DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT/GMO/XO/96J-3

ARTIFICIAL INTELLIGENCE: CAN EXPERT SYSTEMS TECHNOLOGY ENHANCE AIR MOBILITY COMMAND (AMC) DEPLOYMENTS?

GRADUATE RESEARCH PAPER

Joseph A. Jackson, Major, USAF

AFIT/GMO/XO/96J-3

19960617 141

AFIT/GMO/XO/96J-3

ARTIFICIAL INTELLIGENCE: CAN EXPERT SYSTEMS TECHNOLOGY ENHANCE AIR MOBILITY COMMAND (AMC) DEPLOYMENTS?

GRADUATE RESEARCH PAPER

Presented to the Faculty of the Graduate School of
Logistics and Acquisition Management
of the Air Force Institute of Technology
Air University
in Partial Fulfillment of the Requirements for the
Degree of Master of Air Mobility

Joseph A. Jackson, B.S.
Major, USAF

May 1996

Approved for public release; distribution unlimited

Acknowledgments

I would like to thank the Air Mobility Warfare Center and the AFIT staff for making the Advanced Study of Air Mobility (ASAM) program possible. Combining these two assets to provide a professional development course that also incorporates a master's degree, specifically designed for mobility, has made me better prepared to take on the challenges facing today's military.

As usual, the family tends to give the most and complain the least when projects stretch on into the wee hours of the morning. I can not thank my wife, Kay, enough for the understanding and encouragement she gave me to finish this paper and the ASAM program.

Joseph A. Jackson

Table of Contents

Page

List of Figures

Abstract

With the military restructured to react via rapid response from CONUS as opposed to forward presence overseas, Air Mobility Command has been charged to provide the Global Reach needed to project power abroad. The Air Mobility Operations Group (AMOG) structure and the Integral Tanker Unit Deployment (ITUD) concept, used to deploy the initial airlift and air refueling assets forward, were studied to determine the viability of using expert systems technology to aid in matching austere location demands with the proper mix of personnel, equipment and supplies. The paper starts with a thorough background of artificial intelligence and gives two rudimentary examples of working expert systems to show how they differ from conventional software programs. An overview of the AMOG and ITUD then identifies the intended beneficiaries of such a system. Alternative sources for fielding an expert system, ranging from buying off-the-shelf, commercial software to custom designing a system to fit this specific scenario, were then explored. Though this paper does not develop an actual working expert system, it does provide a foundation and offer suggested avenues to further work towards that goal.

ARTIFICIAL INTELLIGENCE: CAN EXPERT SYSTEMS TECHNOLOGY ENHANCE AIR MOBILITY COMMAND (AMC) DEPLOYMENTS?

I. Introduction

General Issues

Though it sounds rather futuristic, the concept of artificial intelligence is not new to the military. On the contrary, it was the military that, over the last fifty years, pioneered research and actual development of smart systems (Ball, 1994:xiii, 3) that work on the principles of artificial intelligence. The lessons of World War II and, in particular, the events at Pearl Harbor convinced President Truman and other important American leaders that our antiquated defense structure had to be reorganized to ensure adequate unified military command and control of our forces. The resulting National Security Act of 1947 coupled with the American tendency to seek technological solutions to its problems sparked much of this fifty years of work centering around the military's command, control, communication and intelligence (C^3I) apparatus (Ball, 1994:3; Wolk, 1975:10-18).

As with many other technologies developed by the military, artificial intelligence applications have migrated to the civilian sector to be used in the business arena. Private industries have taken the basic framework and, because of quantum leaps in computer technology over the last few years, were able to develop advanced smart systems that are both productive and economical (Ball, 1994:32-34; Winston and Prendergast, 1984:1-13). Expert systems, normally the logic generating subsystems of the larger and more encompassing smart systems, are being commercially produced as stand alone modules for smaller applications where most of the variables can be accounted for in the form of databases and user defined inputs (Ball, 1994:12-13). Tens of thousands of operational expert systems are used either as part of an overall smart system or as a stand alone package dealing with a more narrowly defined issue (Ralston and Reilly, 1993:538). Because of this, an entire category of commercial software development has emerged that offers these systems ready-to-use right off-the-shelf or custom designed to fit a user's specific needs. Though this technology has potential application in nearly every facet of the military (Ball, 1994:204), this paper addresses its use in deploying Air Mobility Command assets around the world.

Problem Statement

Air Mobility Command often provides the first cadre to arrive at austere locations when the United States decides to project forces forward. In this role, Air Mobility Command is not fully exploiting available technologies that could ensure the proper personnel and equipment are deployed forward to match the ever increasing and seemingly endless combinations of possible operating environments. With current technology, a user level expert system could be developed to enhance the decision making process for determining the right mix of personnel, equipment and supplies needed to sustain this initial cadre in these unfamiliar locations for a given time frame.

New force structuring has altered the way the United States employs its military machine. As America pulls its forces home, it no longer fields large standing armies prepositioned throughout the world ready to respond at a moments notice. To make up the difference, Air Mobility Command is chartered to provide the rapid Global Reach capability needed to project Global Power. However, part of the pull back led to the withdrawal of the once robust, in-place, world-wide infrastructure needed to support this effort. In its place strategies exist to deploy temporary frameworks forward when the needs arise.

Given ample warning time and an abundance of knowledgeable and experienced decision makers at the unit level, the current heuristic or judgmental methods used to the determine personnel, equipment and supplies needed for the forward location is sufficient. However, imminent crisis combined with a shrinking human experience base may render this approach incapable of providing the best package sizing requirements for Air Mobility Command force employments. A well designed expert system can capture current decision making processes to preserve the human knowledge base and provide continuity for the future (Ralston and Reilly, 1993:537).

Importance of Research

The Air Force and the entire Department of Defense have embraced the quality concept, making it the foundation of every organization. Consequently, decision making has become influenced by quality methods. At the very heart of quality is the idea of continual process improvement. In theory there never is a single best solution or method of accomplishing a goal or task; any procedure that is designed can be improved upon. From a quality perspective it is mandated that an organization improve its processes on a continual basis.

Improvement, however, comes with a price tag and it is important to remember that the benefit gained must outweigh the cost incurred if an initiative is to truly be considered an improvement. For time, money, and effort to be invested in making something better, the ultimate outcome should eventually pay for itself, be it in dollar form or through mission effectiveness. If this is not the case, there is no sense expending the resources in the first place. From the benefits side, two specific examples help illustrate that there is room for improvement.

The author first encountered a problem in this particular area while deploying to Pisa, Italy, as an operations officer for Operation Deny Flight on the unit's first Integral Tanker Unit Deployment (ITUD). Though the KC-135 had been deployed around the world many times before, this was the first time the entire tanker squadron would be temporarily relocated to a forward location without the support of a Tanker Task Force. Squadrons deploying under the new ITUD concept would now be responsible for almost all maintenance, command and control, and support needs.

When it was time to size up the package, heads were scratched for a while and it was quickly realized that a serious lack of experience at the tanker unit level existed in dealing with this kind of operation. The number of airplanes, crews and maintenance personnel needed were mandated from headquarters, but the rest was for the unit to decide. Many questions were asked, like, "What will the electrical needs be? Will converters be required and, if so, how many will be needed? Is the water potable or does that need to be provided as well?"

Luckily, with about three months advance warning, ample time existed to find most of the answers to these questions. Even so, after arriving it was discovered that many things were left behind and others were brought in insufficient quantities. On the fortunate side, again, the location was such that nearly all of the overlooked needs were readily available from the local economy. These two commodities, time and favorable location, were not available in the next example.

In the fall of 1994, an air refueling squadron received orders to deploy to an austere location with a very short lead time. Within hours of the directive, they found themselves launching to near bare-base conditions with only the equipment and personnel they could carry on board their own airplanes; no outside airlift assets were available. Though tons of cargo and personnel were carried, this unit found itself ill-equipped for what awaited them. Upon landing it was discovered that no Material Handling Equipment (MHE), transportation vehicles, useable jet fuel, nor sleeping quarters within fifty miles were available. The latter presented a critical challenge since the deploying unit had not been directed nor were they in a position to foresee the need to bring tents, sleeping bags, or blankets (Hunter, 1995:1-4).

The above example may appear to be an extreme, but the changing size and shape of our military, coupled with the changing world order, means that this scenario will probably be repeated if a method is not found to quickly determine the needs for a variety of possible scenarios. Downsizing means a smaller but more mobile force is needed to cope with future conflicts. Therefore, rapid deployment is and will continue to be the response mechanism to cope with world conditions requiring our presence. The result of continuing with the status quo is either a failed mission because a unit shows up unprepared and ill equipped or we lose the critical element of speed in delivering our forces to a hot spot.

After experiencing the first example and reading about the second, the author looked to the experts in world-wide employment of U.S. forces, the traditional airlifters. They have been doing these types of operations for many years. As early as the China-Burma-India operation in World War II, U.S. airlift has been deployed to austere and remote regions to project an American presence (Tunner, 1964). Discovering that the airlift portion of Air Mobility still uses a better but purely heuristic approach to determine their particular needs sparked the idea that a system incorporating new technologies might prove beneficial. Shortfalls of an approach relying almost entirely on having the advantage of corporate knowledge were demonstrated in the previous two examples.

Background

On June 1, 1992, an historic event occurred for the United States Air Force. On this date Air Mobility Command was formed by melding Military Airlift Command (MAC) with the tanker forces from Strategic Air Command (SAC) (Gant, 1993:v). Initial shock and turmoil was caused by the necessity for the United States to update its force structure to cope with the challenges ushered in with the new world order abroad and the ever increasing budgetary constraints at home.

Though in a painful and clumsy manner at first, the new command continues to deliver airlift and air refueling to all parts of the world, only now, on the strategic scale, it is done in unison with one another. Combining airlift and air refueling assets provides the force extension required to project Global Reach and Global Power. This conscientious decision has become an integral part of our overall military doctrine as the United States pulled its reduced military further back within the confines of its borders (Cirafici, 1995:2; Mathews and Leland, 1995:41). Adding aerial refueling to the airlift picture gives military forces the legs required to reach all points on the globe from starting points within the continental United States (Mathews and Leland, 1995:42-45).

Prior to this union, separate methods for employing airlift and air refueling assets worldwide were developed by their respective commands. Staging from austere locations always played a part in the airlift mission and has recently taken on even greater importance for the reasons stated above. In the past, refueling assets were rotated overseas to augment Tanker Task Forces, but rarely operated from remote locations without full support provided by the Tanker Task Force and almost never operated out of truly austere environments. After the two commands were combined in response to the overall force structure changes, the airlift and the air refueling communities continued to exclusively adapt and improve their processes for worldwide deployment. Currently, air refueling assets deploy using the Integral Tanker Unit Deployment (ITUD) concept and airlift assets stage using the Air Mobility Operations Group (AMOG) concept, both of which will be explained in the next chapter.

