The DRAIR Advisor: A Knowledge-Based System for Materiel Deficiency Analysis

Brian L. Robey, Pamela K. Fink, Ph.D., Sanjeev Venkatesan, and Carol L. Redfield, Ph.D.

Southwest Research Institute
6220 Culebra Road
San Antonio, Texas 78238-5166
brobey@swri.edu

Jerry W. Ferguson

U.S. Air Force Materiel Command
Oklahoma Air Logistics Center
Tinker AFB, OK 73145
jferguso.ocdis01.af.mil

Abstract

Southwest Research Institute (SwRI) and the U.S. Air Force Materiel Command designed and developed an automated system for the preparation of Deficiency Report Analysis Information Reports (DRAIRs). The DRAIR provides Air Force engineers with an analysis of an aircraft item's performance history that includes maintenance, supply, and cost. The DRAIR also recommends improvements for a deficient materiel or aircraft part. The successful design, development, and deployment of the DRAIR Advisor system by applying a combination of knowledge-based system and database management techniques are the subject of this paper.

Introduction

When a problem occurs with an Air Force aircraft part in the field, flight-line personnel prepare a materiel deficiency report (MDR) that describes the problem encountered. Engineers and equipment specialists responsible for the troublesome part, or end item, review the MDR to identify the possible cause(s) of failure. In the past, engineers and equipment specialists have turned to operations research (OR) analysts to assist in item performance analysis. This analysis is usually time consuming, personnel intensive, and requires information from many Air Force data systems. At the Oklahoma City Air Logistics Center, located at Tinker Air Force Base, data collection and analysis require two person-days. This analysis is summarized by an OR analyst in a written document called the Deficiency Report Analysis Information Report, or DRAIR. This document describes an item's performance history including maintenance, supply, and cost. The DRAIR also contains an analysis section and an actions recommended section that suggest performance improvements for the part. To produce a DRAIR, an OR analyst must draw on expertise about acceptable aircraft item performance. This expertise resides among OR analysts, engineers, and equipment specialists. An example DRAIR is illustrated in Figure 1.

To reduce preparation time and produce higher quality DRAIRs, a knowledge-based system was proposed and funded by the Air Force Materiel Command. This automated system had several specific objectives. The primary objective was to reduce the overall time required to produce a DRAIR. To meet this objective, it was necessary to house the data on one central computer accessible to all OR analysts. Another objective was to standardize the DRAIR and make it directly available to the personnel who request it, namely engineers, equipment specialists, and item managers. This would reduce demands on the OR analysts and provide additional time for them to address more complex analysis problems. Further, with the turn over of personnel in the military and the aging of the aircraft fleet, another objective was to capture expertise from personnel who are most knowledgeable about specific aircraft systems and federal stock classes and to make this expertise available to less experienced individuals in the field.

October 7, 1992
REPORT NUMBER: 6255
PREPARED BY: Jerry Ferguson
ORGANIZATION: OC-ALC/TIETP
PHONE: DSN 336-5015

Deficiency Report Analysis Information Report (DRAIR)

SUB-NSN	APPL	WUC	NOUN	UNIT COST
1280012287258	B001B	55AAA	PNL, CITS	$200,248.12
1280012287258	B001B	76GHE	PACU, CP1591/ASQ184	$200,248.12
1280012630536	B001B	55CAA	COMP, CONT, CP1691	$178,679.58
1280012630536	B001B	76GHE	PACU, CP1591/ASQ184	$178,679.58
1280012630536	B001B	55AAA	PNL, CITS	$178,679.58

SOURCE DATA: The data used to prepare this report came from the following sources: 1) Product Performance Subsystem (G099), 2) Supportability analysis Forecasting Evaluation (SAFE), 3) Flying Hours (G099), 4) MICAP Hours (D165B), and 5) VAMOSC (D160B).

MAINTENANCE DATA (D056): A total of 175 inherent failures occurred between JUL 1991 and JUN 1992, which translates into a Mean Time Between Maintenance Type-1 (MTBM-1) of 162 hours. There were no aborts reported. The MTBM-1 trend shows a decrease of 4.9 hours per month. A total of 332 maintenance actions resulted in a MTBM-Total of 85 hours. The percentage of inherent failures to total maintenance actions is 52.7%. The retest OK rate (42%) exceeds 8%. The predominant reason for maintenance actions, also called a how malfunction (HOW MAL), is "No Defect," which accounts for 32%.

Maintenance manhours included 47 scheduled manhours (0.9%), 1857 unscheduled manhours (35.0%), and 3401 shop manhours (64.1%). This results in 18.72 maintenance manhours/100 flying hours.

A MDS/WUC/base significance test was performed. The results, which can be found on page 12 of the CPA, showed that the following MDS/WUC/base combination(s) had failure rates greater than expected:

MDS	WUC	Base
B001B	55CAA	GRANDFORKS AFB ND
B001B	76GHE	ELLSWORTH AFB SD

The part number(s) that consumed a significant number of off equipment manhours were:

6960200-120
560-0000-004

SUPPORT COSTS: The average monthly support costs equal $27,830.08, which translates into a cost per operating hour of $11.78. Further details are available in the yearly support cost summary found on page 3 of the CPA.

SUPPORTABILITY ANALYSIS FORECASTING EVALUATION (SAFE): The Mean Time Between Demand (MTBD) of 680 hours for the two-year period ending MAR 92, increased 8.2% from the previous two-year period. At the base level, the number of spares required is zero. There are no base assets actually on hand. At the depot level, the total number of assets required is 10. The depot reports no assets on hand. There are no unserviceable spares. Over the two-year period there were no condemns. There are no active procurement actions for this item. Further details may be found in the SAFE report.

MICAP (Mission Capable): An item causing an aircraft to fail to meet its mission requirements accumulates MICAP hours. From JUL 91 to JUN 92, a total of 170 MICAP hours were reported, representing two active MICAP incidents. The number of aircraft having reduced mission capability over this time period is 0.0195.

ANALYSIS: In comparison to other B1 components, the reliability is poor, maintainability is high, supply is inadequate, and MICAPS are normal. The reliability could be causing a high demand on supply. The reliability of the listed B1 items appears to be improving. The plastic push buttons on the CITS panel (55AAA) are regularly broken. The item "55AAA" is one of the worst performing items in terms of MTBM-1 in the past. Some components in the 76 system on the B1 have been known to be a problem in the past. The electronic warfare (76) system has been disconnected on some of the B1s. There has been a considerable increase in demand for the items from supply, and the demand is relatively high. The Retest OK for the B1 items should improve with the implementation of new software that alleviates false alarms. The NSNs starting with 12 are in the Fire control federal stock group which contains fire control systems, radar, and sonar.

ACTIONS RECOMMENDED: Further investigate why these items have poor reliability and maintainability. The performance of these items is at such a level that management attention may be desirable. An engineering analysis may be necessary to determine if a MIP should be opened. An alternative would be to investigate the applicable T.O.'s for accuracy. Ensure maintenance workers are properly trained in repairing the equipment. Determine if the reliability is actually causing the increase in demand. Replace the push buttons if needed. The increased demand may necessitate additional procurements. Determine why the increase in demand exists. Research the cause of the high Retest OK rate.

