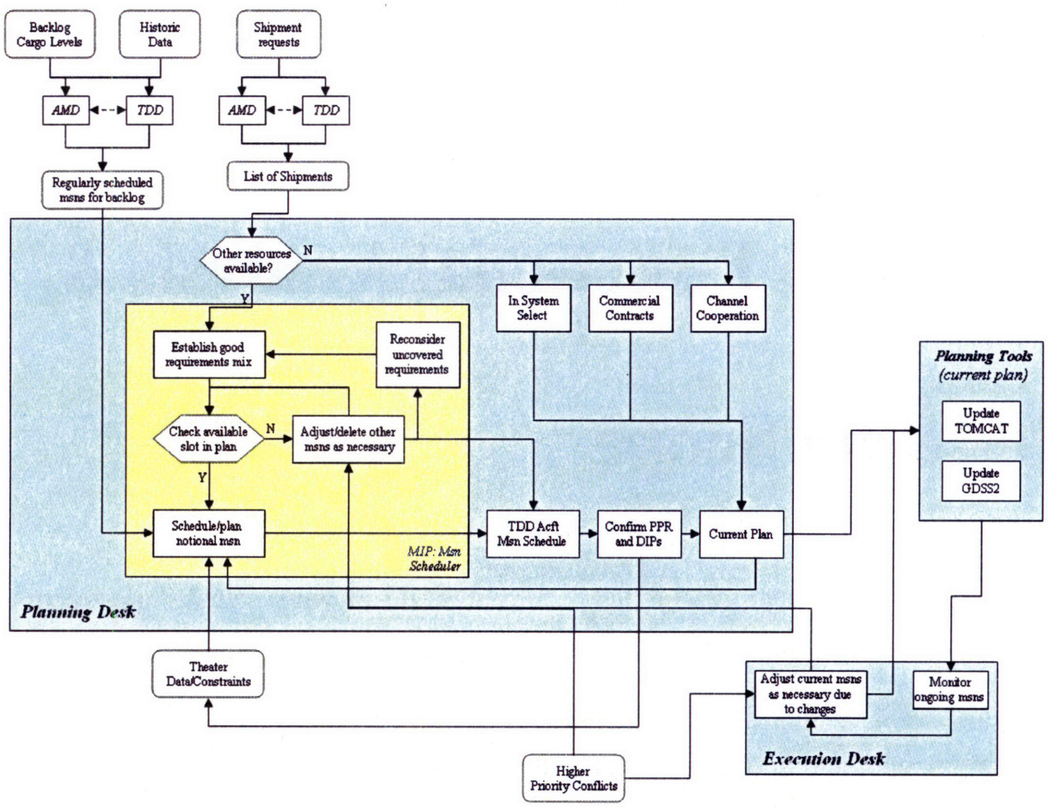

Figure 3-1: Current TDD Mission Scheduling Process with Model Overlay

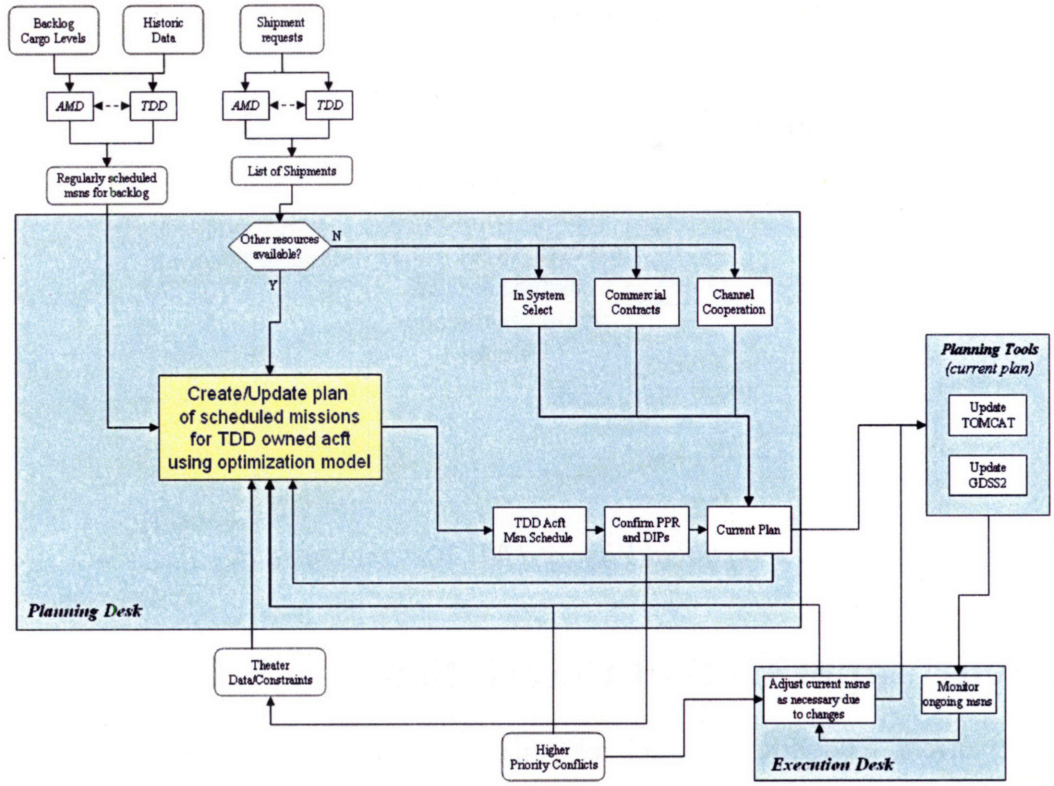

Figure 3-2: Improved TDD Mission Scheduling Process

If we consider our optimization model a functional "black box," we can look at the inputs and outputs for the model. *Figure 3-3* illustrates a simplified view of the model in a similar layout as the entire process illustrated in *Figure 3-2*.

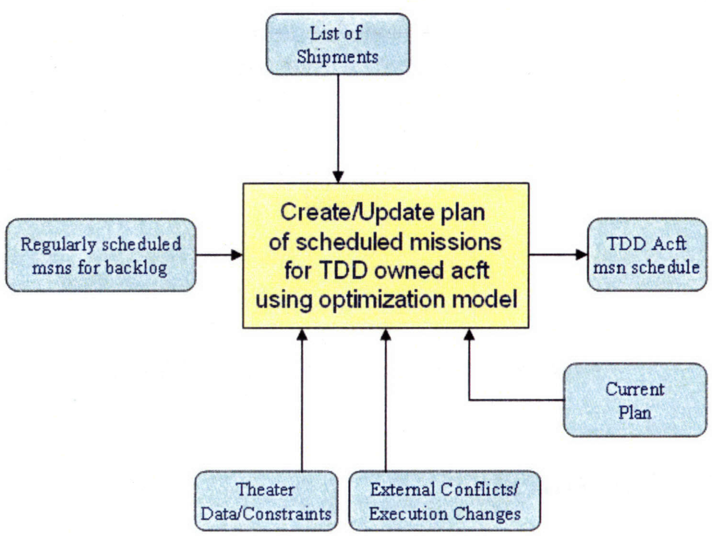

Figure 3-3: Model Inputs and Outputs Based on Scheduling Process

We can further simplify our functional view.

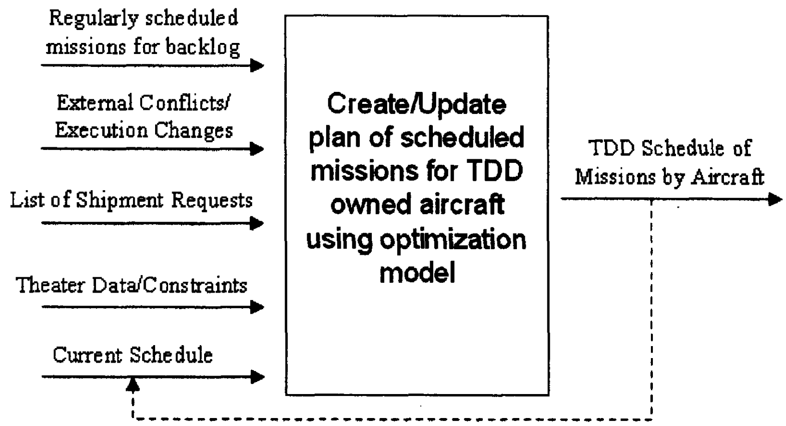

Figure 3-4: Functional View of TDD Mission Scheduling Process

3.2 Model Input/Output Description

Our current modeling does not take into account all relevant inputs from the TDD scheduling process, but considers the inputs that make up the core problem of scheduling organic aircraft. The model takes a set of shipment requests and theater specific data and constraints as inputs. It then outputs a schedule of missions selected to deliver those shipments in order to maximize the objective. We can further define the inputs and outputs to the model.

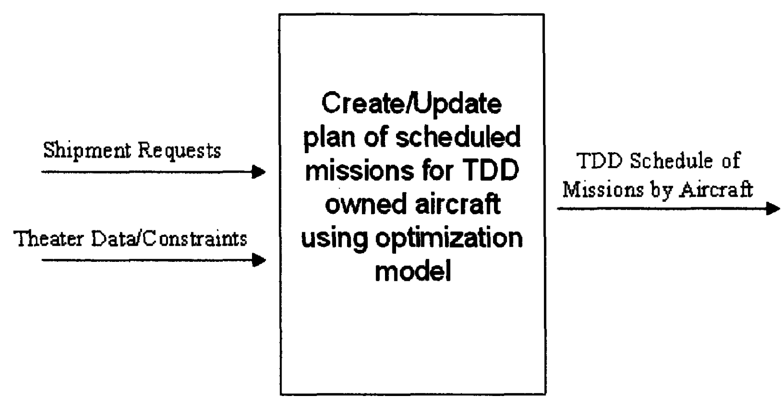

Figure 3-5: Functional View of Model

3.2.1 Shipment Requests

Each individual shipment request contains several attributes: origin/destination, priority, size, weight, time in system (TIS), and a delivery time window. Every origin and destination corresponds to one of twenty aerial ports in the Middle East. We reference each aerial port by its International Civil Aviation Organization (ICAO) airport code.

The priority of a shipment corresponds to its importance. For our model we assume priority levels of one to three, three being the highest priority. Priority 2 and 3 correspond to medium and high priority cargo, respectively. This cargo is important to deliver and has a defined time window. Priority 1 cargo corresponds to low priority cargo that is considered 'backlog' cargo. This cargo does not have a defined time window.

The size of a shipment is defined by the number of pallets and total weight. All pallets have a standard length (108 in) and width (88 in), but their height varies depending on the actual cargo. For our model, we are interested in the C-17, whose cargo capacity is determined by a maximum number of pallets (18) and a maximum weight (170,900 lbs). For simplicity in our model, we only consider size and not weight. It would not be difficult to add a constraint limiting aircraft loads by weight.

Because some of the shipment requests represent passengers, we must convert the number of passengers to pallet positions required. AMC uses a conversion table that estimates the number of pallet positions required for different ranges of passengers.

The *time in system* (TIS) corresponds to the time at which a shipment request is made known to TDD and arrives in their system to be considered in the planning process. The delivery time window consists of an Available Load Day (ALD) and a Required Delivery Day (RDD). All three time metrics are measured in days. Therefore, the time windows are large compared to the amount of time necessary to complete a single mission because no regular mission can be longer than 18 hours, based on the normal crew duty day limit.

3.2.2 Theater Data/Constraints

There are several pieces of information concerning the theater of interest that must be captured in our model. The flight time matrix has flight times between every port in the theater (Table 3-1). These flight times are based on estimates derived using a flight time calculator in GDSS.

	ICAO	Port	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	ORAA	Al Asad	0.0	0.4	0.5	0.6	0.8	0.9	1.2	0.6	0.6	0.9	0.8	1.2	1.6	2.4	1.7	1.8	4.6	0.6	3.1	3.3
2	ORAT	Al Taq	0.4	0.0	0.6	0.4	0.8	0.6	0.9	0.6	0.5	0.7	0.9	1.1	1.5	2.1	1.8	1.8	4.5	1.3	3.1	2.9
3	ORBD	Balad	0.5	0.6	0.0	0.5	0.7	0.5	1.2	0.6	0.5	0.8	0.8	1.2	1.4	2.3	1.8	1.9	4.4	1.5	2.6	3.6
4	ORBI	Baghdad	0.6	0.4	0.5	0.0	0.8	0.8	0.9	0.7	0.6	0.7	0.5	1.1	1.4	2.2	1.9	1.9	4.4	1.4	2.3	3.5
5	ORBM	Mosul	0.8	0.8	0.7	0.8	0.0	0.4	1.3	0.3	0.5	1.1	0.8	1.5	1.6	2.6	1.3	2.1	4.6	1.8	3.3	4.2
6	ORKK	Kirkuk	0.9	0.6	0.5	0.8	0.4	0.0	1.1	0.5	0.4	1.0	0.0	1.4	1.4	2.4	1.4	1.8	4.4	1.9	3.1	3.5
7	ORMM	Basrah	1.2	0.9	1.2	0.9	1.3	1.1	0.0	1.5	0.0	0.4	0.7	0.5	0.5	1.5	0.0	1.1	3.5	1.2	2.1	2.5
8	ORQW	Q West	0.6	0.6	0.6	0.7	0.3	0.5	1.5	0.0	0.4	1.1	1.0	1.5	1.4	2.4	1.4	1.7	4.4	1.6	2.9	4.1
9	ORSH	Al Sahra	0.6	0.5	0.5	0.6	0.5	0.4	1.0	0.4	0.0	0.9	0.4	1.3	1.4	2.1	1.5	1.6	4.4	1.7	3.0	4.0
10	ORTL	Tallil	0.9	0.7	0.8	0.7	1.1	1.0	0.4	1.1	0.9	0.0	0.4	0.6	0.7	1.7	2.1	1.4	3.7	1.0	2.0	3.3
11	ORU1	Al Kut	0.8	0.9	0.8	0.5	0.8	0.0	0.7	1.0	0.4	0.4	0.0	0.9	0.8	1.7	0.0	1.1	3.8	1.6	2.1	3.3
12	OKAS	Al Salem	1.2	1.1	1.2	1.1	1.5	1.4	0.5	1.5	1.3	0.6	0.9	0.0	0.3	1.2	2.7	1.0	3.1	2.1	1.7	2.5
13	OKBK	KCIA	1.6	1.5	1.4	1.4	1.6	1.4	0.5	1.4	1.4	0.7	0.8	0.3	0.0	1.2	2.7	1.1	3.0	2.3	1.6	2.5
14	OTBH	Al Udeid	2.4	2.1	2.3	2.2	2.6	2.4	1.5	2.4	2.1	1.7	1.7	1.2	1.2	0.0	3.3	0.7	3.6	2.5	0.8	2.0
15	LTAG	Incirlik	1.7	1.8	1.8	1.9	1.3	1.4	0.0	1.4	1.5	2.1	0.0	2.7	2.7	3.3	0.0	2.3	4.8	0.8	4.1	5.3
16	OBBI	Bahrain	1.8	1.8	1.9	1.9	2.1	1.8	1.1	1.7	1.6	1.4	1.1	1.0	1.1	0.7	2.3	0.0	3.7	2.6	1.1	2.4
17	HDAM	Djibouti	4.6	4.5	4.4	4.4	4.6	4.4	3.5	4.4	4.4	3.7	3.8	3.1	3.0	3.6	4.8	3.7	0.0	5.2	4.1	2.7
18	OJAM	Amman	0.6	1.3	1.5	1.4	1.8	1.9	1.2	1.6	1.7	1.0	1.6	2.1	2.3	2.5	0.8	2.6	5.2	0.0	3.0	2.9
19	OMAM	Al Dafra	3.1	3.1	2.6	2.3	3.3	3.1	2.1	2.9	3.0	2.0	2.1	1.7	1.6	0.8	4.1	1.0	4.1	3.0	0.0	1.5
20	OOTH	Tumrait	3.3	2.9	3.6	3.5	4.2	3.5	2.5	4.1	4.0	3.3	3.3	2.5	2.5	2.0	5.3	2.4	2.0	2.9	1.5	0.0

Table 3-1: Flight Times for Relevant Ports (hours)

Each aerial port has several attributes. Every port is designated as in or out of the combat zone. This affects mission planning, refueling capabilities, and turn times. Mission planners will typically not plan a mission with several landings at ports within the combat zone. Combat locations do not have refueling capabilities. The estimated turn time in a combat zone is approximately 1 hour 45 minutes, while an aircraft outside the combat zone will take a total time of 2 hours 15 minutes without refueling or 3 hours 15 minutes with refueling (Table 3-2).