How different are these two missions? In the purest sense, air refueling is nothing more than delivering a load, liquid cargo in this case, from point A to point B. Point B for the air refueling mission is usually at an altitude of several thousand feet, but generically speaking, it can be considered a cargo run. The exact mission parameters are different with varied doctrine guiding air refueling versus airlift, but certain basic functions remain the same. More importantly, personnel supporting either mission at a deployed location need transportation, messing facilities, sleeping accommodations, command and control elements, and many other common basic needs. Because of the overlap, this paper considers Air Mobility Command assets involved in ITUD and AMOG operations to be the primary users of a proposed decision making tool based on expert systems technology.

Questions to be Resolved

Although the modern concepts of artificial intelligence have been around for some fifty years, this is a field which only a select group of people really know and understand. Unless one is involved in the actual development or is a direct user of the software employing artificial intelligence, the whole concept is probably very foreign. The first major question to be resolved, then, is *what exactly is an expert system*? To answer this, one must understand how the expert system fits in to the broader field of artificial intelligence. Also, one must investigate how the expert system, itself works, at least from a rudimentary standpoint.

Once a basic understanding is obtained and it is desired to employ an expert system, the next logical question is *who, specifically, would be the primary user of such a system*? Knowing how the AMOG and ITUD are structured and operate is essential in capturing and translating the critical human decision making to a software package.

After knowing what an expert system can do and identifying the potential user, the focus turns to *what critical information and elements are needed to build a specific system that will meet the customers' needs*? To accurately mimic human decision making, the thought processes currently used to accomplish this task must be studied and understood. The required and available input and output parameters needed for such a system must also be determined.

Finally, the last question remaining is *what costs will be incurred and what benefits can be realized by implementing such a system?* As mentioned earlier, simply expending resources for the sake of expending resources does not make sense; some payback should come as a result. It is also important to keep in mind that pure dollar costs and benefits, alone, are only one facet in the equation. Hidden costs and benefits, not immediately apparent on the surface, can be substantial.

Overview of Research

An actual working expert system will not be developed due to the limited scope of this paper. Even if the author wanted to, he does not possess the vast technical knowledge and professional expertise required to do so. Rather, this chapter attempted to identify a need while the rest of the paper will try to present a case supporting the viability of an expert system to enhance decision making for Air Mobility Command deployments by answering the investigative questions presented above. The paper can also serve as a starting point by those more skilled in this area should it ever be decided that further work is warranted.

The next chapter will familiarize the reader with the basic concepts of artificial intelligence by defining the unique terminology used in this field. To give the reader an idea of how they actually work, two specific rudimentary expert systems will then be exemplified. Chapter II will conclude with a review of some of the work, performed to this date, that is pertinent to the topic of this paper.

Chapter III will begin with a brief overview of the AMOG and ITUD to identify the intended Air Mobility Command users targeted by this paper that could benefit from an expert system. For an expert system to be developed, the primary user must be identified and its internal workings understood. Knowing how the AMOG and ITUD concepts operate is important to understanding why, when, and how they would use an expert system.

The ultimate goal of an expert system is to accurately mimic actual human decision making. In this vein, Chapter III will also specifically show how the decisions of what to bring on AMOG and ITUD deployments are currently being made. This is accomplished through direct observation and interviews with those who actually make such decisions. Working backwards, the ultimate desired outcome or output of the decision making process is presented, followed by a discussion of the required and available inputs needed to reach that output.

Chapter IV briefly looks at the alternatives available, ranging from stand alone, off-the-shelf, commercial software to custom designing software fitting this specific scenario. This chapter, in no way, endorses a particular product or makes solicitations from any vendor. Rather it is presented to show the reader that the technology is, in deed, available to field a working and viable expert system.

Chapter V concludes the paper with a summary of the key points and discusses the types of costs and benefits that can be expected if an expert system were to be developed. This chapter will also serves as a launch pad to further research by pointing toward potential areas that should be explored further or in more detail.

With the growing pains of forming a new command nearly overcome, the task at hand is to fine tune the procedures to improve Air Mobility Command's ability to provide formidable Global Reach. Harnessing available and affordable technology can prove invaluable in achieving this goal.

II. Literature Review

Terminology

When discussing any field of study, one must speak the language just to understand even the very basic concepts. Getting past square one is impossible without knowing and differentiating the specific terminology used for that particular field: the realm of artificial intelligence is no different. This section will lay the foundation for understanding the rest of the paper by defining and describing key words and phrases used in this paper and those that will be found should the reader decide to explore this topic further.

As is common in other fields, writers and experts from within as well as from outside this area often confuse the associated terminology by intentionally using different words to say the same thing or by inadvertently using one word when another is more appropriate. This overlapping of terms blurs and obscures meanings, sometimes making it difficult to research this topic. It is also easy to get lost within this realm, not being able to see the forest through the trees, by confusing what level in the artificial intelligence scheme one is actually dealing with. To avoid this confusion, the most commonly accepted definitions and interpretations will be given and presented starting from the macro level, working down to the micro level. Referring to the taxonomy in Figure 1 during the following discussion will help keep the reader oriented.

First it must be understood that artificial intelligence refers to the overall subject area. In broad terms, artificial intelligence refers to the ability of an object to do the same kinds of tasks and reasoning that a thinking human could do (Britannca, 1993:605). Research in this area has gone in two major directions: one being the psychological and physiological research into the nature of human thought, and the other being the technological development of increasingly sophisticated computer systems (Britannica,1993:605). This paper concentrates on the latter, where computer systems are able to perform beyond the point of simply following a series of programming steps.

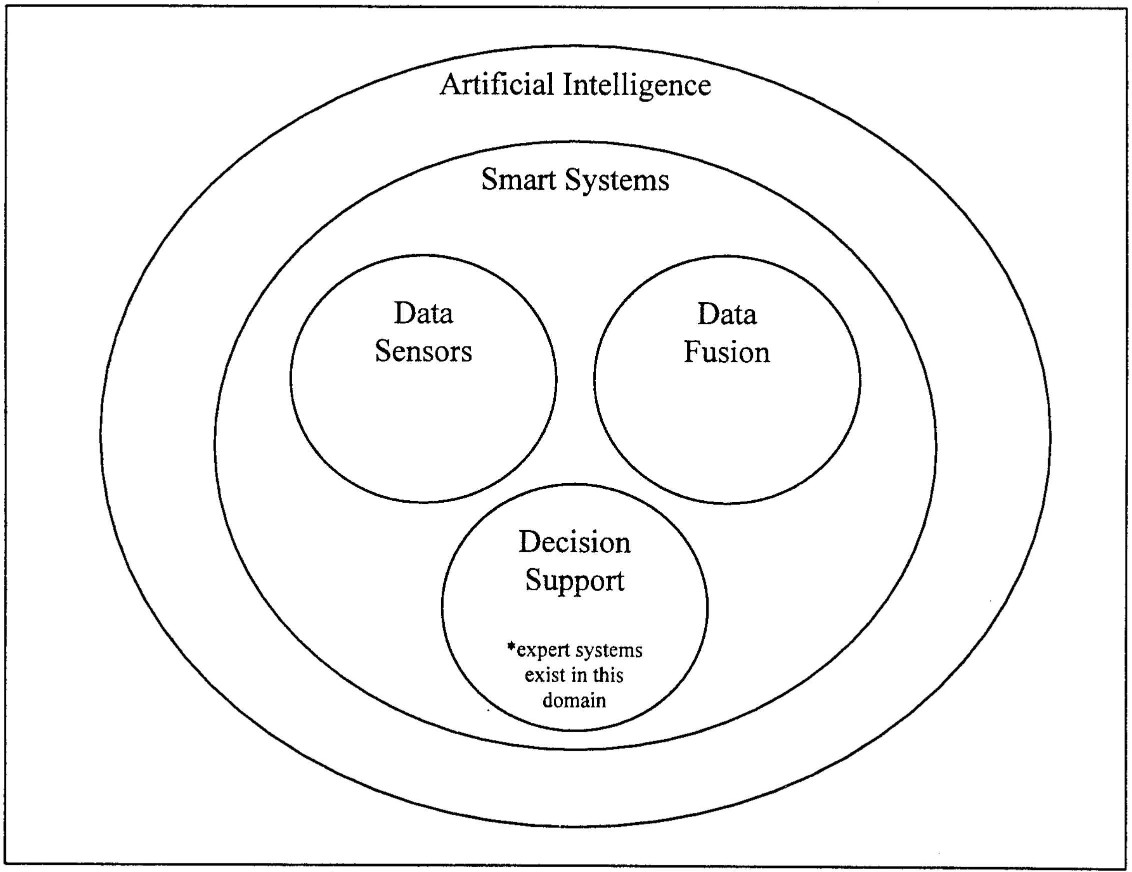

Figure 1. Realm of Artificial Intelligence

Within the broad, overall heading of artificial intelligence lies the concept of smart systems (Ball, 1994:4). As mentioned in the first chapter, Americans are enthralled with technological solutions to problems and have steadily tried to enhance the human operators' abilities to interface with the complex systems and processes that have spawned from the industrial revolution. Smart systems try to improve upon this human interface by extending and sometimes entirely replacing human decision making and actions in some of the many sub-parts to the overall mission or task . These sub-systems draw inputs by directly monitoring the operating environment, by drawing information from a stored databases, or by asking for information from human operators ultimately controlling the entire system. Obviously these smart systems must also be able to draw conclusions, make intelligent decisions, and take appropriate actions based on the input supplied to it (Ball, 1994:4). In general, smart systems are considered to exhibit the characteristics listed below.

(1) Coherent Knowledge (Ball, 1994:5) - Making up parts of the whole, sometimes many smart systems are working together to accomplish the overall task. Coherency ensures that the data used by each smart system is consistent and accessible to the others. Rather than a series of individual, proprietary databases, each module conforms to a standard format so that data can be shared across the entire system.

(2) Integrated Access (Ball, 1994:5) - Integrated access attacks the sharing of data from the opposite angle. Not only should smart systems present and store knowledge and data in a form that is useable to other parts of the system, but the smart system should also know when and where to retrieve this information. Integrated access can be thought of as the pulling of data versus the pushing of data described in the previous characteristic.

(3) Decision Analysis (Ball, 1994:5-6) - Most conventional software programs work like magic black boxes: the program is simply executed to derive some desired output with no presentation of how the results are actually obtained. Smart systems are capable of going a step further. Just as a subordinate should be able to explain, to a superior, the reasoning behind a decision, so should a smart system be able to provide its logic to a user on demand. Providing insight to decision making criteria gives smart systems a level of credibility not garnered by black box technology.