Figure 1. An Example DRAIR produced with the DRAIR Advisor system

The Application of Artificial Intelligence to DRAIR Generation

Before attempting to apply artificial intelligence (AI)-based techniques to the automation of DRAIR generation, the Air Force tried a more conventional approach. This approach involved the use of word processing tools to provide an exhaustive, pre-defined structure for preparing DRAIRs. This format was essentially a fill-in-the-blank facility that resulted in very terse, difficult-to-read documents. And, the preparation of DRAIRs was still time consuming because the data collection and analysis tasks were not addressed. Because editing of the DRAIR was always necessary, each analyst maintained private fill-in-the-blank forms. Thus, DRAIRs were still not standardized.

The use of AI-based techniques in the generation of DRAIRs was appropriate for a number of reasons. First, based on the Air Force's previous experience in trying to automate DRAIR generation, it had become apparent that more conventional approaches were not sufficient. The generation of a DRAIR was not only an issue of data reporting. Rather, it involved the analysis and interpretation of that data with respect to specific domain knowledge about the aircraft item(s) in question. This expertise on aircraft items and their performance existed in a combination of individuals, including the OR analysts, item managers, engineers, and equipment specialists. The OR analysts usually had a high-level appreciation of the problems, while the other experts often supplied the more detailed information about items and the significance of particular failures. The use of AI-based techniques were also appropriate for this application because of the ill-defined nature of the ideal DRAIR document and the DRAIR generation process at the beginning of the system development. The iterative prototyping development methodology that is part of AI-based development was very useful in being able to show and discuss the design and development of the DRAIR Advisor system as it progressed. In addition, through this iterative development process, the OR analysts and other potential users were able to become familiar with the system early on in the work and were, therefore, more comfortable with the system when it was delivered for installation and final testing.

The particular AI-based technique used in the development of the DRAIR Advisor system was knowledge-based systems and, in particular, production rule-based knowledge representation techniques. The AI-based techniques were used in addition to a number of other, conventional software development techniques which included database design and access, and text processing/document generation. The AI-based techniques provided the intelligence for dynamic data query, analysis, interpretation, and text generation. The more conventional techniques provided the actual data access, or input to the system, and document preparation, or output from the system.

A rule-based approach was used in the representation of the expertise in the DRAIR Advisor system because the generation of a DRAIR involves the problem solving tasks of data analysis and data interpretation. As discussed above, originally to produce a DRAIR a human would collect data from a variety of databases and then analyze and interpret this data based on a knowledge of the weapon systems and components in question. The knowledge used by the human experts to perform this task tended to be very high level and heuristically-oriented. As a result, the experts were inclined to talk in terms of rules-of-thumb. They would, for example, describe their reasoning by using phrases such as "if the data for the given item is of this particular form or has this range of values, then I would conclude that . . ." The knowledge they used was not highly detailed, such as would be the case if they were performing model-based reasoning about the functionality of the item in question. Furthermore, the experts did not refer to their knowledge in terms of specific previous experiences, such as "the last time I saw this, . . ." The knowledge was experiential in nature, but high level and generalized, rather than detailed and oriented towards specific example cases. Thus, production rules were a natural way of representing the knowledge that needed to be captured for performing DRAIR generation.

Not only were rules an appropriate format for representing the knowledge to be captured for DRAIR generation, they were also appropriate from the perspective of level-of-granularity. That is, a rule could basically correlate to a generalized situation. When refinement of the knowledge base was performed, often all that was needed to correct/improve the system behavior was the modification of the conditions for which a given rule would apply or the addition of a new rule that covered an entirely new situation. No situations arose where the knowledge could not readily and easily be captured using the rule-based representation paradigm.

Though other knowledge representation techniques would have sufficed for implementing the DRAIR Advisor system, the fit of the production rule paradigm both from a representational as well as a level-of-granularity perspective allowed for much simpler and straightforward implementation and modification of the system. During development we were not required to "work around the tool" or to try to permute what was relayed to us from the experts into a different representation in order to codify the knowledge into the system.

Application Description

The DRAIR Advisor system is a multi-user/multi-job, knowledge-based software package for automatically analyzing the performance history of aircraft end items, a process previously performed manually. The DRAIR Advisor is hosted on a dedicated IBM RS 6000/930 computer running the AIX operating system with five gigabytes of data storage and 64 megabytes of memory. This computer is located at Tinker Air Force Base. All data is maintained in a single Unify 2000 database. Because the DRAIR Advisor system required the generation and maintenance of a large database, as well as the ability to query, analyze, and report on this database, a variety of software development tools were used to implement the system. These included the C language, UNIX shell scripts, Unify 2000 Relational Database Management System, Unify's Structured Query Language (SQL), Unify's RPT Report Writer Language, and the C Language Integrated Production System (CLIPS) knowledge-based system development tool (Giarratano 1990). The DRAIR Advisor system consists of 7,000 lines of C code, 1,600 lines of UNIX Shell script, 520 lines of SQL script, 143 lines of Unify Report Writing (RPT) script, and 603 CLIPS rules.

A key aspect of the DRAIR Advisor system is the high degree of integration between the use of conventional software techniques and AI-based techniques. The Unify 2000 database management tool was used to meet Air Force requirements for compatibility with existing systems. CLIPS was selected as the knowledge-based system development tool because of its ability to readily integrate with more conventional software tools, including C, database management tools, and the operating system. CLIPS was also free to the government, had no per-user licensing fee, and could generate fully compiled, executable modules. It also had a very powerful pattern matching syntax, an important capability for the task of data analysis and interpretation. Few knowledge-based system development tools could meet these capabilities in 1988, when system development began.

The DRAIR is the primary product of the DRAIR Advisor system. Two additional standard data reports, called the Cost Performance Analysis (CPA) and a partial Supportability Analysis Forecasting Evaluation (SAFE), can be obtained in conjunction with, or independent of, the generation of a DRAIR. These two reports provide additional data on the reliability, maintainability, and supportability history for an aircraft item in a conventional database reporting format. The DRAIR, on the other hand, is an English text report that describes the maintenance and supply history for one or more aircraft end items. The report is typically two pages in length. An example DRAIR produced with the DRAIR Advisor is shown in Figure 1. The top of the first page is a header that contains the date, report (or job) number, user's name, organization, and telephone number. Below the header is a line by line listing of items analyzed in the report. Each listing contains the national stock number (NSN), the application (the aircraft or "Mission Design Series" (MDS)), the work unit code (WUC), stock number noun, and the unit cost. The DRAIR contains seven main sections as follows:

SOURCE DATA: Describes the data sources used to prepare the report.

MAINTENANCE DATA: Provides a sentence by sentence description of the maintenance history for the item(s). This section discusses failures, reliability, aborts, mean time between maintenance (MTBM), manhours expended, predominant how malfunctions (HOWMALs), and any significant MDS/WUC/base combinations (due to high failure rates).