Most ports have an unlimited capacity in terms of landings allowed per day. Maximum on ground (MOG) constraints, however, do impact daily mission scheduling. The MOG capacity of a port defines how many aircraft can by on the tarmac at any one time. Mission planners must take this into consideration and stagger mission arrivals when necessary.

ICAO	Airport	Combat	Fuel?	TurnTime	w/fuel
ORAA	Alasad	y	n	1.75	
ORAT	Al_Taq	y	n	1.75	
ORBD	Balad	y	n	1.75	
ORBI	Baghdad	y	n	1.75	
ORBM	Mosul	y	n	1.75	
ORKK	Kirkuk	y	n	1.75	
ORMM	Basrah	y	n	1.75	
ORQW	Q-West	y	n	1.75	
ORSH	Al_Sahra	y	n	1.75	
ORTL	Tallil	y	n	1.75	
ORU1	Al_Kut	y	n	1.75	
OKAS	Al_Salem	n	y	2.25	3.25
OKBK	KCIA	n	y	2.25	3.25
OTBH	AL-Udeid	n	y	2.25	3.25
LTAG	Incirlik	n	y	2.25	3.25
OBBI	Bahrain	n	n	2.25	
HDAM	Djibouti	n	y	2.25	3.25
OJAM	Amman	n	y	2.25	3.25
OMAM	Al Dafra	n	y	2.25	3.25
OOTH	Thumrait	n	n	2.25	

Table 3-2: Aerial Port Information Table

3.2.3 Output

The output from our model consists of a schedule of missions. Each mission consists of several attributes: cargo, aircraft route, duration, start time, and value. The cargo consists of the set of shipments that are covered by a given mission. The aircraft route defines the sequence of ports that an aircraft visits during the mission. Mission duration is the total time to complete a given mission from the moment an aircraft departs its home port to the time it returns to the home port and is prepared to start another mission. This time includes both the flight time between ports and the turn time at each port visited. Mission start time defines when an aircraft is scheduled to leave the home port of OTBH. The mission value is defined as the sum of the value of all shipments covered by a mission. Each shipment's value is determined by multiplying its priority by size, measured in pallets.

3.3 Modeling Scope

Several aspects of the TDD scenarios had to be created to use as data within our models. We spent time discussing different issues with TDD leadership to make our assumptions and scenarios more realistic.

3.3.1 Assumptions

To simplify the initial model, we make the following assumptions:

- no transshipment or throughput;
- aircraft configurations not considered;
- reduced priority set;
- aircraft capacity measured by pallet positions;
- normal crew duty day restrictions;
- all cargo converted to estimated pallets;
- missions begin and end at home port of OBTH; and
- no long term scheduled missions to carry backlog cargo.

Neither transshipment nor throughput is allowed in the current model. Throughput is defined as any cargo that takes intermediate stops to get to its required destination. For example, cargo is required over O/D pair A-C, but flies on a route over A-B-C. Transshipment is defined as cargo moves on multiple aircraft to get to its proper destination. Again, for cargo required over O/D pair A-C, it can travel A-B and get dropped off at B. Another aircraft can then pickup the cargo at B and deliver it B-C.

There are multiple configurations for an aircraft in which extra seats can be carried at the expense of pallet positions. We are not concerned with which configuration is used, but must make sure that the cargo carried can fit onto one of the given configurations. The actual configuration used is decided during execution.

We reduce the number of priorities down to three. Operationally, priorities range from one to ten, but the complexity of different priority levels can be captured in the reduced set of three priority levels. The higher two priorities represent specific requests with ALD/RDD. Lower priority represents general backlog cargo that does not have a specific ALD/RDD. We want to try and deliver all cargo, but it is more important to deliver higher priority cargo within the given time windows and then deliver as much low priority cargo as possible to reduce overall cargo levels at each port.

Aircraft capacity can be measured by volume and weight. We currently only measure aircraft capacity by pallet positions. Eighteen pallets can fit inside a C-17. Later, we should also consider weight.

The crew duty day (CDD) for the C-17 is 16 hrs for a normal crew and 24 hrs for an augmented crew. An augmented crew consists of two normal crews and can therefore extend the CDD. We are currently only considering normal crews.

Cargo consists of two types: pallets (PAL) and passengers (PAX). While we currently use both types of cargo, we convert passengers to an estimated number of pallets using a conversion table from AMC. Later, we should try to keep the cargo in its original form instead of converting everything into PAL because these are only estimates.

We assume that all missions begin and end at the home port. This is not always the case because some missions end before they get all the way to the home port. With the fleet of C-17s based in Qatar, they require a long positioning leg to satisfy requirements in Kuwait or Iraq. Missions can end in Kuwait, but a crew must be pre-positioned there to take the aircraft on its next mission.

Some scheduled missions do not have specific cargo assigned, but rather carry available backlog cargo on predetermined routes. These missions are scheduled far in advance based on historical data. We are not considering missions of this type as input to our model.

Several of these assumptions can be relaxed or at least revisited in later models.

3.3.2 Data

One main aspect of data generation was cargo. Because the actual demand for cargo was not available, we used historic haulage data to represent our demand.

3.3.2.1 Historic TDD Haulage Data

Using historic data from May to October of 2006 for all TDD scheduled missions, we began to filter the data into something usable for our modeling efforts. Our first step was to filter out all missions using the B-747, IL-76, AN-124 or C-5 aircraft because we are only considering scheduling the C-17s. The AN-124, B-747, and C-18 flew a trivial number of mission legs with all fewer than five legs each over the entire six month period. The IL-76 did fly a significant percentage of all TDD missions, as illustrated in the table below, but planning these commercial contracted aircraft is outside the scope of this thesis.

	C-17	IL-76
Total	7463	1408
Days Used	184	181
Avg per Day	40.56	7.78
StDev per Day	6.96	1.27

Table 3-3: Flight Legs by Aircraft Type: May-Oct 06

Next, we filtered the data by mission ID code. All TDD missions fly using *YT* in the mission ID code. Out of the 7463 mission legs flown by the C-17, we removed 161 not using the *YT* designation in the mission ID.

To limit the scope of our modeling, we did not consider all ports visited by TDD aircraft. We filtered the data by port location, noticing that TDD flew to more than 50 ports over the six month period. Some of these ports, however, were visited few times. We eliminated all ports with less than 10 missions for the entire period, limiting our ports to 25. Because we assumed all missions must begin and end at OTBH, we did not consider any of the flights into and out of Afghanistan, eliminating the corresponding five ports.

We chose two specific sets of historic hauled cargo to build scenarios. The first set is from a week's worth of cargo hauled in August 2006 and the second is a week's worth of cargo in September 2006. According to TDD leadership, August represented an average workload and September was an above average workload for TDD.

Table 3-4 and *Table 3-5* show pallets demanded by port of origin and broken down by color representing different destinations. The total number of pallets in the August scenario is 2,095 compared to 2,277 in September, representing a 9% increase. Taking a quick look at the demand data, we can see that a majority of the cargo flows through either OKAS, ORBD, or OTBH.

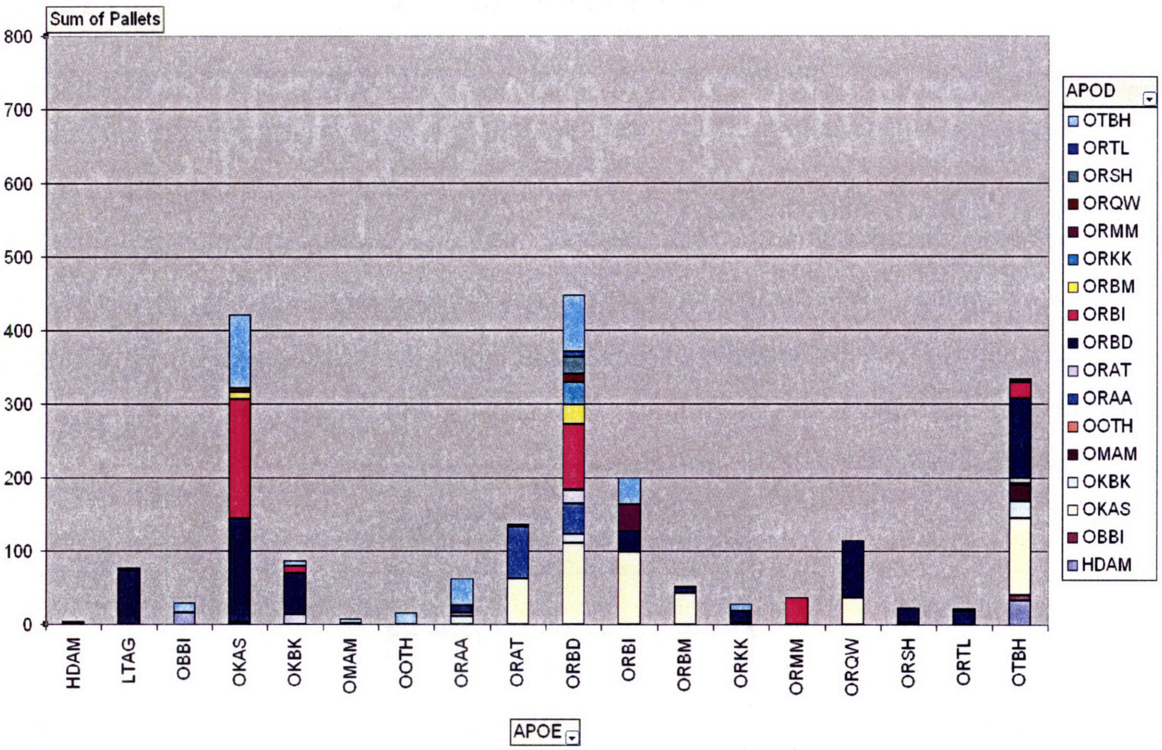

Table 3-4: Demanded Pallets from APOE to APOD for August Scenario

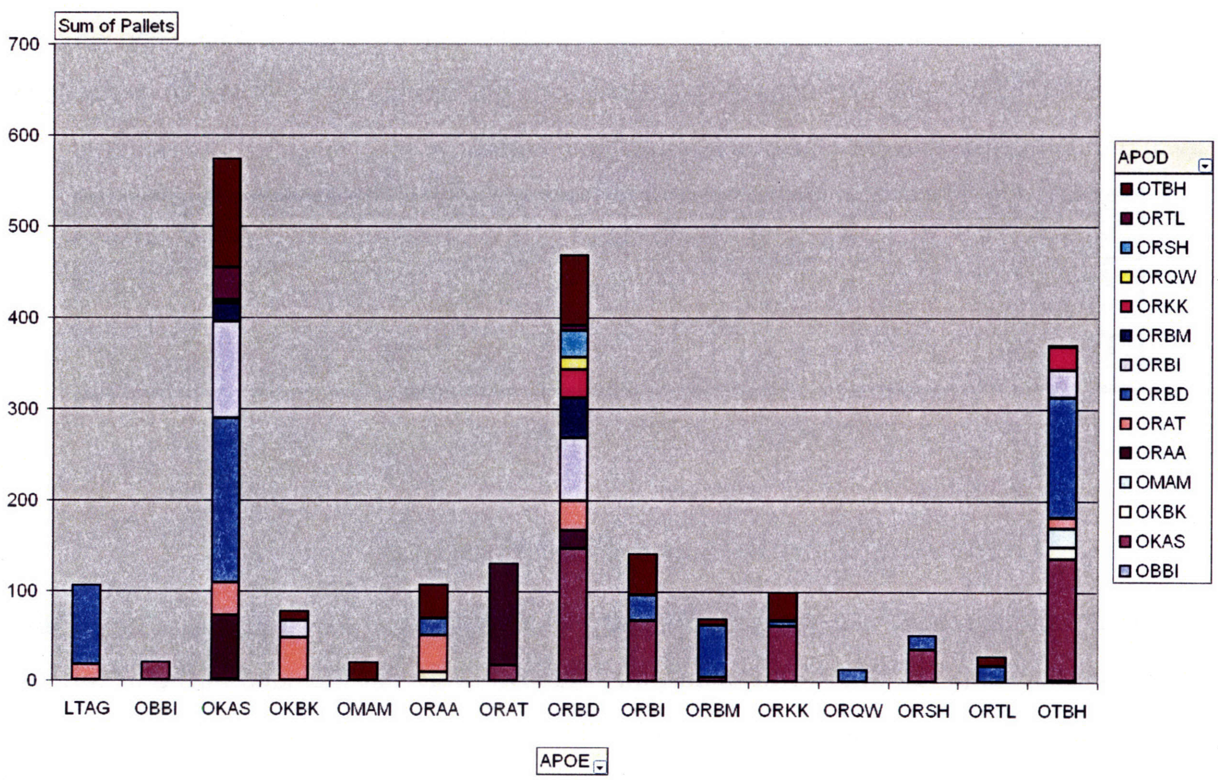

Table 3-5: Demanded Pallets from APOE to APOD for September Scenario

3.3.2.2 Converting Passenger Requirements to Estimated Pallet Positions

One issue with our data was that cargo came in two types, both passengers (PAX) and pallets (PAL). For our modeling efforts, we converted all PAX requirements to estimated PAL positions. As shown in the following table, the C-17 has several configurations that can be used to carry differing amounts of both PAX and PAL. When carrying all pallets, the C-17 uses a double row of pallets. When carrying a significant number of passengers, the pallets can be placed in a single row down the center of the aircraft, allowing for passengers to sit on fold down seats located on each side of the cargo hold. For carrying only passengers, a center row of seats can be installed on the C-17, greatly improving passenger capacity.

Config	PAX	PAL
1	0	18
2	54	11
3	189	4

Table 3-6: C-17 Configurations with PAX and PAL Limits

Several combinations of passengers and pallets can be used. When carrying both passengers and pallets, planners use tables that estimate the number of pallet positions needed for a given number of passengers, illustrated in the table below.

PAL	1	2	3	4	5	7	8	9	10	18
PAX Range	1	5	25	33	43	55	61	71	79	85
	4	24	32	42	54	60	70	78	84	189

Table 3-7: C-17 Passenger (PAX) to Pallet (PAL) Conversion Estimates

3.3.2.3 Generating Priorities and Time Windows

For each shipment, we generated priority levels and time windows based on feedback from TDD leadership. Both priorities and time windows were not included in the historic data and were not available without accessing classified data. We validated our assumptions to generate these fields with TDD, confirming with leadership that our assumptions provided a reasonable sampling of the type of planning done within TDD.

TDD deals with priority levels ranging from 1-10. To simplify our model, but still maintain the operational complexity of different priority levels, we reduced the number of priority levels from ten to three. We assigned a priority level of one, two, or three to each requirement, with three being the highest. These priorities were generated using a discrete uniform distribution, U[1,3]. Priority 1 represents backlog cargo in the system. Priorities 2 and 3 represent higher priority requirements with specific time windows.

For the Time in System (TIS), Available Load Day (ALD), and Required Delivery Day (RDD), we generated each using historical data, the following table, and equations *(3.1)*.