(4) Decision Execution (Ball, 1994:6) - Monitoring is a major function of most smart systems where the computer, through attached sensors, alerts a user when an abnormal condition arises somewhere in the system. Traditionally the smart system not only alerts the user, but also recommends the corrective action. Recent technology is allowing the smart system more autonomy by letting it, rather than the human operator, initiate the appropriate response and simply advising the operator of its actions. Technology is further pushing towards 'smart sensors' that can more quickly perform corrective action at the source of the problem and relay the information to the parent system.

(5) Acquired Knowledge (Ball, 1994:6) - For the human brain, knowledge is acquired through personal experiences and through witnessing or hearing of the experiences of others. Based on the results of previous and somewhat similar episodes, the brain arrives at newer and better decisions as the number of stored episodes increases. In an attempt to model the human brain, smart systems are envisioned with ever increasingly advanced programming logic and neural networks that span the entire

apparatus. Sharing experiences can make the individual smart systems intelligent to the entire system's capabilities and limitations while smarter logic will hopefully allow them to 'think better'. Common gaming schemes have already been developed that allow a computer to play at higher skill levels by drawing upon the results of previous games, demonstrating at least a limited ability to acquire knowledge.

(6) Non-Domain-Specific Problem Resolution (Ball, 1994:6) - A further extension of intelligence allows the human brain to use knowledge and experience gained in one area to reason and solve problems in a completely different one. This truly domain-crossing technology is yet to be developed for the computerized brain, but smart systems are capable of handling some of the problems not specifically anticipated by the developer. This is important since the environment that smart systems are required to operate in makes it impossible for a programmer to foresee all possible scenarios and outcomes.

Going one layer deeper in the artificial intelligence taxonomy, the six characteristics detailed above hint at the three main components or topics integral to a smart system: they are data sensors, data fusion, and decision support (Ball, 1994:3, 7-10, 45; DeJesus, 1995:63). Data sensors are to smart systems what eyes, ears, noses, mouths, and fingers are to humans (Ball, 1994:80). Like the five senses that these body parts provide, data sensors monitor and collect data, while data fusion technology is employed to transmit, categorize, sort, and store this information for immediate or later use by the system (Ball, 1994:112-117). DeJesus calls this process data acquisition and equates it to the functioning of the human nervous system (DeJesus, 1995:63).

These two topics, data sensors and data fusion, are much more involved than presented above, but it is within the third and final domain of smart systems that the expert system exists. Therefore, more time will be spent discussing the subject of decision support. According to Ball, "Any hardware/software system that provides information about its operation to the user of that system is a decision support system." (Ball, 1994:32) The goal for a decision support system, then, is to paint the clearest picture possible of the problem and potential solutions to a given scenario to be observed by any potential user. As discussed earlier, this aspect is what separates smart systems from 'black box' technologies that keep the human operator in the dark. The pictorial view of typical decision support architecture is presented in figure 2. This architecture

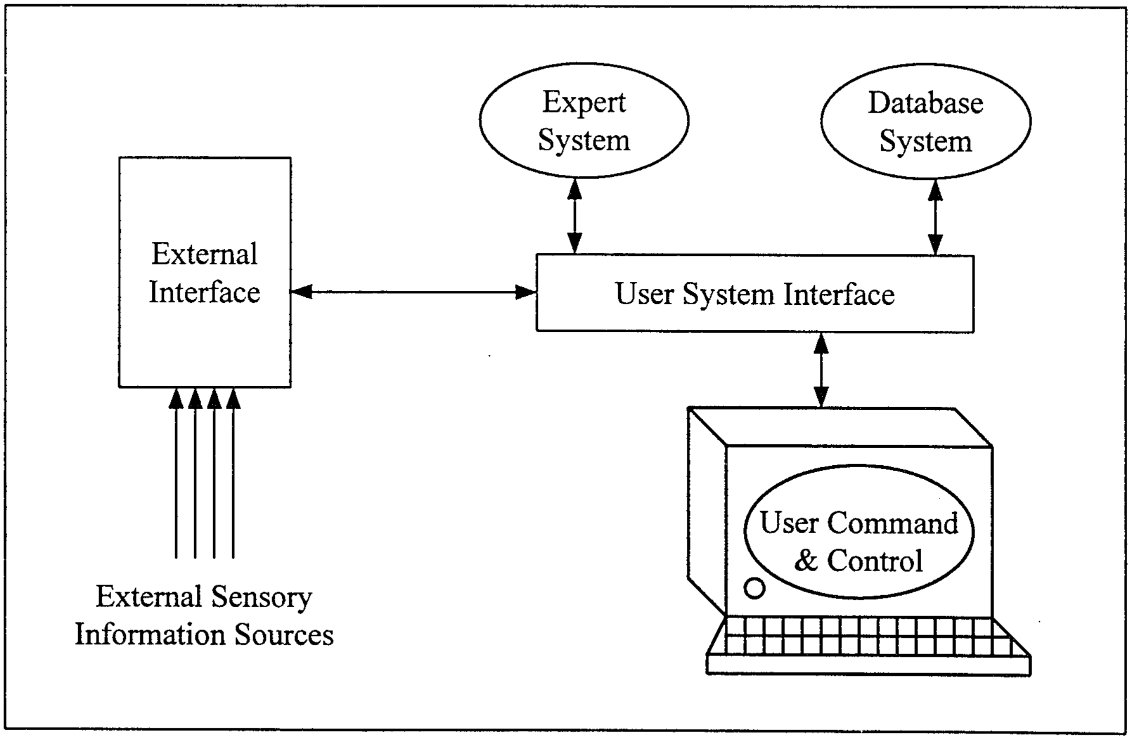

Figure 2. Typical Decision Support Architecture (Ball, 1994:35)

presents a picture of reality for the human user by translating data and decisions derived by the expert system in computer code to a form intelligible to the operator (Ball, 1994:34).

Another way of observing the topic of decision support is to review the evolution of one aspect of the information age. Ball summarizes this evolution in his book "Network Management with Smart Systems." First, electronic data processing (EDP) was used to collect data about processes. This was followed by the management information system (MIS) that organized EDP output in forms that could be used by management to aid in their decision making. The next step centered on management science and operations research (MS/OR) in an effort to aid management in sorting and making decisions from the growing mounds of highly structured MIS reports. Finally, decision support systems (DSS) take a quantum leap and offer managers a way of finding solutions with less structured or incomplete information (Ball, 1994:36). It is the expert system, the major component of a decision support system, that allows this problem solving ability.

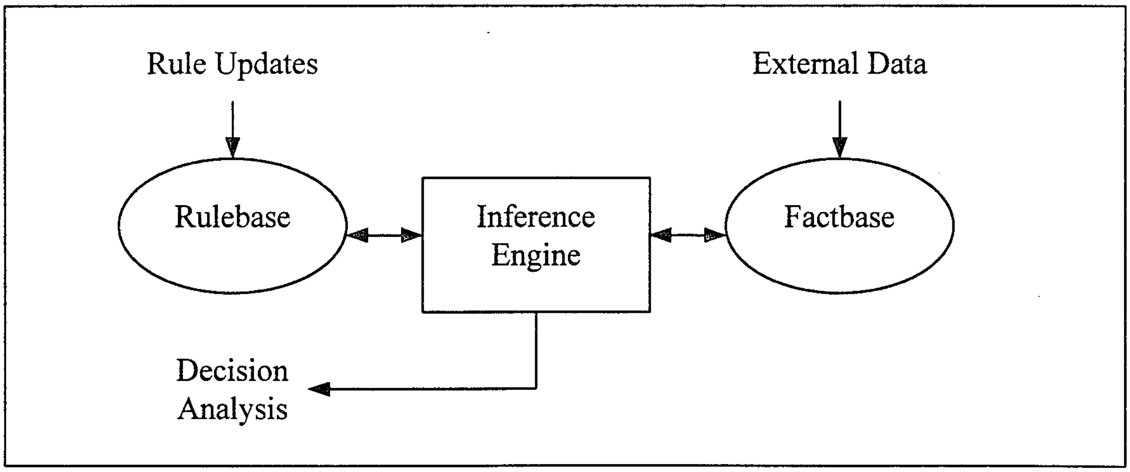

Figure 3. Rudimentary Expert System (Ball, 1994:37)

Figure 3 shows the three main parts of a rudimentary expert system; the inference engine, the rulebase, and the factbase (Ball, 1994:204-205). Ralston and Reilly lump the rulebase and the factbase together, categorizing the expert system as consisting of two principal parts; the inference engine and the knowledge base, but the idea is still the same (Ralston and Reilly, 1993:538). For the purposes of this paper, however, the former grouping consisting of three components is more appropriate.

The inference engine (Ball, 1994:37) is the actual software logic that makes decisions. It accomplishes its task through software, but goes beyond simply following a series of sequential program steps. Fuzzy logic gives the software a more human quality by allowing the expert system to distinguish shades of meanings like the difference between warm and hot or early and late (Gould, 1995:81; Ralston and Reilly, 1993:539). Expert systems differ from conventional software in another way as well: they can let the operator know how confident they are in their various decisions. This will be demonstrated later in the chapter. Six of the more popular decision making techniques used by expert systems are listed below.

(1) Production rules (Ball, 1994:183; Ralston and Reilly, 1993:538), better known as if-then-else rules, are the oldest and most common logic tool used and is familiar to even the most amateur programmer. If a given condition exists, then a preprogrammed response will result. If the condition does not exist an alternative action (or else) will result. This alternative action can be another preprogrammed response or a branch to further investigative routines.

(2) Networks (Ball, 1994:184) attempt to place objects and occurrences in a hierarchical order of importance and match decisions and outcomes with the respective rankings of the instances.

(3) Frames (Ball, 1994:184; Ralston and Reilly, 1993:538) capture an object or occurrence in its entirety and stores them for comparison against future occurrences. This is analogous to a still picture of an event like a child blowing out birthday candles. Not only is the act of blowing out the candles captured, but so is the cake, the table it sits on, the expressions of the partygoers sitting around the table, and a host of others. Any of these attributes might possibly yield a clue or aid in making a decision in the future.

(4) Scripts (Ball, 1994:184) are simply sequential playbacks of canned actions that are useful when the same procedure or response is employed each time a particular scenario is encountered. This would be similar to strictly following a checklist or a recipe to ensure consistent outcomes.

(5) Blackboards (Ball, 1994:184; Ralston and Reilly, 1993:539) in the real world are used to share ideas within a group of people. In the computer world, this same philosophy is used to let different expert systems analyze the same problem or scenario.

(6) The direct approach (Ball, 1994:186) models and compares events through analogies. When an event can not be matched exactly or close enough through the other means mentioned above, comparing it to an analogous situation can provide the perspective needed to make an intelligent decision.