SUPPORT COSTS: Presents the average monthly support costs and cost per operating hour.

SAFE: Provides a sentence by sentence description of the supply history for the item(s). Parameters discussed include mean time between demand (MTBD), the number and condition of depot and base assets, condemnations, and plans to purchase.

MICAP: Discusses the number of MICAP (Mission Capable) hours and incidents. An item that causes an aircraft to fail to meet its mission requirements accumulates MICAP hours.

ANALYSIS: Provides an overall analysis, based on the data contained in the DRAIR, of the item's performance. This section presents both the good and the poor aspects (if any) for the item.

ACTIONS RECOMMENDED: Suggests courses of actions to correct any problems (if any) with the item(s).

This document structure existed, to some degree, before development of the DRAIR Advisor began. During DRAIR Advisor development, the structure and content of the document was formalized and codified based on input from the domain experts.

The DRAIR Advisor system uses a knowledge base that interprets the data stored in a large, mixed source database, generates recommendations concerning the item(s) based on this interpretation, and generates the text that constitutes the DRAIR document. The knowledge needed to perform the data analysis and interpretation, as well as write the report was obtained from experienced OR analysts, engineers, equipment specialists, and item managers on aircraft and federal stock class parts primarily managed at the Oklahoma City Air Logistics Center. However, the system can also handle aircraft not managed at Oklahoma City. The overall system architecture of the DRAIR Advisor appears in Figure 2. The system consists of four main components: the DRAIR database, the database maintenance facility, the user interface, and the DRAIR generation module. Each of these components is discussed in further detail below.

DRAIR Database

The DRAIR Advisor system can be characterized as a tool to support "database decision making." That is, it supports a decision making process that relies heavily on a large amount of data. As indicated above, OR analysts originally had to access a number of different, independent databases to generate specific data reports from which they could obtain the data they needed for analyzing and interpreting the status of a particular aircraft end item. The five databases accessed are illustrated in Figure 2, and are maintenance (D056), supply (D041), flying hours and sorties (G033), mission capable hours (D165), and support costs (VAMOSC). Most of these databases are older, Fortran- or Cobol-based applications that were not designed for integration with each other. In order to put all of this data onto a single machine and into a single database, a relational database design was developed that could incorporate all required data fields. Because different pieces of data were used as keys for the different databases, tables containing cross references were needed to provide a way to correlate data from one system with data from another. The resulting database, once loaded, is approximately 1.6 gigabytes in size. The database is composed of 21 tables and over 200 fields. Due to the large size of the database, search speed and retrieval were an issue. A faster data access methodology based on binary search techniques, called B-trees, was incorporated into the database design. In addition, to increase the speed of access to the data, the DRAIR Advisor system was hosted on a dedicated computer with five gigabytes of disk space. This amount of space was necessary to accommodate pre-processing of data prior to loading in the DRAIR database. The platform selected was the IBM RS 6000/930 RISC-based computer which provided not only the necessary disk storage, but also the performance needed to support the DRAIR Advisor system. The Unify 2000 relational database management system was used to implement the DRAIR database.

Database Maintenance Facility

The database maintenance facility provides a means for the individuals who use and maintain the DRAIR Advisor system at Tinker Air Force Base to update the database. Various data stored in the DRAIR Advisor system are updated, as a minimum, on a quarterly basis. In order for the DRAIR Advisor system to have the most

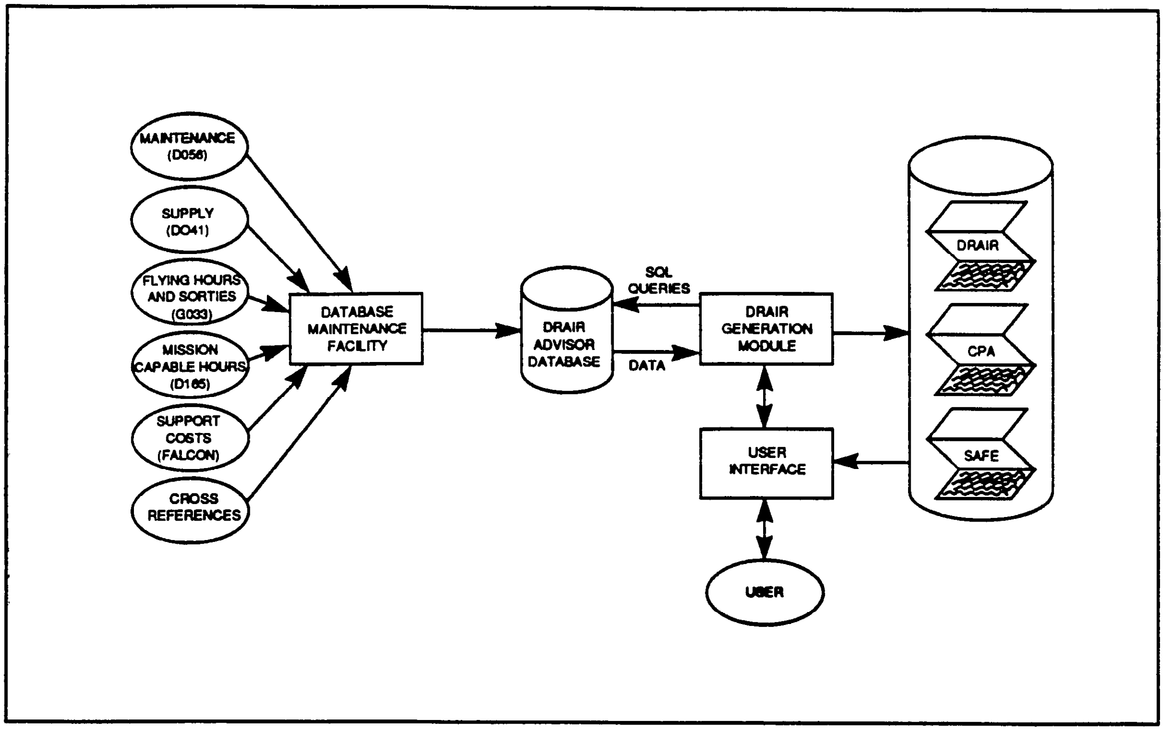

Figure 2. DRAIR Advisor system architecture

accurate and up-to-date data, software tools were developed to assist the database administrator with maintaining the timeliness of data in the DRAIR database. The database maintenance software, written in UNIX shell scripts, automates creating, deleting, and updating table definitions. Further, the software reads incoming data tapes (e.g., for D056) and loads the data into the appropriate tables.

User Interface and Database Query

The DRAIR Advisor user interface has two primary functions: to obtain the identifiers about the item or items for which the user has requested a DRAIR, and to present the resulting DRAIR to the user. All of this information is textual in nature. These item identifiers are the National Stock Number (NSN), the Mission Design Series (MDS), and the Work Unit Code (WUC). Unfortunately, the user does not always know all of this information. Thus, the system is designed to allow the user to enter either (1) the NSN, (2) the MDS and WUC, or (3) the NSN, MDS, and WUC for the item(s) under investigation. In cases (1) and (2) the system finds the missing input identifier or identifiers as appropriate. In addition, to provide flexibility for the user, the MDS and WUC can contain wildcards. Wildcards are special characters that can represent one or more unknown (or unspecified) characters in the MDS and WUC. The use of wildcards simplifies data entry for the user.

