

2

AFWAL-TR-88-2096

INTERFACE II - ADVANCED DIAGNOSTIC SOFTWARE




L.R. LaPierre, P.A. Costen, & D.L. Doel
GE Aircraft Engines
Cincinnati, Ohio 45215
R.P. Shah
GE Corporate Research & Development
Schenectady, New York 12301

December 14, 1988

Final Report for Period July 1985 - July 1988

Approved for Public Release, Distribution is Unlimited



AIR FORCE WRIGHT AERONAUTICAL LABORATORIES
AERO PROPULSION LABORATORY (AFWAL/POTA)
AIR FORCE SYSTEMS COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OHIO 45433-6563

AD-A204 527

NOTICE

When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely Government-related procurement, the United States Government incurs no responsibility nor any obligation whatsoever. The fact that the Government may have formulated or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication or otherwise as in any manner construed as licensing the holder or any other person or corporation; or as conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

This report has been reviewed by the Office of Public Affairs (ASD/CPA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

SEYMOUR G. PRZYTYLO
Project Engineer
FOR THE COMMANDER


LESTER L. SMALL, TAC
Engine Integration & Assessment Branch


ROBERT E. HENDERSON
Deputy for Technology
Turbine Engine Division

If your address has changed, if you wish to be removed from our mailing list, or if the addressee is no longer employed by your organization please notify AFWAL/POTA, Wright-Patterson AFB, OH 45433-6563 to help us maintain a current mailing list.

Copies of this report should not be returned unless return is required by security considerations, contractual obligations, or notice on a specific document.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED			1b. RESTRICTIVE MARKINGS										
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for Public Release; Distribution is unlimited.										
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE													
4. PERFORMING ORGANIZATION REPORT NUMBER(S)			5. MONITORING ORGANIZATION REPORT NUMBER(S) AFWAL-TR-88-2096										
6a. NAME OF PERFORMING ORGANIZATION GE Aircraft Engines		6b. OFFICE SYMBOL (If applicable)		7a. NAME OF MONITORING ORGANIZATION Aero Propulsion Laboratory (AFWAL/POTA) Air Force Wright Aeronautical Laboratories									
6c. ADDRESS (City, State and ZIP Code) One Neumann Way Evendale, Ohio 45215-6301		7b. ADDRESS (City, State and ZIP Code) Wright-Patterson Air Force Base, Ohio 45433-6563											
8a. NAME OF FUNDING/SPONSORING ORGANIZATION Air Force Wright Aeronautical Laboratories; Aeropropulsion Laboratory		8b. OFFICE SYMBOL (If applicable) AFWAL/POTA		9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F33657-85-C-2131									
8c. ADDRESS (City, State and ZIP Code) Wright-Patterson Air Force Base, Ohio 45433-6563		10. SOURCE OF FUNDING NOS.											
		<table border="1"><tr><td>PROGRAM ELEMENT NO.</td><td>PROJECT NO.</td><td>TASK NO.</td><td>WORK UNIT NO.</td></tr><tr><td>63202F</td><td>668A</td><td>01</td><td>36</td></tr></table>				PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.	63202F	668A	01	36
PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT NO.										
63202F	668A	01	36										
11. TITLE (Include Security Classification) INTERFACE II - Advanced Diagnostic Software													
12. PERSONAL AUTHOR(S) Lee R. LaPierre, Peter A. Costen, David L. Doel (GE Aircraft Engines); and Rasik P. Shah (GE Corporate Research and Development)													
13a. TYPE OF REPORT Final		13b. TIME COVERED FROM 85 July to 88 July		14. DATE OF REPORT (Yr., Mo., Day) 1988, December 14									
				15. PAGE COUNT 167									
16. SUPPLEMENTARY NOTATION													
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)										
FIELD	GROUP	SUB. GR.											
21	05		(See Reverse)										
19. ABSTRACT (Continue on reverse if necessary and identify by block number) (See Reverse)													
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS ☐			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED										
22a. NAME OF RESPONSIBLE INDIVIDUAL Stephen J. Przybylko			22b. TELEPHONE NUMBER (Include Area Code) (513) 255-6690		22c. OFFICE SYMBOL AFWAL/POTA								

18. Subject Terms:

Artificial Intelligence
CD-ROM
CEMS (Comprehensive Engine Management System)
Condition Monitoring
Diagnostics
Expert Systems
Gas Turbine Engines
Jet Engines
Kalman Filter
Knowledge-Based Systems
Reliability and Maintainability
TEMS (Turbine Engine Monitoring System)

19. Abstract:

JET-X is an expert system for the diagnosis and maintenance of the TF34-100 jet engine as installed on the USAF A-10A aircraft. JET-X uses input supplied by the TEMS (turbine engine monitoring system) installed on the airplane and, in addition, uses data retrieved from the CEMS (comprehensive engine management system) data base that is a part of the ground computer support system. Both of these monitoring systems generate alarms which are the starting points for JET-X analyses. As part of an earlier contract, diagnostic troubleshooting procedures had been established for resolving each of these alarms. These diagnostic procedures have been embedded into the JET-X system, generally in a much expanded form. In addition, a number of "help" facilities have been developed for JET-X to assist the novice diagnostician. JET-X has been designed to be both a flight-line diagnostic tool and a training aid. Although the system is not complete at this point, a first good prototype has been developed, and some field experience has been gained in order to evaluate the system and to guide future enhancements. This report describes the development of the JET-X system, its features and limitations, results of the field trial, and gives conclusions and recommendations for future work. JET-X has been developed as a part of the INTERFACE IIL Contract (F33657-85-C-2131), sponsored by AFWAL/POTA (of the Air Force Systems Command).

Table of Contents

Section		Page
1.0	Introduction/Program Overview	1
1.1	Program Objectives	2
1.2	Program Timing and Milestones	2
1.3	JET-X Implementation	3
2.0	Results and Conclusions	4
3.0	Recommendations	7
4.0	A-10/TF34-100 Overview	8
4.1	TF34-100 Engine Description	8
4.2	A-10 Aircraft Description	11
5.0	TF34 Engine Monitoring System	12
5.1	TEMS Overview	12
5.2	CEMS IV Overview	12
5.3	Diagnostic Maintenance Procedures with TEMS and CEMS IV	12
6.0	JET-X Overview	15
6.1	JET-X Components	15
6.2	JET-X Interface with CEMS IV	16
7.0	JET-X System Design	19
7.1	JET-X System Architecture	19
7.2	Modular Knowledge-Base Construction	25
7.3	JET-X Knowledge Implementation	25
7.4	Human Factors Issues	27
7.5	Report Generation	28
8.0	JET-X Diagnostics	31
8.1	Symptom-Fault Approach	31
8.2	User's Role in JET-X Diagnostics	31
8.3	JET-X Diagnostic Scope	32
8.4	TEMS and CEMS IV Alarm Coverage	34
8.5	GEN-X Module Types	39



Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
Dist	Original
A-1	

Table of Contents (Continued)

Section	Page
8.6	Modular Knowledge-Base Implementation 39
8.7	Knowledge-Base Execution 40
9.0	Using JET-X 49
9.1	GEN-X End-User Interface Description 49
9.1.1	Procedure Window 49
9.1.2	History Window 49
9.1.3	Explanation Window 49
9.1.4	Pop-Up Menu 51
9.1.5	Title 51
9.1.6	Footer Menu 51
9.2	Interface with External Programs 51
9.3	JET-X Session Features 52
9.3.1	User-Name Entry 52
9.3.2	Activity Selection 52
9.3.3	Diagnose Versus TEMPER 52
9.3.4	Session Type 52
9.3.5	Input ESN (Engine Serial Number) 53
9.3.6	Input Alarms 53
9.3.7	Guided Versus Symptom Troubleshooting 53
9.4	Sample Session: Guided Troubleshooting 53
9.5	Sample Session: Symptom Troubleshooting 77
10.0	JET-X External Support Routines 87
10.1	JET-X Retrieval System 87
10.1.1	JRS Data Flow and Modules 94
10.1.2	JRS Interfaces 94
10.1.3	Role of JRS in Future Expert Systems 97
10.2	Setfacts 97
10.3	Other Support Routines 97
11.0	JET-X Field Evaluation Summary 102
11.1	Test Description 102
11.2	JET-X Diagnostic Performance 102
11.3	System Evaluation 104
11.4	Overall System Assessment 106
12.0	Integration of TEMPER with JET-X 107
12.1	Description of Current TEMPER Algorithm 107
12.2	Need for Improvement in TEMPER Gas-Path Analysis 108

Table of Contents (Concluded)

Section	Page
12.3 Application of TEMPER to the TF34 Engine	109
12.4 Comments on Idealized Integration of TEMPER with JET-X	117
12.5 JET-X TEMPER Interface	119
 13.0 References	 122
 Appendixes	
A AI (Artificial Intelligence) Tools for Diagnostics and GEN-X Overview	A-1
B Knowledge Source(s) and Extraction	B-1
C Background on Gas-Path-Analysis Tools and TEMPER	C-1
C.1 Gas-Path-Analysis Tools	C-1
C.2 Turbine Engine Module Performance Estimation Routine (TEMPER)	C-4
C.3 General Description	C-4
C.4 On-Wing TEMPER Development	C-15
C.5 References	C-25
Nomenclature	C-26

List of Illustrations

Figure		Page
1.	TF34-100 Engine Cut-Away.	9
2.	TF34-100 Trim Curve.	10
3.	Sample Troubleshooting Procedure for "Low Performance" (CP1) Alarm.	13
4.	JET-X Equipment Setup.	17
5.	TEMS/CEMS IV with JET-X.	18
6.	JET-X Knowledge-Base Architecture - High Level Functions.	20
7.	JET-X Knowledge-Base Architecture - Diagnostic Functions.	21
8.	JET-X Introduction Screen.	22
9.	JET-X Glossary Screen.	22
10.	JET-X Symptoms Selection Screen.	24
11.	Example of "If/Then" Rule Table in Development Environment.	26
12.	JET-X Reports Selection Screen.	29
13.	JET-X Engine Serial Number Selection Screen for Reports.	29
14.	Sample JET-X Engine Troubleshooting Record.	30
15.	Sample: T.O. No. 1A-10A-2-71MS-5 Troubleshooting Algorithm.	37
16.	Module CP3TEST.IFT.	43
17.	Module GASPTDMG.IFT.	44
18.	Module VIBANAL.DEC.	45
19.	Module WRMTLTST.IFT.	46
20.	Module WRMTLCHK.DEC.	46

List of Illustrations (Continued)

Figure		Page
21.	Module MTLSIGN.AOR.	47
22.	Module MTLFLAGT.DEC.	47
23.	GEN-X End-User Screen.	50
24.	Frame No. 2 for JET-X Guided Troubleshooting Example.	54
25.	Frame No. 6 for JET-X Guided Troubleshooting Example.	55
26.	Frame No. 7 for JET-X Guided Troubleshooting Example.	56
27.	Frame No. 8 for JET-X Guided Troubleshooting Example.	57
28.	Frame No. 14 for JET-X Guided Troubleshooting Example.	58
29.	Frame No. 15 for JET-X Guided Troubleshooting Example.	59
30.	Frame No. 15(a) for JET-X Guided Troubleshooting Example.	60
31.	Frame No. 15(b) for JET-X Guided Troubleshooting Example.	61
32.	Frame No. 16 for JET-X Guided Troubleshooting Example.	62
33.	Frame No. 17 for JET-X Guided Troubleshooting Example.	63
34.	Frame No. 18 for JET-X Guided Troubleshooting Example.	64
35.	Frame No. 19 for JET-X Guided Troubleshooting Example.	65
36.	Frame No. 21 for JET-X Guided Troubleshooting Example.	66
37.	Frame No. 22 for JET-X Guided Troubleshooting Example.	67
38.	Frame No. 23 for JET-X Guided Troubleshooting Example.	68
39.	Frame No. 24 for JET-X Guided Troubleshooting Example.	69
40.	Frame No. 25 for JET-X Guided Troubleshooting Example.	70

List of Illustrations (Continued)

Figure		Page
41.	Frame No. 26 for JET-X Guided Troubleshooting Example.	71
42.	Frame No. 27 for JET-X Guided Troubleshooting Example.	72
43.	Frame No. 28 for JET-X Guided Troubleshooting Example.	73
44.	Frame No. 29 for JET-X Guided Troubleshooting Example.	74
45.	Frame No. 29(a) for JET-X Guided Troubleshooting Example.	75
46.	Frame No. 30 for JET-X Guided Troubleshooting Example.	76
47.	Frame No. 2 for JET-X Symptom Troubleshooting Example.	77
48.	Frame No. 6 for JET-X Symptom Troubleshooting Example.	78
49.	Frame No. 7 for JET-X Symptom Troubleshooting Example.	79
50.	Frame No. 8 for JET-X Symptom Troubleshooting Example.	80
51.	Frame No. 10 for JET-X Symptom Troubleshooting Example.	81
52.	Frame No. 11 for JET-X Symptom Troubleshooting Example.	82
53.	Frame No. 12 for JET-X Symptom Troubleshooting Example.	83
54.	Frame No. 13 for JET-X Symptom Troubleshooting Example.	84
55.	Frame No. 15 for JET-X Symptom Troubleshooting Example.	85
56.	Frame No. 17 for JET-X Symptom Troubleshooting Example.	86
57.	Sample Engine Cross Section Schematic in JRS.	90
58.	Sample CEMS IV Tabular Display in JRS.	91
59.	Sample CEMS IV Plot Display in JRS.	92
60.	Sample Borescope Photograph in JRS.	93

List of Illustrations (Concluded)

Figure		Page
61.	Data Flow and Module Breakdown for JRS.	95
62.	Illustration of JRS Page Window and Control Panel.	96
63.	Input of TEMS and CEMS IV Alarms, Using Setfacts.	98
64.	Sample of TF34 TEMS Data.	110
65.	TF34 TEMPER Response Analysis.	114
66.	JET-X Screen for "Raw" TEMPER Output.	120
67.	JET-X Screen for "Filtered" TEMPER Output.	120
68.	Maintenance Recommendations, Based on TEMPER.	121

List of Tables

Table		Page
1.	TEMS and CEMS IV Alarms for TF34.	35
2.	Knowledge-Base Execution for a Sampling of CEMS IV Performance Alarms.	41
3.	Video Displays Available in JRS.	88
4.	JET-X External "C" Programs.	100
5.	Field Evaluation of JET-X Diagnostics.	103
6.	TF34 Gas-Path-Analysis Measurements.	111
7.	Standard Deviations for TF34 TEMPER.	113

1.0 INTRODUCTION/PROGRAM OVERVIEW

With the increased complexity of jet engines and the increased awareness of the impact of maintenance problems on the life cycle cost of weapon systems, there has come an increased demand for effective diagnostic tools for engine maintenance. This demand is fueled by the large number of CND's ("cannot duplicate's") and RTOK's ("retest ok's") that are experienced in ordinary service. Studies have suggested that repair often consists of the serial replacement of suspected components until the functionality of the engine is restored. Each removed part must undergo subsequent testing and possible repair before it can be returned to service. There are programs underway to address this difficulty (such as, the GIMADS Contract); however, it is clear that more effective diagnostic tools will positively impact the situation.

The traditional approach to maintenance diagnostics (which, heretofore, has consisted of the use of a number of tools, each of which addresses only a part of the problem) has contributed to this unfortunate situation. The diagnostician is asked to juggle a number of different results in his or her head in order to achieve a unified diagnosis. Frequently, the key "connections" are not in any of the tools and must be learned through experience. This is not necessarily correctable; however once a connection is learned by one diagnostician, there should be a means to propagate this information to others.