Priority	PAX	PAL
1	20+	20+
2	U[1,2]	4
3	U[1,2]	4

Table 3-8: RDD Generation based on Priority and Cargo Type

The start of each time window, the ALD, was set as the day the cargo was actually delivered based on historical data. The end of each time window, the RDD, was then set based on our assumptions outlined in Table 3-8.

As mentioned previously, the Priority 1 cargo represents backlog and does not have a specific time window. Because our scenarios only cover a single week's worth of cargo, the Priority 1 cargo effectively has no limiting time window. Based on inputs from TDD, we modeled higher priority PAX requirements with a time window of either one or two days and PAL requirements with a time window of four days. We assigned the TIS for each requirement between two and four days. Each TIS was set by subtracting a randomly generated value using a discrete uniform distribution, U[2,4], from the ALD.

$$\begin{aligned} ALD &= HistoricDeliveryDate \\ RDD &= ALD + F[priority, type] \\ TIS &= ALD - U[2,4] \end{aligned} \tag{3.1}$$

3.4 Problem Classification

One key aspect of building a model is understanding how our problem relates to similar problems that have been solved in the past. Our problem is an extension to the classic *traveling salesman problem* (TSP) known as the *pickup and delivery problem with time windows* (PDPTW).

3.4.1 Define TSP, VRP, PDPTW

The TSP is a classic optimization problem, defined as follows: "Given an undirected graph $G=(N,E)$ with n nodes, and costs c_e for every edge $e \in E$, the goal is to find a tour (a cycle that visits all nodes) of minimum cost" [19]. In the last half century, significant progress has been made to solve larger and larger instances of the TSP and other similar problems. Currently, no polynomial time approximate algorithm exists for the TSP and other *NP-hard* problems [19]. The core of our problem can be viewed as a TSP, assuming a single aircraft that had infinite range and capacity.

We are interested in a problem that adds complexity to the classic TSP, namely:

- multiple vehicles;
- capacitated vehicles;
- tour duration limits;
- time windows on each load; and
- precedence constraints.

First, we are trying to build tours for multiple vehicles (aircraft) at the same time. For m vehicles this amounts to solving m TSPs simultaneously. This is commonly referred to as the *vehicle routing problem* (VRP). The limitations on our vehicles and nodes make it a *capacitated vehicle routing problem with time windows* (CVRPTW). Lastly, we are not concerned with visiting single nodes, but rather pairs of nodes (ports) that correspond to the pickup and delivery of each load. This is known as a *pickup and delivery problem* (PDP).

3.4.2 Dynamic Nature of Mission Planning

As mentioned in Chapter 2, TDD is responsible for not only planning missions, but also monitoring execution. Each day follows a similar routine. TDD personnel monitor all missions currently in execution, while others are responsible for updating the current plan with any necessary changes and additional missions to deliver newly available cargo.

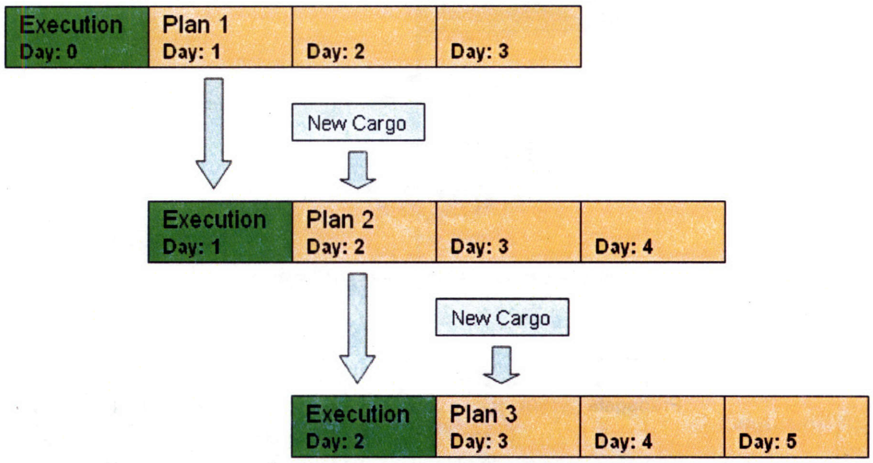

Figure 3-6: Daily Re-Planning within TDD

In general simplistic terms, most planning merely adds missions to the schedule to cover new cargo being available in the system (*Figure 3-6*). The cargo arrives periodically and this is known as a *time-based trigger*. The planning process, however, can rapidly become more complicated (*Figure 3-7*) and require immediate re-planning in which currently planned missions must be moved, adjusted, or deleted. For example, if high priority cargo arrived in the system, other lower priority missions might need to be delayed in response. Sometimes events occur that require immediate re-planning to adjust the schedule. These *events-based triggers* occur when there is a delay in execution or other unforeseen change, such as an aircraft going down for maintenance.

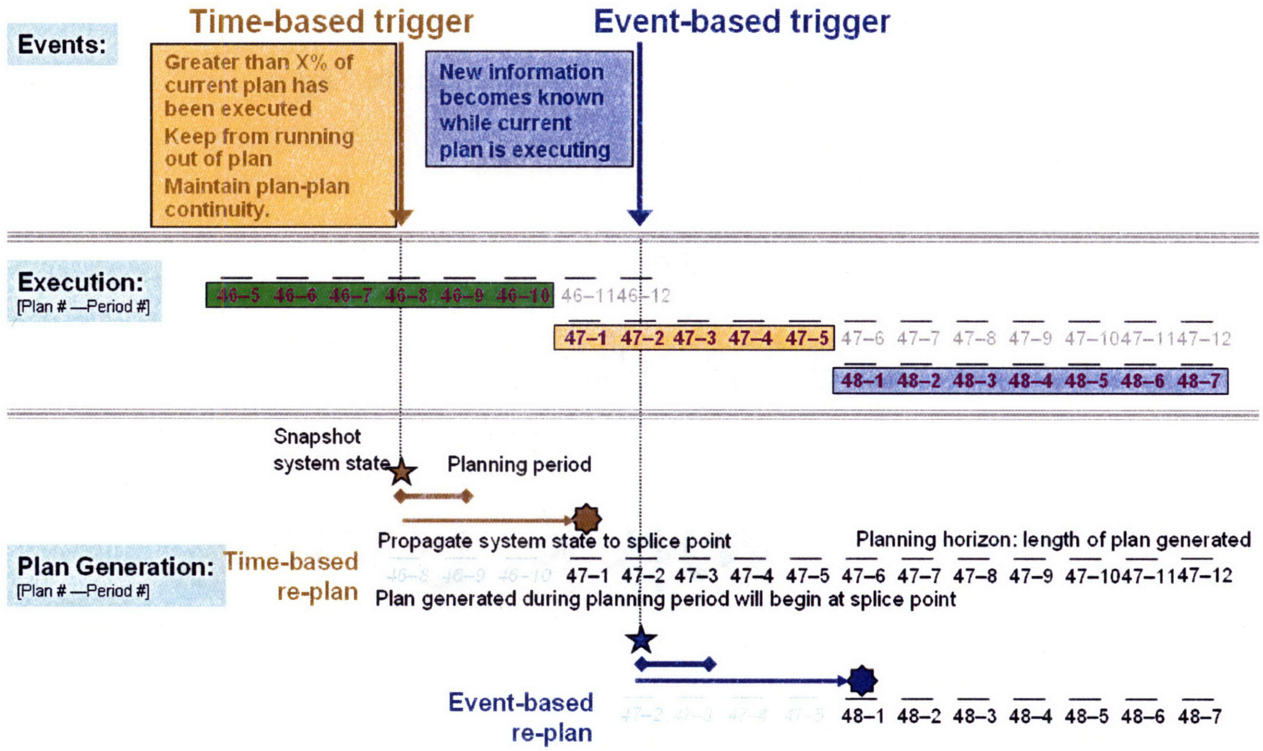

Figure 3-7: Planning while Executing (Re-planning)

3.4.3 Missions as Composite Variables

Conventional approaches to similar problems without time windows have included arc and path based formulations. We have decided on an approach that uses a different structure to define our variables. We refer to our variables as *composite variables* in that each one represents more information than the decision to choose a given arc, path, single commodity, and vehicle [15]. Each *composite variable* represents an entire mission. We define a mission as a unique combination of shipments and specific route to cover that set of shipments.

The number of variables considered grows exponentially with the number of shipments considered for each mission. For our problem, the number of all combinations of shipments considered does get large, but does not become intractable because most missions carry four or fewer shipments. The following table illustrates a histogram of shipments contained on each mission, based on historical data. We see that over half the missions only deliver one shipment, and over 90% of the missions deliver four shipments or fewer.

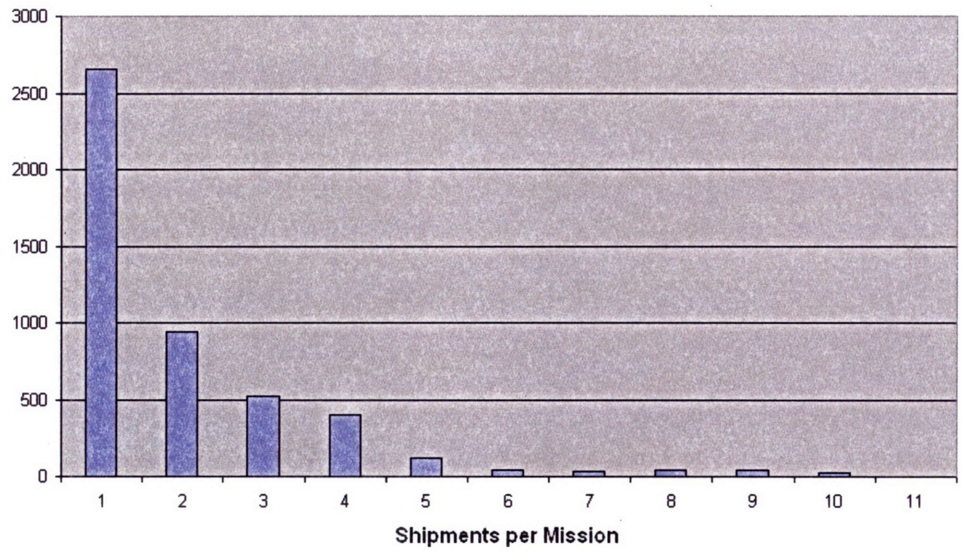

Table 3-9: Histogram of Shipments per Mission from May – Oct 06

Our variable choice requires a significant amount of preprocessing before solving our problem. For each feasible shipment combination, we build a good route to cover the set of shipments. However, this then simplifies our model by only having to assign whole missions and their start times on the schedule. We will further explain our modeling techniques in the following chapter.

[This Page Intentionally Left Blank]

Chapter 4 Modeling

The previous chapter introduced a functional model and discussed the data and assumptions used to implement that model. This chapter has three objectives:

1. Describe the our modeling approach;
2. Introduce a MIP formulation and greedy heuristic to assign missions to the schedule; and
3. Discuss improvements and extensions to our algorithm.

4.1 Theater Airlift Model

We introduced the input and output to our model in *Section 3.2*and now we will describe how our 'black box' works. As shown in the previous chapter, we again illustrate the overall model in *Figure 4-1*.

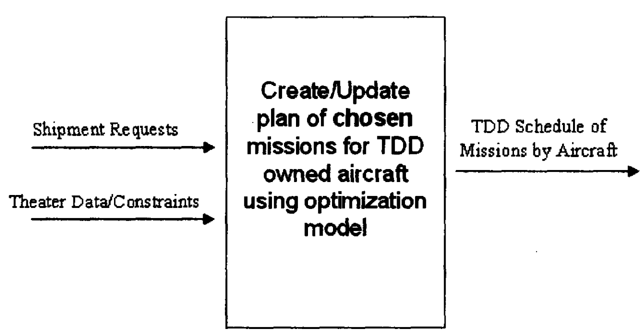

Figure 4-1: Functional View of Overall Model

As described in *Section 3.4.3*, missions are composed of a set of shipments and a specific route to cover those shipments. The output from our model assigns missions to a specific aircraft at a given time.

We break down the overall model into its main functions, as illustrated in *Figure 4-2*. If we were to consider the entire list of shipments for scheduling at once, the problem becomes intractable. Because the entire list of shipments is too big to consider at one time, we break the list of shipments apart and solve a piece at a time. For each subset of shipments, we build a set of feasible missions to serve as our variables and run an optimization model to select the missions. The optimization approach solves a mixed integer program (MIP), selecting missions that satisfy the constraints and maximize the objective function. The last step is to update our master schedule and shipment list to reflect the newly chosen missions. The algorithm exits the loop when no shipments remain to be scheduled.

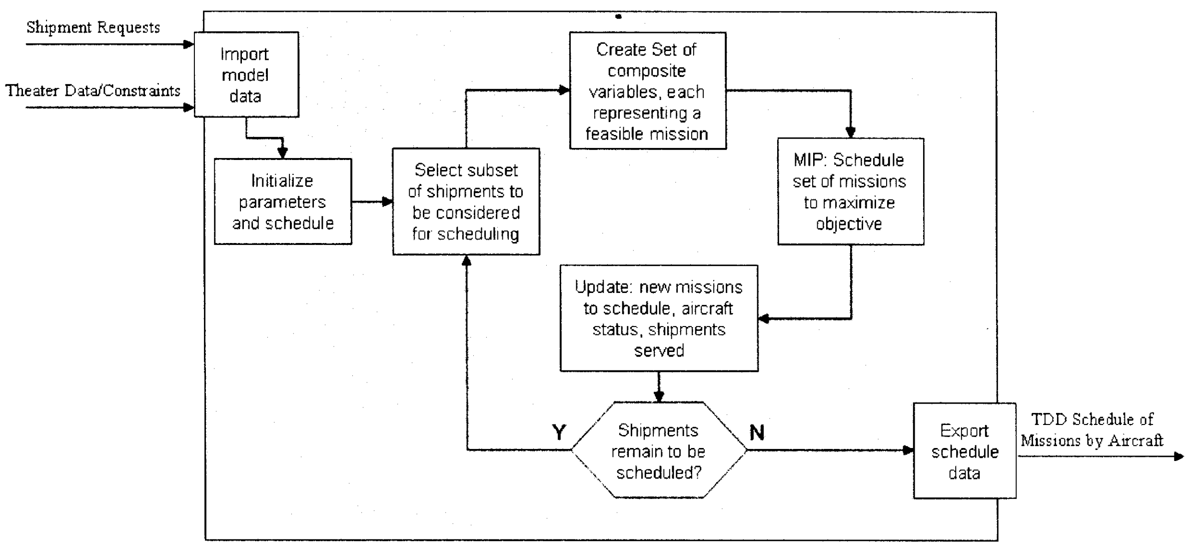

Figure 4-2: Functional Components of Overall Model

4.1.1 **Initialization**

The first step of our model, after reading in the necessary data, is to initialize the parameters and state variables. Some parameters deal with operational constraints, while other parameters adjust model run times and termination conditions.

The main state variable that must be initialized is the master schedule. The schedule contains all relevant mission data, described later in *Section 4.1.3*, as well as start times for each mission, broken down by each individual aircraft on the schedule. Each aircraft's available start time is determined by the end of the last mission scheduled. We start all scenarios with no planned missions and all aircraft available; however, it is not difficult to start the algorithm with some missions already on the schedule or currently in progress.

The important parameters in our model are detailed in the following sections.

Operational Parameters

The operational parameters are:

- Crew Duty Day (CDD);
- Number of Aircraft Available;
- Number of Periods;
- Size of Cargo Sets (Maximum number of Shipments per Mission); and
- Excess Slack Time between Missions.

For our modeling, we only consider normal CDD limits. The number of periods in our algorithm refers to the number of missions that can be scheduled for an aircraft within the MIP. A period is defined as block of time required to complete a mission. When setting the number of periods to consider in the MIP, the number of missions assigned can be no greater than the number of periods. By increasing the number of periods we greatly increase the size of the MIP. The maximum size of *cargo sets* limits the number of shipments considered for a mission. Slack time is additional time available for an aircraft to stay at the home port. This allows aircraft to wait before the next mission is scheduled.