A rule-based approach was used to generate the complex database queries required to access data based on missing input identifiers and wildcards. When a user provides only the NSN, the system will generate queries to search the database for corresponding MDS/WUC combinations. When a user provides MDS/WUC combinations (including wildcards), the system will search the database tables for all corresponding NSNs and subsequently all MDS/WUC combinations for these NSNs. Additional queries to the database are generated by the DRAIR Advisor system to extract information for further analysis by the knowledge-based modules for preparation of the DRAIR, as well as for the CPA and SAFE reports.

The key driver in the user interface design was the requirement that the system be accessible over a dial-in modem. Thus, the interface had to be character-based and keyboard-driven. It was implemented in the C language and with ASCII display codes to accommodate different types of terminals. The interface consists primarily of a series of menus that guides the user through the input of the few pieces of information needed by the system. These menus were also designed to allow the user to request different reports. The design provides the user with instructions and examples for data entry. Error checking is performed on all text values entered by the user. Examples of the first three screens of the DRAIR Advisor, in which the user is queried for input, appear in Figure 3.

After a user has entered the required input for a given request, the system can be exited or another request can be made. This permits batch processing of report requests, since complex reports require processing times of up to an hour. For each request, an electronic mail message is sent to the user that informs whether the request was completed successfully and if so, where the reports have been stored. If the reports are generated as requested, the mail message contains instructions for printing or viewing the report files.

DRAIR Generation Module

The DRAIR generation module is the heart of the intelligence in the DRAIR Advisor system. It takes the data obtained from the dynamically generated database queries concerning end items to be investigated, and then uses a set of rule bases to analyze, interpret, and report the results. Thus, the module's primary input is a set of data about the item(s) in which the user is interested and its primary output is the English language, textual information report on the status of the item(s), namely the DRAIR.

The knowledge contained in the DRAIR Advisor system consists of both general knowledge about how to analyze and interpret aircraft end item data, as well as more specific knowledge about specific aircraft and federal stock classes. It also has knowledge concerning how a DRAIR is structured and what it should contain. The overall organization of the DRAIR Generation Module appears in Figure 4. It consists of several components, some oriented towards generation of the DRAIR document structure, including the DRAIR Main Template Generator and the Analysis and Actions Recommended Generator, and others oriented towards the generation of the information to be contained in the DRAIR, including the General DRAIR Knowledge Base and the Specialized DRAIR Knowledge Base.

Based on a user's input, the DRAIR Advisor system first dynamically generates the appropriate SQL scripts needed to obtain the data from the DRAIR Advisor Database as discussed in the User Interface and Database Query section, above. The data obtained is then directed to the DRAIR Main Template Generator which uses the data to generate and write the DRAIR header, including the list of item identifiers and the source data, as well as the maintenance data, support costs, supply data (SAFE), and mission capable (MICAP) sections of the DRAIR. The General DRAIR Knowledge Base contains rules about how to interpret data concerning reliability,