In a military setting, some problems may be exacerbated by a lack of experience. At a location such as Barksdale AFB in Louisiana, there is substantial experience available for troubleshooting jet engines. Maintenance people at Barksdale generally have several years of experience with the same engine; thus, they know how to recognize and correct virtually all problems. At the other extreme, there are bases such as Suwon AFB in South Korea where, typically, the length of stay for a mechanic is of the order of 1 year. Clearly, this does not allow for the development of sufficient experience to be effective in jet engine diagnosis. (Some of the more difficult diagnostic problems may not occur as often as once per year.)

Fortunately, the field of AI (artificial intelligence) appears to have developed a solution to this problem. The solution referred to is called the knowledge-based system or, more popularly, the expert system. Knowledge-based systems are differentiated from other, conventional computer software in that they separate the instructions (or rules of operation) from the inference mechanism. This facilitates the development and maintenance of decision-based, procedural codes relative to conventional programming languages such as FORTRAN. Knowledge-based systems have been applied to a number of types of problems; however, the most consistent area of success has been in diagnostics. This should not be surprising, since the knowledge-based system facilitates the coding of the connections that were referred to above.

This report describes the first stage of the development of a diagnostic expert system that will be used to diagnose TF34 engine problems. The system has been developed by a team of people working for GEAE (GE Aircraft Engines) and is specifically aimed at the TF34-100 engine

on the A-10A aircraft. The knowledge-based system which is described herein has been dubbed JET-X (for Jet Engine Troubleshooting eXpert). This work has been sponsored by AFWAL/POTA (a division of the Air Force Systems Command) under Contract No. F33657-85-C-2131.

1.1 Program Objectives

The overall objective of the JET-X Program was to explore the application of expert system technology to aircraft gas turbine engine diagnostics. It was decided that the best opportunity for successful demonstration of this objective was to address the analysis of airborne acquired monitoring system data on the ground (postflight). Often, effective analysis of this data can be beyond the capabilities of engine mechanics, especially novices; therefore, this approach was considered meaningful.

Because of this, it was necessary to identify a GE engine which was in the U.S. Air Force inventory that had a mature engine monitoring system and a substantial experience base that would provide a solid foundation for development of an expert system. The TF34-100 engine on the A-10 aircraft met these criteria. The TF34 engine is equipped with Northrop Electronics Division's TEMS (Turbine Engine Monitoring System). Ground display and processing of data is provided by the CEMS IV (Comprehensive Engine Management System, Increment IV), developed under contract to the USAF by Systems Control Technology, Inc.

In addition to the above-stated objective, several related goals were also essential elements of the program; these include:

- Investigate the application of the GE TEMPER (Turbine Engine Module Performance Estimation Routine) to a military engine. TEMPER is a data analysis technique for identifying faulty or deteriorated engine components, particularly those related to gas path performance. Currently all TEMPER applications are on large commercial engines.
- Develop a video "help" facility to augment the capabilities of the expert system. Identifying the specific type of information to include and the appropriate hardware for such a system were key subobjectives to this goal.
- Allow the domain expert(s) to actually build the knowledge base for this expert system. This was considered the most effective method for developing a comprehensive and user-oriented knowledge base. However, inherent in this approach is the need to provide an expert system tool that could be easily used by a novice.

1.2 Program Timing and Milestones

The JET-X Program ran from September 1985 to July 1988. Key program milestones are summarized as:

March 1986	Preliminary design review held at WPAFB
April 1987	Critical design review and laboratory demonstration of the developed system

January 1988	Set up the JET-X system at Barksdale AFB, LA to initiate the JET-X field evaluation
April 1988	Field demonstration and presentation of the completed system at Barksdale
June 1988	Complete the JET-X field demonstration
August 1988	Submit JET-X final report for USAF review.

1.3 JET-X Implementation

JET-X is implemented as a PC-based expert systems application. As previously stated, it is applied to ground-based troubleshooting of TF34-100 engines on the USAF A-10 aircraft. JET-X provides fault-isolation capabilities to supplement the current engine fault detection system. JET-X incorporates domain specific knowledge from existing troubleshooting manuals that has been augmented and enhanced by a team consisting of a factory engineer with a background in engine performance analysis and a field technical representative with considerable on-site troubleshooting experience. The expert system shell used for this application is GEN-X. This report is intended to give a comprehensive overview of the diagnostic scope and design objectives of JET-X, as well as the development issues and what procedures were employed to produce a powerful and flexible tool for use by the mechanic.

An expert system like JET-X, for the troubleshooting of military aircraft engines, is an extension of the existing maintenance practices. Users of military and commercial jet engines are relying increasingly on computer-based engine monitoring systems to provide information to direct maintenance activities. This approach allows engines to be serviced "on condition" or when actually needed rather than according to a preset schedule. The aircraft- or engine-mounted portion of the monitoring system examines engine data in real-time; abnormal or out-of-limit conditions are flagged, and in most cases, a data snapshot is recorded. In addition, routinely acquired trend data, stored in a ground-based computer, can be analyzed for anomalous changes to produce trend alerts. JET-X provides systematic analysis techniques for these data.

2.0 RESULTS AND CONCLUSIONS

1. The development of expert (or knowledge-based) systems for jet engine diagnosis and repair can be productive in at least two ways. First, the resulting system can be utilized by a novice technician to solve problems that would normally require the assistance of a more experienced mechanic. Second, the more seasoned diagnostician may be helped on unusual problems which have not been previously experienced.
2. When properly designed, the expert system is also useful as a training device. The novice mechanic can gain experience by using the system to diagnose problems that have been encountered in the past, or even imagined problems. To be successful as a training tool, the expert system should include a variety of help facilities that permit exploration at greater depth.
3. The expert system is neither "expert" nor "intelligent," as might otherwise be suggested by the name of the technology. The diagnostic techniques are largely procedural with decision points that respond to available data. The system will cope only with problems that it has been programmed to handle. The ability to solve unanticipated problems is beyond the current state-of-the-art. (Reference 1 offers convincing arguments that intelligent behavior from computers is unattainable.)
4. The information that is contained in the expert system could be presented in a Tech Order; the benefit of the expert system is that the indexing is done automatically by the computer. An expert system can comfortably accommodate procedures that would be very complex on the printed page.
5. The techniques that are used to fault-isolate jet engine problems are primarily developed through experience in the field. This may partly be a result of inadequate attention during the design phase, but another cause is that many of the field problems are unanticipated during development. Thus, the mechanic must develop procedures in the field to fault-isolate problems that were not expected during the design. Many engine problems are solved in new designs to be replaced by others.
6. Tools, such as GEN-X, for developing expert systems facilitate the programming of procedural knowledge. This makes them particularly effective for jet engine diagnosis, because of the need to update the tool based on field experience. The use of conventional programming languages, such as FORTRAN, for these efforts would be far more difficult. The use of Tech Orders for communicating this type of information is also very awkward. In the field demonstration, it was possible to modify the knowledge base to accept a new problem-solving technique in about 10 minutes.
7. If circumstances permit the expert to build the knowledge base, as was the case with this program, several benefits can result. An intimacy between the expert and his knowledge base can result, which may yield a tool with greater technical depth. By working constantly on the problem, the expert may be able to recall

greater detail and penetrate problems to a greater depth than might otherwise be the case if the interview approach were used to extract troubleshooting methods.

8. In this climate, there is a compatible role for the expert system specialist to guide the expert in shaping the knowledge base into an efficient design. The design of the architecture and the modularization of the knowledge is an appropriate project for one who is familiar with other expert systems. Ideally, there will be a close working relationship between the expert and the knowledge-base system designer.
9. One of the most difficult areas in building a knowledge-based system is the assurance that all logic paths have been adequately covered. To some extent, this difficulty traces to the fact that it is nearly impossible to reproduce a moderately complex knowledge base on paper. The knowledge-base structure is similar to a tree, and this problem is equivalent to assuring that every twig is properly constructed. Although a modular design helps to solve this problem, additional features of expert system shells that address this issue would be greatly appreciated.
10. One of the advantages of the expert system is that it can be fun to use, provided it is properly designed. A system that is fun to use will get used and will meet its objectives; this suggests that considerable emphasis needs to be placed on the end user interface. Features such as color graphics and well-designed, lively displays can add to the fun of using the system. The evaluation of the end user interface is an important element of the field testing of the system and should not be overlooked.
11. One of the difficulties in designing an acceptable end user interface is the variety of desires for the system. An obvious difference is between novice and expert users; the novice will want to explore the details, while the expert will want to bypass them. The novice will want considerable help, and the expert will believe that he/she doesn't need it. The system must accommodate these differences while assuring that the job is also performed correctly. Perhaps individual configuration selections could be helpful in addressing these issues.
12. A more subtle issue regarding the interface to the user is giving the user a proper role in the decision process. As an example, JET-X involves the user in evaluating the features of the trended data (steps, scatter, etc.). Automated techniques for determining these features in trended data are not always as effective as the human eye. Giving the user responsibility for these decisions can help to retain a sense of ownership for the outcome. Nevertheless, in the field trial, the absence of automated data feature extraction was criticized.
13. During the field evaluation, JET-X was able to diagnose several hard to solve problems with multiple symptoms. Often, key symptoms that identify the root cause are accompanied by secondary symptoms that can confuse diagnosis; however, JET-X looks primarily for the key symptoms and avoids the confusion associated with secondary symptoms. This capability is, of course, a function of

the expert experience utilized in building the knowledge base. Where experience was lacking, JET-X was not able to make an accurate diagnosis.

14. The JRS (JET-X Retrieval System) graphical interface established as part of JET-X to display borescope photographs, samples of previous TEMS data frames or CEMS plots, etc. is a valuable analog to the diagnostic expert system. Ultimately, this feature could form the bridge between diagnosis and repair by allowing reference to the repair manual procedures. The emerging development of optical storage devices should substantially increase the amount of information that can be available to these systems.
15. The JET-X hardware requirement, involving two AT-class personal computers with a conventional color monitor and a high resolution black and white monitor, in addition to the independent CEMS terminal, was excessive. It should be possible to combine the software pieces into a single integrated software system that could reside on a single platform. The developing OS2 or WINDOWS packages should be able to support the resultant, integrated system.
16. The marriage of Kalman-filter-based gas path analysis techniques to expert systems technology should result in a more effective gas path analysis tool. This capability was not demonstrated as part of the JET-X Program because of the limited gas path instrumentation on the TF34 and because of the absence of a direct link to the CEMS data base management system. For future engines, such as the ATFE (Advanced Tactical Fighter Engine), which will have adequate sensor sets, this development should be pursued.
17. During the field evaluation, several specific comments surfaced regarding the JET-X end user interface. These comments are given in Section 11 of this report; in general, they relate to details of the JET-X system that should be improved if the system is to be converted to a production application.

3.0 RECOMMENDATIONS

Experience gained with the prototype JET-X system has demonstrated that an expert system can provide valuable assistance in the interpretation of monitoring system data for maintenance purposes. Based on this experience, the following recommendations are made for the design of a production version of the JET-X system:

1. All ground-station features should be combined into a single platform. The expert system should be linked to the CEMS IV data base so that alarm data, as well as numerical data, can be accessed without user intervention. The CEMS display, the JRS monitor, and expert system screen should be integrated into a single device; a "windows" environment offers the technology to accomplish this.
2. While the knowledge base of the prototype system is comprehensive, a production system should provide a completely enhanced knowledge base. It is recommended that a survey of accumulated troubleshooting experience with TEMS and CEMS on the TF34 engine be conducted and knowledge-base procedures be updated accordingly.
3. An important element of any follow-on program is the design of the end user interface. To acquire data to support this effort, extensive field testing of the prototype system is recommended to gather comments on end user interface improvements.
4. A second engine model should be included as part of a follow-on program in order to gain an appreciation of the elements of the knowledge base that can be carried forward to new applications versus those which are engine specific.

In addition, the following recommendations apply to the extension of the technology studied by this contract.

1. The linking of Kalman-filter-based gas path analysis technology to knowledge-based systems technology should be pursued using an engine with adequate gas path instrumentation.
2. The linking of expert systems technology to digital image storage technology, such as was accomplished in JRS, should be pursued for application to integrating jet engine diagnostic and repair procedures.

Finally, it is recommended that anyone considering the development of an expert system should consider the following advice:

- Expert systems should be built by the experts
- Shells should permit this type of development environment
- Experts in knowledge-based systems design should be employed to design the overall architecture and to assist in the detailed structural design of the system.

4.0 A-10/TF34-100 OVERVIEW

4.1 TF34-100 Engine Description

Designed and manufactured by GEAE of Lynn, Massachusetts, the TF34-100 is a high bypass, front fan, gas turbine engine with a bypass ratio of approximately 6 to 1. The -100 Engine, developed under contract to the USAF, is a derivative of the TF34-400 engine built for the U.S. Navy and used on the ASW (anti-submarine warfare) S-3A Viking aircraft. The -100 Engine powers the Air Force A-10 attack aircraft.

The -100 Engine (as well as the -400) employs a single-stage fan which operates at an optimum pressure ratio of 1.5 to 1, and is driven by a 4-stage LPT (low pressure turbine). A 14-stage axial compressor, powered by a 2-stage HPT (high pressure turbine), incorporates variable position stators and provides a nominal 14.5 to 1 pressure ratio. Combustion occurs in a through-flow, full annular combustor fed by 18 fuel injectors. An engine-mounted gearbox driven by a PTO (power takeoff) shaft from the core compressor supplies approximately 225 horsepower extraction capability for hydraulic, electrical, and fuel pumping power requirements. The lubrication system is a self-contained, closed-circuit, pressurized, dry sump system designed to supply oil for lubrication and cooling to the necessary engine components during engine operation. Fuel scheduling is provided by a hydromechanical control, which employs an electrical subsystem for temperature limiting (fuel flow modulation) during maximum power operation.

The two rotor systems are supported and contained by a total of seven main bearings, located in three separate sumps. Two of the seven are ball bearings, which absorb the axial loads of the fan and compressor respectively; the other five bearings are roller type. Oil leakage from each of the three sumps is prevented by air-pressurized radial carbon seals. Three major structural frames on the engine support the main engine bearings in which the high and low pressure rotors turn. Struts of the front, combustion, and exhaust frames are used for lubrication system service and scavenge lines, as well as various air lines. Figure 1 is a cut-away view of the TF34-100 Engine.

Roughly 85 percent of the thrust produced by the TF34 results from air accelerated through the fan nozzle; the remaining thrust comes from core engine flow. Since fan flow has the dominant effect on engine thrust, it is a good measure of the overall health of the engine (thrust producing ability at a given turbine outlet temperature). However, since fan flow cannot be easily measured on-wing, fan speed is used in the field as an indicator of flow or thrust. For this reason "fan speed trim margin" is the parameter utilized in the field to monitor engine performance; this is calculated as the difference between actual engine speed and a minimum allowable fan speed, determined from a curve in the T.O. (Technical Order) Manual as a function of OAT (outside air temperature). This curve, known as the "trim curve," is presented in Figure 2.




Figure 1. TF34-100 Engine Cut-Away.




Figure 2. TF34-100 Trim Curve.