Derived Parameters

The derived parameters are:

- Schedule Time, with

  $ScheduleTime = NumPeriods \cdot (CDD + HomePort\_TurnTime + SlackTime)$; and

- *Working List* Maximum Length, with

  $WorkingList\_Length = periods \cdot aircraft \cdot CS\_size \cdot factor$.

The schedule time defines the amount of time considered for mission scheduling beyond an aircraft's current availability. This time is based on several parameters. For example, if we consider the number of periods equal to two, the excess slack time equal to three, the CDD equal to 18, and a home port turn time of 3.25 hours, as described in *Section 2.2.2.2*, then $ScheduleTime = 2 \cdot (18 + 3.25 + 3) = 48.5$ hours.

The maximum length of the working list is determined by the number of shipments that can be scheduled during an iteration, multiplied by a factor larger than one. This gives us a set of shipments to consider during each iteration, larger than the number we can actually schedule. For example, if we consider 9 aircraft, 2 periods, 3 shipments per mission, we are able to schedule 54 shipments. By using a factor of 1.25, then we set the maximum working list length to be 68.

Run Time Parameter

The run time parameter is the Integer Program (IP) – Linear Program (LP) Relative Gap. This parameter is specific to the MIP, solved using CPLEX version 10.1.0. The IP-LP relative gap refers to the relative tolerance between the relaxed LP solution and the IP solution. For example, if the IP-LP relative gap is set to 0.05, CPLEX will stop when a feasible integer solution is within 5% of optimal.

4.1.2 Selecting a Subset of Shipments

The next step of the algorithm is deciding which shipments to consider in the current iteration, illustrated in *Figure 4-3*. One of the parameters of our model is the maximum length of the list of shipments, called the *working list,* considered at each iteration. This list is ordered based on a shipment's time window and value. Based on all shipments remaining to be delivered, we add a shipment to the *working list* if it meets certain criteria, namely:

- At least one aircraft available on the schedule after shipment *available load day* (ALD);
- All shipments with earlier *required delivery day* (RDD) already included on *working list;* and
- Current *working list* less than maximum length.

We continue to add shipments to the *working list* until we reach the maximum length or no other shipments exist that meet our criteria. The values are then updated based on the *value function* we define within the algorithm. As the deadline of a shipment approaches, it is more important to select that shipment to be delivered in upcoming missions, reflected by higher values. The criteria for adjusting the shipment's value can be adjusted by changing the *value function* used to define a shipment's value.

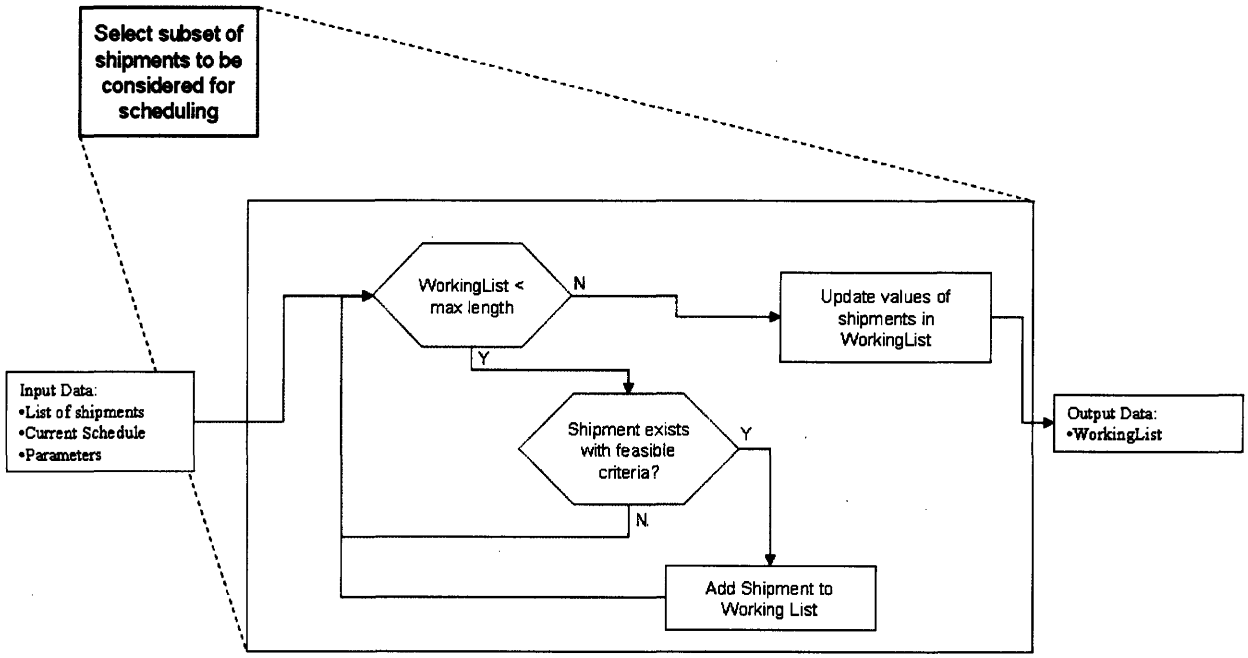

Figure 4-3: Functional Description of Working List Selection

During initial experiments, we used a simple combination of a shipment's size multiplied by its priority to determine its value. Because this value function did not take into consideration the amount of time left until the shipment was due, our algorithm did not deliver several of the shipments within the required time window. We have currently defined two functions to calculate the updated value for shipments, taking into account the time until the RDD of a shipment.

The first value function *(4.1)* does not consider size or priority of the shipment and defines a shipment's value by its position in the working list multiplied by a factor of 10. We define the following parameters:

n Length of the working list;

i Position of shipment in list from *0..n-1*; and

s_i Shipment *i*, with attributes: origin, destination, size, priority, TIS, ALD, and RDD as defined in *Section 3.2.1.*

$$s_{i,value} = 10 \cdot (n-i) \tag{4.1}$$

Using the value function *(4.1)*, the last shipment in the list has a value of 10, with each shipment above it increasing by 10. Because the list is ordered according to shipment RDD, this corresponds with higher values the earlier a shipment is due.

Equation *(4.2)* below involves a step function based on the RDD of each shipment plus a combination of the shipment's size and priority. We define the following parameters:

e - Earliest time an aircraft is available on schedule;

α - Value added if $e - s_{i,RDD} \leq 1$ (day) after first available aircraft, with

$$\alpha \succ \beta + s_{i,size} \cdot s_{i,priority} \quad \forall i \in WorkingList;$$

β - Value added if $2 \leq e - s_{i,RDD} \leq 1$ (day) after first available aircraft, with

$$\beta \succ s_{i,size} \cdot s_{i,priority} \qquad \forall i \in WorkingList.$$

$$s_{i,value} = \begin{cases} \alpha + s_{i,size} \cdot s_{i,priority} & e - s_{i,RDD} \leq 1 \\ \beta + s_{i,size} \cdot s_{i,priority} & 1 < e - s_{i,RDD} \leq 2 \\ s_{i,size} \cdot s_{i,priority} & 2 < e - s_{i,RDD} \end{cases} \tag{4.2}$$

The selection of the specific value function used within our algorithm almost surely results in different missions and shipments scheduled. Other alternative value functions should be explored in the future.

4.1.3 Composite Variable Generation: Feasible Mission List

After selecting the subset of shipments to be considered and defining their value, we generate composite variables, each representing a different feasible mission, illustrated in *Figure 4-4*. We base each mission on a unique combination of shipments. By limiting the maximum number of shipments that can be considered in any mission, we limit the combinatorial explosion that will otherwise ensue if we considered all possible combinations of shipments. This limitation on the number of shipments in any given mission is a reasonable assumption based on historical operational data, as mentioned in *Section 3.4.3,* which shows that most missions contain 4 or fewer shipments.

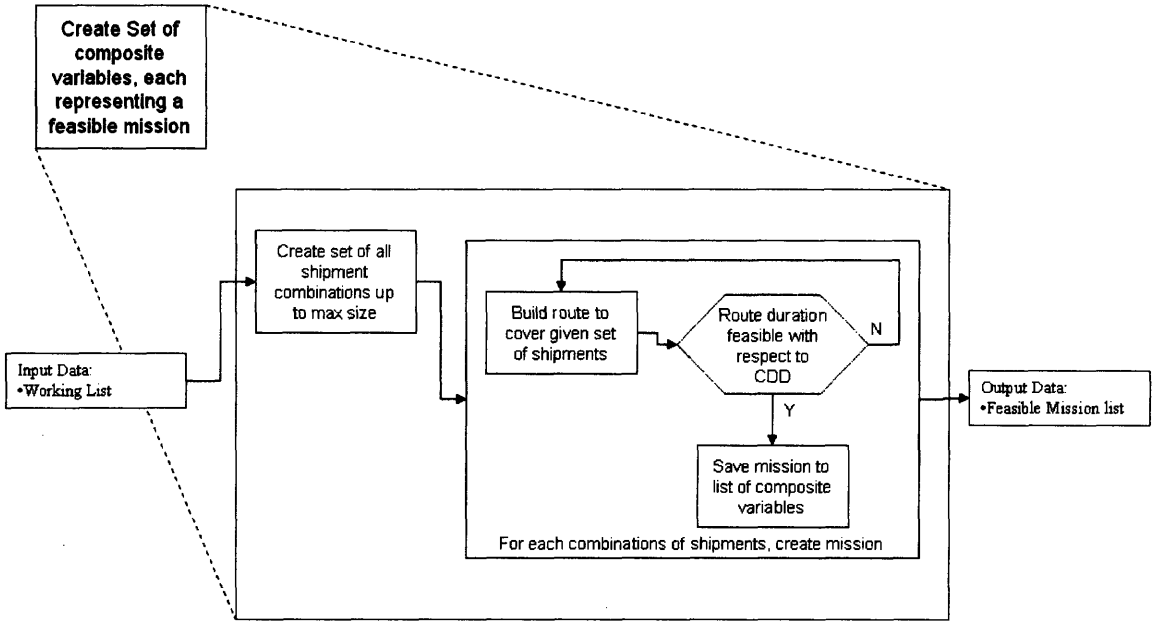

Figure 4-4: Functional Description of Composite Variable Generation

Based on the current *working list* of shipments, we build a set containing all combinations of shipments, referred to as a *cargo set*. For each *cargo set* we then build a route that covers those shipments. We illustrate route building in *Figure 4-5*.

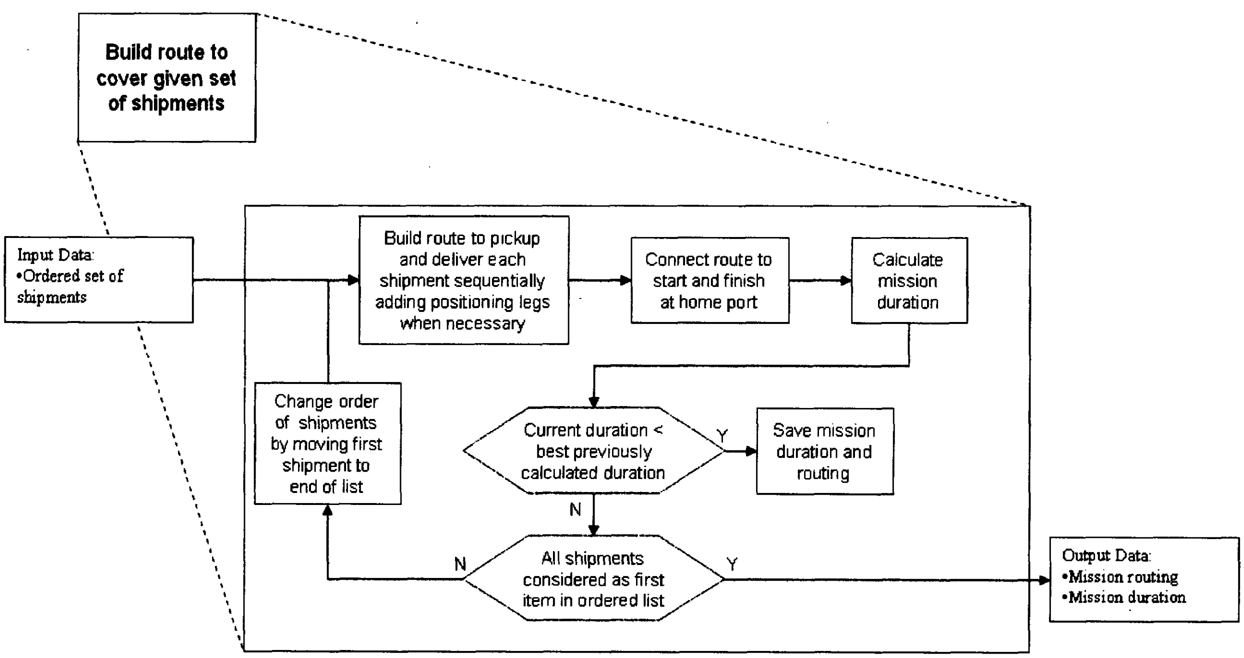

Figure 4-5: Building a Route for a Given Set of Shipments

While we do consider different orderings of shipments while finding the shortest route, one shipment must be dropped off before the next shipment is picked up. The final route chosen is suboptimal because we do not consider all permutations of shipment sequencing and because we are only considering one shipment on an aircraft at a time. We can build improved routes by implementing route optimization models, which are discussed later in this chapter (*Section 4.3.1*).

After building the routing for a given combination of shipments, we check to see if that route is feasible in terms of mission duration. Given normal CDD restrictions, crews cannot fly missions over 18 hours in length. Therefore, any mission whose duration is greater than CDD limits is not a feasible composite variable.

Each mission that is saved as a composite variable has several attributes, namely:

- Route;
- Duration;
- Shipments covered;
- Value;
- ALD; and
- RDD.

Based on the route selected, we calculate the estimated mission duration by summing the associated flight time and port turn time estimates. Based on the shipments covered in a given mission, we calculate the mission value, ALD, and RDD. The value gained by executing a mission is based on the sum of all shipment values. The ALD and RDD for each mission represent the latest ALD and earliest RDD that cover all shipments within a mission.

There are no built in checks for feasibility before an entire mission is built. Therefore, missions are generated for each possible *cargo set*, defined and afterwards checked for feasibility, in terms of mission duration. If we keep track of infeasible *cargo sets*, we can build all feasible missions without having to consider the routing necessary for each cargo set. We will discuss ways to determine a mission's feasibility before actually calculating the routing in *Section 4.3.2*.

4.1.4 MIP: Assigning Missions to Aircraft Schedules

After we have generated a set of feasible missions, we must select which missions to assign to our schedule. Each mission already contains the necessary information of aircraft routing and which shipments to deliver. Within our constraints, we assign missions that maximize the value and number of shipments delivered while minimizing the number of aircraft used, total flight time, and mission finish times.

For our formulation, we use the set of missions generated as described in Section 4.1.3. We add the *null* mission to this set of missions. The *null* mission is defined as a mission having zero value, zero duration, an unlimited time window, and no shipments. This provides the algorithm with the option of choosing no mission for an aircraft.

We define *periods* in the following formulation as "slots" on the schedule in which a mission can be assigned and do not relate to real time. Periods are defined in such a manner that only one mission can be assigned to any period. In our scenarios, missions are between 10-20 hours, so each period represents an arbitrary length of ~20 hours.

We remind the reader that a shipment's ALD and RDD refer to the *available load day* and *required delivery day*. Because our formulation works in hours, we convert these shipment time window limits. For all time windows referenced in days, the time refers to 2400 at end of the given day. For example, if we start day one at zero hours, a shipment with and ALD of five and a RDD of 7 relates to a time window of [120,168] in hours.