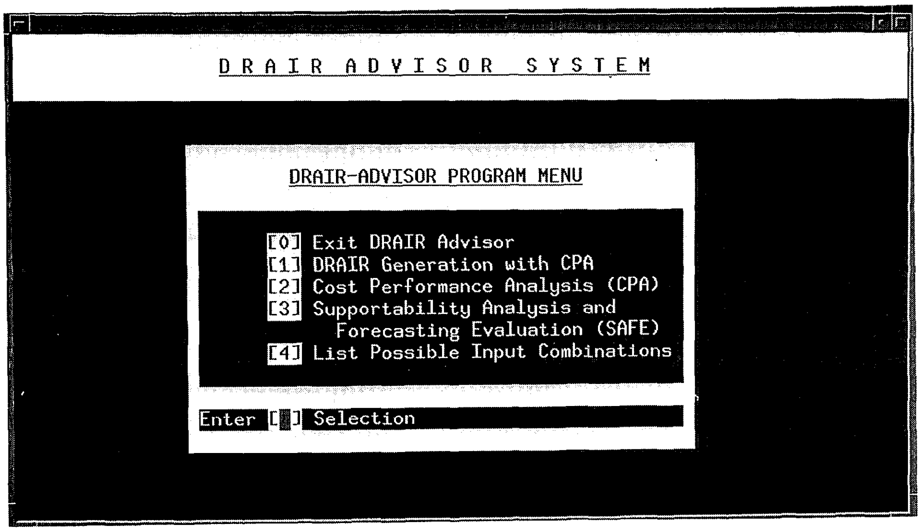

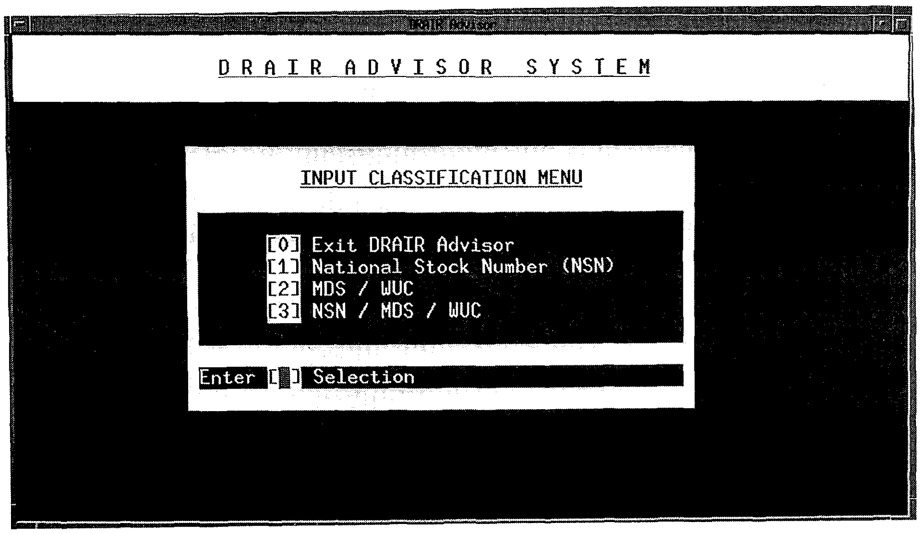

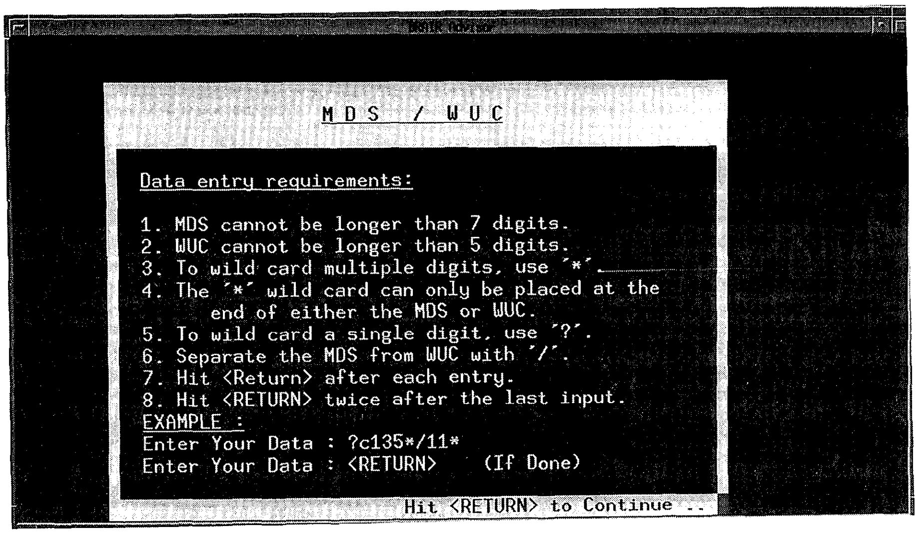

Figure 3. Example DRAIR Advisor user interface screens

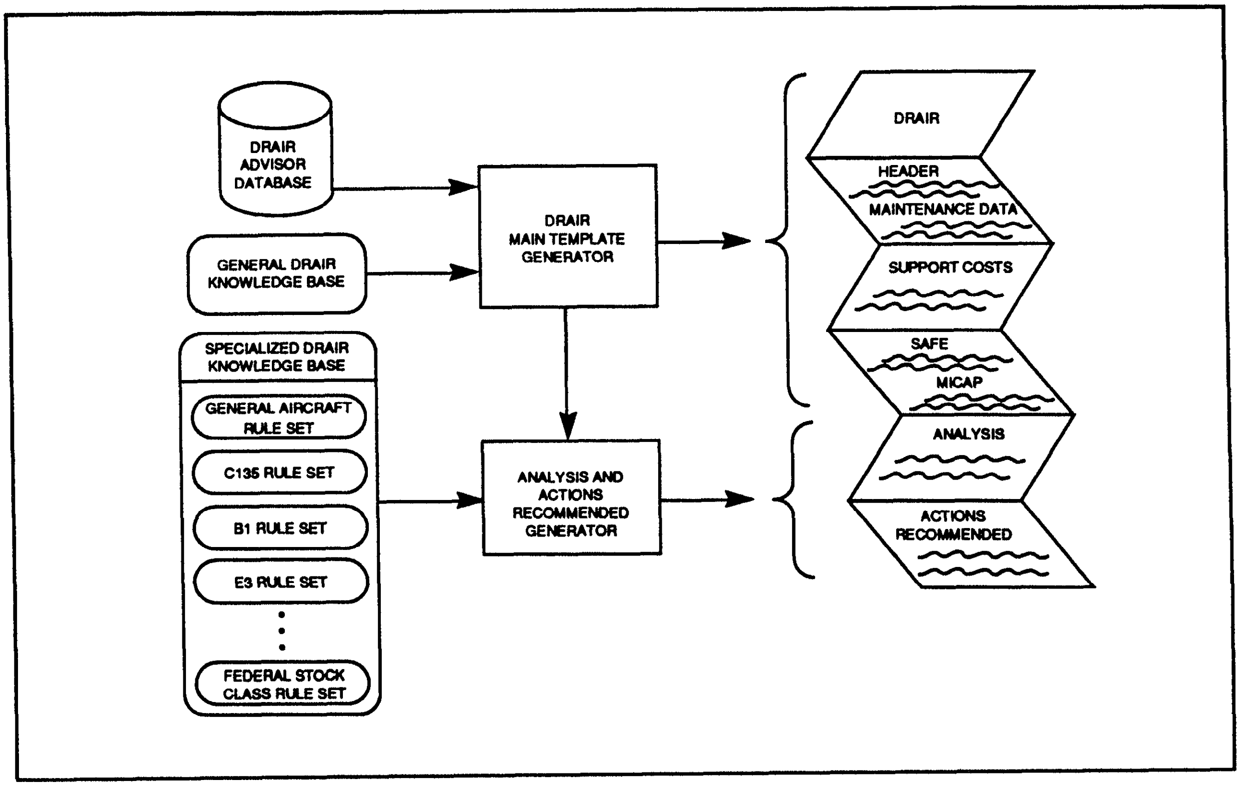

Figure 4. DRAIR Generation Module

maintainability, supply, and MICAP factors such as number of failures, MTBM (mean time between maintenance), how malfunction (HOW MAL) codes, RETEST OKs (i.e., items for which no problems are found during testing), maintenance manhours per flying hour, cannibalizations, MTBD (mean time between demand), and number of MICAP hours and incidents.

Once the DRAIR Main Template Generator has analyzed the data and generated the header and the appropriate sentences in each of the first four sections of the DRAIR, the Analysis and Actions Recommended Generator is called to generate these last two sections of the DRAIR. This module uses the Specialized DRAIR Knowledge Base which contains rule sets to handle specific knowledge about selected aircraft, such as the B-1 bomber, C-135 cargo, and E-3 (AWACS), and selected federal stock classes such as 1650 (hydraulics), 2995 (miscellaneous aircraft engine accessories), 4820 (valves), and 6615 (autopilot and gyroscopes). A general aircraft rule set is called when none of these aircraft are under consideration (e.g., the user is analyzing the F16) or when any combination of these three aircraft are under consideration.

The aircraft rule sets are written in the CLIPS language. The rules represent the knowledge of engineers, equipment specialists, and OR analysts who are experts on the particular aircraft. These rule sets use the results of the analysis that appear in the first four sections of the DRAIR. The source code for each of these rule sets is readable and understandable, and is divided into two parts: general analysis rules about the aircraft in question and specific rules about items or systems within each of the aircraft modules. The first set of rules analyzes the reliability, maintainability, and supply of the item for that particular aircraft. The analysis differs slightly for each aircraft. The result of these analyses is a set of statements that is placed in the DRAIR about the normality of these aspects of the end item. The second section of rules deals with the systems contained on the aircraft. For example, if the items are from a system that has been known to be a problem in the past, a statement about that knowledge is written to the file to be printed in the DRAIR report. Each of the systems on the aircraft has at least one piece of information about it in the rule base. Finally, there are rules about specific items (i.e., by national stock number).

Although each of the specific aircraft modules has different evaluations for the meaning of such terms as "poor reliability," each module calculates the reliability in a similar manner by considering the same types of values. The mean time between maintenance (MTBM) value determines if the reliability is poor or good. The maintenance manhours per flying hour determines if the maintainability is poor or good. To determine if supply is at an adequate or inadequate level, the overall assets are compared to the overall requirements. To determine if MICAP factors are acceptable or unacceptable, the fleet size, number of MICAP incidents, and number of MICAP hours may all be considered. Typically, if there is a significant trend, a statement is made about the MTBM trend or about the failure trend (reliability is increasing or decreasing). Each of the modules addresses the issue of a high number of aborts, RETEST OKs, condemnations, and cannibalizations (i.e., the borrowing of parts from other aircraft). An analysis is performed on whether there is a shortage of supply at the base or depot. The reasons for any shortages are also determined, if possible. In each of the modules, a secondary analysis is performed to relate maintenance with supply, flying hours with supply and maintenance, and other appropriate combinations.