As engine performance deteriorates over time, the fan speed margin will drop, indicating lower available thrust. Performance can be recovered through maintenance, or fan speed margin can be increased by raising the limiting turbine temperature maintained by the control system. This limiting temperature can be raised by an adjustment on the engine; however, this governing temperature can only be elevated to a redline value before maintenance must be carried out if performance is to be regained. The process of uptrimming, or simply "trimming" the engine, refers to the process of adjusting the governing turbine temperature in the control system in order to achieve a fan speed above the minimum required by the trim curve.

4.2 A-10 Aircraft Description

Dubbed the "Tank Killer," the A-10 is the primary CAS (close air support) aircraft of the U.S. Air Force. Powered by two TF34-GE-100 high bypass turbine engines, its main weapon is the GE GAU-8/A 30-mm Gatling gun, firing up to 4000 rounds per minute. In addition, the aircraft can carry up to 16,000 pounds of external weapons on 11 hard point pylons.

5.0 TF34 ENGINE MONITORING SYSTEM

The TEMS (Turbine Engine Monitoring System) and the CEMS IV (Comprehensive Engine Management System, Increment IV) provide all engine monitoring functions for the TF34-100 and were developed to support the OCM (on-condition maintenance) or RCM (reliability centered maintenance) concepts.

TEMS is a stand-alone, on-board condition monitoring system; whereas, CEMS IV software is mainframe resident and is a data storage, display, and diagnostic support tool. Together, these provide the user with limit-exceedance and event data, parametric trends, trend deviation alarms, and the correlation of performance, spectrometric oil analysis, maintenance, and life-limit data enabling a comprehensive assessment of engine "health."

5.1 TEMS Overview

The TEMS continuously monitors specific engine and airframe parameters during operation. Automatic recording of data is triggered by detection of any 1 of 33 limit-exceedance conditions. Data are also captured for trending at two preprogrammed flight conditions: takeoff and cruise. The EPU (electronic processing unit) hosts system and event detection software and stores event data. The UDU (umbilical display unit), located in the nose wheelwell of the aircraft, displays the engine status at the end of each flight and serves as the interface for data retrieval from the EPU. Using either of the two data retrieval/storage units, flight data can be transferred into the ground station (a personal computer), which can provide a printed copy of each data record. Limit-exceedance data from TEMS can then be examined immediately, while the takeoff data (which is utilized for trending), must be transferred into the CEMS IV data base for trend evaluation.

5.2 CEMS IV Overview

The CEMS IV provides mass data storage in a base level mainframe computer, accepting TEMS performance, oil analysis, maintenance, and component life-limit data. In addition to providing "user-friendly" data displays, abnormal parameter trends are automatically identified and alarmed. Statistical analysis and comparison, trend rate calculations, and correlation of various data types by means of plot or table format are available for comprehensive evaluation of engine trend and event alarms.

5.3 Diagnostic Maintenance Procedures with TEMS and CEMS IV

In order to conduct consistent and comprehensive analysis of the TEMS detected limit-exceedance events and CEMS IV generated trend alarms, diagnostic troubleshooting procedures were developed. The procedures are constructed in logic-tree format; branching is based on user response to questions regarding various data displays available through CEMS IV. Figure 3 shows an example of one of these troubleshooting procedures for a "low performance" alarm.




Figure 3. Sample Troubleshooting Procedure for "Low Performance" (CP1) Alarm.

Development of these original procedures was based on the limited data base generated by 27 TEMS-equipped A-10's operated at Barksdale AFB, LA over a 3-year period. After testing, evaluation, and modification, the procedures were incorporated into a USAF Technical Order (T.O. No. 1A-10A-2-71MS-5, Reference 2). The current version of this T.O. is 3 years old; update and correction has proven to be very slow and difficult. Because it is a publication, the depth of the procedures is limited, as are the examples, explanations, and computational capabilities available to the user.

These diagnostic procedures are intended for use with the CEMS IV data terminal. For each CEMS IV alarm or TEMS event, a specific troubleshooting algorithm is available which guides the user in calling up appropriate data displays on the terminal and posing questions about symptoms that may be present in the data. The objective of this session is to take advantage of all available TEMS/CEMS IV data to help isolate the cause of the alarm before actually doing testing or maintenance on the engine.

While the JET-X knowledge base builds on the methods included in the original procedures, significant enhancements were incorporated to take advantage of the expert system software environment and to reflect more recent troubleshooting experience.

6.0 JET-X OVERVIEW

Expert systems are well-suited for use with modern engine monitoring systems, since the examination of recorded data is often required as part of the fault-isolation process. Depending on the education and skill level of the mechanic, this can be a non-trivial task. JET-X provides the TF34 engine technician with a systematic method of analyzing operational (TEMS) and trend (CEMS IV) alarms using all available data in the ground station. In many instances, the procedure for complete evaluation of an alarm may be too complicated for inclusion in a manual; the expert system provides an excellent medium for dispensing complex procedures in a nonthreatening format.

In actual use, JET-X guides the operator in calling up relevant graphical displays on the ground station terminal and queries him about specific symptoms that may be evident in the display. In its present version, JET-X does not interface electronically with ground station data. A typical session begins with the user (operator) entering all alarms for a given engine. Based on these alarms, the user is directed to call up data displays on the CEMS ground station. JET-X queries the operator relative to these displays, who then makes judgements concerning the data and enters corresponding responses into JET-X by means of menu prompts. In the event that the operator lacks sufficient experience to interpret the CEMS IV displays and make accurate judgements, extensive on-line video help facilities, consisting of example cases from archival CEMS IV data, are provided. The operator's answers are translated by JET-X into facts with either a true or false value depending upon the response. The combination of facts resulting from a session are then examined within JET-X to determine the appropriate maintenance recommendations.

6.1 JET-X Components

JET-X is a composite set of software that is resident in two separate Zenith Z248 microcomputers each configured with 640-K bytes of RAM (random access memory). One of these devices hosts GEN-X, the JET-X knowledge base, and all external programs and files which support non-GEN-X functions. The other unit supports the JET-X retrieval system (JRS) which provides video help displays to assist the user during the troubleshooting session. The first unit employs a CGA color monitor, while JRS utilizes a high resolution black and white monitor.

Software elements comprising JET-X include the following:

- GEN-X - This is the GE-developed expert system shell software which was used to build JET-X. The knowledge base was built using the development mode and is run in the end user display mode. In the end user environment, the details of knowledge base construction are transparent to the operator.
- Knowledge Base - Consists of all the GEN-X modules developed to perform TF34 diagnostics and control the actual troubleshooting process. These modules

cannot be accessed independent of GEN-X and are unique to the JET-X application.

- JRS Executive Program - Controls the display of auxiliary video displays on the JRS monitor, by providing "handshaking" functions between the knowledge base and the host software on the JRS computer.
- Bit Maps - All video help displays are stored on the JRS computer in the form of bit maps, and are called from the knowledge base by the executive software.
- Microsoft Windows - This is the host software environment utilized on the JRS computer for display of the video help images.
- External Support Programs - Numerous "C" language programs, which are called from the knowledge base (GEN-X) and are external to GEN-X, provide capabilities that were considered desirable but not within current GEN-X capability. These features, described in greater detail in Section 10.0, perform reporting, display, user input, and numerical calculation functions.

6.2 JET-X Interface with CEMS IV

Since no link currently exists between JET-X and the CEMS IV engine data base, information can only be passed to the knowledge base through the user. For this reason, the JET-X system physically sits beside the CEMS IV ground station terminal, through which all available engine data is accessed. Figure 4 is a photograph of the actual JET-X field setup tested at Barksdale AFB, LA.

Figure 5 illustrates the overall data flow in the TEMS/CEMS IV monitoring system and demonstrates where the JET-X system fits in. As discussed previously, all data acquired by TEMS and CEMS IV is available for display on the CEMS IV terminal. When running JET-X, the user is directed to input specific commands into the CEMS IV keyboard which pulls up required data displays onto the screen. Ideally, these commands would be initiated from the knowledge base directly to CEMS IV; however, establishing this electronic link was beyond the program scope.




Figure 4. JET-X Equipment Setup.




Figure 5. TEMS/CEMS IV With JET-X.

7.0 JET-X SYSTEM DESIGN

In developing an expert system, we know that the bulk of the effort is involved in extracting the domain expertise and translating this expert knowledge into appropriate logic modules. A parallel activity of significance is the design of the approach to be used for integrating these logic modules. Approximately 30 percent of the JET-X project effort required was addressed to knowledge-base design issues. The design or architecture of the knowledge base includes knowledge-base structure, control of the inference process, and execution efficiency. A poor overall system design can greatly diminish the effectiveness of the expert system as well as the users' perception of the tool.

This section describes the JET-X architecture, its principal features, and the human factor considerations reflected in its design. The troubleshooting examples described in Section 9.0 will illustrate the features discussed here in more specific terms.

7.1 JET-X System Architecture

A high-level diagram of the JET-X knowledge-base architecture is presented in Figures 6 and 7. The uppermost levels of this hierarchical design are controller modules; while lower level modules contain the actual troubleshooting logic of the system.

At the control level in Figure 6, the "Log-In" module handles the opening transactions with the user; it records the user's name, time, and date, and opens the session files for entering information. The next step, "Activity Selection," offers three major choices: "Troubleshoot, Reports, and Introduction," as well as a "Help" option for explaining each of these functions.

Selecting Introduction produces the screen view of Figure 8, which provides a mini-training session that is useful to individuals unfamiliar with TEMS and CEMS IV, as well as with JET-X. Introduction offers an overview of TEMS and CEMS IV, provides a glossary of TEMS and CEMS IV acronyms that are used in JET-X, offers step-by-step procedures for interpreting various CEMS IV plots, and presents a brief explanation of how JET-X operates. An example of the "glossary" feature is shown in Figure 9.

Reports gives the operator access to records of past JET-X troubleshooting sessions. Using this feature, past experience with engines can be examined; the capability to delete engine records is also available. The Reports section supports generation of printed records of diagnostic sessions in various formats, as explained in detail in Section 7.5.

Selecting Troubleshoot presents the user with two menu picks: "Diagnose" or "TEMPER." Diagnose enters the JET-X troubleshooting procedure for the TF34 engine; whereas, the selection of TEMPER provides access to a simulation of the TEMPER interface which is built into JET-X. TEMPER and the JET-X TEMPER interface simulation are described in detail in Section 12.0.




Figure 6. JET-X Knowledge-Base Architecture - High Level Functions.




Figure 7. JET-X Knowledge-Base Architecture - Diagnostic Functions.




Figure 8. JET-X Introduction Screen.




Figure 9. JET-X Glossary Screen.

In the main troubleshooting path (Diagnose), the user is asked to indicate whether this is a "Work" or "Training" session; this selection determines whether the engine troubleshooting records are to be saved at the end of the session. In the Work mode, the specified engine number is used by JET-X to first check whether any prior maintenance recommendations were carried out satisfactorily before proceeding further.

The main troubleshooting activity begins by displaying a list of all 53 TEMS and CEMS IV alarms in logical groups; the user is then asked to indicate the specific alarms occurring on the engine of interest. One of these alarm categories contains the three performance-related alarms, which indicate an abnormal trend condition. The procedures to establish root causes of these alarms are extensive; JET-X addresses 12 symptoms or conditions that can help to isolate the cause of the performance alarms (see Section 8.0 for further details).

For these three complex performance alarms, the operator can select either of two modes of troubleshooting: "Guided" or "Symptom." In the Guided method, JET-X governs the sequence of troubleshooting steps by leading the user through all relevant portions of the knowledge base. This search sequence is defined by the expert/knowledge base developer, and varies depending on other actual CEMS/TEMS alarms indicated. In operation, JET-X displays questions and prompts the user to input a response at each step from a menu of possible choices. This method of diagnosis is well-suited for training, as well as for actual fault analysis. The alternate Symptom method of troubleshooting is useful for the more experienced diagnostician and is applied only to those faults which have very extensive, complex troubleshooting procedures (currently, only the three performance alarms). In the Symptom mode, the user selects the likely cause for the engine alarm, and JET-X guides the search for the root cause of the indicated alarm from this educated guess.

The functional relationship of the Guided and Symptom methods is presented in Figure 7. The diagnostics and fault-isolation procedures in the Guided troubleshooting method ensure that all 53 alarms and expected combinations are addressed to cover all possible troubleshooting scenarios. Covering all alarms combinations with a "brute-force" technique would require an extremely large number of rules with an adverse impact on the efficiency of search. Instead, a novel technique was conceived and implemented in JET-X to ensure that all alarms and their relevant combinations are addressed, while requiring a minimum number of rules, resulting in very efficient execution. The strategy of evaluating alarm combinations is a major strength of JET-X. While some alarm combinations would be recognized by most TF34 troubleshooters in the field, others require considerable experience to relate. Incorporating these kinds of relationships into an understandable Tech Order format is difficult and, currently, is not done.

In a typical Guided session, alarms are input, and control is passed to the "guided troubleshoot control" section of Figure 7. Input alarms are examined to determine the order and combination in which they will be evaluated. This portion of evaluation takes place in the three alarm sort levels. Additional search control logic transfers inference to the applicable troubleshooting procedures and steps. The "Alarm Modules" block of Figure 7 represents the bulk of the

JET-X knowledge base. It includes logic modules for the detailed troubleshooting procedures for various alarms and contains rules to combine various data symptoms for generating recommendations. It also can display session progress and results and can provide detailed explanations for its maintenance recommendations.

When JET-X is troubleshooting an engine with multiple, unrelated alarms under the Guided mode, the first alarm will be evaluated following the steps described above. After reviewing the maintenance recommendation made for this alarm, the technician may choose to continue with the analysis of the other alarms, restart to troubleshoot another engine, or exit from the system. If he chooses to continue, the next alarm(s) will be evaluated in a similar manner. During the troubleshooting session, the operator is provided various "help" facilities in the form of explanations or video images (by means of the JRS).

The Symptom's method of troubleshooting is also functionally blocked in Figure 7. It is structured to accommodate the more seasoned troubleshooter by allowing him/her to circumvent certain control portions of the guided mode and jump directly to the analysis of his/her "hunch." In the JET-X implementation, the Symptom's method is available for the analysis of the CEMS IV performance alarms. For the analysis of the indicated performance alarm, entry to the Symptom's method would provide the user with a list of all possible faults associated with that alarm (Figure 10). Utilizing one's knowledge, together with background information for the engine, the user may select one of the listed faults, and JET-X will provide step-by-step procedures to verify that fault. After verification, the user may select another fault for evaluation, switch to the Guided method, or exit from the system. This mode of troubleshooting can be tailored to suit any skill level.




Figure 10. JET-X Symptoms Selection Screen.

7.2 Modular Knowledge-Base Construction

Since JET-X is built using GEN-X, the knowledge base is modular. As a result, specific troubleshooting procedures or searches in JET-X are contained in individual logic modules. For example, the logic for "rising vibration trend" analysis is contained in a logic module named "vibanal.dec", where the suffix ".dec" denotes it to be a "decision-tree" module (see Appendix A for discussion on GEN-X). The modular structure of the JET-X knowledge base offers several advantages for the developer and the maintainer of the knowledge base. New modules can be added to accomplish additional troubleshooting tasks and then linked into the existing architecture using the backplane commands of GEN-X. Changes in existing modules are made by altering only the affected module. If modules are kept relatively small in size during the development, the effort involved in updating them will be trivial. For jet engine applications of expert systems, the ability to easily alter and add procedures is important, since troubleshooting experience expands as service time on the engine fleet increases.