The following notation describes our formulation:

Sets

A set of all aircraft $a \in A$. Each aircraft has an available start time.

B set of all aerial ports $i \in B$

K set of all cargo shipments $k \in K$. Each shipment is indexed by $\{i,j,p,w\}$, indicating that cargo must move from aerial port i in B to aerial port j in B, has p priority, and has size w measured in pallets.

M set of all missions $m \in M$ where M is the set of feasible missions with respect to operational constraints plus the *null* mission. Each mission has an associated value, duration, and time window, input as data. The shipments covered by each mission m are defined by the indicator variables δ_m^k equal to 1 if mission $m \in M$ covers shipment $k \in K$, and equal to 0 otherwise. The routing information for each mission is not relevant for our MIP.

P set of all periods $p \in P$ from *1..n*.

Decision Variables

x_{map}: equal to 1 if mission $m \in M$ is assigned to aircraft $a \in A$ during period $p \in P$ and 0 otherwise;

B_{ap} start time of the mission assigned to aircraft $a \in A$, during period $p \in P$;

y number of shipments delivered;

z total required flight time; and

$usedAcft_a$ equal to 1 if any mission is assigned to aircraft $a \in A$ and 0 otherwise.

Data

$value_m$ value measured by cargo moved of assigning any given mission $m \in M$;

$time_m$ duration of mission $m \in M$;

ALD_m ALD of mission $m \in M$ defined as latest *available load day* for all shipments covered;

RDD_m RDD of mission $m \in M$ defined by earliest required *delivery day* for all shipments covered; and

s_a Available start time of aircraft $a \in A$.

Parameters

scenLength length of current scenario (hrs); and

ϕ, φ, γ Non-negative coefficients for different terms of the objective function that can be adjusted to change the overall objective function.

Indicator Variables

δ_m^k: equal to 1 if mission $m \in M$ covers cargo $k \in K$ and 0 otherwise; and

u_a: equal to 1 if aircraft $a \in A$ is assigned at least 1 mission and 0 otherwise.

Formulation

$$Max \sum_{m\in M} \sum_{a\in A} \sum_{p\in P} \left(value_m \cdot x_{map}\right) + \phi \cdot y - \gamma \cdot z - \sum_{a\in A} \left(usedAcft_a + \varphi \cdot B_{ap|p=n}\right) \tag{4.3}$$

s.t.
$$\sum_{m\in M} \sum_{a\in A} \sum_{p\in P} \delta_m^k \cdot x_{map} \leq 1 \qquad \forall k \in K \tag{4.4}$$

$$\sum_{m\in M} \sum_{p\in P} x_{map} \cdot time_m \leq scenLength \qquad \forall a \in A \tag{4.5}$$

$$B_{an} + \sum_{m\in M} \left(x_{man} \cdot time_m\right) \leq s_a + scenLength \qquad \forall a \in A \tag{4.6}$$

$$\sum_{m\in M} x_{map} \leq 1 \qquad \forall a \in A, p \in P \tag{4.7}$$

$$\sum_{m\in M} x_{map} \geq \sum_{m\in M} x_{maq} \qquad \forall a \in A, p = 1,\ldots,n-1; q = p+1 \tag{4.8}$$

$$\sum_{m\in M} x_{map} \cdot value_m \geq \sum_{m\in M} x_{maq} \cdot value_m \qquad \forall a \in A, p = 1,\ldots,n-1; q = p+1 \tag{4.9}$$

$$\sum_{m\in M \backslash m=null} x_{map} \leq usedAcft_a \qquad \forall a \in A, p \in P \tag{4.10}$$

$$B_{aq} \geq B_{ap} + \sum_{m\in M} x_{map} \cdot time_m \qquad \forall a \in A, p = 1,\ldots,n-1; q = p+1 \tag{4.11}$$

$$B_{ap} \leq \sum_{m\in M} x_{map} \cdot \left(24 \cdot RDD_m - time_m\right) \qquad \forall a \in A, p \in P \tag{4.12}$$

$$B_{ap} \geq 24 \cdot \sum_{m\in M} x_{map} \cdot ALD_m \qquad \forall a \in A, p \in P \tag{4.13}$$

$$B_{a1} \geq s_a \qquad \forall a \in A \tag{4.14}$$

$$\sum_{k\in K} \sum_{m\in M} \sum_{a\in A} \sum_{p\in P} \delta_m^k \cdot x_{map} \geq y \tag{4.15}$$

$$\sum_{k\in K} \sum_{m\in M} \sum_{a\in A} \sum_{p\in P} time_m \cdot x_{map} \leq z \tag{4.16}$$

$$x_{map}, usedAcft_a \in \{0,1\} \qquad \forall m \in M \backslash null, a \in A, p \in P \tag{4.17}$$

$$B_{ap}, y, z \in \Re + \tag{4.18}$$

The objective of our formulation maximizes the number and value of shipments delivered while minimizing flight time, aircraft used, and mission start time. Constraint *(4.4)* requires that each shipment must be covered by at most one mission. Constraint *(4.5)* forces the total amount of time assigned to each aircraft to be less than the amount of time in the scenario. The mission assigned during the last period must finish before the end of the scenario *(4.6)*. At most one mission is assigned to an aircraft per period *(4.7)*. Constraints *(4.8)* and *(4.9)* specify the order of mission assignments. For each aircraft, missions must be assigned to the earliest open period *(4.8)* and a mission with higher values must be assigned before other missions *(4.9)*. An aircraft is considered used if it is assigned at least one mission *(4.10)*. For each aircraft, the start time for each assigned mission must be greater than the start time for the mission assigned in the previous period plus the duration of the previous mission *(4.11)*. Constraint *(4.12)* ensures that an assigned mission's finish time is before the RDD of the mission cargo and *(4.13)* ensures that an assigned mission's start time is after the ALD of the mission cargo. All missions must start after each aircraft is available *(4.14)*. Both mission assignment variables and aircraft usage variables are binary *(4.17)*. Variables B, y, and z are all real positive numbers *(4.18)*. Constraint *(4.15)* with the objective function forces the variable representing shipments delivered to be no greater than the actual number of shipments delivered. Constraint *(4.16)* with the objective function ensures that the variable representing total flight time equals the actual flight time.

This MIP takes a set of missions as input with only aggregate information concerning their time and value. It outputs the mission assignment/schedule for a set of aircraft over a short time horizon of one to two days.

A current limitation of our model is that each feasible mission does not consider port operating hours or MOG constraints. Because the assignment problem does not currently consider any information 'inside' each mission, it cannot check port operating hours or MOG for each scheduled mission. For most circumstances, MOG conflicts can be handled during execution by shifting a mission forward or backward on the schedule to avoid congestion. However, hard constraints that limit the number of landings during congested periods are helpful to planners. As mentioned in *Section 2.2.2.2*, TDD must reference a global MOG tool when scheduling a mission because MOG capacity is not only affected by TDD missions, but by all incoming aircraft at a port. MOG constraints can be implemented into the MIP that determines the scheduling of missions by taking input data from the MOG tool.

Currently, the problem considers a static assignment with no uncertainty. Operationally, the problem is a rolling horizon with uncertainty. Due to uncertainty, strategies need to be implemented to make the plan more flexible to change. One example would be to leave at least a set number of aircraft open during each period.

4.1.5 Updating Schedule and Master List of Shipments

Following the selection of missions to be added to the schedule, we update both the master schedule and shipment list, as shown in *Figure 4-6.*

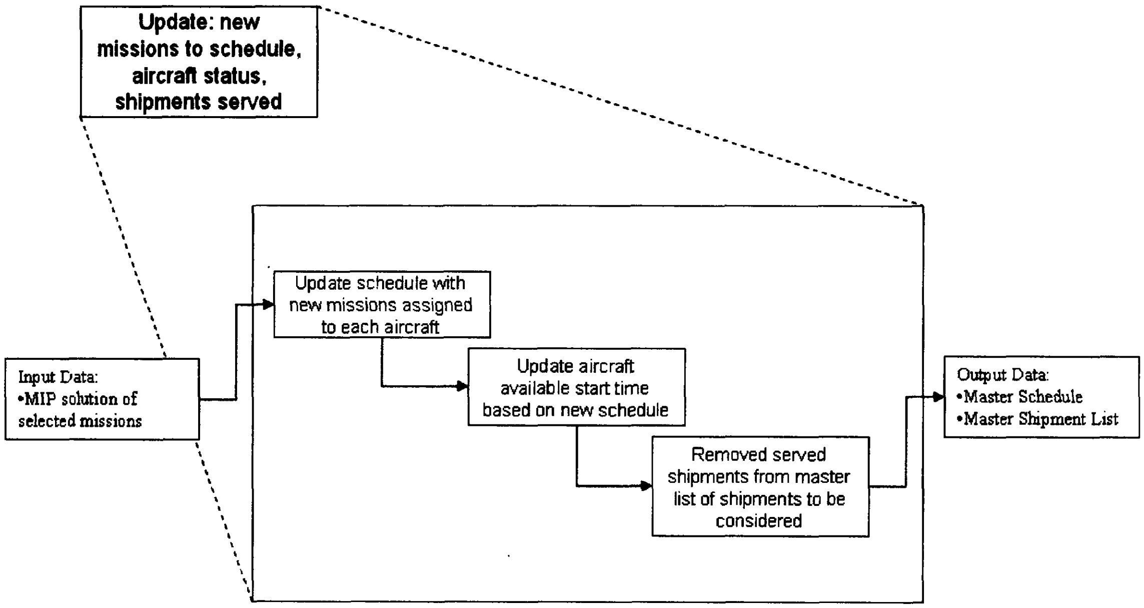

Figure 4-6: Function Description of Updating Master Problem

The master schedule contains all necessary mission information as well as the times for each mission flown for each individual aircraft. As missions are added to the schedule, the master shipment list is updated to reflect shipments that have already been served so they are not considered during the following iterations. The model stops when no shipments remain to be considered for scheduling. If additional shipments remain, the updated schedule and shipment list is then used during the next iteration of our model to build TDD missions.

4.2 Greedy Heuristic

We have also created a greedy heuristic for comparison with our model. While there are several ways to construct a greedy algorithm, we decided to use the same idea of composite mission variables. Instead of using a MIP to select and schedule missions, we use a greedy heuristic. Another approach is to use a greedy heuristic in the creation of mission variables by ordering the set of shipments according to their value and assigning the highest valued set of available shipments. *Figure 4-7* illustrates how our greedy algorithm uses the same construct to create mission variables, but then assigns them to the schedule individually based on value. This approach more accurately reflects the current operational process.

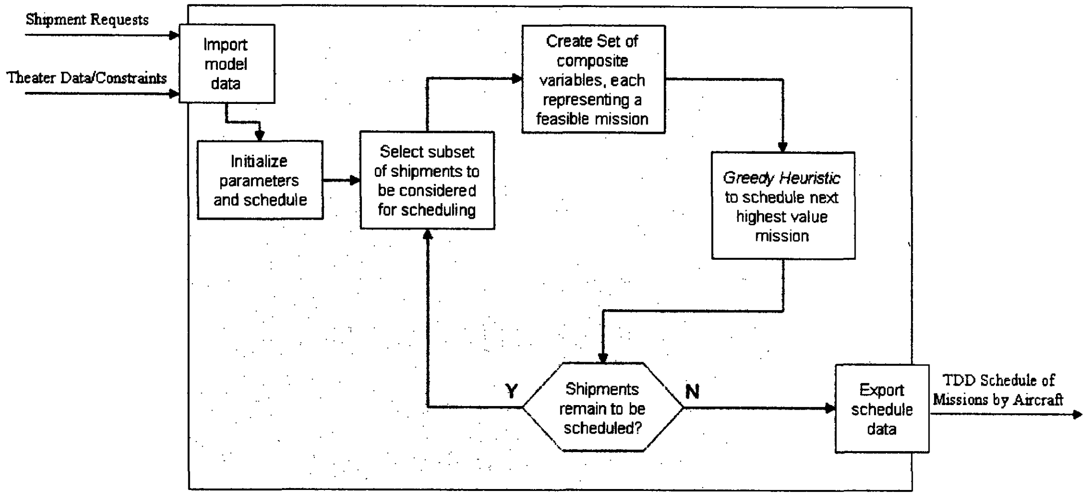

Figure 4-7: Overall Model Using Greedy Heuristic

A functional description of our greedy heuristic is illustrated in *Figure 4-8*. The greedy heuristic does not consider adding multiple missions to the schedule at once like the MIP. Instead, the highest valued mission, with shipments not yet covered, is selected. We still incorporate the idea of a *working list* of shipments to limit the number of composite mission variables created.

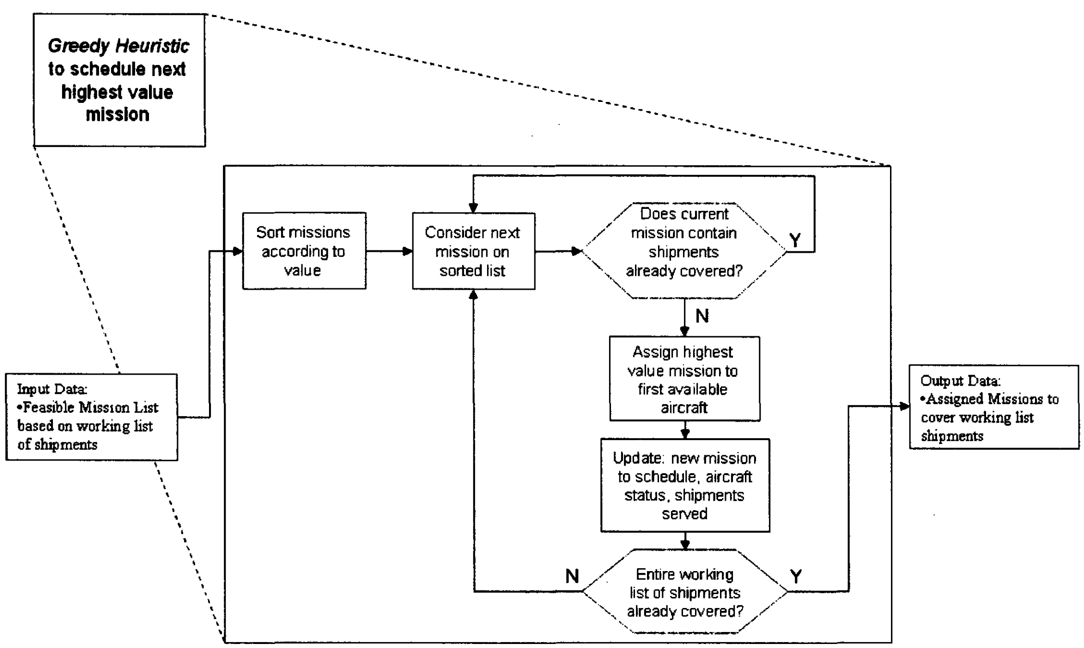

Figure 4-8: Functional Description of Greedy Heuristic

4.3 Possible Improvements

There are several areas in our model where extensions and improvements can be applied in the future. The following section discusses some of these issues.

4.3.1 Building a Route for a Mission

For each mission, we must generate a tour covering the set of shipments in that mission. In our model, we use a heuristic to find the minimum cost sequencing for a given set of shipments to serve with a single aircraft. We can formulate this problem as the classic minimum cost TSP. The network flow formulation for the TSP describes our example of theater airlift.

We assume that missions must begin and end at the home port. Each shipment must be less than the aircraft capacity and there is at most one shipment for each origin/destination pair. This formulation assumes a non-empty set of shipments to be delivered, intermediate stops before delivering a shipment are not allowed, and that the tour visits does not cycle through the home port more than once. If a set of shipments exists such that the solution shuttles back and forth between the home port and another destination to cover a set of shipments, then those shipments can be split up such into multiple sets that can be solved sequentially.