After the appropriate aircraft rule set has been executed, the federal stock class (FSC) knowledge base is used as the final analysis. It produces statements that are placed in the analysis and actions recommended sections of the DRAIR report. The knowledge base uses the results not only of the main DRAIR analysis but also of the aircraft modules. There is a rule for each of the federal stock groups (the first two digits of a national stock number) that produces a description statement about that group. These rules are fired only when all of the items in the DRAIR report are contained in one of the classes. The classes for which the system currently contains specific knowledge are 1650 (hydraulics), 2995 (miscellaneous aircraft engine accessories), 4820 (valves), and 6615 (autopilot and gyroscopes).

DRAIR Advisor System Innovations

Development of the DRAIR Advisor system required the novel integration of various software development technologies. A knowledge-based system was combined with conventional programs to access a large database, perform data analysis, and interpret information. Specifically, innovativeness is present in the use of a rule base to dynamically generate complex database queries based on user input and in the use of a knowledge base to produce a textual report which varies depending on the analysis and interpretation of the data found. Particularly innovative is the accessibility of this system to users throughout the U.S. Air Force. These innovations are discussed in the following subsections.

Dynamic Database Query Using a Rule Base

The generation of a DRAIR depends on the acquisition of specific data from the DRAIR database about the aircraft item(s) in question. Though the type and source of data needed is known in advance by the system, the specific qualifiers for finding the data is not known until the user provides a request. The fact that a user can provide different inputs depending on what is known, as well as the fact that wildcards can be used to describe the desired data, complicates further the problem of acquiring the relevant data. To solve this problem, a rule-based approach was taken to generate the complex database queries required to access the required data. Rules written in CLIPS provided increased flexibility and a powerful pattern matching capability that permitted the handling of wildcarded input as well as the possibility of multiple data input. The rule base was capable of handling queries in which the user only knew one or two of the input identifiers (i.e., NSN, WUC, or MDS). A first-level set of rules was designed to dynamically construct queries to obtain any user-unknown identifiers from the database. A second-level set of rules then uses the results of the first-level queries to build additional queries that actually access the relevant maintenance and supply data.

The dynamically generated queries often become complex and unwieldy, depending on user input. Typically, 21 database tables are accessed to obtain approximately 200 different fields of data required for DRAIR, CPA, and SAFE report generation. Because of the limitations in the Unify 2000 SQL regarding the number of nested queries and tables that can be selected in a single query, it was necessary to dynamically create database views during program execution. A view is a collection of tables and fields that together represent a virtual table. The queries required to create these views are also constructed by a set of CLIPS rules.

Knowledge-Based Report Generation

The DRAIR document is an English text-based report that provides information concerning the status of one or more aircraft end items. The structure of the report and its contents is essentially a set of expertise acquired from the OR analysts that is embedded in a rule base of the DRAIR Advisor system. This knowledge base is called by a conventional C language program after the required data has been retrieved from the database. Data is made available to the knowledge base through the reading of files that contain the results of the database queries and other high level analyses. At the time the DRAIR Advisor system development began, most knowledge-based system development tools were not capable of reading standard text files. One reason the CLIPS tool was chosen was because of its ability to read and write UNIX files.

Different sentences are included in the DRAIR depending on the results of the data analysis and interpretation. In the knowledge base, rules are grouped in sets according to sentence purpose. For example, there is a set of rules for interpreting and describing the reliability of an item. What is said about reliability depends upon various factors (e.g., mean time between maintenance) and their combinations of possible values. During execution, the knowledge base determines what sentences are appropriate, modifies them based on the data, and writes them to a file. The file is written in the UNIX troff format. A C language program executes troff with this file in order to prepare the formatted report. This file is editable by the OR analyst should changes be necessary.

User Accessibility

The potentially large user base demanded that the DRAIR Advisor system be easily accessible from a variety of locations and access methods. In the past, the use of AI technology required specialized software and hardware. Commercially available AI-based software packages often require expensive development and user runtime licenses and large amounts of memory and disk space. Our goal was to overcome these limitations and successfully deploy a large, AI-based application that required little or no hardware and software investment by users. Experience has shown that in order for a new system to gain acceptance, it should be easily accessible and integrate well with existing software environments (i.e., accessible from computer hardware and software already on the user's desk).

The development team overcame these limitations through careful selection and application of both AI and conventional programming techniques. An AI-based tool was chosen that is very powerful and inexpensive. A development license for CLIPS is approximately $200 (free to the government) and is capable of generating a C-based executable. CLIPS also includes royalty-free use of the resulting runtime software. Users can access the system over DoD Internet or via telephone modem. The only software required for PC users is a telecommunications package such as Procomm. All processing is done by the IBM RS 6000 computer at the Oklahoma City Air Logistics Center. Reports can be down-loaded and printed at user sites since they are simply ASCII text files.

Application Use and Payoff

The DRAIR Advisor system has been well received by the targeted users. The system is used by OR analysts, engineers, equipment specialists, and logisticians in the Air Force Materiel Command. Currently the system is used by personnel at the Oklahoma City and San Antonio Air Logistics Centers (ALCs) as well as at Ellsworth Air Force Base in South Dakota. At Oklahoma City alone, approximately 25 DRAIRs are produced per month by 15 different users. Other Air Force bases, including Warner Robins ALC (Georgia), Sacramento ALC (California), and Ogden ALC (Utah), are expected to begin using the DRAIR Advisor in the near future. The combined user population of these bases alone would exceed 2,000.

Acceptance and usage of the DRAIR Advisor has been positive, in part, due to system accessibility. No specialized hardware or software licenses are required. Users can dial in to the system via a modem or use a workstation or PC connected to the DoD Internet. At Oklahoma City ALC, users typically use PCs that are connected to the base-wide local area network (LAN). As part of system deployment, the IBM RS 6000 computer was connected to the LAN and access to the DRAIR Advisor was made a menu-selectable option on the main computer at the Oklahoma City ALC.

Accessibility of the system by engineers and equipment specialists has reduced significantly the time that the OR analysts at Oklahoma City spend on DRAIR generation. Generation of a DRAIR requires only a few minutes of time to enter the necessary data; most reports are processed by the computer within an hour. The OR analysts are now able to spend additional time on other job functions (which have always been in their job descriptions). Even though the number of OR analysts in the office at the Oklahoma City ALC has been reduced, more DRAIRs are now being produced. Reports produced by the DRAIR Advisor system are also standardized. Reporting preferences among the analysts were consolidated into a single DRAIR format. In addition, potential errors and oversights by analysts have been eliminated through the use of the automated system.