The modular structure of knowledge base also enabled easy implementation of the Symptom's troubleshooting method, since the same logic modules used for troubleshooting in the Guided mode were directly employed. Finally, by loading only the logic module needed for the current diagnostic step, a large-scale application (several hundred rules) can run efficiently on a microcomputer.

7.3 JET-X Knowledge Implementation

The overall strategy employed by JET-X is to establish the truth value (either "True" or "False") of globally available facts that represent symptoms or conditions essential to isolating the cause of an engine alarm or alarms. The procedural knowledge of the expert for searching out and verifying these symptoms is implemented in decision trees. Once these truth values have been established, maintenance recommendations are generated by executing "If/Then" tables which consist of rules relating combinations of these facts. Because JET-X is able to relate a much larger number of facts than would be possible in a technical manual, the troubleshooting capability of the system offers a significant enhancement over traditional troubleshooting techniques. However, the value of an expert system for this type of application is dependent upon the number of unique parameters (that is, facts) from which symptoms can be obtained. For jet engine monitoring systems, this number is directly related to the number of parameters that are measured and/or calculated. If only a few are recorded, the rules relating these parameters to specific engine faults may also be few in number and easy to learn; consequently, the complexity and expense of employing expert system technology would be misplaced.

In addition to the procedural knowledge of how to interpret and identify symptoms in the engine data, another aspect of the expert experience included in the JET-X knowledge base is the frequency with which faults actually occur. Those faults known to cause a specific symptom most frequently are examined first; this strategy avoids having to examine more data than is usually needed before focusing on the real cause of an engine alarm. The issue of which

examinations to perform first, based on the complexity or cost of performing the test, is moot for JET-X because, by design intent, JET-X only tests data that is immediately available on the CEMS IV ground station terminal. Consequently, all tests are essentially of equal difficulty and cost.

The control strategy implemented in JET-X is handled by utilizing the If/Then rule tables. An example of a rule table is given in Figure 11. The rules contained in these tables are executed as needed during inference, but their contents (as seen in Figure 11), are not displayed on the screen in the end user mode. Rules imbedded in these tables establish relationships between TEMS and CEMS IV alarms leading to specific follow-up actions, simulating the expert's capability to recognize symptom patterns characteristic of a specific engine fault.

Rule	Condition	Action
FFG TND > LMT	5000	100
vib trend	B 5000	100
rising p5p1 tnd	B 5000	100
sign perf loss	B 5000	100
MFTS TND < LMT	5000	100
valid stop	B 5000	100
LOW THRUST	5000	100
valid perf prb	B 5000	20
FFG > LMT	5000	100
STALL	5000	100
valid stall	B 5000	20
vib trend-ind	B	100
gas path dng		175
gunfire check	B	175

Figure 11. Example of "If/Then" Rule Table in Development Environment.

From a human factors' viewpoint, it is undesirable to repeat a question that has already been answered earlier in the troubleshooting session. Without proper care, a modular knowledge base can suffer this drawback, especially when several modules are involved in a single session. Utilizing global nodes (a GEN-X feature which allows the same question to be shared by more than one decision tree), knowledge obtained in one procedure is passed to another, avoiding the need to repeat a question. Once a response to a global question has been given in one decision tree, the answer will be applied automatically if the same question occurs in another tree.

7.4 Human Factors Issues

During the JET-X development, a number of concerns were encountered regarding the end-user's reaction to various JET-X features and methods. Consequently, many alterations were made to functions and architecture to address these human factor issues. Developers concluded that 25 to 30 percent of the labor used to build the JET-X knowledge base was directed towards these "creature comforts."

The original JET-X design approach could leave the user feeling "swept away" because decisions and recommendations seemed to come from nowhere; the impression communicated to the user was that of a spectator rather than a participant in the process. Consequently, conveying the sense of being "in control" and "informed" during troubleshooting became a critical human factor design problem. The user's sense of involvement is significantly enhanced by obtaining frequent feedback on conclusions reached and knowing what segment of the troubleshooting task is about to be addressed. During longer sessions, the communication of what has been accomplished and what remains to be done becomes important to this sense of control. Also, the option to display a summary of the troubleshooting session while still in the session has been incorporated.

Another feature contributing to a user's sense of control is the ability to exit when desired, to avoid the uncomfortable feeling of being "trapped." While it is impractical to provide a "graceful" exit from JET-X at all times, the architectural design was modified to include frequent opportunities to "Exit" or "Restart" a session. When either of these options is selected, a summary of the troubleshooting session will be displayed, and appropriate records will be updated and stored for a work session.

Military jet engine mechanics and troubleshooters possess a wide range of experience and skill. Being able to accommodate users with a variety of capabilities was another critical human design issue. An expert system aimed at too narrow an experience level would likely waste the more experienced users' skills and therefore be rejected by them. The modular format of the knowledge base enabled implementation of a powerful method of bridging the user experience range. For potentially complex performance alarms, both the Guided and Symptom methods of troubleshooting are available to the diagnostician. The inexperienced user can select the Guided approach, which will lead through all possible data investigations relevant to the alarm being worked. The user who may already have a hunch of what the problem is can choose the Symptom method and proceed directly to that portion of the knowledge base addressing this suspicion.

While some technicians pursue their responsibilities only out of duty, many are interested in understanding the engine - how it works, and how it responds when not operating properly. In order to provide helpful, as well as insightful information to meet the needs of this latter group, two features were employed. First, questions asked by JET-X which may not obviously relate to the problem under investigation have an associated explanation that can be accessed by menu pick. Second, maintenance recommendations made by JET-X that involve complex

relationships between symptoms are fully explained in the extended help facility (Section 10.4), which is available if the user desires. These provisions serve not only to accommodate curiosity, but are excellent teaching tools as well.

An expert system designed to be used by a diagnostician should not leave the user feeling replaced; a proper balance must be achieved between including him/her in the troubleshooting process and employing analysis that may be beyond his/her skill and experience. This trade-off is best met when the user finds the tool interesting and fun to use. Within the JET-X architecture, allowing the user to make the judgements about data symptoms was the right approach. If the system directly accessed information in the database, analyzed it, developed a recommendation, and then asked the user for his concurrence, he could be left feeling "out of the loop." In general, the analytic methods of testing data for trends or other symptoms are no more accurate than the eye; consequently, no loss in performance is borne by having the operator input user judgements directly.

The actual format and appearance of the screen the user sees can either be an enticement or deterrent to system use. JET-X utilizes a color monitor, which provides universal appeal. No effort was made to alter the split-screen format of the conventional GEN-X end user interface, as it was considered adequate for a proof-of-concept project. If desired, minor modifications in the end user display format could be readily made to further enhance user receptivity.

Another element contributing to user acceptance is the ability of the expert system to deal with tough issues and "think" of things the user might not. An expert system which is shallow and of no real help when the troubleshooting is difficult will not be used at all. Employing experts to build JET-X enabled many difficult and unusual troubleshooting practices to be included in the system. Often these less-traveled paths did not occur to the builders on their first construction of a procedure but, rather, were "jogged" out of memory by constant contact and thought about the knowledge base.

7.5 Report Generation

If the user selects Reports as a menu pick when entering JET-X, he/she will be confronted by two options: "Engine Records" and "Delete Engine." The Engine Records feature allows the user to obtain on the screen a summary of all engine records in JET-X, or a detailed diagnostic record for the specific engine indicated by the user. Two types of summary reports are offered: sorted by engine number or sorted by record date (Figure 12). The sorted summary records are displayed by session date in reverse chronological order; the user can indicate the specific engine records for display (Figure 13). Figure 14 is an example of an engine troubleshooting record.

The Delete Engine function is provided so that the user can delete diagnostic records from JET-X for engines taken out of service. All records in JET-X will be deleted for those engines specified by the user for deletion.




Figure 12. JET-X Reports Selection Screen.




Figure 13. JET-X Engine Serial Number Selection Screen for Reports.

```

SESSION IDENTIFICATION:
session number: 0054
date: 04-01-88 12:00
user name: Costen
engine number: 5947
action: Guided

"CENS/TEMS ALARMS EVALUATED"
-----
145 Check for excessive compressor
    bleed or horsepower extraction
167 Troubleshoot engine per I.O.
    based on all findings

CENS/TEMS ALARMS - INPUT:
Level 2 Flameout

CENS/TEMS ALARMS - EVALUATED:
Flameout

DIAGNOSTIC PATH MARKERS:

MAINTENANCE RECOMMENDATIONS:
-----
NOTE - Following recommendations
       are for:
           Hit 'P' to print this information
           Hit 'C' to continue on to next session or 'Q' to quit.

```

Figure 14. Sample JET-X Engine Troubleshooting Record.

8.0 JET-X DIAGNOSTICS

8.1 Symptom-Fault Approach

Current expert system technology recognizes two fundamental approaches for addressing diagnostic problems. While terminology may vary, these two methods can be classified as either "model-based" or "symptom-fault;" although some hybridization of the two is possible. As its name implies, model-based reasoning depends on a functional (and possibly physical as well) model or representation of the system or device to be diagnosed. When abnormal operation of the device is detected, the model will identify which components could contribute to the observed malfunction. Subsequent isolation of the faulty component or subsystem is performed by testing data within the device and comparing it with the model; this testing procedure would be automatically optimized for efficiency.

The success of this approach is dependent upon the accuracy and completeness of the model and the ability to make measurements of data within the system for isolation testing. This technique is most amenable to electrical systems where the laws governing performance are well known; for complex mechanical systems, such as a gas turbine engine, the modeling task is prohibitively intricate. Consequently, only simple models are practical, which in many instances would not be adequate for a reasonable measure of fault-isolation. Also, the issue of data measurement for comparison with the model becomes very difficult and costly for a jet engine. For these reasons, the symptom-fault technique was considered the only viable approach for JET-X.

Rather than depending on a functional model of the system, the association of specific symptoms with corresponding faults is the foundation of this alternate diagnostic approach. The association of symptom and fault is generally derived from experience and, thus, requires an expert or experts with first-hand knowledge of the relationships between possible component failures and how they manifest themselves. Having sufficient experience to relate symptoms to all possible component failure modes is the obvious drawback of this method. While coverage of all possible failures is perhaps impossible, coverage can be quite complete if the experience base used to develop diagnostic rules is sufficiently mature. Symptom-fault diagnostics suffers its greatest handicap when applied to new engines with little field experience.

JET-X diagnostics is based on years of experience in applying available data to analyze the various TEMS and CEMS IV alarms. This experience is embedded in the knowledge base in modular question and answer sessions which are used to identify clues in the data that may help to isolate the cause of the alarm.

8.2 User's Role in JET-X Diagnostics

JET-X depends heavily on the user to accomplish its diagnostic functions. Specifically, JET-X directs the user to call up various data displays on the CEMS IV terminal and queries

the user about parameter levels and conditions that may exist in the data. Appropriate responses to these questions are available in the menu, which is displayed concurrently with the question. In making a menu selection, the user is actually setting a JET-X fact true or false (invisible to the user).

Eliminating the user from this process was beyond the scope of JET-X. This capability would have required JET-X to interface directly with the CEMS IV data base and examine data for specific symptoms. Establishing a data link with CEMS IV would have required considerable outside contractor support, for which there was no provision in the program. In addition to the program issues, the automated analysis of data with no user-participation introduces human factors concerns as well as technical issues.

Technically, the evaluation of trend data using statistical techniques can, at times, be a poor substitute for the eye. If automated data analysis could have been employed, JET-X would have been computationally intensive at the expense of other features. A more subtle problem to be confronted, if this approach were to be used, is the role the user should actually play in the diagnostic process with an expert system. Performing all data evaluation computationally without user participation, would eliminate the user's contribution to the process. While this may be advantageous in some environments, it suppresses human involvement and concern. The process of actively engaging the user in decision making, actually trains the user over a period of time, from the guidance supplied by JET-X. Consequently, the operator can become more proficient and, potentially, more competent. This type of involvement would probably be more applicable to novice technicians; seasoned troubleshooters may actually prefer a more streamlined process. Fostering positive human involvement can be a valuable by-product of the expert system.

8.3 JET-X Diagnostic Scope

Usually, troubleshooting an aircraft gas turbine with an engine monitoring system is a two-step process. The monitoring system provides information which may isolate a fault directly, but in general, this is not the case. More often, available data will narrow or direct the scope of the troubleshooting that must be completed on the engine itself. This on-engine troubleshooting can consist of special inspections, removal and testing of components, or removal and replacement of components. By design, JET-X addresses only that portion of troubleshooting dealing with the interpretation of monitoring system data.

This approach has the advantage that for a given engine, JET-X troubleshooting can be concluded during a single session, since all relevant data is available in the CEMS IV terminal. This alleviates the need to return to the terminal at some later time with new information; having to gather information and return to JET-X hours, or even days later, was not considered a viable approach to engine diagnostics. When insufficient data exists to enable JET-X to isolate a fault, the recommendation produced will direct the user to a specific T.O. procedure, in order to complete the troubleshooting task. If JET-X could be carried planeside for engine troubleshooting, the value of having all phases of troubleshooting built in would be obvious.

The result of this definition of scope directly impacts the diagnostic penetration or fault-isolation capability of JET-X. For those engine faults for which there is very little applicable data, JET-X diagnostics will necessarily be shallow. Conversely, where much data is available (as in the case of the three CEMS performance alarms), penetration can be excellent.

The knowledge base for the three CEMS IV performance alarms consists of three separate subcategories. The first section guides the user in searching the CEMS IV data base for symptoms of degraded gas path performance. The gas path refers principally to the gas generating hardware in the engine: the high pressure compressor, the combustor, and high pressure turbine. The health of these components has a very strong influence on the overall performance (thrust and fuel consumption) of the engine. JET-X searches for the following symptoms which can be evidence of gas path damage or deterioration:

- Rising Core Vibration Trend
- High In-Flight Core Vibrations
- Engine Stall
- Rising Turbine Pressure Ratio
- Sudden Loss in Overall Performance
- In-Flight (TEMS) Low Thrust.

Checking for maintenance-generated performance changes is the second dimension of the CEMS performance alarm knowledge base. While maintenance would not be expected to induce a performance loss, real performance losses can follow certain maintenance tasks, even when they are properly performed. Improperly executed maintenance as well as inaccurate reporting practices can also result in CEMS IV performance alarms. The key used to identify this possibility is a maintenance action occurring immediately prior to a trend performance drop. If one occurred, the type of maintenance is identified which provides a clue to the probable cause of the performance change.

A search for miscellaneous symptoms comprises the last portion of CEMS performance alarm analysis. Some of these symptoms are evidence of true change; others only result in a drop in calculated (by TEMS) performance, while the actual engine performance remains unchanged. The following symptoms fall into this third knowledge base category.

- TT2 Error
- P0 Error
- Error Caused by TEMS Performance Extrapolation
- VG Tracking Changes
- Aircraft Deployment (Takeoff Location Seems to Affect Performance).

One of the issues addressed by expert diagnostic systems is the prioritization of the troubleshooting process on the basis of probability and cost. Where isolation of a faulty component must be achieved by testing or replacement of suspect components, the selection of which test to perform first will be a function of the probability of the component being at fault (utilizing a priori information) and the cost of accomplishing the test or replacement. Because JET-X troubleshooting only addresses monitoring system data, the issue of probability and cost of evaluation becomes moot, since all available data can be examined with equal ease and cost. For events where experience indicates that certain causes are more likely than others, the JET-X architecture has been designed to search for symptoms of these causes first. The user may want to continue looking at more data; otherwise, the session can be terminated.