Sets

N set of all nodes containing a source node, a sink node, and nodes representing an aerial port corresponding to either the origin or destination of a shipment. The source node has supply of +1 aircraft and the sink node has supply of -1 aircraft, and all other nodes have supply of 0.

A set of all arcs (i,j) such that $i, j \in N$. Each arc represents the flight leg between port i and port j.

A_s set of all arcs $(i, j) \in A$ such that each (i,j) pair corresponds to an origin and destination of a shipment to be delivered

$O(i)$ set of all $i \in N$ such that an arc exists going out from node i to node j.

$I(i)$ set of all $i \in N$ such that an arc exists going in from node j to node i.

S collection of all proper subsets $s \in S$ of nodes N such that each set s contains at least two nodes

Decision Variables

f_{ij} flow along arc (i,j) representing the number of times a flight leg from port i to port j is used

Data

t_{ij} total time to fly between arc (i,j) and turn the aircraft at the destination such that it is ready to begin another flight leg

o source node

d sink node

Formulation

$$Min \sum_{(i,j)\in A} t_{ij} \cdot f_{ij} \tag{4.19}$$

$$\text{s.t.} \quad \sum_{j\in I(i)} f_{ji} - \sum_{j\in O(i)} f_{ij} = 0 \qquad \forall i \in N \tag{4.20}$$

$$\sum_{j\in O(o)} f_{oj} = 1 \tag{4.21}$$

$$\sum_{i\in I(d)} f_{id} = 1 \tag{4.22}$$

$$f_{ij} = 1 \qquad \forall (i,j) \in A_s \tag{4.23}$$

$$\sum_{i\in s} \sum_{j\in N\backslash s} f_{ij} + \sum_{i\in N\backslash s} \sum_{j\in s} f_{ij} >= 2 \qquad \forall s \in S \tag{4.24}$$

$$\{0,1\} \qquad \forall (i,j) \in A \tag{4.25}$$

Objective *(4.19)* minimizes the necessary time for a tour. Constraint *(4.20)* maintains conservation of flow through the network. Constraints *(4.21)* and *(4.22)* enforce that the aircraft must flow in from the source node to one of the shipment origin nodes and out from one of the shipment destinations to the sink node. Constraint *(4.23)* enforces that the aircraft must cover the given set of arcs corresponding to shipments. Constraint *(4.24)* ensures we eliminate subtours. A drawback of this formulation is that there are exponentially many of these subtour constraints. We can solve without these constraints. If a subtour is generated, we construct the necessary constraint to eliminate the subtour and resolve. We repeat adding necessary constraints until the solution is feasible. Finally, aircraft flows over flight legs must be binary *(4.25)*.

The solution to this problem gives the set of arcs, or flight legs, contained in the minimum cost tour for an aircraft to cover a set of shipments. Because the solution is a flow along a simple cycle, we can find the sequence order by flow decomposition [21], which in this case is simply tracing arcs forming the cycle starting at the home port and picking each arc in the solution in an order that connects each one and whose final destination is the home port.

4.3.2 Building Set of Feasible Missions without Generating all possible Cargo Sets

One way to improve our model is to reconsider the brute-force method of considering all *cargo sets* when building missions. Checks on cargo set feasibility can be based on the *minimal forbidden sets* (MFS) instead of repeating all necessary calculations to determine feasibility for each cargo set generated. The MFS is the collection of cargo sets such that each set is infeasible due to resource limitations and any subset is feasible [19]. After finding a cargo set to be infeasible, we save it to the MFS if it meets the aforementioned criteria. Any time we run across another cargo set with any subset already included in the MFS, we immediately reject the current cargo set as infeasible.

4.3.3 Other Improvement Strategies

One area in which different strategies can be implemented is mission generation. The default is to consider sets of cargo combinations up to a maximum size. Additional constraints can force shipments to have the same priority if considered as a cargo combination. This would create high/low valued missions in terms of cargo. Then, if a mission has to be cancelled, the low priority cargo missions are the first to be cancelled.

Different objectives can be considered in both the working list generation and the optimization. For the working list, different value functions for shipments can be considered based on proximity to due date. For the MIP objective function, one variant is to maximize shipments delivered, regardless of priority. Other variants include placing different weights on each priority level.

Mission details can be considered in the assignment problem. This introduces a large amount of information into the assignment problem, but can ensure that MOG and port operating hour constraints are satisfied.

4.4 Alternative Approaches

4.4.1 Building a Smaller Set of Missions to Consider

Considering all shipments at once is not tractable and creates a combinatorial explosion in the number of (mission) variables to consider in the MIP. One approach is to filter the sample space of missions by solving an IP over subsets of all available shipments to determine good candidate missions. We can then save the 'good' candidate missions from each run and use this set of 'good' candidate missions as the domain from which to choose the missions scheduled over the entire time horizon.

This approach can use similar preprocessing steps as those described earlier to build missions covering subsets of a set of shipments. An IP formulation of the 'knapsack problem,' using a given set of mission variables built from subsets of shipments, can be solved to select a set of candidate missions to add to the sample space for the actual problem. By running this formulation over different, overlapping subsets of all shipments, we form the sample space of candidate missions from which to select.

The following formulation solves for the best set of candidate missions to add to the master problem sample space. Each mission contains the information of aircraft routing and which shipments to deliver. We select candidate missions subject to limits on the amount of available aircraft hours to cover a given set of shipments. The objective maximizes priority shipments delivered while minimizing mission duration.

Sets

B set of all aerial ports i in B

K set of all cargo shipments $k \in K$. Each shipment is indexed by $\{i,j,p,w\}$, indicating that cargo must move from aerial port i in B to aerial port j in B, has p priority, and has size w measured in pallets.

M set of all missions $m \in M$ where M is the set of feasible missions with respect to operational constraints. Each mission is indexed by $\{R,n\}$, indicating that a mission contains a unique shipment combination R and uses tour n. The tour used for each mission is the shortest feasible path to cover all of the shipments in a given mission.

Decision Variables

x_m: equal to 1 if a mission is selected and 0 otherwise.

Data

value$_m$ value based on amount of cargo moved of assigning any given mission $m \in M$;

duration$_m$ length of mission $m \in M$;

ρ maximum number of candidate missions to select; and

ϕ, φ Non-negative coefficients for different terms of the objective function that can be adjusted to change the overall objective function.

Indicator Variables

δ_m^k: equal to 1 if mission $m \in M$ covers cargo $k \in K$ and 0 otherwise.

Formulation

$$Max \sum_{m \in M} (\phi \cdot value_m - \varphi \cdot duration_m) \cdot x_m \tag{4.26}$$

$$\text{s.t.} \quad \sum_{m \in M} x_m \leq \rho \tag{4.27}$$

$$x \in \{0,1\} \tag{4.28}$$

The objective is to selects the number of missions. The missions selected must be fewer than the number of candidate missions desired to be added to the sample space of the master problem *(4.27)* and mission variable is binary *(4.28)*. We solve this problem by ordering the candidate missions according to their objective value and then use a greedy selection of non-negative to select up to ρ missions.

4.4.2 PDPTW Formulation

Another approach is to consider a network formulation based on the pickup and delivery problem with time windows (PDPTW). Recent models developed by Ropke, Cordeau, and Laporte have solved instances of the PDPTW with up to eight vehicles and 96 requests [19].

Their problem contained similar constraints dealing with vehicle capacity, pairing and precedence based on origin/delivery of shipments, and time windows. Their algorithm uses a branch-and-cut method. Additional inequalities are implemented to strength the bounds on the formulation.

Using this formulation likely is not be an efficient way to solve our scenarios, which have over 250 requests and at least nine vehicles. However, it can possibly be implemented on subsets of shipments similar to how we have run each iteration using a *working list* of shipments.

Chapter 5 Results and Analysis

The previous chapter described the modeling techniques used to implement our algorithm. This chapter has three objectives:

1. Define the metrics and objective function(s) used to analyze our models;
2. Discuss the parameter settings used in experimentation; and
3. Describe experiments performed and discuss the results.

5.1 Metrics

As mentioned in *Section 2.3.1*, TDD is concerned with the *efficiency* and *effectiveness* of their operations. To gauge the performance of our models in both areas, we will use several metrics.

5.1.1.1 Efficiency

We measure efficiency from the perspective of TDD. By decreasing the values of the following metrics while delivering a given set of shipments, TDD is able to better utilize their main assets, the C-17s, to either carry additional cargo or be tasked for other important assignments. The metrics of interest to us are:

- Total flight time;

- Average aircraft finish time;
- Number of flight legs; and
- Number of positioning legs.

The metrics listed above give us several ways to measure the efficiency of our solution. Less *flight time* in a solution translates to a more efficient use of aircraft. *Aircraft finish time* is based on the time at which an aircraft completes all scheduled missions and the average aircraft finish time shows when aircraft should be available for other uses, such as operating future missions. We note that the average aircraft finish time might not vary significantly, even if all other metrics seem to suggest a more efficient solution. The aircraft finish time is likely constrained by certain shipments that have later availability dates.

We also look at the *flight legs* used in a solution. A *positioning leg* is defined as any flight leg in a mission in which no shipments are carried. The aircraft is only being positioned to pick up shipments on following flight legs or returning to its home port. Solutions with fewer positioning legs provide less flight time spent without carrying any cargo and increase flight time available for missions.

We have not included the aircraft capacity utilization rate in our list of metrics because it does not vary in our modeling approach. For each mission, shipments must be delivered before another shipment is picked up. Because only one shipment is on an aircraft at any time, the capacity utilization rate is the same for any solution as long as all shipments are delivered. Kiel-doesn't utilization get affected by the amount of empty repositioning that goes on? If extensions to our model are applied that include multiple shipments on an aircraft at any one time, then it would make sense to also track the aircraft capacity utilization as a measure of efficiency.

5.1.1.2 Effectiveness

We measure effectiveness from the perspective of the war fighter requesting airlift via TDD. If TDD is unable to deliver all requested shipments, it will probably result in further delays and the need to find some alternate means for delivery. The effectiveness metrics of interest to us are:

- Percent of available shipments delivered; and
- Hold time for moved shipments.

Measures of effectiveness include the *percent of available shipments moved* and the hold time for these shipments. The *hold time for moved shipments* within the necessary time windows is the primary measure of effectiveness. In our models, if all aircraft are scheduled beyond a shipment's time window, the shipment is removed from the list of shipments to be considered and shown as not delivered in the final solution.

If all shipments are delivered, we can compare the hold time for each shipment. We define hold time as the time starting when a shipment is available to the time it is delivered to its destination. A shipment's hold time includes both the wait time until a shipment is picked up and the transit time after it is picked up until the shipment reaches its destination. We calculate the hold time for all shipments and also broken down based on shipment priority.

5.1.1.3 Additional Metrics

Planners are concerned not only with the efficiency and effectiveness of plans, but also time necessary to actually build plans, as well as the chance that a plan will have to be adjusted in the future. This leads to the following additional performance metrics:

- Algorithm run times;
- Slack time between missions; and
- Mission duration.

We also include additional metrics that measure computational performance and solution robustness. *Algorithm run times* measure the amount of time it takes to solve our model using a Dual Core 3.6/3.59 GHz Intel Pentium 4 processor with 1 GB of RAM.

Slack time between missions and *mission duration* are two metrics that can be used to examine the robustness of a solution. Solutions with longer average slack times between scheduled missions are less likely to need to be rescheduled due to unforeseen circumstances that cause missions to be delayed. If the average mission duration is near the crew duty day (CDD) limit, a delay can cause the crew to violate the CDD and warrant the use of an augmented crew that has an extended CDD.

5.2 Objective Function

The objective function used in the MIP drives which feasible missions are added to the schedule at each iteration of our solution algorithm. This objective, as shown again in *Equation 5.1*, maximizes the total value of shipments delivered as well as the number of shipments delivered, while minimizing total flight time, aircraft used, and mission end time.

$$Max \sum_{m \in M} \sum_{a \in A} \sum_{p \in P} \left(value_m \cdot x_{map}\right) + (\phi \cdot y) - (\gamma \cdot z) - \sum_{a \in A} \left(usedAcft_a + \left(\varphi \cdot B_{ap|p=n}\right)\right)$$

Equation 5.1: Objective function of MIP Mission Scheduler

The direct benefit of assigning a specific mission to an aircraft and period, as decision variable *x*, is defined by $value_m$. The $value_m$ for each mission *m* is recalculated every iteration using the *value function* and is dependent on the shipments contained within the mission. As described in *Section 4.1.2,* the value of each shipment is dependent on the time remaining until the end of the shipment's time window as well as the shipment's priority and size.

Variables *y* and *z* are derived variables based on the missions assigned. Decision variable *y* defines how many shipments are being delivered. Decision variable *z* defines the total flight time required.

The decision variable $usedAcft_a$ is defined to be one if any mission is assigned to a specific aircraft *a*. Therefore, the model will choose to use as few aircraft as possible if not all aircraft are needed. By minimizing mission start time, $B_{ap|p=n}$, missions are all assigned as soon as possible on the schedule.

We set the objective function so that delivering high-valued shipments and delivering more shipments overall is much more important than efficiency , as measured by total flight hours. However, we can modify objective function coefficients to reflect different priorities and levels of importance for each term of the objective function.

5.3 Parameters

As described in *Section 4.1.1*, there are model parameters that are set at initialization. The purpose of this section is to examine the effects of changing some of these parameters. In our testing, we change *max periods scheduled per iteration (parameter P3), max cargo set size* *(parameter P4), working list percent of possible shipments (parameter P6),* and *working list value function (parameter P7).* The effects due to these changes are summarized in the following sections. We change each parameter individually and compare results against the results from the baseline setting shown in *Table 5-1*. We compare both operational metrics and computational performance.

Parameters	Baseline
P1: CDD	18 (hrs)
P2: Aircraf Available	9
P3: Max Periods Shceduled per Iteration	1
P4: Max Cargo Set Size	4
P5: Excess Slack Time Available	2 (hrs)
P6: Working List % of Possible Shipments	125%
P7: Working List Value Function	**FCT-2**
P8: IP-LP Relative Gap	0%

Table 5-1: Model Parameter Descriptions and Baseline Settings

5.4 Greedy Heuristic Evaluated on Historic Operational Data

We define **SCENARIO-1** as the shipment data for a week of TDD historical operational data from August 2007, as described in *Section 3.3.2*. Using baseline settings, we ran our greedy heuristic **GRD-1** on **SCENARIO-1**. We see from the results, shown in *Table 5-2*, that all shipments are delivered and the mean finish time for aircraft is during the seventh day of the scenario.

Efficiency		Greedy-2
Total Flight Time	(hrs)	607.15
Mean Aircraft Finish Day	(days)	7.61
Flight Legs		437
Positioning Legs		165
Effectiveness		
Priority 3		
Percent Delivered		100%
Mean Delay	(hrs)	10.65
Priority 2 Shipments		
Percent Delivered		100%
Mean Delay	(hrs)	9.97
Priority 1 Shipments		
Percent Delivered		100%
Mean Delay	(hrs)	17.12
Additional Metrics		
Run Time	(sec)	316.20
Mean Mission Duration	(hrs)	13.94

Table 5-2: Greedy Heuristic Overall Metrics

In *Figure 5-1* we show the distribution of hold time by each shipment. The lines represent the upper and lower quartiles of responses and the boxes represent the inner quartile, with the median show as an asterisk.