Management of the B-1 bomber at the Oklahoma City ALC relies heavily on the DRAIR Advisor system to provide accurate, up-to-date reports on weapon system performance. Based on the analysis and actions recommended sections in the DRAIR, courses of action have been determined for aircraft items. Furthermore, decisions have been made regarding the purchase of spare parts. Another reason for system success has been the confidence users have in the information provided by the DRAIR Advisor system. DRAIRs are often produced by product improvement working group (PIWG) members for discussion at quarterly meetings. Also, the DRAIR can be used by engineers, equipment specialists, and item managers to identify deficiencies in aircraft parts before they become a concern. Active participation by end-users during system development and evaluation contributed to this acceptance of the system. With new users at other bases accessing the system, these system benefits are anticipated for other weapon systems as well.

In order to obtain funding to develop the DRAIR Advisor, the Air Force had to justify the required time and money quantifiably. An automated system for DRAIR generation is expected to provide a cost avoidance of approximately $120,000 per year and to save 1,900 person-hours per year. Amortization of the system began in October of 1991 (with installation of a working version) and will be complete in September of 1995. Though the system will not pay for itself quantifiably for another two years, an immeasurable number of immediate, long-lasting qualifiable benefits such as those described above already have been realized.

Application Development

The DRAIR Advisor system was developed in two major phases, an initial prototyping phase that provided proof-of-concept as well as a limited working system, and a full-scale development phase that expanded the working prototype into a complete DRAIR Advisor system. The overall cost of development was approximately $500,000 and took place over a period of three and a half years. The development team for both phases consisted of Southwest Research Institute and the U.S. Air Force Technology & Industrial Support Directorate at Tinker Air Force Base. The Institute was responsible for overall system design and implementation. The Air Force provided all of the domain expertise as well as considerable input on system design and functionality. In addition, the Air Force provided extensive input on each iterative delivery of the system and had considerable responsibility in the installation and testing of the system before final delivery.

The prototype was developed in approximately one and a half person-years of effort over a period of nine months (Robey et al. 1990). It was completed in early 1989. Prior to development of the prototype, preparation of a DRAIR required approximately three-person days of effort that included accessing many computer systems. The prototype reduced this to a few hours and proved the potential for an automated DRAIR generation system. However, because the initial effort was only a prototype, the DRAIR Advisor had limitations that restricted widespread usage and access. For example, the database only contained maintenance data (D056 Product Performance System) for aircraft maintained at Oklahoma Air Logistics Center. Further, the prototype system required all three known aircraft identifiers - the mission design series (MDS), the work unit code (WUC), and the national stock number (NSN). Often all three identifiers were not available to the individual wishing to generate a DRAIR. The prototype system was developed on a DEC VAX 8650 with limited storage and memory; only one user could use the prototype system at a time. In addition, the computer was shared with a number of other large database applications which also limited the system speed and the availability of disk space. However, the success of the prototype led to a demand for expansion and full-scale development of the DRAIR Advisor.

Full-scale development of the system began in mid-1990 and was completed in January of 1992. Approximately four person-years of effort were required. The full-scale development effort for the DRAIR Advisor system focused on five areas: (1) database expansion, (2) user input, (3) knowledge base expansion and additions, (4) multiple-user access, and (5) overall system performance. The prototype DRAIR Advisor database was expanded to include maintenance data for all five Air Force Air Logistics Centers. Software tools were written to assist the database administrator in maintaining the DRAIR database. User input was simplified by including cross-referencing data for MDS/WUC to NSN in the database. This permitted users to enter either the MDS/WUC combination or the NSN or all three identifiers. In addition, expansion of the user interface allowed users to wildcard their input for MDS and WUC. The knowledge base that interprets the data and suggests recommendations was expanded to contain expertise from engineers and equipment specialists on a set of pre-specified aircraft and federal stock classes. This allowed for more detailed analyses and recommendations on certain classes of MDS/WUC inputs while maintaining the ability provided in the original prototype knowledge base to reason in a general sense about other aircraft and federal stock classes. Extensive knowledge engineering was performed to acquire, code, test, and refine this knowledge. To accommodate future system growth, a modular approach was taken and a methodology was developed to allow simple expansion of the knowledge base to include additional aircraft and federal stock classes. Multiple-user and multiple-session capabilities were implemented by assigning and tracking unique job identification numbers. Overall system performance was improved by rehosting the system on a dedicated IBM RS 6000/930 computer running AIX with five gigabytes of disk space and 64 megabytes of memory.

During both the initial prototyping and the full-scale development efforts, the DRAIR Advisor system development approach proceeded through the five stages of knowledge-based system development: problem identification, conceptualization, formalization, implementation, and testing (Buchanan & Shortliffe 1984). These five stages were repeated several times, resulting in a number of intermediate deliveries of the system to the domain experts and selected end users. These intermediate deliveries allowed the users to clearly see the system and provide concrete feedback concerning system design, functionality, and performance. It also provided a means of iteratively testing the system for correct and reasonable behavior. This knowledge-based system development approach proved very powerful, as it permitted highly modular development of the software.

The primary mode of knowledge acquisition employed was interviews. However, printed resources (e.g., Air Force technical manuals) were also reviewed. Based on an initial formalization of the knowledge required to analyze aircraft end item performance, a rule-based development environment was chosen. This approach allowed relatively straightforward representation of the knowledge obtained from experts. Experts were selected based on their experience, ability to articulate knowledge, and personal interest in the project.

A key aspect to the success of the DRAIR Advisor system was that the Air Force had a highly motivated, proficient champion for the project from the start of the initial prototype through full-scale development, and beyond into fielding and maintenance. In addition, success of the DRAIR Advisor required more than just an experienced capability in artificial intelligence. Development and deployment of the system depended heavily on an interdisciplinary team of individuals knowledgeable in artificial intelligence specifically, as well as in database design, software engineering, and computer science in general. It also required dedicated, open-minded, forward-thinking experts in the domain of application that were willing to provide time and input throughout the development process.

Deployment of the DRAIR Advisor

The full-scale DRAIR Advisor system was deployed on an IBM RS 6000/930 computer running AIX (IBM's version of UNIX) and the Unify 2000 relational database management system. The Air Force officially designated the DRAIR Advisor as the G050 system. All software is licensed for up to 16 simultaneous users. The IBM RS 6000 is connected to the Air Logistics Center (ALC) LAN and to the DoD Internet.

Users access the system over the Air Logistics Center LAN, DoD Internet, or telephone modem. A menu selection for the DRAIR Advisor was added to the central computer at Oklahoma City Air Logistics Center. This central computer is connected to the DoD Internet which supports remote login. Minimal training is required to access and run the DRAIR Advisor. New users are able to run the system with little or no assistance since the user interface consists of only five menus and a maximum of three data input types.

The DRAIR Advisor was officially deployed on January 31, 1992. However, the iterative development process allowed the system to be operational starting in October, 1991. During this four-month operational period, domain experts and end-users contributed significantly to knowledge base verification and refinement and to user interface design. Verification, or confirming that the report output is as intended, was simple since the domain experts were closely involved in knowledge base development and available to review results throughout system development. Validation of the knowledge base was accomplished by allowing potential users (i.e., equipment specialists, engineers, and item managers) not involved with system development to run the system, obtain reports, and provide comments.