8.4 TEMS and CEMS IV Alarm Coverage

JET-X diagnostics is focused on providing the user with the tools to analyze each of the possible alarms that can be generated by the TEMS airborne monitoring system and the CEMS IV ground station. Table 1 lists all 54 of these alarms. The diagnostic procedures that were developed for JET-X are based on troubleshooting techniques that had been previously created by GEAE under contract to the USAF (SA-ALC, Kelly AFB). These initial algorithms, currently included in USAF Technical Order No. 1A-10A-2-71MS-5, consist of single page (8.5 by 11 inch) troubleshooting trees for each TEMS/CEMS IV alarm; Figure 15 is an example of one of these procedures. When completed (1984), these trees contained the latest experience available for applying available data to the analysis of engine alarms. Because most of the possible alarms had never occurred, the analysis techniques built into many of these trees were unproven. Although additional experience since initial release has been acquired, this cannot be easily incorporated into the original T.O. algorithms; consequently, they have undergone only minor revision.

Initially, JET-X was designed to evaluate 3 of the total 54 alarms. These three alarms were CEMS IV performance trend alarms: one identified sudden steps in trended engine performance, another flagged a gradual deterioration in engine performance, and the last alerted maintenance personnel to an unacceptable level of absolute (trended) performance. Several reasons supported the selection of these three alarms for the JET-X knowledge base. First, low performance is one of the prime maintenance drivers for the TF34-100 engine. In most cases, low performance is simply caused by the accumulation of contaminants from the A-10's gun on the core compressor, and washing the engine usually solves this problem. However, being able to quickly identify when this is not the cause of low performance is important from a flight safety and aircraft availability viewpoint. Second, many data symptoms have been identified that were not available for the original work which could be applied to the validation of performance alarms and isolation of cause. The incorporation of these additional symptoms into alarm troubleshooting represented a significant sophistication in troubleshooting methodology not amenable to T.O. format.

Table 1. TEMS and CEMS IV Alarms for TF34.

Alarm Name	TEMS or CEMS	Alarm Description
Performance Alarms		
NFTR MRV Below Limit	C	Most Recent Value of Takeoff Performance is Low
NFTR Forecast Below Limit	C	Takeoff Performance Trend Forecasted Low
FT5 Trend Below Limit	C	Step Loss in Trended Takeoff Performance
Low Thrust (In-Flight)	T	In-Flight Low Performance
Miscellaneous Alarms		
Low Idle Speed	T	Idle Speed Below Limit
Level 2 Flameout	T	Engine Flameout (Core Speed < 56%)
Level 2 NG Rollback	T	Core Speed Below Schedule
Level 1 T5 Shift	T	T5 is Higher/Lower than Trim Value
Engine Stall	T	Core Compressor Stall
Engine Stall Out of Envelope	T	Stall: Operation Out of Envelope
Level 1 Slow Start	T	Excessively Slow or Hung-Start
Stall Max Above Limit	C	Number of Detected Stall Pulses Limit
Vibration Alarms		
NF VIBS - Level 1	T	High Fan Frequency Vibs at One Pick-Up
NF VIBS - Level 2	T	High Fan Frequency Vibs at Two Pick-Ups
NG VIBS - Level 1	T	High Core Frequency Vibs at One Pick-Up
NG VIBS - Level 2	T	High Core Frequency Vibs at Two Pick-Ups
Rising Vibration Trend	C	Rising Core Frequency Vibration Trend
Vibration Trend Step	C	Step Rise in Core Frequency Vib Trend
Wear-Metal Alarms		
Rise in Wear-Metal - Iron	C	Rising Concentration of Iron in Oil
Rise in Wear-Metal - Silver	C	Rising Concentration of Silver in Oil
Rise in Wear-Metal - Nickel	C	Rising Concentration of Nickel in Oil
Rise in Wear-Metal - Chromium	C	Rising Concentration of Chromium in Oil
Rise in Wear-Metal - Copper	C	Rising Concentration of Copper in Oil
Rise in Wear-Metal - Aluminum	C	Rising Concentration of Aluminum in Oil
Rise in Wear-Metal - Silicon	C	Rising Concentration of Silicon in Oil
Rise in Wear-Metal - Titanium	C	Rising Concentration of Titanium in Oil
Rising Iron (Fe) Trend	C	Step Change in Iron Oil Concentration
VG System Alarms		
Level 1 VG Closed	T	VG's Off-Schedule (Closed Side)
Level 2 VG Open	T	VG's Off-Schedule (Open Side)
VG Schedule Shift - Open	C	Step Change in VG Tracking Trend - Open
VG Schedule Shift - Closed	C	Step Change in VG Tracking Trend - Closed

Table 1. TEMS and CEMS IV Alarms for TF34 (Concluded).

Alarm Name	TEMS or CEMS	Alarm Description
Fluctuations Alarms		
Level 1 Sensor Flux - NG	T	Abnormal Core Speed Fluctuation
Level 1 Sensor Flux - NF	T	Abnormal Fan Speed Fluctuation
Level 1 Sensor Flux - ITT	T	Abnormal ITT (T5) Fluctuation
Level 1 Sensor Flux - WF	T	Abnormal Fuel Flow Fluctuation
Level 2 Engine Flux	T	Abnormal Fluctuation in Two or More
Lube System Alarms		
Level 1 POIL > 95 psig	T	High Oil Pressure
Level 2 Oil psi < Schedule	T	Low Oil Pressure
Level 2 Oil psi Flux	T	Abnormal Fluctuation in Oil Pressure
Level 1 Fuel Filter Delta P	T	High Fuel Filter Pressure Drop
POIL Trend Above Limit	C	Rising Oil Pressure Trend
OIL Trend Below Limit	C	Falling Oil Pressure Trend
COIL.MRV Above Limit	C	Cumulative Oil Added Above Limit
TOIL.MRV Above Limit	C	Time Since Oil Change Exceeded
Overspeed/Overtemp Alarms		
Level 1 ITT < 890, No WF Override	T	Overtemp With T5 Control
Level 2 ITT < 890, WF Override	T	Overtemp With T5 Control Disabled
Level 2 ITT < 945 or 1000, No WF Override	T	Overtemp With T5 Control
Level 2 ITT > 945 or 1000, WF Override	T	Overtemp With T5 Control Disabled
Level 2 Hot Start	T	Maximum Starting Temperature Exceeded
Level 1 NG < 99.4	T	Core Overspeed Above 99.4%
Level 2 NG < 102	T	Core Overspeed Above 102%
Level 1 NF < 98	T	Fan Overspeed Above 98%
Level 1 NF < 94.5	T	Fan Overspeed Above 94.5%
Level 2 NF < 99.7	T	Fan Overspeed Above 99.7%




Figure 15. Sample: T.O. No. 1A-10A-2-71MS-5 Troubleshooting Algorithm.

Using the primitive diagnostic algorithms as a start, each of the three CEMS performance alarms was revised and expanded. Specific enhancements address the following symptoms/conditions:

- Concurrent In-Flight Stall Event (TEMS)
- Concurrent In-Flight Low Performance Event (TEMS)
- Concurrent In-Flight Vibration Event (TEMS)
- Trend Data Check for Possible Transition Liner Collapse
- Correlation of Wear-Metal Debris with Specific Component
- Trend Data Check for Possible VG Tracking Changes
- Soft Sensor or Data Errors.

While it is difficult to quantify the complexity and contribution of this added analytical capability, one simple measure is the quantity of rules. Using an assumption of three decision-tree nodes per rule, the initial tree for each of the CEMS performance alarms is composed of about 10 rules; the JET-X versions use between 120 and 140 rules. Besides searching for a greater number of potential symptoms, much of the expanded diagnostic capability of JET-X results from its ability to examine combinations of symptoms to generate unique troubleshooting recommendations. This additional diagnostic "penetration" is a function of the amount of data or information that can be applied to the analysis of a problem.

Once the knowledge base for the initial 3 CEMS IV performance alarms had been developed, procedures for evaluating the remaining 51 alarms were added. This unplanned extension of JET-X capability was undertaken in order to make the system a more effective prototype of a production system. Besides providing diagnostic methods for each of the TEMS and CEMS IV alarms, integrating all alarms required that the JET-X architecture address the added complexity required to accommodate all alarms.

Initially, the other 51 alarms were added exactly as found in the T.O. However, modifications to these procedures were also introduced, but not on the same scale as those undertaken for the original three CEMS IV alarms. These enhancements fell into the following three categories:

- Incorporate experience acquired since the original algorithm was created
- Automate the answering of questions using information already in the knowledge base
- Include calculation procedures for analyzing fuel-control-related problems.

Modifications to 15 of the alarms were of the first and/or third type, which added to the diagnostic coverage of the procedure. Another 10 trees were enhanced by the second category of

change, automatic answering. In summary, a total of 28 of the original 51 diagnostic procedures were modified to varying degrees of complexity. Adding the remaining 51 alarms to JET-X enhances their use and accessibility, making JET-X more complete. However, most of the above enhancements (except automatic answering) could also be included in the original T.O. format, although implementation would be less convenient for the user.

An additional benefit was achieved by including diagnostic procedures for all 54 alarms in the knowledge base. Often combinations of alarms are indicative of specific failures, requiring little, if any, additional user-intervention to identify their cause. JET-X searches for related alarm combinations before proceeding with the analysis of individual alarms. This feature contributes a significant dimension to TEMS/CEMS IV troubleshooting technology that is not well-suited to a T.O.

8.5 GEN-X Module Types

A GEN-X knowledge base is built using three module types for representing rules and procedures, all of which are created in a graphical format available in the GEN-X development mode. These three modules are Decision (DEC) Trees, If/Then (IFT) Tables, and "And/Or" (AOR) Trees. Modules can be linked together using the GEN-X "back plane" command language. As previously discussed, the modular nature of GEN-X enables separate troubleshooting methods and techniques to be built into individual modules.

Decision trees, containing procedural knowledge, are used to guide the user through the evaluation of TEMS and CEMS IV data. The user examines data displays on the CEMS IV terminal and makes judgements concerning the presence of symptoms. In general, a yes or no response is made by the user (by means of menu selection) on the GEN-X end-user interface. Depending on the question and corresponding response, facts are set true or false as the session proceeds. Once the search for symptoms is complete, acquired facts (symptoms) are "examined" within If/Then Tables inferenced in a forward-chaining (event-driven) mode. When rules relating facts in the tables are fired, during inference, maintenance recommendations are generated.

If/Then Tables are also employed in the forward-chaining mode to search for related alarm combinations, which also trigger maintenance recommendations when rules are fired. Control of the overall troubleshooting scheme is carried out by If/Then Tables in the forward-chaining mode, while control of the more complex performance troubleshooting procedures is governed by If/Then Tables inferenced in backward-chaining (goal-driven) mode.

8.6 Modular Knowledge-Base Implementation

As stated, a natural by-product of the modular knowledge base is that specific troubleshooting procedures or techniques for examining TEMS and CEMS IV data can be built into a single module or combination of modules. For the three CEMS IV performance alarms, separate modules perform the following functions:

- Vibration Analysis
- Performance Step Assessment
- P5P3 Trend Evaluation
- Wear-Metal Trend Identification
- Wear-Metal Alarm Validation
- Stall Validation
- TEMS Performance Alarm Validation
- Maintenance-Induced Performance Changes
- Data-Error-Induced Performance Changes
- VG Tracking Changes
- Gun Round Check
- Performance Step Alarm Validation.

In addition, numerous modules support these functions by providing help facilities to assist data interpretation, feedback to the user on troubleshooting status, and control functions.

The additional 51 CEMS IV and TEMS alarm procedures are less developed than are the CEMS IV performance alarms and each was, therefore, represented by one or two decision tree modules.

The modular knowledge base enables simple and rapid modification of troubleshooting procedures. For a symptom-fault approach to troubleshooting, the knowledge base can be continually evolving as experience accrues. Consequently, the ability to easily update procedures is essential. This feature was exercised frequently during development because experience from more than one source was added to the knowledge base.

8.7 Knowledge-Base Execution

In order to illustrate the GEN-X inference process, Table 2 presents a step-by-step description of the execution of a small subset of the knowledge base for a CEMS IV performance alarm. Since only a few of these steps are actually seen by the JET-X operator, this provides an insight into all the "behind-the-scenes" functions performed. An explanation of the column titles for this table follows.

Step Number: Sequential index of the functions.

Module Name: Name of the GEN-X module in which the step occurs.

Text Displayed to User and/or Function Performed: Summary of the actual text seen in the end-user procedure window and/or description of the inference processes

Table 2. Knowledge-Base Execution for a Sampling of CEMS IV Performance Alarms.

STEP #	MODULE NAME	TEST DISPLAYED TO USER AND/OR FUNCTION PERFORMED	MENU SELECTION	BACK PLANE FUNCTIONS
1	CP3TEST.IFT	No text. Goal mode inference seeks to determine value of the fact "gas path damage"	N/A	TS: call gasptdg3.ift g* PS: none
2	GASPTDG3.IFT	No text. Goal mode inference seeks to determine the value of the fact "vib trend"	N/A	TS: call vibanal.dec* PS: none
3	VIBANAL.DEC	"Enter C10 on the CEMS terminal"	Continue	
4	VIBANAL.DEC	"Does the most recent point (or points) on the vibration trend plot show any evidence of being higher than previous vibration points"	Yes* No Example Explain Fact	TS: none PS: none
5	VIBANAL.DEC	"Could data scatter be the cause of the apparent rise in trended vibrations?"	Yes No* Scatter Explain Fact	TS: none PS: none
6	VIBANAL.DEC	No text. A return to the calling module (gasptdg3.ift) is executed. Prior to so doing the fact "vib trend" is set true and a short description of results is written to a file outside GEN-X.	N/A	TS: none PS: set "vib trend" true* append flags.rec "Rising vibration trend identified" + return*
7	GASPTDG3.IFT	No text. "vib trend" is set true.	N/A	TS: call gasptdg3.ift g PS: none
8	GASPTDG3.IFT	No text. Rule # 2 fires, fact "vib trend" is set true. Execute back plane command.	N/A	TS: none PS: t: call wrmtlst.ift f*
9	WRMTLST.IFT	No text. Forward chain through all rules; Rule # 6 fires, setting fact "wearmetal chk" false and execute back plane commands.	N/A	TS: none PS: f: call wrmtlchk.dec* t: call mtplot.dec t: call vfyntalm.ift g call mtlalign.aor
10	WRMTLCHK.DEC	"Because rising vibrations have been identified it is advisable to check for rising wearmetal trends. Enter C10 on the CEMS terminal to view wearmetal plots."	Continue	TS: none PS: none
11	WRMTLCHK.DEC	"Examine the wearmetal plots and identify any metals that appear to be rising beyond normal scatter. Enter rising wearmetals on the screen to follow"	Continue	TS: none PS: none
12	WRMTLCHK.DEC	No text. Back plane command calls an external program (setfacts.exe)	N/A	TS: Default Continue* PS: run setfacts -d-metals.def* -o-metals.dat update metals.dat

* Indicates Item Selected by User to Proceed, Which Activates Next Incremental Step Listed

Table 2. Knowledge-Base Execution for a Sampling of CEMS IV Performance Alarms (Concluded).