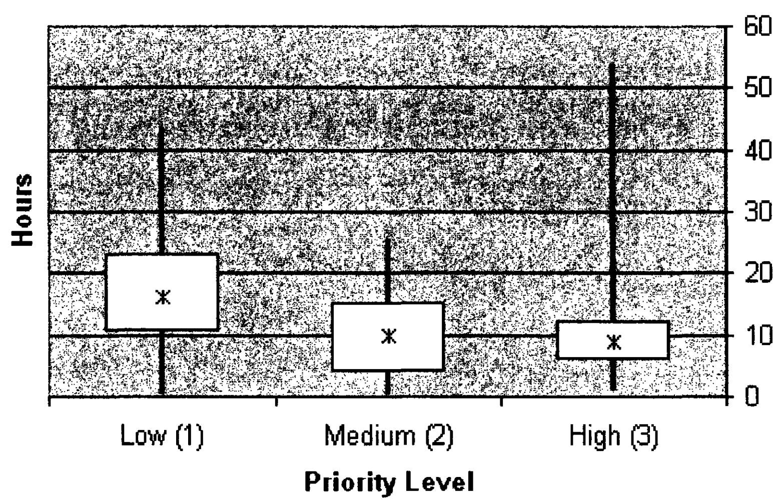

Figure 5-1: Boxplots for each Shipment in Hours by Priority Level

5.5 Greedy Heuristic vs. Optimized Model

Hypothesis: The optimization-based model will improve operational metrics.

Results: Comparing the greedy heuristic (**GRD-1**) to the optimized model (**OPT-1**), we see improvements in both the shipment hold time and flight times. *Figure 5-2* shows that shipment hold time improved for all priority levels with the greatest improvement in terms of mean hold time of the lowest priority shipments. The maximum hold time for high priority shipments dropped from 54 hours to 25 hours. All shipments were delivered in both cases.

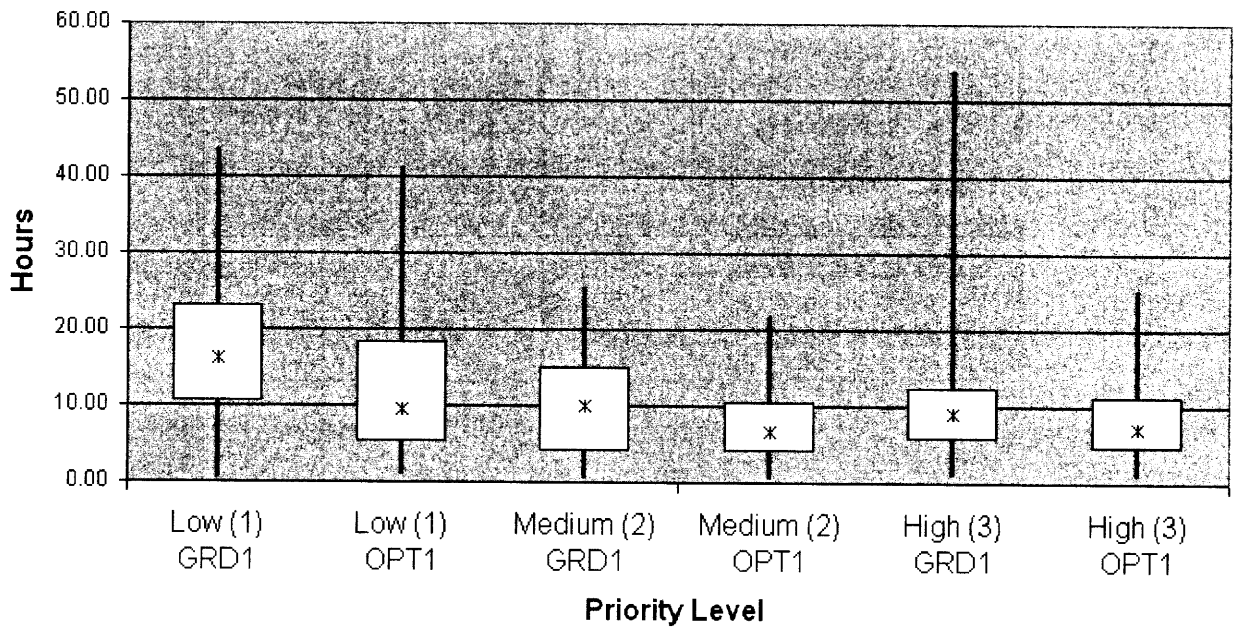

Figure 5-2: Comparison of GRD-1 and OPT-1 Shipment Hold Time by Priority Level

A more significant result than the decrease in shipment hold time is the improvement in efficiency metrics for **OPT-1** compared to **GRD-1**. Using **OPT-1** we find a 20% decrease in the number of flight legs and the corresponding total flight time. Most of this improvement is due to the large decrease in required positioning legs.

Efficiency		OPT-1	GRD-1
Total Flight Time	(hrs)	480.78	607.15
Mean Aircraft Finish Day	(days)	7.45	7.61
Flight Legs		347	437
Positioning Legs		75	165
Effectiveness			
Priority 3			
Percent Delivered		100%	100%
Mean Delay	(hrs)	8.20	10.65
Priority 2 Shipments			
Percent Delivered		100%	100%
Mean Delay	(hrs)	7.59	9.97
Priority 1 Shipments			
Percent Delivered		100%	100%
Mean Delay	(hrs)	12.22	17.12
Additional Metrics			
Run Time	(sec)	976.46	316.20
Mean Mission Duration	(hrs)	11.83	13.94
Mean Mission Slack Time	(hrs)	5.55	2.31

Table 5-3: OPT-1 vs. GRD-1 Overall Metrics

5.6 Adjusting Optimization Model Parameters

In the following sections, we discuss the different results using **OPT-1** with different parameter settings.

5.6.1 Available Aircraft (P2)

Hypothesis: As this parameter for number of available aircraft is decreased, effectiveness metrics will get worse while efficiency metrics should stay relatively constant.

Results: The decrease in the number of aircraft available had a significant effect on both sets of metrics. As expected, shipment hold times for all priority levels increased with fewer aircraft available (*Figure 5-3*). Because we put more value on higher priority shipments, the high priority shipments were not as strongly affected as lower priority shipments.

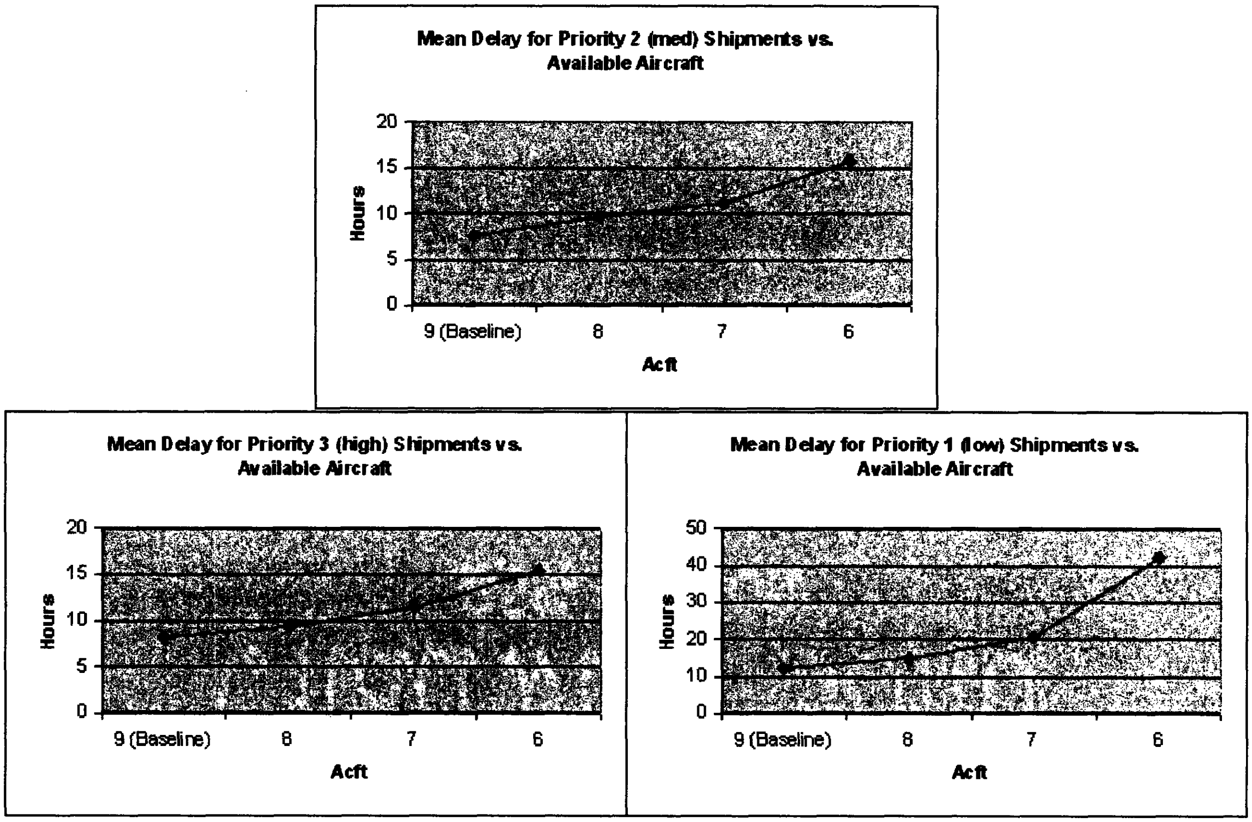

Figure 5-3: Effect of Available Aircraft on Hold Time by Shipment Priority

An unexpected result was the non-monotonicity in total flight time with the decrease in available aircraft (*Figure 5-4*). While there was less than a 1% change between nine and seven aircraft, total flight time for six aircraft increased 3.5% compared to the baseline. The slight initial decrease makes sense due to the baseline model having enough slack in the schedule to plan missions that are more efficient, but deliver some shipments sooner than otherwise possible.

The increase in flight time for six aircraft is mainly caused by more shipments being constrained by their RDD. This further constrains the missions that must be chosen to cover all shipments within their time windows.

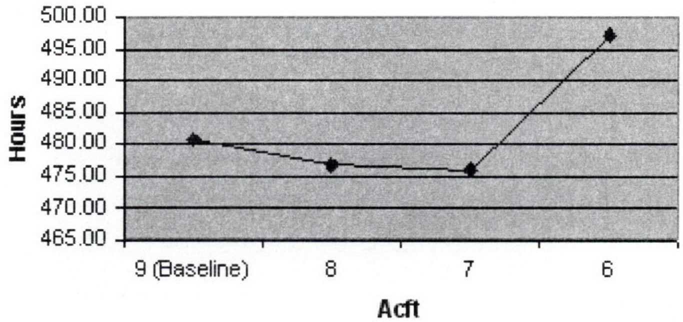

Figure 5-4: Total Flight Time vs. Available Aircraft

Using six aircraft, our model was not able to deliver all shipments within their respective time windows. As shown in *Figure 5-5*, the mean slack time between missions significantly decreased. In real-time execution, this change could cause significant problems as mission delays affect subsequently scheduled missions. With little slack in the schedule, more re-planning is necessary and some missions are cancelled all together. There is a constant tradeoff in terms of reduced slack time caused by increased aircraft utilization versus increased slack time to increase overall schedule robustness. The most important metric is successfully delivering all required shipments.

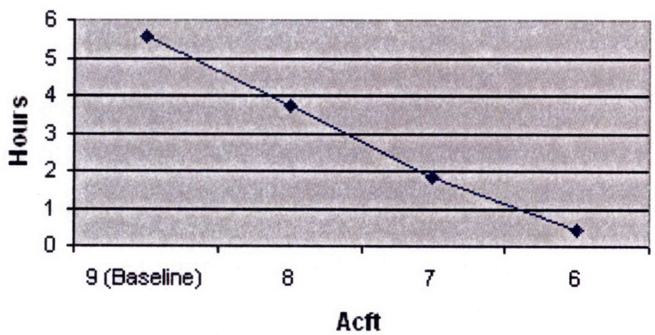

Figure 5-5: Mission Slack vs. Available Aircraft

5.6.2 Max Cargo Set Size: Number of shipments considered per mission (P4)

Hypothesis: As this control parameter is increased, more efficient composite variable missions can be built, improving operational metrics while significantly increasing run times due to combinatorial complexity.

Results: We were unable to examine the results of our model setting the max cargo set size (P4) above four. When setting P4 to five, our computer ran out of memory during the first iteration while trying to construct all possible missions during composite variable mission generation, only producing the number of variables generated.

We can examine the number of composite variable missions generated setting P4 from two to five. As shown in *Figure 5-6*, the number of possible missions grows exponentially while the number of feasible missions grows at a slower rate. As P4 is increased, fewer missions are feasible due to mission duration restrictions. Of the 3.5 million missions generated when P4 is set to five, only 72,000 are feasible (*Table 5-4*).

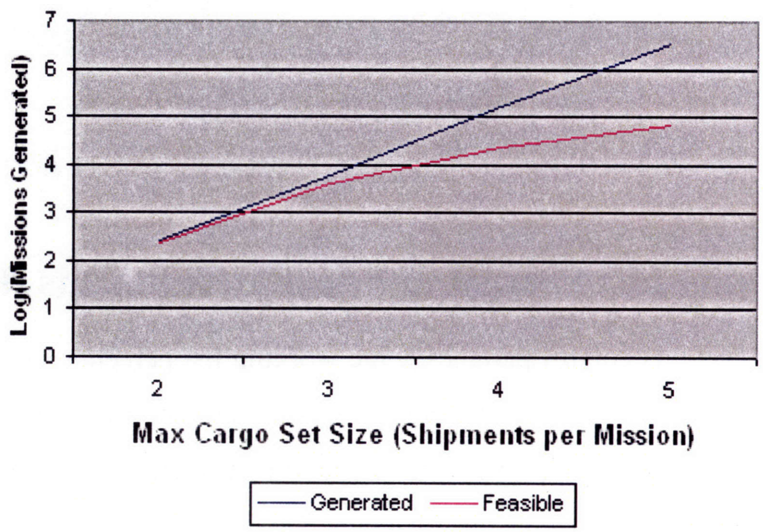

Figure 5-6: Logarithmic Comparison of Generated Mission Variables vs. Feasible Mission Variables using different Max Cargo Size (P4) during Single Iteration

	Total Msns	Feasible Msns
2	253	247
3	6017	4200
4	164220	23700
5	3505050	72000

Table 5-4: Generated vs. Feasible Missions using different Max Cargo Size (P4) during Single Iteration

As expected, operational metrics improved as P4 increased. The improvement stemmed from missions containing more shipments, which resulted in far fewer required positioning legs to deliver all shipments (Figure 5-7).

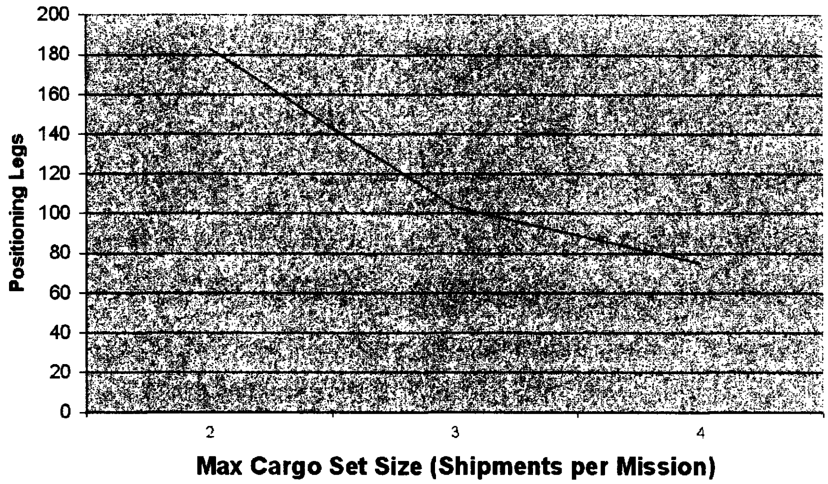

Figure 5-7: Positioning Legs vs. Max Cargo Set Size (P4)

5.6.3 Working List Percent of Possible Shipments (P6)

Hypothesis: As this control parameter, the working list percent of possible shipments, is increased, there are more shipments and therefore missions to choose from during each iteration of the MIP, translating into an improvement in operational metrics while increasing run times due to a larger sample space.