Again referring to Figure 3, it can be seen that the rulebase and the factbase interact with the inference engine as it formulates decisions. As its name indicates, the rulebase contains the rules or laws that define and govern the scenario for which the expert system was designed (Ball, 1994:37). In a military application this could be operating instructions, flying regulations, standard operating procedures, or possibly international law. The factbase provides the facts or knowns about the particular problem being processed (Ball, 1994:37-38). This information is usually in the form of a database that receives its information externally from sensors (as described earlier) or from manual inputs. Examples could include airplane cargo and range information, the number of pallets to be loaded, or simply the intended destination and its attributes.

Two Rudimentary Expert Systems

Now that the foundation is laid for understanding what an expert system is, it is time to look at two specific examples. These examples are somewhat dated and very rudimentary in nature but are used in the private sector and provide a good representation of how today's even more powerful applications perform.

The first example is a program that diagnoses automobile problems by asking the user a number of investigative questions. Figure 4 is a printout of what the user would actually see while executing this program. The expert system analyzes the user-supplied responses, asks further questions based on these responses, and finally gives a number of diagnoses and treatments. A Certainty Factor (CF) tells the user how confident the system is in deriving its solution ranging from very unsure (0) to being certain (1). In this case, the expert system places almost no confidence in its first conclusion that the problem is a discharged battery. The system claims its last two conclusions, however, are

```
AUTO CONSULTATION PROGRAM VERSION TWO

This program will analyze your car problems by asking questions about the functioning
of specific areas of your automobile. Please type in one of the words that follows the
three dots in each question. Any other input will produce erroneous results.

*****QUESTION 1*****
Gasoline odor is...none normal strong? strong

*****QUESTION 2*****
Headlights are...dim normal? dim

*****QUESTION 3*****
Fuel filter is...clogged not-clogged? not-clogged

*****QUESTION 4*****
Battery cables are...loose corroded normal? loose

*****QUESTION 5*****
Starter cranking is...no-crank slow-crank normal grind? normal

*****QUESTION 6*****
Outdoor temperature is? 45

*****QUESTION 7*****
Gas gauge reading is...empty not-empty? not-empty

CONCLUSIONS:
*****CF 0.000000e + 000*****
Diagnosis is that the battery is discharged.
Treatment is to charge or replace battery.

*****CF 6.000000e - 001*****
Diagnosis is that car is flooded.
Treatment is wait ten minutes then try starting
or depress accelerator to the floor while starting.

*****CF 6.000000e - 001*****
Diagnosis is that the battery cables are in disrepair.
Treatment is to clean and tighten battery cables.
```

Figure 4. Auto Consultant Program (Ralston and Reilly, 1993:1085)

correct with sixty percent confidence that either of these treatments will solve the user's problem (Ralston and Reilly, 1993:1084-1085).

The second example presented is the expert system known as MYCIN, a medical program that diagnoses infectious diseases. This example is presented in greater detail in order to expand upon and tie into the knowledge and concepts given in the first part of this chapter. It will also demonstrate the characteristics that distinguish a true expert system from a program that simply executes lines of code.

Before launching directly into MYCIN, it is valuable to understand why and how such a system was developed. Three main factors allowed MYCIN to come into being; the overuse and misuse of antibiotics, the uneven dispersion of medical expertise across the nation and the world, and the limitation that MYCIN would operate in a relatively small domain (Winston and Pendergast, 1984:29-30). The final factor is worth repeating; MYCIN was not designed to solve all medical problems. Rather, like its human counterparts, MYCIN is an expert in diagnosing infectious diseases and is not intended to become the Jetsons' version of the old country doctor. MYCIN captures the highly judgmental, inexact cognitive process of associating symptoms and diseases.

Figure 5 exemplifies a typical interview conducted by this expert system as it seeks knowledge on the patient's condition. Notice that the system designers were careful to build the program so that it would resemble the existing consultation process that doctors are already accustomed to (Winston and Pendergast, 1984:31). In this example the doctor did not know the identity of ORGANISM-1, so MYCIN asked further probing questions to try to narrow down the possibilities.

```
----------PATIENT-248----------
1)      Patient's Name: (first-last)
**      Dick H.
2)      Sex:
**      MALE
3)      Age:
**      52
4)      Have you been able to obtain positive microbiological
        information about a possible infection of Dick H.?
**      YES
----------INFECTION-1----------
5)      What is the infection?
**      ENDARTERITIS
6)      ...
7)      ...
8)      ...
The first significant organism from this blood culture
(CULTURE-1) will be referred to as:
----------ORGANISM-1----------
9)      Enter the identity of ORGANISM-1:
**      Unknown
10)     Is ORGANISM-1 a rod or coccus (etc.):
**      ROD
11)     The gram stain of ORGANISM-1:
**      GRAMNEG
12)     Have there been POSITIVE cultures yielding organisms
        about which you will NOT be seeking advice?
**      YES
```

Figure 5. MYCIN Expert System (Winston and Pendergast, 1984:31)

After many questions along these same lines, the doctor finally gets a diagnosis. This example provided the expert advice shown in Figure 6, where MYCIN narrowed the possibilities to six probable infections. Like a medical expert in this field, MYCIN presents all known infections that have a convincing probability of being the culprit.

```
INFECTION-1 is ENDARTERITIS with BACTEREMIA
<Item 1>     E. COLI [ORGANISM-1]
<Item 2>     SALMONELLA (species unknown) [ORGANISM-1]
<Item 3>     KLEBSIELLA-PNEUMONIAE [ORGANISM-1]
<Item 4>     PSEUDOMONAS-AERUGINOSA [ORGANISM-1]
<Item 5>     ENTEROBACTER [ORGANISM-1]
<Item 6>     PROTEUS-NON-MIRABILIS [ORGANISM-1]
```

Figure 6. MYCIN's Diagnosis (Winston and Pendergast, 1984:32)

The dilemma now becomes how to treat all six possible infections without overdosing the patient with antibiotics. This is where the expert part of an expert system really gets started. MYCIN goes into its therapy stage and prescribes two drugs that can safely defeat all six infections. See Figure 7.

```
[Rec 1] My preferred therapy recommendation is as follows:
        In order to cover for Items <1 3 4 5 6 >:
                Give:  GENTAMICIN
                Dose:  128 mg (1.7 mg/kg) q8h IV (or IM) for 10 days
                Comments:    Modify dose in renal failure
        In order to cover for Item <2>:
                Give:  CHLORAMPHENICOL
                Dose:  563 mg (7.5 mg/kg) q6h for 14 days
                Comments:    Monitor patient's white count

Do you wish to see the next choice therapy?
**      NO
```

Figure 7. MYCIN's Therapy (Winston and Pendergast, 1984:32)

Earlier, the integral parts of an expert system were discussed and presented in Figure 3. MYCIN's inference engine draws upon its rulebase of nearly 500 rules to analyze the factbase (supplied by the user) in order to come up with its recommendations.

Figure 8 shows MYCIN's rule number 27, in this case being a simple if-then inferential rule. Within this rule MYCIN gives hints at fuzzy logic by demonstrating its ability to see options on a continuum rather than simply yes and no or on and off; MYCIN suggests a 0.7 chance at the organisms identity.

Rule 27

If
1) the gram stain of the organism is gram negative, and
2) the morphology of the organism is rod, and
3) the aerobicity of the organism is anaerobic,

Then There is suggestive evidence (.7) that the identity of the organism is Bacteroides.

Figure 8. MYCIN's Rule 27 (Winston and Pendergast, 1984:34)

The inference engine uses these nearly 500 self-contained rules in much the same way a human expert would in coming up with its conclusions. In a process called backward chaining, MYCIN starts with trying to achieve its goal of identifying the unknown organism and moves backwards through its vast rulebase. Like a physician remembering training from medical school or knowledge gained through experience, MYCIN first recalls all rules stored in its memory that attempt to identify an organism, like rule number 27. Then it must move another layer deeper by searching for all rules that deal with gram stains. This process of moving backwards continues until MYCIN can go no further, at which point MYCIN asks the physician user questions for even more information (Winston and Pendergast, 1984:35).

Though MYCIN works well on this backward chaining process, a related method that operates in a reverse manner has also been implemented in other expert systems. This process is called forward chaining where the system starts at a known position and reasons forward to arrive at a conclusion (Ralston and Reilly, 1993:536). Forward chaining systems are better suited at making predictions in contrast to MYCIN's diagnostic prowess using backward chaining (Gould, 1995:80-81).

```
14)    Is the patient's illness with ORGANISM-1 a hospital-acquired infection?
**     Why

[2.0]  This will aid in determining the identity of ORGANISM-1.
       It has already been established that
              [2.1]   the morphology of ORGANISM-1 is rod
              [2.2]   the gram stain of ORGANISM-1 is gramneg
              [2.3]   the aerobicity of ORGANISM-1 is facul

       Therefore, if
              [2.4]   the infection with ORGANISM-1 was acquired while the
              patient was hospitalized
       then
                      there is weakly suggestive evidence (.2) that the identity
                      of ORGANISM-1 is pseudomonas [RULE050]
```

Figure 9. Insight Into Logic (Winston and Pendergast, 1984:35-36)

To truly grasp the superior ability of an expert system, MYCIN must be observed one step further. This system can work like a black box, being totally transparent to the operator, but MYCIN's real strength lies in its ability to give the user insight into its logic. In this example on line 14 of the interview, the system asks the physician a question about where the infection may have been acquired; this dialogue is given in

Figure 9. Here the physician is curious as to 'why' this information is relevant. As if conferring with a colleague, MYCIN answers the physician quite eloquently.

When the interview and consultation are complete, the physician can ask additional questions for clarification. In Figure 10, the physician wants to know how MYCIN arrived at one of its conclusions. This example demonstrates that, in addition to making a conclusion, the expert system can also store and keep track of how it was derived. Not only is this valuable as an audit tool having legal ramifications, but MYCIN can actually take on the role of teacher.

```
** HOW DID YOU DECIDE THAT ORGANISM-1 WAS AN E. COLI?
I used rule 084 to conclude that the identity of
        ORGANISM-1 is e.coli.  This gave a cumulative
        c.f. of (.51).
I used rule 003 to conclude that the identity of
        ORGANISM-1 is e.coli.  This gave a cumulative
        c.f. of (.43).
```

Figure 10. Further Logic (Winston and Pendergast, 1984:36)

It must be pointed out that the expert systems in the previous two examples can not totally replace the human experts they are modeled after. To the contrary, expert systems are generally designed and employed to compliment the human, especially where a shortfall of expertise exists. It should be clear that these two expert systems were intended to be utilized by someone with a good, solid working knowledge in the respective field. Feeling threatened by the development of an expert system can seriously prejudice the decision to explore this option and must be overcome in order to reap the benefits such a system might yield.