STEP #	MODULE NAME	TEST DISPLAYED TO USER AND/OR FUNCTION PERFORMED	MENU SELECTION	BACK PLANE FUNCTIONS
13	SETFACTS.EXE (external pgm)	User indicates which wearmetals (if any) have a rising trend: iron, silver, nickel, chromium, none of the above. After entry return to wrmtlchk.dec.	Iron Silver Nickel Chromium None of the Above*	N/A
14	WRMTLCHK.DEC	No text. Execute back plane command which loads facts set by the setfacts.exe program into the knowledge base.	N/A	TS: default Continue PS: run setfacts -d-metals.def -o-metals.dat update metals.dat *
15	WRMTLCHK.DEC	No text. Execute the back plane command "return", which returns to calling module wrmtlst.ift.	N/A	TS: none PS: return *
16	WRMTLTST.IFT	No text. Execute the next back plane command (call mtlsgn.aor).	N/A	TS: none PS: f: call wrmtlchk.dec t: call mtlplot.dec t: call vlymtalm.ift g call mtlsgn.aor *
17	MTLSIGN.AOR	No text. Forward chaining sets the fact "wearmetal sign" false and in so doing executes the PS back plane command "call mtlflagf.dec".	N/A	TS: none PS: t: call mtlflagf.dec f: call mtlflagf.dec *
18	MTLFLAGF.DEC	"No abnormal wearmetal signatures were identified by this part of the analysis". Execute back plane command to write a message to an external file (flags.rec).	Continue	TS: none PS: append flags.rec "No abnormal wearmetal signatures" *
19	MTLFLAGF.DEC	No text. Return to the calling module (mtlsgn.aor).	N/A	TS: none PS: return *
20	MTLSIGN.AOR	No text. Return to the calling module (wrmtlst.ift). Note: the return is implicit.	N/A	TS: none PS: t: call mtlflagf.dec f: call mtlflagf.dec
21	WRMTLTST.IFT	No text. Return to the calling module (gasptd3.ift). Note: the return is implicit.	N/A	TS: none PS: f: call wrmtlchk.dec t: call mtlplot.dec t: call vlymtalm.ift g call mtlsgn.aor
22	GASPTD3.IFT	No text. All conclusion facts have been set (step # 8). Return to calling module (cp3test.ift).	N/A	TS: none PS: none


* Indicates Item Selected by User to Proceed, Which Activates Next Incremental Step Listed

which are transparent to the end user. Text displayed to the user is in quotes. All steps for which the phrase "No Text" appears, perform transparent functions only.

Menu Selection: Display of the menu selections available to the operator at this step (if no menu is presented at a given step, N/A appears). The selection made at each step is indicated with an asterisk.

Back Plane Functions: List of all the functions which are present in the back plane of a fact in an If/Then Table, or of a node in a decision tree. Back plane functions execute in one of two modes, either to set the value of a fact or node, or else after the value of the fact has been set. In the table, functions of the first type are preceded with TS (to set) and the latter with PS (post set). The command(s) executed at each step is denoted by an asterisk.

Seven modules are accessed in this simulation: three If/Then Tables, three Decision Trees, and one And/Or Tree. Figures 16 through 22 show the front plane views of each of these modules (available only to the developer; the end user does not see this structure). In general only decision tree node text is displayed to the operator; the setting of facts true/false, inferencing of rule tables, and the execution of most back plane functions are not seen by the end user.



Fact	B	Value	Cost	1	2	3	4	5	6
Gas Path Damage	B	0.5000	100	T	F				
Maintenance Cause Check	B	0.5000	100			T	F		
Symptom Cause	B	0.5000	100					T	F
GPD Exit	B			<T	<F				
Maintenance Cause Exit	B					<T	<F		
Symptom Exit	B							<T	<F

Figure 16. Module CP3TEST.IFT.

Each step from Table 2 is located on the corresponding module diagram, showing the fact or node to which the step most directly applies. A brief description of each of the seven modules accessed in this example is included below.

CP3TEST.IFT: This table controls the order of execution of the three categories of symptoms for troubleshooting a step change in trended engine performance. This module is inferenced in the goal or backward chaining mode and is invisible to the operator.

Fact	B	Value	Cost	1	2	3	4	5	6	7	8	9	10	11	12	13
Gunfire Check	B	0.5000	100	.	.	.	.	.	.	.	.	.	.	.	.	.
FFG Tnd > Lmt		0.5000	100	T	F					F	F	F	F	F	F	F
Vib Trend	B	0.5000	100		T					F	F	F	F	F	F	F
Rising PSP3 Tnd	B	0.5000	100			T				F	F	F	F	F	F	F
Low Thrust		0.5000	100				T			F	T	F	T	F	T	
Valid Perf Prb	B	0.5000	20				T				F		F		F	
FFG> Lmt		0.5000	100						T	F	F	F	F	F	F	
Stall		0.5000	100					T		F	F	F	F	T	T	
Valid Stall	B	0.5000	100					T						F	F	
Vib Tnd-Ind	B		100	>T	>T											>F
Gas Path Dmg			175	>T	>T	>T	>T	>T	>T	>F	>F	>F	>F	>F	>F	



Figure 17. Module GASPTDMG.IFT.

Fact	B Value	Cost	1	2	3	4	5	6
FE FCST > Lmt	0.5000	100	T					F
AG FCST > Lmt	0.5000	100		T				F
NI FCST > Lmt	0.5000	100			T			F
CR FCST > Lmt	0.5000	100				T		F
FE Tnd > Lmt	0.5000	100					T	F
Wearmetal Check B			>T	>T	>T	>T	>T	>F

9
16
21

Figure 19. Module WRMTLTST.IFT.




Figure 20. Module WRMTLCHK.DEC.




Figure 21. Module MTLSIGNAOR.




Figure 22. Module MTLFLAGT.DEC.

GASPTDG3.IET: The search for all gas path symptoms that may be related to a step change in trended performance is controlled by this module. In addition, if vibrations are detected, this module will direct an examination of wear-metal data for evidence of secondary damage induced by the vibrations. This table is inferenced in the backward chaining mode and is also user-transparent.

VIBANAL.DEC: A decision tree which guides the search for a rising vibration trend signature in the CEMS IV data.

WRMTLTST.IET: This table checks for the presence of any wear-metal trend alarms in a forward chaining mode and is called when a vibration trend has been identified; its functions are unseen by the operator. If no alarms are present, the user will be directed to a module (**WRMTLCHK.DEC**) which will guide him in searching for wear-metal trends.

WRMTLCHK.DEC: In the case of rising vibration trends, if no CEMS wear-metal trend alarms occurred, this decision tree is called and directs the examination of wear-metal trend plots for any evidence of increasing trends.

MTLSIGN.AOR: Using a logical "or," this And/Or Tree "manifolds" facts that were set during the evaluation of CEMS IV wear-metal trend plots. If any wear-metal(s) were found to be rising, a single representative fact is set true, otherwise it is set false. Like rule tables, the functions of this module are transparent to the user.

MTLFLAGE.DEC: A two-node decision tree which communicates to the operator the fact that no wear-metal trends were identified. The objective of this and similar trees is to provide the user with feedback on symptom findings and session status.

The process represented in Table 1 is for a step loss in trended engine performance; the same alarm is utilized in the sample end-user session documented in Section 9.0. In an actual JET-X session, other nondiagnostic steps would precede the functions entered in Table 1; Step No. 1 is the beginning of the actual diagnostic process.

9.0 USING JET-X

In order to convey how JET-X is utilized and the user's perspective, a sample JET-X troubleshooting session is presented. An attempt has been made to be reasonably complete in the presentation of the actual screens viewed by the user during the troubleshooting process. All communication between the knowledge base and the user are through the GEN-X end-user interface.

9.1 GEN-X End-User Interface Description

The GEN-X end-user screen format consists of two side-by-side rectangular windows (the developer works in a different environment when constructing a GEN-X knowledge base). Figure 23 illustrates the features of this end-user interface, and the function of each of these display areas is summarized below.

9.1.1 Procedure Window

The left side of the screen is the procedure window which contains the current communication between the knowledge base and the end user. This communication might be in the form of a question which requires the user to choose a response from the pop-up menu, a step in a multi-stepped procedure, or a single statement passing information to the user on the status of the troubleshooting session. On the actual monitor, the window is blue and the question or current statement is highlighted within a cyan-colored window.

9.1.2 History Window

The right-hand zone of the screen is the history window and contains an upward scrolling record of questions and statements that have appeared in the procedure window, along with the responses made by the user. In general, the statement or question in the history window is a cryptic abbreviation of the wording that appeared in the procedure window.

What appears in the history window is under the control of the developer; some transactions between the knowledge base and the user are insignificant and therefore, are kept from appearing in the history window by the designer. The user's response to each procedure window entry is found immediately to the right of the summary statement in the history window, under the screen column titled "Value." Responses to procedure window questions other than by menu selection by the operator (for example, supplied by other modules, external programs, etc.) will also appear in the history window.

9.1.3 Explanation Window

Questions appearing in the procedure window are intentionally concise to avoid wearying the frequent JET-X user with more text than needed to perform an operation. However, for the novice, many of these questions may seem to be cryptic or incomplete. For each entry in the procedure window, the developer is able to provide additional text to support the procedure window question or statement. This information, when available, is displayed in the explanation window when the user selects "Help" or "Explain Fact" from the pop-up menu.




Figure 23. GEN-X End-User Screen.

9.1.4 Pop-Up Menu

All possible responses to a procedure window entry are contained in the pop-up menu which appears as an overlay on the history window. If a question is in the procedure window, the pop-up menu will contain a minimum of two possible responses, generally "Yes" and "No." Statements (or steps in a procedure) have only a single menu option which, in JET-X, is typically "Continue."

Multiple menu selections also are presented when the user has a choice of how he/she would prefer to proceed; menu selections at such a fork in JET-X might be "Search, Restart, Exit," etc. At any time the pop-up window can be toggled on and off, using the "+" key. Menu selections are made by moving a cyan-colored cursor over the desired menu selection, using the Up and Down arrow keys, and then depressing the "Enter" key. For a single-item menu, only Enter is necessary.

In addition to these major display areas, two other regions on the end-user screen offer information or control options to the user. These are described in the following paragraphs.

9.1.5 Title

To the right of the word "Title," the name or descriptor of the knowledge-base module currently being processed may appear. The words that appear here are under the control of the developer and are intended to give the user some idea of what is being executed, as well as to provide the developer with a flag to identify the current module. This latter function is of value when troubleshooting knowledge-base errors or malfunctions.

9.1.6 Footer Menu

At the bottom of the GEN-X end-user screen is a selection of overall control functions available to the user. In general these will either terminate the current session, initiate the current session, or terminate the session and restart another. The number preceding each selection corresponds to the keyboard function key (F1, F2, etc.) that will activate the associated feature. Equivalent functions built into the JET-X knowledge base are available to the user at appropriate forks in the troubleshooting process.

Session record-keeping functions are executed when termination is initiated by means of the JET-X functions (which are activated through the pop-up menu); this is not the case when the footer menu is used. Footer menu functions abort the current session and bypass all JET-X close-out activities.

9.2 Interface with External Programs

Numerous, external (to GEN-X), "C" language programs are part of JET-X. These present the user with interface formats which, with one exception, appear very much like the GEN-X end-user screen.

Support programs will present the user with two, blue, rectangular-format windows identical to the GEN-X procedure and history windows. This design was intentional to avoid distracting screen changes when leaving and returning to the GEN-X knowledge base. Subtle differences do exist, however. Absent from the support program screens will be: all GEN-X identification information at the top of the screen; Title, History, and Value headers; and the Footer menu.

Functional interaction with these external programs is also different from GEN-X, although each individual program has unique interface requirements. Table 4 (contained in Section 10.0) summarizes the external programs. In some cases external programs require no user interaction, but serve to display information only. Section 10.0 provides further discussion of these external programs.

9.3 JET-X Session Features

Most JET-X troubleshooting sessions will have several features in common. These general characteristics are described as follows.

9.3.1 User-Name Entry

The first interaction the user has with JET-X occurs in a support program outside of GEN-X which asks the user to enter his/her name. A list of users appears on the right-hand screen; selection is made using the Up and Down arrow keys and the Enter key. The user's name is included in various JET-X reports.

9.3.2 Activity Selection

Besides Troubleshoot, two other top level functions are available in JET-X. Selecting "Reports" allows the user to view summaries of past troubleshooting sessions, delete records no longer needed, and examine a catalog of available historical information. The selection of "Introduction" provides several training features, including: an overview of JET-X, an explanation of TEMS and CEMS IV, a glossary of parameters and terms unique to TEMS and CEMS IV, and a feature to enable browsing through the video support facilities included in JRS.

9.3.3 Diagnose Versus TEMPER

If Troubleshoot is selected as the activity, the user is then asked to indicate if he/she wishes to perform troubleshooting using the JET-X knowledge base ("Diagnose") or walk through a simulated diagnostic session with TEMPER. More information on the TEMPER option is given in Section 12.0 of this report.

9.3.4 Session Type

JET-X next asks whether the purpose of the session is Work or Training. If "Work" is selected, a permanent session record will be created, containing the actual session data. No records are retained following a Training session, which as its name implies, is used primarily for training and demonstration.

9.3.5 Input ESN (Engine Serial Number)

If Work is selected as the session type, then JET-X will ask the user to input a 4-digit ESN (engine serial number), which is used in the session records. The last 4-digits of the TF34 ESN are to be input at this point. This function is performed by an external program.

9.3.6 Input Alarms

The user is presented with a multiscreen display of all possible TEMS and CEMS IV alarms and is asked to select those that pertain to the engine being diagnosed. Alarm entry is external to GEN-X and is performed by the "Setfacts" program.

9.3.7 Guided Versus Symptom Troubleshooting

If at least one of the CEMS IV performance alarms was entered, the user will be given the option of analyzing this alarm in JET-X by one of two methods, "Guided" or "Symptom." JET-X will guide the analysis of the alarm when the Guided method is used. For Symptom troubleshooting, the user selects the particular investigation the user feels is appropriate.

JET-X troubleshooting is interactive; the operator will be requested to call up relevant CEMS IV displays and answer questions regarding the data on the screen. The depth of a particular JET-X session is dependent on the amount of data available in CEMS IV that can be applied to troubleshooting a specific alarm.

9.4 Sample Session: Guided Troubleshooting


The following sample JET-X session (Figures 24 through 46) illustrates the Guided approach to analyzing a step change in trended performance ($\text{NFT5 TND} < \text{LMT}$), which is one of the three CEMS IV performance trend alarms for which JET-X troubleshooting logic is well-developed. The intent of this example is to demonstrate the nature of a JET-X session, rather than to completely document the troubleshooting logic for a given alarm.

In this example, the rapidly falling engine performance is due to damage to the core compressor, caused by land material breaking free of the rotor and producing airfoil damage. The symptom that will enable identification of this malfunction is a rising core vibration trend that results from the increasing imbalance of the compressor rotor.

For each screen which is part of this sample session, the following four pieces of information are presented:

- Frame Number - Indicates the sequence number of the frame being displayed for an actual JET-X session. When the numbers are nonsequential, frames have been skipped.
- Preceding Screens - Since not all of the end-user screens that would be part of the real JET-X session appear in the example, this section summarizes the content and purpose of the screens which have been skipped.

- Current Screen - Provides an explanation of the current screen and the options available to the user at that point.
- User Selection - Indicates the choice (or menu selection) made at the current node for this example. Where applicable, other options that were not selected are explained.




Frame No. 2:

Preceding Screen(s) - The user was asked to select his/her name from a list of possible JET-X users.

Current Screen - Introduction to JET-X; the user can select the desired JET-X function: TROUBLESHOOT, REPORTS, INTRODUCTION, HELP.