Overall system administration was a key issue during deployment of the DRAIR Advisor. Administration is necessary for user accounts, the database, and the operating system. A key project team member from the Air Force assumed system administrator responsibilities. A complete set of documentation was prepared to assist the system administrator. These documents included a user's guide, a database administrator's guide, a programmer's reference guide, and a source code listing.

System Maintenance

System maintenance was addressed during the development of the DRAIR Advisor. In order to maintain up-to-date reporting capabilities, maintenance is required for the database and the knowledge base. The management of user accounts is also essential. For maintenance of the database and user accounts, the system administrator uses software developed specifically for these tasks. Modifications to the knowledge base can also be made by the system administrator, who was the primary domain expert and assisted in knowledge base development. Maintenance approaches for the database, knowledge base, and user accounts are described below.

Database Maintenance

The system administrator maintains the DRAIR database with assistance from the database maintenance facility described above. Database tables are updated monthly or quarterly, depending on source (e.g., D056, SAFE, MICAPs etc.). The timeliness of reports depends on efficient, regular updating of the database. With this software, the system administrator has the flexibility to modify the database design to support data format changes or expansion needs.

The data stored in the DRAIR database changes over time, with recent maintenance activities being added and older ones being deleted. This data updating is performed at various intervals depending on the source. The data has historically been obtained by the Air Force via memorandums of agreement with the supplying agencies. Data is typically transferred on nine-track magnetic tape media. In most cases, the data exists in multiple volumes of tape media. The system administrator is responsible for updating the DRAIR database as necessary to maintain the required data type overlap for successful DRAIR Advisor system operation. For example, if the database maintains two years worth of SAFE data, the D056 data in the database must reside within those two years.

Menu-driven software tools were written using UNIX shell scripts, the C language, and UNIFY 2000 SQL scripts to read the data from tape media, process the data, and load it into database tables. Due to variations in data format on the tapes, the software was designed to assist in selecting the best format to use when reading a tape. The software allows batch processing for loading of data, as this task is time consuming. The software also allows the system manager to modify, backup, and update the database and the data dictionary. These capabilities are essential to accommodate changes in data types and formats provided by the supplying agencies. Once the data is available electronically over the Air Force computer network, the plan is to obtain the update data over the network and the use of data tapes will not be as vital to the maintenance of the DRAIR Advisor system.

Knowledge Base Maintenance

The DRAIR Advisor knowledge base was designed for maintenance by the principal domain experts (i.e., the OR analysts). The modular design permits the domain experts to maintain, update, and enhance the aircraft and federal stock class knowledge. In fact, within the first three months of deployment, the domain experts had successfully modified the E-3 aircraft knowledge base to include changes in analysis criteria. In the future, the Air Force plans to include specific knowledge about additional aircraft.

A knowledge acquisition methodology was developed specifically for obtaining expertise in federal stock classes. A set of template rules were designed for use in knowledge base expansion. The DRAIR Advisor system is expected to be expanded to include detailed knowledge on all 400 federal stock classes. Selected domain experts were trained in the methodology through active participation in interviews, rule generation, knowledge base modification, and compilation. An example rule is shown in Figure 5. The six basic steps to the methodology are listed in Figure 6. Verification and validation of newly added knowledge is easy for the domain experts since testing of new rules can be accomplished without modification to the other system components. Further, verification is simple because the expert is the one who placed the knowledge into the system.

```
;-------------------------------------------------------------;
;       Any problem with an item in this FSC                  ;
;-------------------------------------------------------------;

(defrule fsc_problem_.... "values match aircraft modules"

  (nsn ?nsn &: (eq 1 (str-index "..." ?nsn)))
  ;;; delete unwanted values below
  (or (reliability low|very_low|very_very_low)
          (maintainability high|marginal)
          (supply inadequate|not_ok)
          (micaps unacceptable|high|marginally_high)
          (aborts ?abort &: (> ?abort 0))
  )
  (not (comment fsc_problem_....))

=>

  (assert (comment fsc_problem_....))
  (printout action "Problems with the ... items should be investigated further.")
)
```

Figure 5. Template rule for the DRAIR Advisor knowledge base

(1) Query the DRAIR database using this query:

```
SET CURRENT SCHEMA TO drair_data;
select unique nsnnoun, count(*) from falcon
where fcnsn shlike '....*' group by nsnnoun
into 'fsc-subclass.txt';
```

In place of .... put the 4 digits of the federal stock class (FSC). You can put this query in a file such as fsc-subclass.sql and enter the command:

```
SQL fsc-subclass.sql
```

The resulting "fsc-subclass.txt" file will have all NSN nouns that are in the FSC.

(2) Use this list to define subclasses by using the nouns directly or grouping some of them together into a subclass. Check the H6-1 for descriptions of sample nouns, if possible, to use as descriptions in some analysis statements.

(3) Run an example CPA on as many subclasses as possible to find typical HOWMALs.

(4) Locate and interview the engineers and equipment specialists who are most familiar with the whole FSC. Use example questions found in the system documentation. Attempt to find the best communicative and willing person to use as the primary expert and for feedback of the resulting DRAIR reports.

(5) Fill in the template rules in the "fsc.clp" file with the rules and put the new rules in numerical order of FSC.

(6) Test, produce DRAIR reports, get feedback and revise the code (edit the "fsc.clp" file and compile the source code). See Section 5 of the Programmers Reference Guide for additional information on recompiling source code.

Figure 6. Steps for adding federal stock class knowledge

User Access

Software was written using UNIX shell scripts to assist the system administrator in managing user accounts. The software automates the setup of user accounts to provide access to the DRAIR Advisor system, including database authorizations. For new users, the software sends electronic mail that provides instructions for using the DRAIR Advisor system. Privileges can also be removed from inactive accounts. Electronic mail, in many cases automatically generated, is used for all communication with users.

Conclusions

The work described represents one of the first fielded application of knowledge-based systems technology in the Air Force materiel management environment. Because of widespread user accessibility and enthusiastic acceptance, the DRAIR Advisor system has become one of the most highly recognized, successful programs in artificial intelligence undertaken by the Air Force Materiel Command. The DRAIR Advisor system is used by operations research analysts, engineers, item managers, and equipment specialists to obtain fast, up-to-date, concise reporting on the performance status of aircraft parts. Actual use of the DRAIR Advisor has resulted in both qualitative (e.g., higher quality reporting that affects courses of action) and quantitative (e.g., time and money savings) benefits. The Air Force plans to build additional software systems that will utilize report information obtained from the DRAIR Advisor.

References

Buchanan, B.G. and Shortliffe, E.H. 1984. *Rule-Based Expert Systems*. Reading, MA: Addison-Wesley Publishing Company.

Giarratano, J.C. 1990. *CLIPS Reference Manual*. Houston, TX: National Aeronautics and Space Administration.

Robey, B., Fink, P., Redfield, C., and Ferguson, J. 1990. A Knowledge-Based System for Materiel Deficiency Analysis. In proceedings of the Artificial Intelligence Applications for Military Logistics Symposium and Workshops, 272-277. Williamsburg, VA: American Defense Preparedness Association.