Previous Work

An enormous body of research is available on the broader topics of artificial intelligence and expert systems. Therefore, after defining the general terminology and gaining a good understanding of the capabilities of an expert system, the search was narrowed down to previous military or military-related studies and applications. A three-pronged approach was ensued to include searching Defense Technical Information Center (DTIC) holdings, reports written by the RAND Corporation, and computer systems already developed for the military.

A consistent pattern emerged as a result of this search; it appeared that the topic of artificial intelligence and expert systems technology was hot in the late 1980s, as the bulk of the printed material was produced then. The DTIC search revealed more than eighty theses addressing this area, most of which were written between 1984 and 1990. Similarly, RAND published more than twenty reports with all of these written between 1985 and 1989. Like the dropping of a fad, the printed material from these two sources nearly disappeared as RAND and DTIC papers moved on to other topics.

In particular, there were three theses among the DTIC documents that were related in some regard to the specific objectives of this paper. The first discussed a model for the rapid deployment of armed forces (Tate, 1984). The program DEPLOY, reviewed by Tate, has interactive capabilities but is more a model than an expert system. It also concerns the broader planning issues as opposed to the more execution oriented focus that this paper is targeting. The second thesis reviews a rudimentary system that draws upon stored knowledge to plan aircraft hazardous cargo loads (Sanborn, 1986). Though dated, this paper could provide a piece of the puzzle when it comes to prioritizing cargo versus available airlift. The final related thesis is more recent and brings to light the issue of route planning (Kilpatrick, 1992). Route choice will affect the amount of fuel required, which in turn affects how much cargo can be carried.

If development of a system proposed by this paper is pursued, there are two RAND reports dealing with the management side of fielding this type of project that should be reviewed. The first report warns of a number of risk factors that are still valid today (Bankes, 1985). These risk factors include escalating costs as the size of the project expands, a more difficult task in debugging and verification than in traditional software, and the development of support resources (hardware, software, and skilled personnel) before expert system technology can be applied. The second report, written four years later, echoes the managed-risk theme of the previous report (Kameny and others, 1989). This paper goes into more detail by presenting a step-by-step guideline highlighting managerial and technical risks during each phase of expert system software development.

Kameny and his colleagues also state that expert systems development can be viewed as a software engineering endeavor and, therefore, should be subject to much the same software development methods and standards as prescribed by the Department of Defense (Kameny and others, 1989). This blends in nicely with a review of some of the computer systems that have recently been developed for the military. Due to the explosion in the number of computer applications, the Department of Defense has setup the Defense Information Services Agency (DISA) to set standards for all new DOD computer systems. Now, all new proposed computer systems must go through DISA making it a prime initial source for researching such systems.

In researching existing DOD computer systems, an interesting and pertinent project was found called the TRAnscom Command and Control Evacuation System, coded TRAC2ES for short (DISA, 1995:Vol 4, 9). TRAC2ES is expected to become a complete mission planning system with its primary purpose to manage patient aeromedical evacuation. Its relevance to this paper lies in TRAC2ES ability to optimize the allocation and scheduling of aeromedical evacuation partly by matching Medical Treatment Facility (MTF) resources on one end with the nature of the patients' traumas on the other. If not already aware of how this relates to the deployment of a TALCE or an ITUD, the next chapter's discussion on the process both the TALCE and ITUD use to match resources available to those thought to be required should help the reader make the connection.

III. Discussion of Problem Components

Two Proposed Users

As mentioned earlier, many facets of military operations could be considered for incorporation of some type of expert system. However, this paper takes a look at two closely related areas within Air Mobility Command that could possibly be served by the same or very similar systems. Remember that one of the critical aspects in building an expert system is that it should not attempt to do everything for everybody. Rather, it should, like the human counterpart it is designed to emulate, stick to the field in which it possesses the expertise. The two areas that this paper addresses are the Air Mobility Operations Group (AMOG) and the Integral Tanker Unit Deployment (ITUD) concepts of employing AMC assets.

The AMOG Structure. There are two AMOGs: one at Travis Air Force Base, California and the other at McGuire Air Force Base, New Jersey (AMC, 1995:Ch 2, 7). They were strategically placed on each coast to provide the leadership and expertise needed to support rapid response of air mobility assets around the world (Cirafici, 1995:82). As discussed in Chapter I, the AMOGs are chartered with the task of temporarily beefing up the world-wide enroute support system when the need arises. The AMOGs do not own or operate any airplanes directly, as will later be contrasted with the ITUD concept. Rather, their mission is to set up or fortify a staging location for AMC airplanes and crews transiting the area.

Much of the mission support forces to set up a bare base operation is extracted directly from the AMOGs' subordinate squadrons. The rest of the personnel required are augmented from active duty squadrons and air reserve component (ARC) forces (Cirafici, 1995:10, 82-83). An AMOG consists of six squadrons with an organizational structure depicted in figure 11 (AMOG-621st, 1995).

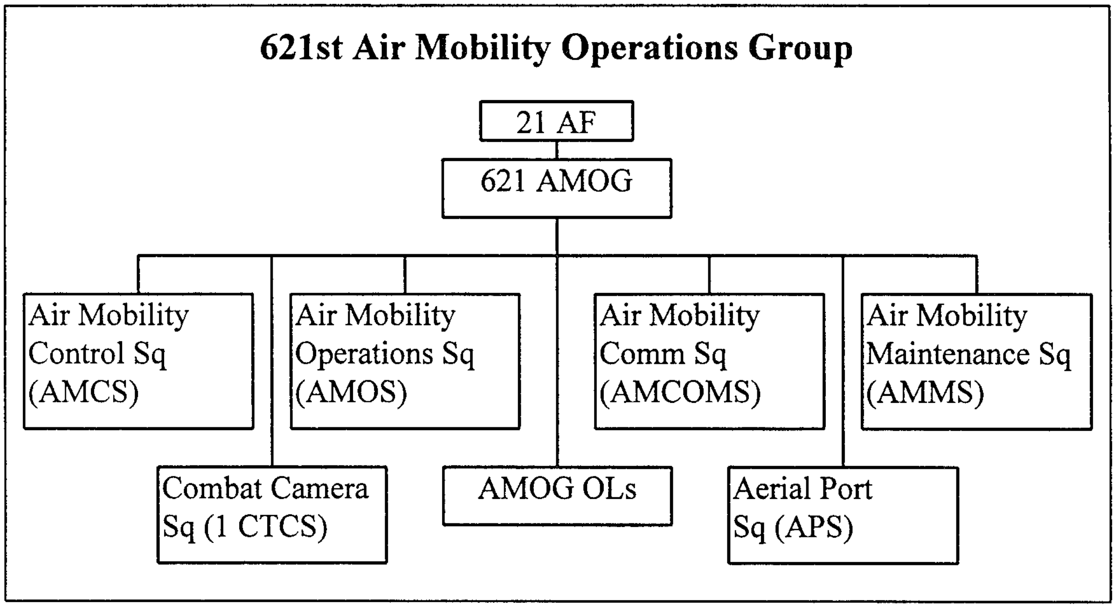

Figure 11. AMOG Organizational Chart (AMOG-621st, 1995)

Neither the AMOG nor its squadrons deploy directly, as their chartered, stateside missions are to train their personnel for deployment not actually do it as units (AMOG-621st, 1995). Rather, the AMOG forms a Tanker Airlift Control Element (TALCE) by piecing parts together from its component squadrons using a building block approach. The result is an organizational structure resembling a typical airlift wing. The TALCE then deploys to provide support for air mobility forces where a full time AMC enroute structure is nonexistent or inadequate (Cirafici, 1995:12-13). Specifically, a TALCE provides command and control, aerial port, logistics, intelligence, combat camera, civil engineering, security, weather, intelligence, and other assets as required (AMC, 1995:Ch 4, 19). One can see why it is important, in most airlift operations, for the TALCE to be the first to arrive on the scene.

If multiple locations are involved, more than one TALCE may be formed and sent to the theater to handle their respective locations. If the operation is large enough an Air Mobility Element (AME), formed in the same fashion, would be dispatched to centralize the control of the TALCEs (Cirafici, 1995:10). This building block approach is shown in Figure 12.

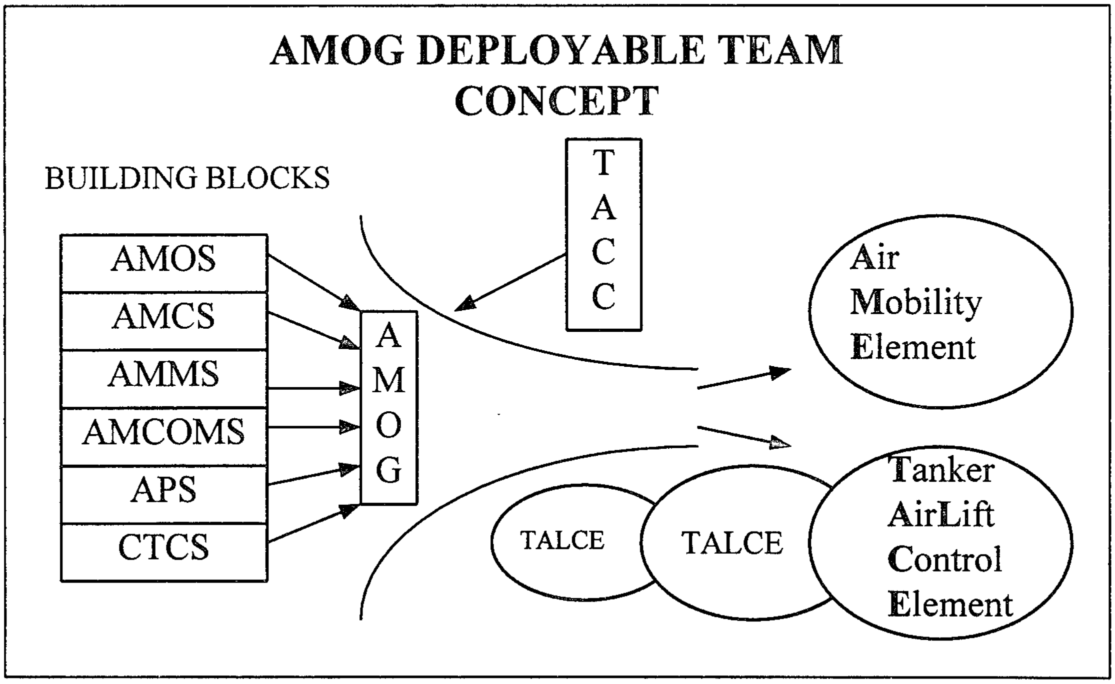

Figure 12. Forming an AME and TALCE (AMOG-621st, 1995)

The ITUD Concept. In contrast to the AMOG which is an organizational structure, the ITUD is a concept of operations. As the name implies (Integral Tanker Unit Deployment), tanker units are tasked to deploy as an aggregate whole. An ITUD tasking comes from Headquarters Air Mobility Command to the tankers at the air refueling wing level (AMC, 1995:Air Refueling MAP, 2). From this point it is up to the individual wings to decide how they will actually fill the tasking. The intent is for integral units, teams that are used to working together on a daily basis, to move their operation, mostly intact, to a forward operating location.