Results: There were not any strong trends in performance as we increased this parameter except for the increase in run time due to the additional model and variable generation complexity at each iteration of the model (*Table 5-5* and *Figure 5-8*).

Efficiency		100%	125% (Baseline)	150%
Total Flight Time	(hrs)	496.62	480.78	486.18
Mean Aircraft Finish Day	(days)	7.49	7.45	7.50
Flight Legs		351	347	346
Positioning Legs		79	75	74
Effectiveness				
Priority 3				
Percent Delivered		100%	100%	100%
Mean Delay	(hrs)	8.22	8.20	8.71
Priority 2 Shipments				
Percent Delivered		100%	100%	100%
Mean Delay	(hrs)	8.01	7.59	8.46
Priority 1 Shipments				
Percent Delivered		100%	100%	100%
Mean Delay	(hrs)	14.47	12.22	11.71
Additional Metrics				
Run Time	(sec)	565.34	976.46	1444.60
Mean Mission Duration	(hrs)	11.77	11.83	11.71
Mean Mission Slack	(hrs)	5.13	5.55	4.97

Table 5-5: General Metrics while Adjusting Working List Percent of Possible Shipments

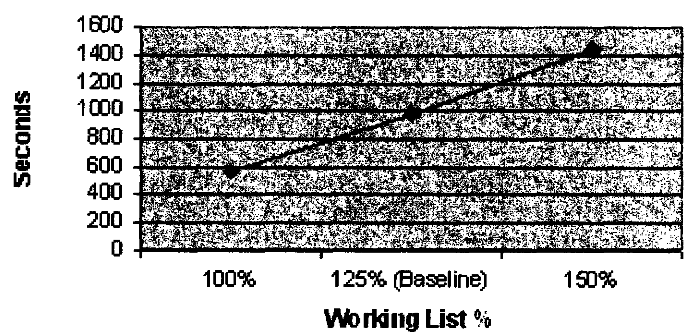

Figure 5-8: Run Time vs. Working List %

Further experimentation with different *value functions* might produce different results when adjusting the working list percent parameter. Using a *value function* with fewer weight on shipments that have an immediate RDD, the model could pick different available shipments if the working list percent parameter is increased above 100%.

5.6.4 Value Function: FCT-1 vs. FCT-2 (P7)

Hypothesis: As the *value function* is changed, there could be differences in the overall solution due to changes in function parameters. In our case, FCT-1 only sets value based on time until a shipment's RDD and does not take into account a shipment's priority. FCT-2 places a much higher weight on shipments due within the next 48 hours and also takes into account a shipment's priority. FCT-1 will result in improved efficiency metrics compared to FCT-2. FCT-2 will result in lower hold time for higher priority shipments compared to FCT-1.

Results: Contrary to our hypothesis, FCT-1 did not result in higher efficiency metrics in terms of flight time or number of flight legs. As shown in *Table 5-6*, the efficiency metrics for FCT-1 were actually slightly worse. Using FCT-2 did indeed lower hold time for high priority shipments, as well as all other shipments, but not a significant amount.

Efficiency		FCT-1	FCT-2
Total Flight Time	(hrs)	490.22	480.78
Mean Aircraft Finish Day	(days)	7.47	7.45
Flight Legs		352	347.00
Positioning Legs		80	75.00
Effectiveness			
Priority 3			
Percent Delivered		100%	100%
Mean Delay	(hrs)	8.46	8.20
Priority 2 Shipments			
Percent Delivered		100%	100%
Mean Delay	(hrs)	7.86	7.59
Priority 1 Shipments			
Percent Delivered		100%	100%
Mean Delay	(hrs)	13.40	12.22
Additional Metrics			
Run Time	(sec)	775.64	976.46
Mean Mission Duration	(hrs)	12.20	11.83
Mean Mission Duration	(hrs)	5.42	5.55

Table 5-6: Value FCT-1 vs. FCT-2

FCT-1 does not place a high premium on how soon a shipment is due and does not consider a shipment's priority level. We hypothesized this would result in more instances in which alternate shipment combinations were considered to build missions. Although we hypothesized that using the two different value functions would have an effect on the results of the model, we did not observe any significant changes.

5.6.5 Max Periods Scheduled per Iteration: 1 vs. 2 periods (P3)

Hypothesis: As the maximum periods scheduled per iteration control parameter is increased, we expect operational metrics to improve and run times to increase due to problem complexity.

Results: We compared setting parameter P3 to one and two periods with seven aircraft and an IP-LP relative tolerance of 1% to improve solution time for the two period case. By changing our settings to multiple periods, we consider scheduling subsequent missions for each aircraft, up to the number of periods, at each iteration of the model. We hypothesized that using multiple periods for our model should improve operational metrics overall, but this was not the case.

The change in P3 demonstrates the trade-offs between shipment hold time, flight time, and overall aircraft finish time. As shown in Table 5-7, using two periods, as opposed to a single period, resulted in more efficient scheduling of aircraft missions, using fewer positioning legs, with a 5.8% reduction in overall flight time.

		Periods	
Efficiency		**2**	**1**
Total Flight Time	(hrs)	524.67	557.12
Mean Aircraft Finish Day	(days)	8.91	8.12
Flight Legs		360	383
Positioning Legs		88	111
Effectiveness			
Priority 3			
Percent Delivered		100%	100%
Mean Delay	(hrs)	13.37	13.68
Priority 2 Shipments			
Percent Delivered		100%	100%
Mean Delay	(hrs)	14.44	12.80
Priority 1 Shipments			
Percent Delivered		100%	100%
Mean Delay	(hrs)	74.24	33.27
Additional Metrics			
Run Time	(sec)	13383.30	78.54
Mean Mission Duration	(hrs)	10.38	11.07
Mean Mission Slack	(hrs)	2.65	0.50

Table 5-7: 2 vs. 1 Period(s) Scheduled per Iteration

However, along with increased efficiency were additional hold times for priority 1 cargo (*Figure 5-9*). Using two periods, the overall aircraft finish time increased by over half a day, representing the amount of time during which another mission could be flown. Also important is the significant difference in run times. Using a single period, the model took a little over a minute to finish, while using two periods the model took nearly 4 hours.

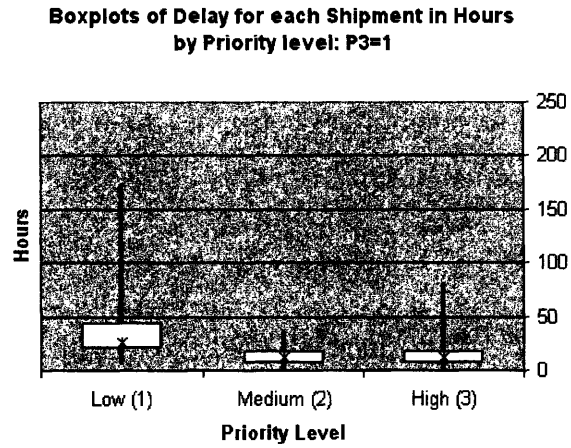

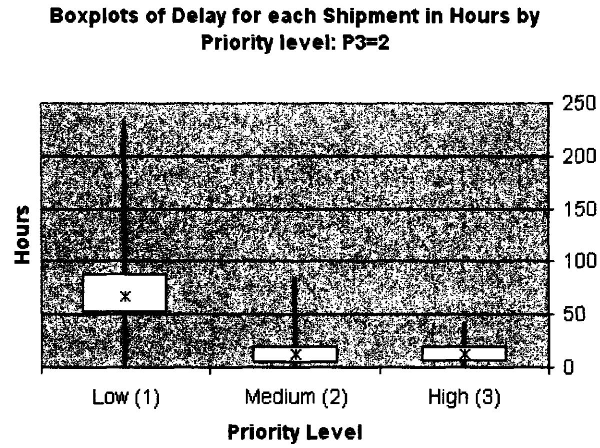

Figure 5-9: Comparison of Shipment Hold Time by Priority Level for P3=2 vs. P3=1

5.7 Summary

In this chapter, we discuss the benefits of using an optimized model versus a greedy heuristic in the mission scheduling process for TDD. The optimized approach provides improvements in both efficiency and effectiveness. We also explore how some of the possible adjustments to the optimized model affected both operational performance and model run times. Operationally, there are advantages to incorporating tools, such as our optimized model, into the planning process. Our model can produce initial results much faster than can be done by hand and with improved operational metrics. Implementing our model gives planners the opportunity to compare different possible solutions and can improve the overall process.

Chapter 6 Summary and Future Work

6.1 Thesis Summary

In this thesis we explore theater airlift operations for CENTCOM involving C-17 aircraft. We describe the current mission planning and execution process used by Theater Direct Delivery (TDD). After discussing important metrics, we introduce a model to improve the mission planning process. This model gives planners an initial solution for the C-17 fleet stationed at Qatar.

We analyze this model, comparing it to a greedy heuristic, and discuss the results due to changes in model parameters. The solution to the optimized model showed improvements in both efficiency and effectiveness metrics compared to the solution generated using a heuristic scheduling approach.

6.2 Future Work

The contributions in this thesis only represent one step in better solving problems such as planning military airlift. We summarize some recommendations for improvements and extensions as future work.

- *Incorporate dynamic shipment arrivals.* For our models, we did not incorporate the time in system (TIS), instead assuming that all shipments were visible in the system at the beginning of the scenario. By incorporating TIS for each shipment and then comparing the performance of operational metrics compared to scenarios with the assumption of perfect information, we could determine the decrease in performance due to uncertainty of upcoming shipments. We could test strategies to improve plans taking into account this uncertainty.
- *Implement ability to re-plan currently scheduled missions due to new information.* When shipments are given as dynamic arrivals into the system, the idea of re-planning based on new information becomes a requirement.
- *Incorporate and evaluate robustness in terms of aircraft availability.* In our models, we explore changing the number of available aircraft, but assume aircraft availability is known a priori and does not change. However, aircraft availability is constantly changing within TDD due to maintenance or higher priority taskings. A better approach might be to incorporate robustness using Bertsimas/Sim [15], or other, approaches to handle uncertainty in the number of available aircraft.
- *Explore improvements in composite variable generation.* Limits in variable generation hampered our model. As mentioned in *Section 4.3.2*, breadth first searches incorporating minimum forbidden sets could improve variable generation time and memory usage by eliminating the need to generate the entire set of possible missions, of which only a small percentage are feasible.
- *Extend planning to include commercial contracts and other organic aircraft in theater.* Our model only looks at TDD scheduled C-17s. This is only a fraction of all airlift scheduled in CENTCOM. Considering scenarios with other organic aircraft, as well as commercial contracts, provides a much more integrated plan. Improvements in modeling techniques are necessary to incorporate more aircraft.

- *Implement additional operational constraints.* Our model did not consider either maximum on ground (MOG) capacity at each port or port operating hours. Implementing these operational constraints into the model provides a solution that better deals with the complexities addressed daily by TDD planners.
- *Improve mission routing and shipment loading.* One way to improve our model is to consider more options during mission variable generation. We describe a simple improvement in *Section 4.3.1*that finds the optimal routing given the same rules for shipment loading. This idea can be expanded into a formulation that fits multiple shipments onto a single aircraft at one time, if possible, and considers transshipment.

Appendix A: Glossary of Terms

ALCT	Airlift Control Team
ALD	Available Load Date
AMC	Air Mobility Command
AMD	Air Mobility Division
APOD	Aerial Port of Debarkation
APOE	Aerial Port of Embarkation
ATEM	Air Tasking and Efficiency Model
CAMPS	Consolidated Air Mobility Planning System
CAOC	CENTCOM Air Operations Center
CDD	Crew Duty Day
CDDOC	CENTCOM Deployment Distribution Operations Center
CENTCOM	(United States) Central Command
DIRLAUTH	Direct Liaison Authorized
EMTF	Expeditionary Mobility Task Force
GATES	Global Air Transportation Execution System
GDSS	Global Decision Support System
ICAO	International Civil Aviation Organization
ITARS	Intratheater Airlift Request System
Msn(s)	mission(s)
MOG	Maximum on Ground
MFS	Minimum Forbidden Set
PAL	pallets
PAX	passengers
RDD	Required Delivery Date
TDD	Theater Direct Delivery
USTRANSCOM	(United States) Transportation Command
TACC	Tanker Airlift Command Center
XOC	Command and Control Directorate (TACC)

XOZ Director of Operations (TACC)

References

[1] Air Force Doctrine Document 2, (Organization and Employment of Aerospace Power). US Air Force Publications, February 2000.

[2] Air Force Doctrine Document 2-6: Air Mobility Operations. US Air Force Publications, March 2006.

[3] Air Force Pamphlet AFPAM 10-1403. 1 June 1997. US Air Force Publications. Dec 2006. <http://www.fas.org/man/dod-101/usaf/docs/afpam10-1403.htm>

[4] Air Mobility Command Fact Sheet. 18th Air Force. Oct 2006. <http://www.amc.af.mil/library/factsheets/factsheet.asp?id=231>

[5] "Doctrinal Implications of the Joint Deployment Operations Center." JWFC Doctrine Pamphlet 8. 10 Feb 2006.

[6] US Military Dictionary: "direct liaison authorized". Feb 2007 <http://usmilitary.about.com/od/glossarytermsd/g/dirliaauth.htm>

[7] Miles, Donna. "Air Force Staff Restructures to Improve Joint Ops, Communication." 30 Jan 2006. American Forces Press Service. Mar 2007. <http://www.defenselink.mil/news/Jan2006/20060130_4052.html>

[8] C-17 Globemaster III. Global Security. Feb 2007. <http://www.globalsecurity.org/military/systems/aircraft/c-17-pics.htm>

[9] C-17 Globemaster III Tactical Transport Aircraft, USA. Air Force Technology. Dec 2006. <http://www.airforce-technology.com/projects/c17/>

[10] Google Earth. Version 3.0.0762. 28 Mar 2007.

[11] Theater Direct Delivery Training Document (draft). Global Mobility Air and Space Operations Detachment. Scott AFB, IL. 2006.

[12] Greenlee, Paul. "C-17 TDD Reachback Cell Position Training." Air Mobility Command, Tanker Airlift Control Center, Presentation (Powerpoint).

[13] Greenlee, Paul. Personal Interview. 25 Aug 2006. TDD Leadership.

[14] Weston, Fredrick. Personal Interview. 25 Aug 2006 TDD Leadership.

[15] A. Armacost, C. Barnhart, and K. Ware, "Composite Variable Formulations for Express Shipment Service Network Design," Transportation Science, Focused issue on freight transportation, Vol. 36, No. 1, pp. 1-20, February 2002.

[16] J. Bridgers. "Planning US CENTCOM Intratheater Airlift Routes." Naval Post-Graduate School. September 2006.

[17] D. Bertsimas and M. Sim. "The price of robustness", Operations Research 52 (2004) 35-53.

[18] D. Bertsimas and M. Sim. "Robust discrete optimization and network flows", Mathematical Programming 98 (2003) 49-71.

[19] D. Bertsimas and J. Tsitsiklis. *Introduction to Linear Optimization.* Belmont, Massachusetts: Athena Scientific, 1997.

[20] Nielsen, Christopher. "Large-Scale Network Design using Composite Variables: An Application to Air Mobility Command's 30-day Channel Route Network." MIT Thesis, 2002.

[21] K. Ravindra, T. Ahuja, J. Magnanti, and B. Orlin. *Nework Flows: Theory, Algorithms, and Applications.* Prentis Hall, Jan 1993.

[22] S, Ropke, J. Cordeau, and G. Laporte. "Models and a Branch-and-cut algorithm for Pickup and Delivery Problems with Time Windows". July 14, 2005.

[23] M. Uetz. "Minimal Forbidden Sets – On the Representation of Resource Constraints in Project Scheduling", *10th DFG-Workshop on Resource–Constrained Project Scheduling*, Technische Universität Berlin, Germany (April 2 – 4, 2001)