User Selection - TROUBLESHOOT; this selection enters the JET-X diagnostic and troubleshooting portion of the knowledge base for analyzing TEMS and CEMS IV alarms.

Figure 24. Frame No. 2 for JET-X Guided Troubleshooting Example.



Frame No. 6:


Preceding Screen(s) - The user was given a choice of troubleshooting methods: DIAGNOSE or TEMPER, DIAGNOSE was selected, which accesses the JET-X knowledge base.

- Next, the choice of WORK or TRAINING was presented; WORK was selected, which triggers the creation of a permanent record of the session.*
- After selecting WORK, the user was then asked to enter a 4-digit engine serial number.*

Current Screen - The user is advised that the next screen will be a menu of all possible TEMS and CEMS IV alarms, on which he/she will enter the alarms for the current engine.

User Selection - Only one menu item available, CONTINUE, which continues the procedure.

Figure 25. Frame No. 6 for JET-X Guided Troubleshooting Example.




Frame No. 7:

Preceding Screen(s) - None.

Current Screen - The TEMS CEMS IV alarm list; all 54 alarms are on this list, which can be scrolled up and down with the arrow keys. The alarms are selected (or deselected) using the left and right arrow keys.

User Selection - "NFT5 TREND BELOW LIMIT."

Figure 26. Frame No. 7 for JET-X Guided Troubleshooting Example.




Frame No. 8:

Preceding Screen(s) - None.

Current Screen - Because a CEMS IV performance alarm was indicated, the user can choose to troubleshoot using either the GUIDED or SYMPTOM method.

User Selection - GUIDED troubleshooting, which will direct the analysis of all possible data symptoms that could be related to alarm that was entered.

Figure 27. Frame No. 8 for JET-X Guided Troubleshooting Example.



Frame No. 14:


Preceding Screen(s) - Verification of the NPTS Trend Below Limit Alarm was accomplished by guiding the user through a check of the actual data on which the alarm was based. For this example, the alarm was found to be valid and troubleshooting will continue.

Current Screen - "Enter C30 ..." directs the operator to enter the command "C30" on the CEMS IV terminal. This will display a trend plot of Core Vibrations (FFG) against EOT (Engine Operating Time) on the CEMS IV terminal.

User Selection - Choices are CONTINUE or HELP; select CONTINUE.

- HELP, if selected, would display an explanation of the instruction to "Enter C30 on the CEMS terminal."

Figure 28. Frame No. 14 for JET-X Guided Troubleshooting Example.



Frame No. 15:


Preceding Screen(s) - None.

Current Screen - The question directs the user to make a judgement about the trend plot (C30) currently displayed on the CEMS terminal. Specifically, he/she is to decide if it appears that an abnormal rise in vibration trend is present.

User Selection - Select EXAMPLE, which will display an example of a rising vibration trend on JRS and provide a brief description on the JET-X terminal of how to identify such a trend.

- HELP would provide a brief explanation of how to identify a rising vibration trend.

Figure 29. Frame No. 15 for JET-X Guided Troubleshooting Example.




Frame No. 15(a):

Preceding Screen(s) - None.

Current Screen - The entire narrative (which comes into the procedure window one step at a time) explaining the example trend plot on JRS is displayed. This is intended as a guide to the novice user who may be unfamiliar with making judgements about TEMS and CEMS IV data. Experienced CEMS troubleshooters would not need to consult the example.

User Selection - Select CONTINUE (only option available).

Figure 30. Frame No. 15(a) for JET-X Guided Troubleshooting Example.




Frame No. 15(b):

Preceding Screen(s) - None.

Current Screen - At the conclusion of viewing the example, the operator is returned to the original question at which the EXAMPLE was selected. The user must now answer the question: "Is a rising vibration trend present in the actual engine data?"

User Selection - Select YES, indicating that a rising vibration trend is observed to exist.

Figure 31. Frame No. 15(b) for JET-X Guided Troubleshooting Example.



Frame No. 16:


Preceding Screen(s) - None.

Current Screen - This question provides an opportunity for the troubleshooter to indicate whether he/she thinks the apparent vibration trend is due to data scatter. Often it may not be clear whether a real trend exists or not; if scatter is a possibility, then a method will be provided to assist in determining if the trend is real.

User Selection - Select HELP, which will provide a brief explanation that will be a guide to identifying data scatter.

- SCATTER would provide a more in-depth example, along with a JRS display illustrating how to judge when scatter is present.

Figure 32. Frame No. 16 for JET-X Guided Troubleshooting Example.




Frame No. 17:

Preceding Screen(s) - None.

Current Screen - This screen is the same as that displayed and designated "Current Screen" in Frame No. 16, but with the explanation displayed above the procedure and history windows.

User Selection - Enter YES, indicating that data scatter could be the cause of the apparent vibration trend.

Figure 33. Frame No. 17 for JET-X Guided Troubleshooting Example.




Frame No. 18:

Preceding Screen(s) - None.

Current Screen - As a first step in the procedure to determine if scatter is present, the operator is directed to enter the command "C61" on the CEMS terminal and respond when the display is on the screen. C61 will display a table containing the vibration trend data.

User Selection - Select YES: the table is displayed on the terminal.

Figure 34. Frame No. 18 for JET-X Guided Troubleshooting Example.




Frame No. 21:

Preceding Screen(s) - A statement informing the operator that the rising vibration trend was confirmed was displayed prior to this current JET-X screen.

Current Screen - Gun-gas ingestion can enter the engines during flight, contaminating the compressor and eventually causing a deterioration in performance. If a large number of rounds were fired on a single flight, a step loss in trended performance can occur.

User Selection - Select NO. This bypasses the determination of whether gun-gas ingestion might be the cause of the reduced performance.

Figure 36. Frame No. 21 for JET-X Guided Troubleshooting Example.



Frame No. 22:


Preceding Screen(s) - None.

Current Screen - This screen is principally informative, indicating what is to follow. Entering "C18" on the CEMS terminal will pull up wear-metal trend plots for scrutiny by the diagnostician.

User Selection - Select CONTINUE to proceed with the analysis of wear-metal alarms.

- HELP would provide a brief explanation of the relationship between vibrations and wear-metals.

Figure 37. Frame No. 22 for JET-X Guided Troubleshooting Example.




Frame No. 23:

Preceding Screen(s) - None.

Current Screen - Instructions directing the user to examine the wear-metal plots for abnormally rising trends. A screen will be presented on which the user can indicate which, if any, of the wear-metals are rising.

User Selection - Select CONTINUE.

Figure 38. Frame No. 23 for JET-X Guided Troubleshooting Example.




Frame No. 24:

Preceding Screen(s) - None.

Current Screen - An external "C" program lists the various wear-metals for which plots are displayed on the CEMS terminal. The operator moves the highlighted block up and down and selects the rising wear-metal(s) using the left/right arrow keys.

User Selection - Select NONE OF THE ABOVE, which indicates that none of the wear-metals were found to be rising.

Figure 39. Frame No. 24 for JET-X Guided Troubleshooting Example.




Frame No. 25:

Preceding Screen(s) - None.

Current Screen - A status statement to confirm what was entered on the screen identified as "Current" in Frame No. 24.

User Selection - CONTINUE is the only option.

Figure 40. Frame No. 25 for JET-X Guided Troubleshooting Example.



Frame No. 26:


Preceding Screen(s) - None.

Current Screen - Another status statement; however, this gives a summary of the first segment of the troubleshooting session, which searches for symptoms of gas path damage (principally the compressor and turbines).

- *For this example, evidence of gas path damage was found (the rising vibration trend); a summary of the session will be displayed.*
- *If no positive gas path damage were found, no session summary would be presented, but rather, analysis would continue on to the next portion of troubleshooting.*

User Selection - CONTINUE is the only option.

Figure 41. Frame No. 26 for JET-X Guided Troubleshooting Example.



Frame No. 27:


Preceding Screen(s) - None.

Current Screen - This is the session summary screen, which provides the user with an overview of the troubleshooting session. Included are five pieces of information: log-in and record-keeping data, alarms input, alarms analyzed thus far, significant symptoms noted during the session, and maintenance recommendations resulting from the session.

- *A borescope recommendation has been made due to the detection of the rising vibration trend.*
- *The Session Summary is produced by an external program.*

User Selection - Hit ENTER to continue. If a print-out of the summary is desired, hit "P" before ENTER.

Figure 42. Frame No. 27 for JET-X Guided Troubleshooting Example.




Frame No. 28:

Preceding Screen(s) - None.

Current Screen - The user is presented with the option of viewing and/or printing a more complete recommendation than that which appeared on the Session Summary that was just presented.

User Selection - Select VIEW/PRINT which will display the extended explanation and provide an option to print it out.

Figure 43. Frame No. 28 for JET-X Guided Troubleshooting Example.



Frame No. 29:

Preceding Screen(s) - None.


Current Screen - The extended recommendation is displayed by one of the external programs. In this case, the recommendation is tied more completely to the symptoms which were entered during the interactive session with the user.

- *The right-hand side of this screen indicates that a video (JRS) display is available to supplement the recommendation.*
- *The additional help is selected with the "H" key.*

User Selection - Select "H" to view the JRS help video.

- *ENTER would return the user to the JET-X knowledge base, thus enabling him/her to continue.*

Figure 44. Frame No. 29 for JET-X Guided Troubleshooting Example.




Frame No. 29(a):

Preceding Screen(s) - None.

Current Screen - This is a borescope photograph indicating the kind of damage that is the likely cause of the rising vibration trend. This photograph appears on the JRS screen, not on JET-X.

User Selection - ENTER returns the user to JET-X.

Figure 45. Frame No. 29(a) for JET-X Guided Troubleshooting Example.



Frame No. 30:

Preceding Screen(s) - None.


Current Screen - The user can select what to do next. If satisfied with the JET-X recommendation he/she may want to EXIT or RESTART with a new engine. END will return the user to the first JET-X screen. If the user would like to continue analyzing the current alarm, he/she would select SEARCH.

User Selection - Being satisfied with the JET-X recommendation, select EXIT. This ends the session and concludes the example!

Figure 46. Frame No. 30 for JET-X Guided Troubleshooting Example.

9.5 Sample Session: Symptom Troubleshooting

The following sample JET-X session (Figures 47 through 56) illustrates the Symptom approach to analyzing a step change in trended performance (NFT5 TND < LMT), the same alarm that was evaluated for the Guided example provided in Section 9.4. The same cause for the problem (compressor rotor land delamination) will be assumed for this example. It will be noted that many of the screen interfaces are the same as for the Guided session. Up until the selection of Symptoms or Guided troubleshooting, the paths are obviously identical. The actual troubleshooting for a rising vibration trend is identical, because the same GEN-X modules are called by both methods.




Frame No. 2:

Preceding Screen(s) - The user was asked to select his or her name from a list of possible JET-X users.

Current Screen - Introduction to JET-X; the user can select the desired JET-X function: TROUBLESHOOT, REPORTS, INTRODUCTION, HELP.

User Selection - TROUBLESHOOT; this selection enters the JET-X diagnostic and troubleshooting portion of the knowledge base for analyzing TEMS and CEMS IV alarms.

Figure 47. Frame No. 2 for JET-X Symptom Troubleshooting Example.



Frame No. 6:


Preceding Screen(s) - The user was given a choice of troubleshooting methods: DIAGNOSE or TEMPER; DIAGNOSE was selected, which accesses the JET-X knowledge base.

- *Next, the choice of WORK or TRAINING was presented; WORK was selected, which triggers the creation of a permanent record of the session.*
- *After selecting WORK, the user was then asked to enter a 4-digit engine serial number.*

Current Screen - The user is advised that the next screen will be a menu of all possible TEMS and CEMS IV alarms, on which he or she will enter the alarms for the current engine.

User Selection - Only one menu item (CONTINUE) is available, which continues the procedure.

Figure 48. Frame No. 6 for JET-X Symptom Troubleshooting Example.



Frame No. 7:

Preceding Screen(s) - None.

Current Screen - The TEMS and CEMS IV alarm list; all 54 alarms are on this list, which can be scrolled up and down with the arrow keys. The alarms are selected (or deselected) using the left and right arrow keys.

User Selection - "NFT5 TREND BELOW LIMIT."

Figure 49. Frame No. 7 for JET-X Symptom Troubleshooting Example.

Title: Select Activity

History

Value

CEMS performance alarms are indicated on this engine. These alarms can be analyzed by specific symptoms or using a guided method similar to all other alarms.

Do you want to
follow a GUIDED trouble
shooting procedure
or
use Specific SYMPTOMS for
trouble shooting?

(If SYMPTOMS is selected this
feature will execute first)

Troubleshoot / Reports /
Introduction? Troubleshoot
Troubleshooting method? Diagnose
Type of session? Training
CEMS & TEMS alarms entered Continue
Any CEMS perf alarms? Yes

MENU
Guided
Symptoms
Help

1:Reset 2:Replay 9:Execute 10:Stop

Frame No. 8:


Preceding Screen(s) - None.

Current Screen - Because a CEMS IV performance alarm was indicated, the user can choose to troubleshoot using either the GUIDED or SYMPTOM method.

User Selection - SYMPTOMS troubleshooting.

- This will allow the user to search the data for specific symptoms that he/she may consider the most likely cause of the alarm(s).

Figure 50. Frame No. 8 for JET-X Symptom Troubleshooting Example.



Frame No. 10:

Preceding Screen(s) - After selecting SYMPTOMS, the user was asked whether he/she would like to view an explanation of the SYMPTOMS method of troubleshooting. YES or NO options were available; for this example, NO was selected.


- If YES had been chosen, an explanation of how the SYMPTOMS method works would have been presented.

Current Screen - A menu of possible symptoms is presented, each of which could be related to a cause of low trended performance. The user selects one symptom and hits ENTER. This screen is generated by the SETFACTS Program and is external to GEN-X.

User Selection - For this example, the symptom "RISING VIBRATION TREND" is selected.

- Based on experience, the user would select this option if it were suspected that a rising vibration trend accompanied a sudden step loss in engine performance.

Figure 51. Frame No. 10 for JET-X Symptom Troubleshooting Example.



Frame No. 11:

Preceding Screen(s) - None.


Current Screen - "Enter C30 ..." directs the operator to enter the command "C30" on the CEMS IV terminal.

- This will display a trend plot of Core Vibrations (FFG) against EOT (engine operating time) on the CEMS terminal (same as GUIDED, Step No. 14).

User Selection - Choices are CONTINUE or HELP; select CONTINUE.

- HELP, if selected, would display an explanation of the instruction to "Enter C30 on the CEMS terminal."

Figure 52. Frame No. 11 for JET-X Symptom Troubleshooting Example.



Frame No. 12:

Preceding Screen(s) - None.

Current Screen - This question asks the user to make a judgement about the trend plot (C30) currently displayed on the CEMS terminal. Specifically, the user is to decide if an abnormal rise in the vibration trend is present. (This step is the same as Step No. 15 of the GUIDED approach.)

User Selection - Select YES, indicating the presence of an abnormal vibration trend plot.

- *EXAMPLE would display (on JRS) an example of a rising vibration trend and provide a brief description on the JET-X terminal of how to identify such a trend.*
- *HELP would provide a brief explanation of how to identify a rising vibration trend.*

Figure 53. Frame No. 12 for JET-X Symptom Troubleshooting Example.

Title: Select Activity

History

Value

Could data scatter be the cause of the apparent rise in trended vibrations?