As opposed to the airlift mission, where planes and crews rotate in and out as they haul cargo to and from the location, air refueling airplanes and crews are bedded down along with the support infrastructure for forty-five to ninety days at the deployed site. Unlike the AMOGs, air refueling wings tasked with an ITUD deploy their own airplanes to the remote location. The two concepts are similar, however, in that the ITUD provides nearly the same type of support services that the TALCE performs, as described above. The ITUD concept sounds simple enough, but one must understand how this process evolved in order to fully appreciate it.

When the tanker force from SAC joined with the former MAC, it was more than simply a reorganization and name change to Air Mobility Command: it was an entire cultural change for the tanker fleet. It was not the reorganization itself, but the reason for the reorganization that was behind this culture shock. Like the entire military, the very role of the air refueling mission was undergoing a drastic shift in focus. Since the early 1960s, when Strategic Air Command was formed as part of the nuclear triad, the primary mission of the tanker force was to refuel bombers in support of the Single Integrated Operations Plan (SIOP) (Gant, 1993:xi, 32). Pulling or sitting alert put the forces at a high state of readiness to respond and launch quickly in event of a surprise attack from an adversary and became the hallmark of SAC for the next three decades (Gant, 1993:25).

The fall of the Soviet Union caused the nuclear deterrent mission to decrease and brought the tankers off of alert, making them more available as mobility assets (Killingsworth and others, 1994:iii, xi). As a result, today's mission of supporting the deployment and employment of U.S. global power has become the tanker's primary job. A recent RAND study found that "forward-located" tanker assets were critical during the early stages of a deployment, especially before the enroute system has had a chance to expand (Killingsworth and others, 1994:xii).

Prior to the Soviet demise and the RAND study, however, the emerging concept behind the ITUD was already being explored. Operation Desert Storm sparked an Airpower Research Institute (ARI) project to look at the future role of the air refueling force (Gant, 1993). In this paper, Gant proposed a Deployable Refueling Package (DRP) that would be "self-supporting and internally led." This package concept would make it possible to rapidly deploy tanker operations in groups containing their own maintenance and operational support needs. Still bound by the alert commitment, Gant developed the basic DRP as a four-ship fleet of tankers that could be combined with other units' DRPs if the scale of the mission dictated. As the SIOP took a back seat, the DRP concept eventually evolved into the ITUD where an entire unit would deploy. It appears that Gant's work was the forerunner of today's ITUD.

Current Procedures

The speed at which the ITUD concept was incorporated into operations caught the tanker units somewhat unprepared, leaving them without a specific procedure to determine what to bring on the deployment. As pointed out in the example in Chapter I, a brainstorming-type technique was used at first. Eventually, as more experience was gained and the tanker force became more in tune with existing AMC methods, a process resembling that used by the AMOG began to develop. For this reason, only the method currently used by the AMOG will be presented since it adequately represents both of the proposed users. Unless cited otherwise, the information in this section comes from a personal interview with the squadron commander of the Air Mobility Operations Squadron (AMOS) at McGuire Air Force Base, New Jersey (Smith, 1995).

When directed by AMC Headquarters, the Tanker Airlift Control Center (TACC) tasks the AMOG to supply AMC assets using unit type codes (UTCs) as basic building blocks to form a package. These five-character designators identify specific package capabilities and can be found for all Air Force packages in Volume III of the War and Mobilization Plan (AMWC, 1995:55). The building blocks are put together, as needed, to provide command and control, aircraft maintenance, planning and execution, communications, intelligence, weather, and any other mission specific functions for operational and humanitarian missions.(Cirafici, 1995:9).

Using these generic, off-the-shelf UTCs is a good first step, but it does not account for the variety of austere locations we now find AMC assets deploying to. They also do not get down to the specifics or micro level planning required at the unit level, especially in the support areas. In reality, when tasked for an operation, the AMOG "pares and tailors" the tasked UTCs as required to fit the location and mission specifics. Previous experience and rule-of-thumb reasoning is used along with the most recent site survey information, the Summary of Aifield Restrictions, the Airfield Suitability Report, and any other available information to determine what needs to be shipped with the deployment in order to cope with the scenario.

Through experience, it has been determined that certain items are needed on all deployments regardless of the location. Examples include office supplies like pencils, pens, paper, computers, telephones, fax machines, and other similar items. Where legal to do so, the AMOG stocks these items in deployment boxes that are stored and ready to be shipped in a moment's notice. Depleted supplies are replenished when the TALCE returns stateside after the mission is complete so that the boxes are ready to go for the next tasking. Items impractical and against regulation to stock on the shelf, like computers and fax machines, are pre-designated and taken from the individual squadrons when tasked.

Many other items are ready to be deployed if needed and range in size from 40,000-pound Material Handling Equipment (MHE) to relatively small Meals Ready to Eat (MREs). The AMOG must determine what is needed and match this against what is available at the forward location, either as part of the existing infrastructure and host nation support or that which can be purchased from the local economy. This cargo must also be prioritized so that it arrives in the proper sequence; it would be beneficial, for example, to have MHE, used to load and unload the cargo, on one of the first airplanes.

The fact that airlift is not an unlimited resource and not everything that is desired can be taken is another important reason to prioritize. Some of the list is easy to rank order, based upon the scenario, while the rest requires expert judgment. The "old pros" further demonstrate their expertise by recalling similar circumstances from past experiences to set priority to the list and to insert items that are often overlooked by the "rookies".

Basic Framework for the Expert System

Using the above discussion, some meat can now start to be placed on the skeleton structure of the expert system presented in Figure 3 from the previous chapter. Again, this is not an attempt to build a working expert system. Rather, it is presented to help put the pieces together and try to categorize the various inputs used in determining the best mix of equipment and supplies to cope with the deployed location.

Figure 13 has been adapted from Figure 3 to incorporate some specifics presented in this chapter. The rulebase would consist of various flying regulations, Operating

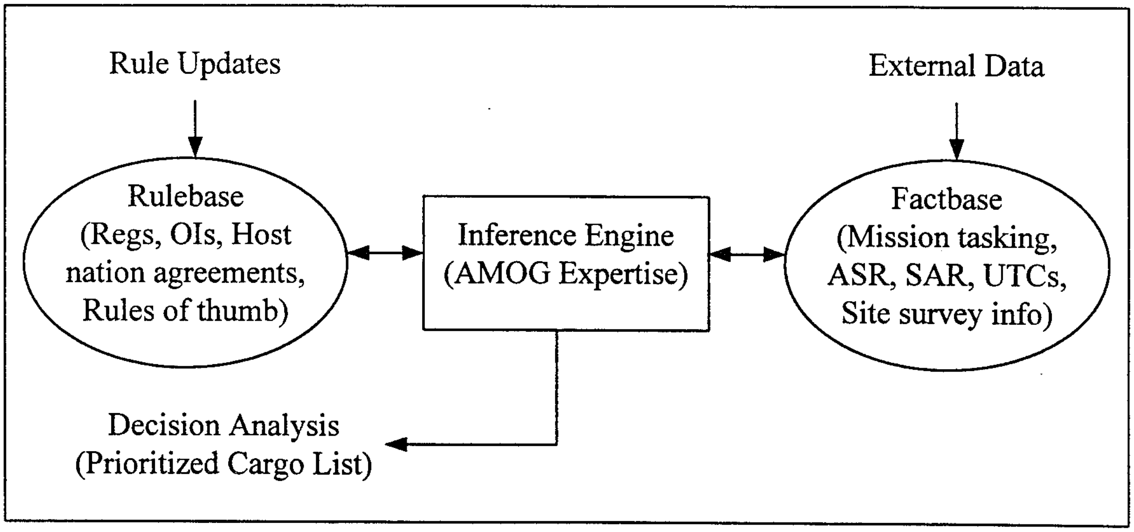

Figure 13. Basic Expert System for the AMOG (adapted from Ball, 1994:37)

Instructions, standard operating procedures, agreements with host nations, airplane specifications to include cargo and range capabilities, rules of thumb, and guidance from higher headquarters. It could also house the rules giving preference to certain types of equipment and supplies, depending on the scenario.

The factbase would include the mission specifics dictated in the tasking such as location, duration, number and type of aircraft allotted for airlift, and the nature of the mission. It would also include information from the Airfield Suitability Report (ASR), the Summary of Airfield Restrictions (SAR), the most current site survey, and any other sources that could indicate what is already available at potential deployment locations.

The inference engine draws information from the rulebase and the factbase as it interacts with the user in order to come up with its solution. This would be an attempt to capture and mimic the existing thought processes of the highly skilled AMOG personnel currently making these decisions.

Figure 13 and the discussion above are, by no means, all-inclusive and it would take much further research to ferret out all available inputs. The next chapter will explore the options available to field a fully functional expert system.

IV. Alternatives

One More Definition

Before discussing the possible alternatives, one more definition must be presented. Though Chapter II has an entire section devoted to definitions, it was more appropriate to wait and explain the art of knowledge engineering here since it is critical in understanding how an expert system is made. Ball provides his working definition of knowledge engineering, stating that, "It is the translation of unstructured, even unarticulated, knowledge into a machine-readable form" (Ball, 1994:188). The following paragraphs will expand upon this definition.

Figure 14 shows the knowledge engineering process that begins with gathering information from the organization desiring the expert system (Ball, 1994:188). Through interviews and observation, the knowledge engineer tries to capture the targeted rare or critical expertise and clone it for later use (Winston and Pendergast, 1984:18). The process ends with the knowledge encoded in such a way that it can be interpreted and manipulated by a computer. The role that knowledge engineering plays, therefore, can be roughly equated to that of a translator in that it bridges the gap by interpreting between the expert or experts and the software programmers who eventually design the actual software shell for the system (Ball, 1994:188).

While bridging this gap, the knowledge engineer must determine the complexity of the problem, the availability of the inputs, the importance of the search speed in finding the problem solution, and the importance of accuracy in the solution. According

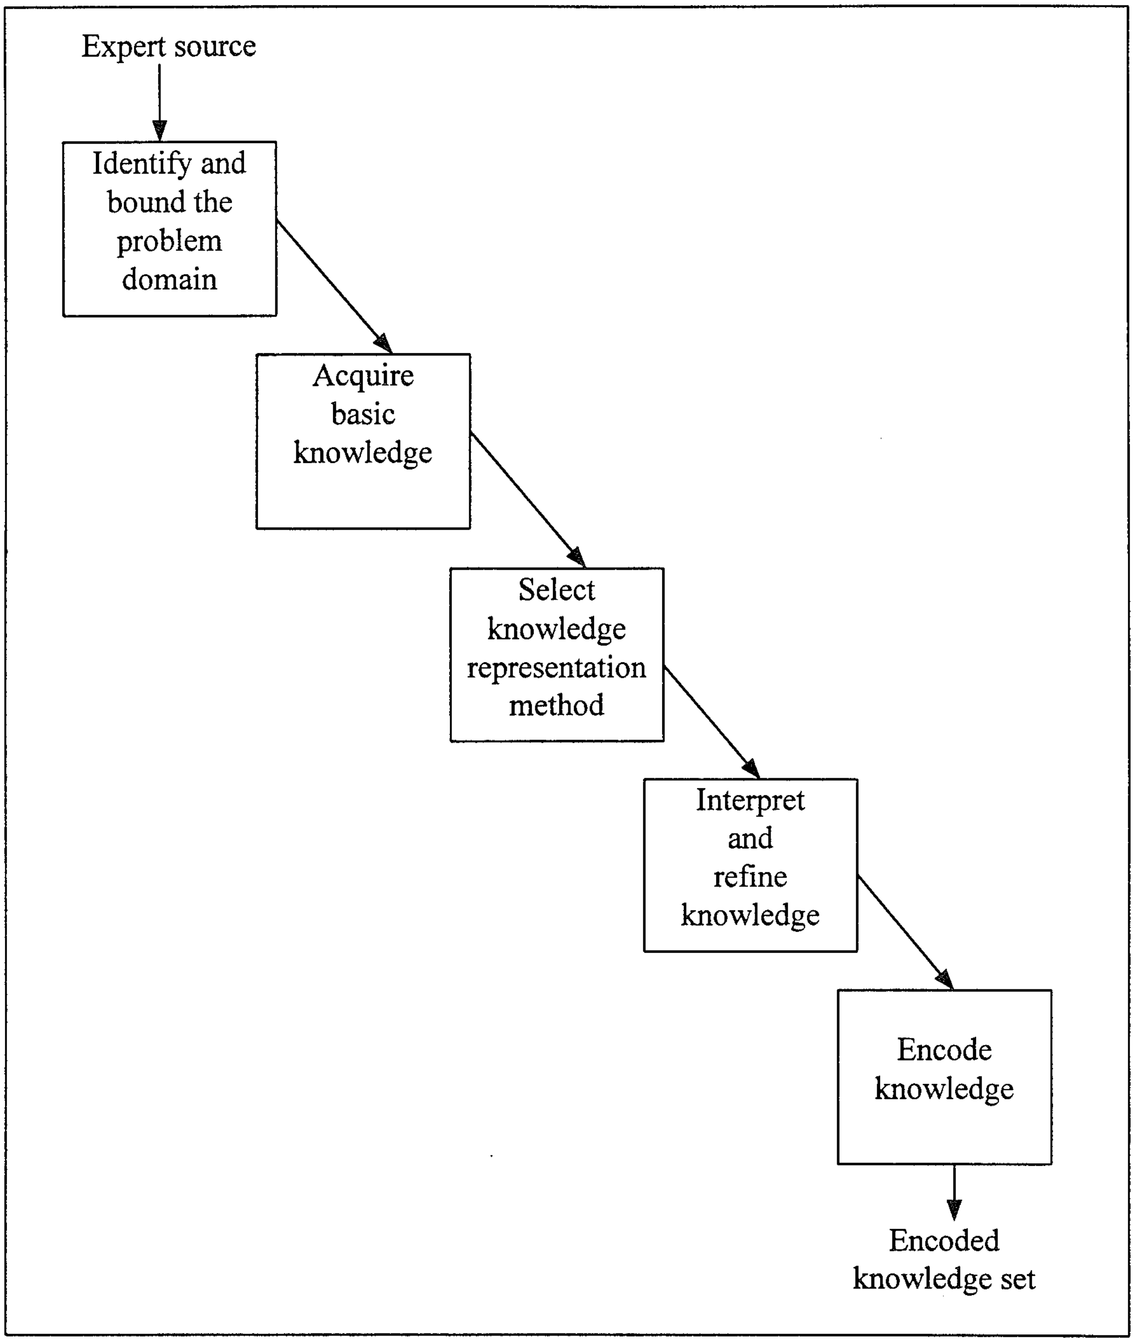

Figure 14. Knowledge Engineering Process (Ball, 1994:189)

to Ball, the knowledge engineer needs the following skills: interpersonal communications skills, the ability to summarize, explanation and justification, assembly of results, recognition of constraints, development of hypothesis, assessment and use of certainty factors, organization and control of situations, and finally, computer technical skills (Ball, 1994:190).

One of the earliest examples of this process, involving an expert maintenance system, can be used to tie this concept together. After more than thirty years of experience, General Electric's chief diesel locomotive maintenance expert was getting ready to retire. The company, fearing the loss of this man's keen ability to detect and accurately diagnose symptoms of failure, frantically sought a way to "pick his brain" in order to "capture" his knowledge. They found that the best method was through multiple interviews where the retiree described all that he knew about the steps and circumstances surrounding the many possible scenarios of diesel locomotive operation, failure, and maintenance. The detailed written records from these interviews were eventually migrated to a computer system to be used by future maintenance personnel (Ball, 1994:204).

Obtaining an Expert System

With the somewhat long-winded definition of knowledge engineering fresh in the reader's mind, the first alternative can be presented. This would involve hiring a professional knowledge engineer to translate "AMOG and ITUD speak" to computer code and having a private software firm turn the results into a workable computer program. The result would be a commercially developed, fully customized expert system that would be built specifically for the purpose espoused in this paper. The discussion of current procedures in Chapter III was an initial, though be it insufficient, attempt by this author to demonstrate where the knowledge engineer might start in going about his business.

This alternative has advantages and disadvantages, as do all options. The main advantage is the obvious focus that the custom approach would bring to the problem. Not only would the Air Force get a professionally designed, custom tailored system, but the scrutiny from the inside as well as the outside could uncover previously unseen or newly developed glitches in AMC's overall deployment strategy. These imperfections could be corrected in conjunction with the development of the expert system. Two major disadvantages are the time and cost involved with custom designing software and the difficulty with technical bugs inherent in a new computer system.

The second option is to seek out programs that have already been developed for commercial use. A number of generic programs, capable of running on a standard PC, promise the ability to adapt to the specific needs of the user, similar to the way a financial program like "Quicken" can be tailored for use by a wide range of businesses. The advantages and disadvantages of this approach are parallel but opposite to those of the custom approach presented above.

Limited financial resources available to this project precluded testing actual working products to verify the claims, but it should be considered if further work is pursued in this direction. Instead, four of the more promising programs are mentioned below and were extracted from The Software Encyclopedia.

(1) "Expert Choice" by Expert Choice Incorporated retails for $495.00 and provides executive decision support for complex problems involving multiple criteria, alternatives, and scenarios (R. R. Bowker Company, 1994:280). It can also explain to the human user why a certain decision was made. One of its suggested uses is in resource allocation which can be related to the problem proposed in this paper.

(2) "Expert Ease" by Dynacomp Incorporated sells for $395.00 and is touted as a decision tree analysis tool (R. R. Bowker Company, 1994:280). It "creates an expert system of your own design, using your own specialized knowledge," to be used by non experts so they can arrive at the same result or decision that the expert would. By simply typing in real world examples, this program claims to create a decision tree based on rules that it derives from the examples.

(3) Dynacomp also offers the less expensive "Expert System Tutorial" which retails for $29.95 (R. R. Bowker Company, 1994:280). This system allows the user to do a "needs analysis" to determine if the procedure they are trying to emulate is amenable to capture by an expert system in the first place. If the problem domain is small enough, this program may be sufficient all by itself.

(4) The final program, selling for $495.00, is called "Expert 87" and is offered by Thinking Software Incorporated (R. R. Bowker Company, 1994:280). Though this program is probably dated, as its title suggests, it still has features that are worth mentioning. It can combine the viewpoints of as many as twenty five experts to generate a system that works on a consensus type basis. Each individual expert's unique perspective, as they all address the same problem, is preserved.

The final option for sourcing an expert system to fit the needs laid down in this paper comes from within the Air Force itself. This should be of no surprise since it was pointed out in the very first part of this paper that the initial work in artificial intelligence was spawned in the military. Efforts to harness technology continue in the military to this day as personnel from within the ranks possess the skills to develop complex computer programs for simulation, war-gaming, and command and control. With this in mind, an AFIT graduate student, working towards a more technical advanced degree, could use this paper as a starting point to perform their thesis in an effort to take the idea to its next step. In addition to summarizing, the next chapter offers suggested paths should someone decide to proceed on this endeavor towards the development of an expert system.

V. Summary and Conclusion

This paper has pointed out an area where technology might be used to enhance air mobility operations. Due to their similarities, the AMOG and the ITUD concepts were targeted as candidates for an expert system to aid in the deployment process. A thorough background of expert systems, in general, and how they fit into the broader scheme of artificial intelligence laid the foundation for understanding the terminology involved in this area. The AMOG and the ITUD were detailed to understand how these concepts are put into practice and how current procedures and inputs needed to implement them might be overlaid on the framework of an expert system. Finally, the options available to obtain an expert system were presented.

Though development of an actual working expert system was beyond the scope of both this paper and its author, an attempt was made to show that the concept is feasible with today's technology. This paper, then, could be used as a platform for further work at various levels of research in the direction of a fielded expert system. Before considering any of the alternatives in Chapter IV, however, a detailed impact study should be performed in an attempt to quantify the need. While this paper cited some examples where an expert system could have improved the deployment process, no primary research was performed to gather hard data showing a favorable cost to benefit ratio. This would be a good starting point for further research.

New projects hinge upon obtaining funding, and obtaining funding comes down to convincing Congress that the benefits of a proposed program outweigh the costs. In a not-for-profit organization like the Air Force, however, it is difficult to accurately assign numerical values to costs and benefits. Especially in the military, weighing pluses and minuses goes deeper than just measuring the immediate outcome derived from an object's use on one side and its price tag on the other. The following paragraphs point to some of the considerations that a researcher trying to quantify the pros and cons would need to think about.

Authors Sproull and Kiesler touch on this subject when they discuss cost and benefit implications for an organization setting up a computer network. In their book, they talk about the second-level effects realized when new technology is introduced. Referring to the typewriter when it was first developed, Sproull and Kiesler point out that its first-level effect was intended to be the ability to produce letters of the quality coming from a printing press with the anticipated customers being clergymen and writers. The unforeseen second-level effect it had on clerical operations in the business world proved to be even more extensive and beneficial. Similarly, few people initially envisioned that the telephone and the personal computer, seen primarily as tools for business, would become standard household features revolutionizing the areas of communications and information (Sproull and Kiesler, 1991:6-7).

Sproull and Kiesler point to other benefits of new technology that are also hard to quantify. First of all, electronic communications and manipulation of data can make routine procedures more efficient and reduce the need for redundancy, while at the same time expanding the amount of information that can be assimilated (Sproull and Kiesler, 1991:141). A second benefit accompanying technological change is the introduction of new ways of thinking and can be summed up by an old Chinese proverb that says, "If we don't change direction, we'll end up where we're headed." In other words, keeping pace with technological change can expand one's horizon, an important attribute to have when organizations are large and have internally diverse activities or when change is great and nonroutine (Sproull and Kiesler, 1991:141-142, 173). Both of these describe the state of affairs for today's military where new ways of thinking, acting, and organizing are critical. Finally, Sproull and Kiesler point to the increased speed and reliability associated with direct access to information (Sproull and Kiesler, 1991:170).

The costs associated with a new computer system are a little more straight forward, but one must look past the first-level direct costs of the equipment and the programming required to develop the system. Other costs include installation, routine maintenance, upgrades, possible organizational changes, and those "risk factors" warned of by the authors of the Rand reports presented near the end of Chapter II. In any event, when looking at cost to benefit relationships, a quote from Sproull and Kiesler can guide one beyond the first-level effects. Though they are speaking of network connections, the concept is the same as they state,

> New connections will be most attractive to organizations committed to employee competence and involvement and to organizational flexibility as ways to achieve and sustain success. Without such commitments, analyses of electronic communication will be dominated by first-level efficiency thinking. (Sproull and Kiesler, 1991:161)

Once a decision is made to field an expert system, research at the next level could be performed by examining the alternatives in this paper (and any others that might be added by further investigation) in more detail. The two RAND studies cited near the end of Chapter II and mentioned again in the previous paragraph should then be reviewed. In addition to identifying risk factors that warn of escalating costs, debugging problems, and additional external resources needed actually field an expert system, the later paper presents a step-by-step guideline highlighting managerial and technical risks during each phase of expert system software development. Also, as previously pointed out, the Defense Information Services Agency (DISA), must be in the loop on the development of any new Department of Defense computer system. The standards they set must be adhered to. DISA would also be the starting point to search for existing and related expert systems that might save duplication of effort, making development easier.

The final level of further research alternative would be the building of a complete and operational expert system. To do this, the developer should consider adhering to the Knowledge Engineering Process shown in Figure 14 in the previous chapter. This process would involve much greater study of the AMOG and the ITUD process to capture the human process to machine code. As pointed out before, the Air Force has computer programmers that are capable of writing powerful and complex programs, but the actual knowledge engineering process may have to be contracted to an outside professional due to the scarcity of expertise in this field.

The military initiated much of the work in the field of artificial intelligence over the last fifty years. Since then the outgrowths of this discipline, expert systems, have migrated to the civilian sector to be used in many commercial applications. Advances in computer technology, coupled with improvements in expert systems technology, make it feasible for Air Mobility Command to field a system to enhance mobility operations.

Bibliography

Air Mobility Command (AMC). Air Mobility Master Plan-1996. AMMP-96. Scott AFB IL: HQ AMC/XP, 2 October 1995.

Air Mobility Operations Group (621st). "Unit Mission Briefing." McGuire AFB NJ, 1995.

Air Mobility Warfare Center (AMWC). Acronym Dictionary (Second Edition). Ft Dix NJ: HQ AMWC/WCOM, February 1995.

Ball, Larry L. Network Management With Smart Systems. New York: McGraw-Hill, Inc., 1994.

Bankes, S.C. Future Military Applications for Knowledge Engineering. Santa Monica CA: The Rand Corporation, July 1985, (N-2102-1-AF)

Cirafici, John L. Airhead Operations - Where AMC Delivers. Maxwell AFB AL: Air University Press, 1995.

DeJesus, Edmund X. "Manufacturing Data," Byte,Vol 10, Issn 12: 63 (December 1995).

Defense Information Systems Agency (DISA). Integration Decision Paper for the Joint Transportation Corporate Information Management Center's Transportation Migration Systems Recommendations. Falls Church VA: 1 February 1995.

Gant, Henry N. Deterrence in the Face of Explosive Peace. Maxwell AFB AL: Air University Press, 1993.

Gould, Lawrence. "If AI Ran the Zoo," Byte,Vol 10, Issn 12: 79-83 (December 1995).

Hunter, John H. "The KC-135 in AMC Contingency Operations: Still Some Rough Edges." Position Paper Submitted in the Air Mobility Operations Course, 95-C, March 1995.

Kameny, I., U. Khan, J. Paul, and D.G. Taylor. Guide for the Management of Expert Systems Development . Santa Monica CA: The RAND Corporation, July 1989, (R-3766).

Killingsworth, Paul, Keith Henry, Laura Melody, and James Stucker. Tankers: Air Mobility for the 1990s. . Santa Monica CA: The RAND Corporation, November 1994, (DRR-913-AF/OSD).

Kilpatrick, Freeman A. An Investigation of Discovery-Based Learning in the Route Planning Domain. MS thesis, AFIT/GCE/ENG/92D-07. School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, 1992 (AD-A259141).

Matthews, James K. and John W. Leland. General Ronald R. Fogleman - An Oral History. Scott AFB IL, 1995.

Ralston, Anthony and Edwin D. Reilly. Encyclopedia of Computer Science (Third Edition). New York: Van Nostrand Reinhold, 1993.

R. R. Bowker Company. The Software Encyclopedia. New Providence NJ: Reed Publishing USA, 1994.

Sanborn, Roderick S. A Knowledge-Based System for Load Planning Hazardous Cargo. MS thesis, AFIT/GCS/ENG/86D-14 School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, 1986 (AD-A178184).

Smith, Gregory. Commander, Air Mobility Operations Squadron, McGuire AFB NJ. Personal interview. 28 December 1995.

Sproull, Lee and Sara Kiesler. Connections: New Ways of Working in the Networked Organization. Cambridge: MIT Press, 1991.

Tate, David O. DEPLOY: An Interactive Goal Programming Model for the Rapid Deployment of Armed Forces. MS thesis, AFIT/GOR/OS/84D-14. School of Systems and Logistics, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, 1984 (AD-A152007).

The New Encyclopedia Britannica-Volume 1 (Fifteenth Edition). Encyclopedia Britannica Incorporated, 1993.

Tunner, William J. Over the Hump. New York: Duell, Sloan, and Pearce, 1964. Reprint by U.S. Government Printing Office, 1991.

Winston, Patrick H. and Karen A. Pendergast. The AI Business: Commercial Uses of Artificial Intelligence. Cambridge: MIT Press, 1984.

Wolk, Herman S. “The Defense Unification Battle, 1947-1950: The Air Force,” Prologue, 7: 10-18 (Spring 1975).

Vita

Major Joseph A. Jackson was born on 10 April 1960 in Belmond, Iowa. He graduated from Kanawha High School in 1978 and entered undergraduate studies at Iowa State University in Ames, Iowa. He graduated with honors, earning a Bachelor of Science degree in Aerospace Engineering in 1982. He received his commission on 15 May 1982 through the Reserve Officer Training Corps program at Iowa State University.

After graduating from Undergraduate Pilot Training, he remained at Vance AFB as a First Assignment Instructor Pilot. His second assignment was at Kadena AB, Japan where he flew the KC-135 and served as flight commander and assistant operations officer. His third assignment took him to Grand Forks AFB where he was the chief of tanker standardization and evaluation branch and chief of wing training. In February 1995, he entered the School of Logistics and Acquisition Management, Air Force Institute of Technology as part of the Advanced Study of Air Mobility (ASAM) program.

Permanent Address: 1112 9th Ave. S.
Clear Lake, IA 50428

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
May 1996

3. REPORT TYPE AND DATES COVERED
Graduate Research Paper

4. TITLE AND SUBTITLE
ARTIFICIAL INTELLIGENCE: CAN EXPERT SYSTEMS TECHNOLOGY ENHANCE AIR MOBILITY COMMAND (AMC) DEPLOYMENTS?

5. FUNDING NUMBERS

6. AUTHOR(S)
Joseph A. Jackson, Major USAF

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Air Force Institute of Technology,
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER
AFIT/GMO/XO/96J-3

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
HQ AMWC/WCOA
FT DIX NJ 08640

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

With the military restructured to react via rapid response from CONUS as opposed to forward presence overseas, Air Mobility Command has been charged to provide the Global Reach needed to project power abroad. The Air Mobility Operations Group (AMOG) structure and the Integral Tanker Unit Deployment (ITUD) concept, used to deploy the initial airlift and air refueling assets forward, were studied to determine the viability of using expert systems technology to aid in matching austere location demands with the proper mix of personnel, equipment and supplies. The paper starts with a thorough background of artificial intelligence and gives two rudimentary examples of working expert systems to show how they differ from conventional software programs. An overview of the AMOG and ITUD then identifies the intended beneficiaries of such a system. Alternative sources for fielding an expert system, ranging from buying off-the-shelf, commercial software to custom designing a system to fit this specific scenario, were then explored. Though this paper does not develop an actual working expert system, it does provide a foundation and offer suggested avenues to further work towards that goal.

14. SUBJECT TERMS
Air Mobility Command, Artificial Intelligence, Computer Programs, Deployment, Expert Systems,

15. NUMBER OF PAGES
61

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT
Unclassified

18. SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

19. SECURITY CLASSIFICATION OF ABSTRACT
Unclassified

20. LIMITATION OF ABSTRACT
UL

AFIT Control Number AFIT/GMO/XO/96J-3

AFIT RESEARCH ASSESSMENT

The purpose of this questionnaire is to determine the potential for current and future applications of AFIT research. **Please return completed questionnaire** to: DEPARTMENT OF THE AIR FORCE, AFIT/LAC BLDG 641, 2950 P STREET, WRIGHT-PATTERSON AFB OH 45433-7765 or e-mail to dvaughan@afit.af.mil or nwiviott@afit.af.mil. Your response is **important.** Thank you.

1. Did this research contribute to a current research project? a. Yes b. No

2. If you answered YES to Question #1, do you believe this research topic is significant enough that it would have been researched (or contracted) by your organization or another agency if AFIT had not researched it? a. Yes b. No

3. The benefits of AFIT research can often be expressed by the equivalent value that your agency received by virtue of AFIT's performing the research. **Please estimate** what this research would have cost in terms of manpower and dollars if it had been accomplished under contract or if it had been done in-house.

Man Years____________ $____________

4. Whether or not you were able to establish an equivalent value for this research (in Question 3), what is your estimate of its significance?

a. Highly Significant b. Significant c. Slightly Significant d. Of No Significance

5. Comments (Please feel free to use a separate sheet for more detailed answers and include it with this form):

____
Name and Grade

____
Organization

____
Position or Title

____
Address