Select SCATTER for help interpreting vibration data scatter.

Troubleshoot / Reports /	
Introduction?	Troubleshoot
Troubleshooting method?	Diagnose
Type of session?	Training
CEMS & TEMS alarms entered	Continue
Any CEMS perf alarms?	Yes
Use Guid	Symptoms
(Index)?	Yes
Want SYM	No
Vibratio	Scatter
availabl	Help
Evidence	Continue
vibration trend?	Yes

1:Reset 2:Replay 9:Execute 10:Stop

Frame No. 13:

*Preceding Screen(s) - None.**Current Screen - This question asks the troubleshooter to indicate whether he/she thinks the apparent vibration trend is due to data scatter or not.*

- Often, it may not be clear whether a real trend exists or not. If scatter is a possible cause for the apparent trend, then selecting SCATTER will provide a method to assist in determining if the trend is real. (This step is the same as GUIDED Step No. 16.)

User Selection - Select NO.

- HELP, would provide a brief explanation that is a guide to identifying data scatter; whereas, SCATTER would give a more in-depth example along with a JRS display illustrating how to judge when scatter is present.

Figure 54. Frame No. 13 for JET-X Symptom Troubleshooting Example.

SESSION IDENTIFICATION:
session number: 0058
date: 04-06-88 12:15
user name: Costen
engine number: 2323
action: Symptoms

CEMS/TEMS ALARMS - INPUT:
NFT5 TND < LMT (CP3)

CEMS/TEMS ALARMS - EVALUATED:

DIAGNOSTIC PATH MARKERS:
Rising vibration trend identified.

MAINTENANCE RECOMENDATIONS:

gp03 Rising core vibration trend:
borescope

Hit 'P' to print this information. Hit Return to Continue

Frame No. 15:

Preceding Screen(s) - The user was told that a summary of the current session will be displayed. The only menu option available was CONTINUE.

Current Screen - This is a session summary of the results of troubleshooting for a vibration symptom. The format is the same as the session summary displayed at the conclusion of the GUIDED example.

- Recommendations made by JET-X when in the SYMPTOMS mode have a different code number than those made in the GUIDED method, even though the recommendations may be the same. The reason for this is that different modules are used by the two methods to combine symptoms to produce recommendations. Since more facts are generally entered during a GUIDED session, these recommendations may be more robust and, therefore, are distinguished from those made during SYMPTOMS troubleshooting.

User Selection - No selection, the user hits ENTER to continue the session.

Figure 55. Frame No. 15 for JET-X Symptom Troubleshooting Example.

Title: Select Activity

History

Value

A summary of the results of the current search using symptoms will be displayed.

Choose Next Step:

SYMPTOM: check another symptom
 GUIDED: guided troubleshooting
 RESTART: enter a New Engine
 EXIT

Troubleshoot / Reports /	
Introduction?	Troubleshoot
Troubleshooting method?	Diagnose
Type of session?	Training
CEMS & TEMS alarms entered	Continue
Any CEMS perf alarms?	Yes
Use Guid	Symptoms
(Index)?	No
Want SYM	Guided
Vibratio	Restart
availabl	Exit
Evidence	Continue
vibration trend?	Yes
Could data scatter cause the trend?	No
Display Results	Continue
Any further display function?	None

1:Reset 2:Replay 9:Execute 10:Stop

Frame No. 17:

Preceding Screen(s) - The user was presented with the option of viewing and/or printing a more complete recommendation than that which appeared on the Session Summary, which was presented in Frame No. 15. NONE was selected for this example. (This process selection is the same as GUIDED Step No. 28.)

Current Screen - The options presented are to:

- EXIT the session
- RESTART troubleshooting with a new engine
- Select another SYMPTOM to examine, or
- Continue the session using GUIDED troubleshooting.

User Selection - EXIT is selected, which terminates the session and stops GEN-X execution.

Figure 56. Frame No. 17 for JET-X Symptom Troubleshooting Example.

10.0 JET-X EXTERNAL SUPPORT ROUTINES

Certain features considered necessary or desirable for the JET-X application were not directly available in GEN-X; consequently, external "C" language programs were written in order to provide these capabilities. These external programs are called from the knowledge base using the back-plane feature of GEN-X. The programs are designed to provide an end-user screen format similar to that provided by GEN-X. (Note that some of these capabilities will be included in future releases of GEN-X, based upon their contribution to the JET-X system.)

10.1 JET-X Retrieval System

The JRS (JET-X Retrieval System) is the most significant feature developed to augment the troubleshooting procedures built into the GEN-X shell. GEN-X permits "Help" text to be displayed during the course of a session to facilitate diagnostics; however, the Help text is only a small part of the information that could prove useful during the diagnostics and repair procedures. Access to diagrams, photographs, and repair manuals may also be extremely valuable. Although these resources are available in printed form (for example, Tech Orders), a better means is needed to access the information. Specifically, it is desirable that the computer automatically find the relevant information for a specific problem, thereby bypassing the need to find the appropriate page of the correct manual. This type of facility is a goal of the "hyper-text" systems that are beginning to appear on the market.

The JRS provides digital (bit-mapped) storage and display of video Help screens relevant to the troubleshooting procedures and maintenance recommendations of JET-X. These Help screens consist of photographs, sample CEMS IV screens, diagrams, and drawings, as listed in Table 3 and illustrated in Figures 57 through 60. These graphical images represent the kind of visual information a troubleshooter may find useful; hence, they are made available at appropriate points during the JET-X diagnostic session.

One of the uses of these displays is to supplement the experience of the user in making decisions and judgements about CEMS IV data. Specifically, when asked by JET-X whether a certain symptom is present in the engine data, JRS will provide an example of that symptom. For example, Figure 59 shows a JRS display that illustrates a performance shift accompanied by a trend in the P5P3 characteristic. Another key use of this feature is the enhancement of maintenance recommendations with additional information in graphic form. For example, photographs (taken through a borescope) of compressor and turbine damage are available when JET-X recommends a borescope inspection of these components (Figure 60). Assistance of this type can be of great value to the inexperienced maintenance technician, who may be unsure of what to look for when performing the required inspection.

For this prototype system, JRS resides on a separate PC (personal computer) equipped with a high resolution black and white graphics monitor. Selected graphical or photographic materials were scanned and converted to digital format and placed on the JRS hard disk using applicable compression techniques to enhance storage efficiency. Calls for a JRS display can

Table 3. Video Displays Available in JRS.

1. Photograph of an A-10 firing its gun in flight; displayed when logging into JET-X
2. Three-dimensional cut-away view of the TF34-100 engine
3. Diagram illustrating the relationship between TEMS and CEMS IV
4. Photograph of TEMS hardware and the A-10 aircraft
5. Borescope photograph of compressor damage
6. CEMS plot presenting scatter interpretation in trended VG tracking data
7. CEMS plot demonstrating a rising vibration trend
8. CEMS table illustrating the validation of performance trend data
9. CEMS table with associated instructions for selecting the appropriate data and performing the P5P3 step change calculation
10. CEMS plots illustrating the difference between a real performance trend and data scatter
11. CEMS data frame with example of a "TEMS Low Thrust" event with the opposite engine at idle power
12. CEMS data frame with example of a TEMS Low Thrust event caused by an excessive extrapolation of performance data
13. CEMS data frame with example of a TEMS Low Thrust event accompanied by a T5-amp lock-out
14. CEMS performance trend plot with maintenance actions indicated
15. Split screen CEMS trend plot showing performance changing as trended VG tracking changes
16. CEMS trend plot depicting a step change in performance
17. CEMS data frame with example of a low thrust event caused by a bad value of TT2
18. CEMS takeoff data frame, with low takeoff performance caused by an erroneous P0, which manifests itself with a high takeoff Mach number
19. CEMS takeoff data frame with low takeoff performance caused by an erroneous TT2 value
20. CEMS takeoff data frame with low takeoff performance caused by an excessive extrapolation of performance from a part-power condition
21. CEMS takeoff data frame with low takeoff performance caused by an excessive extrapolation of performance due to a T5-amp lock-out event

Table 3. Video Displays Available in JRS (Concluded).

22. Schematic diagram of an outer transition liner collapse
23. Borescope photograph illustrating an outer transition liner collapse
24. Borescope photograph showing an outer transition liner collapse and the Stage 3 nozzle
25. Borescope photograph depicting erosion of the land coating on a compressor rotor
26. T.O. diagram of the maneuver envelope of the A-10
27. CEMS vibration trend plot illustrating data scatter interpretation
28. CEMS data frame illustrating an in-flight record with the aircraft slats extended
29. CEMS vibration trend plot showing a step increase in trended vibrations
30. Split screen CEMS trend plot with trended VG and trended NGNF, used to confirm a true shift in VG tracking
31. CEMS trend plot showing how to identify scatter in trended P5P3
32. Split screen CEMS trend plot illustrating falling performance and rising P5P3
33. Example of TEMS high resolution data of a TEMS detected in-flight stall
34. CEMS trend plot illustrating differentiation of a rising iron wear-metal trend from data scatter
35. CEMS trend plot demonstrating the differentiation of a rising silver wear-metal trend from data scatter
36. CEMS trend plot showing differentiation of a rising nickel wear-metal trend from data scatter
37. CEMS trend plot depicting differentiation of a rising aluminum wear-metal trend from data scatter
38. CEMS performance trend plot, with an example of an additional flight, following a low performance takeoff
39. CEMS data table illustrating how to calculate the average performance prior to a low performance takeoff




Figure 57. Sample Engine Cross Section Schematic in JRS.




Figure 58. Sample CEMS IV Tabular Display in JRS.




Figure 59. Sample CEMS IV Plot Display in JRS.




Figure 60. Sample Borescope Photograph in JRS.

be initiated in the knowledge base, either by user option or automatically, and are passed to JRS through serial port communications. JRS software, which runs in the Microsoft Windows environment, provides the following features:

- Communications Between the JET-X PC and JRS
- Interpretation of Commands Received from JET-X
- Display of Scaled Images on the JRS Screen.

10.1.1 JRS Data Flow and Modules

A top level diagram of the JRS data flow is presented in Figure 61. Data enters the system at three places: from JET-X, through the serial port; from the user, by means of a "mouse" attached to the JRS computer; and from the disk where images are stored. Data leaves the JET-X Retrieval System as status information to the JET-X system, or as a bit map of pixels for the monitor. The JRS software functions are allocated to three modules using a "windows" architecture.

The first module is the Communication Manager. This program accepts commands from the JET-X host and performs the necessary "handshaking." Database management functions are provided by a second module called the Table Manager. This program provides the translation from a "key," supplied by the JET-X system, to a set of records which specify document names, page names, and the file names where individual pages are stored. The third module, called the User-Interface Manager, handles the user interfaces and page presentation.

10.1.2 JRS Interfaces

The JRS is designed to act as a context-sensitive, help facility controlled by the JET-X system. Thus, it is the responsibility of the JET-X system to provide the context of the problem; the user, however, is free to explore the information within that context. Accordingly, the JRS has two interface modes: the interface to the JET-X host, and the interface to the end user.

- Host Interface - The JET-X host system provides context to the JRS by issuing one of 13 commands through the serial port of the computer. The commands include: query commands to search for a document or set of documents by a meaningful key; display commands to display a page, or a portion thereof; attribute commands to change the appearance of a page by reversing the image, or a portion of the image, or by overlaying graphics and text.
- User Interface - Once the context is set, the user can choose to explore the document options requested by the host (JET-X) system. Since the user interacts with the JET-X expert system using a keyboard on the host PC, the JRS-user interface is completely mouse-controlled. This eliminates the confusion that would arise if the user had to work with two keyboards.

The JRS-user interface uses two windows, as shown in Figure 62 (also demonstrated in Figures 58 through 60). The top window is called the page window; it is the view port for the currently




Figure 61. Data Flow and Module Breakdown for JRS.

selected page. The second window is called the control panel. Analogous to the dashboard of a car, this panel provides switches to select options and control operation; it also provides status information to show the state of the JRS. The control panel uses a number of standard window controls to give the JRS "point and select" operation. For example, to cause the image tones to be reversed, the user selects the negative box. There are no commands to remember.




Figure 62. Illustration of JRS Page Window and Control Panel.

To examine a document, the user can select document pages in sequence (selection is made by clicking on the "Next Page" button), or can randomly access any other page by double-clicking on its name (listed in the box identified as "Page"). Bookmarks can also be inserted, to allow a quick return to a particular document and page. The user can view an enlargement of any portion of the image by selecting the "Zoom" box and highlighting the area.

10.1.3 Role of JRS in Future Expert Systems

The role of JRS in future versions of JET-X for TF34 engines may be different. In a mature, integrated, diagnostic expert system, the CEMS IV data base would be integrated with the expert system knowledge base. Consequently, a direct access to engine data in the CEMS IV data base, without user intervention, coupled with data interpretation from a diagnostic viewpoint, would probably reduce the need for including CEMS IV data frames in JRS, except perhaps, for training purposes. On the other hand, considerable additional data, such as bore-scope photographs, for interpreting complex situations might be added. Certain unique test sequences may be simulated to provide a "video" test procedure. JRS would play a vital role if the details of repair and maintenance procedures were to be provided, along with the data evaluation procedures of the prototype JET-X.

The computer hardware for an integrated, diagnostic expert system could be integrated into a single (portable, if desired) unit with a text screen and a graphics screen and with sufficient disk space.

10.2 Setfacts

During the development of JET-X, there were several junctures where the user was asked to select one or more options from a list. For example, the entry of TEMS and CEMS IV alarms at the beginning of the troubleshooting session is, in essence, a selection from a list of candidates. Similarly, the choice of path in the Symptom method requires choosing from a list of options. Thus, it was decided that these selection processes would be greatly facilitated if the full selection list could be displayed, enabling the user to select as many items as desired. In the absence of such a facility in GEN-X, an external routine called "Setfacts" was developed.

Using the Setfacts program, all possible TEMS and CEMS IV alarms are listed on a "simulated" GEN-X end-user screen (Figure 63). The user moves a highlighted block up and down the multiscreen list, and selects or deselects alarms by toggling the left- or right-arrow keys. Once selected, these alarms are automatically uploaded into JET-X as facts; those selected are set "true," all others are set "false." Ideally, direct communication with the CEMS IV ground station could permit loading of this information without user intervention; however this option was not possible in JET-X, although it is recommended for any production expert system. The solution employed, Setfacts, was simple and involved no hardware interface problems (or costs), which could have been substantial if a direct data link had been utilized in the current prototype expert system.

10.3 Other Support Routines

Several other support routines were also developed to implement features desired for JET-X. The functions and the method of implementation of these routines in JET-X are described in the following paragraphs.




Figure 63. Input of TEMS and CEMS IV Alarms, Using Setfacts.

Record Keeping

Record keeping was an essential ingredient of the JET-X system that required implementation outside of the GEN-X shell environment. The objective of this feature was to maintain a record of all working (as opposed to training) troubleshooting sessions performed with JET-X. A reporting program was designed to append essential information (such as, the engine serial number, date, who executed the session, etc.) onto a master catalog file. Details about a specific engine troubleshooting session are held in an engine file that can be referenced from the master catalog file. The engine file provides the user with a listing of the TEMS and CEMS IV alarms for which the session was initiated, as well as a summary of the maintenance recommendations made by JET-X.

The record-keeping program automatically maintains only the five most recent files for a given engine serial number. The user can also manually delete any engine files or records in the master catalog by means of the reporting program; this feature could be used when an engine is transferred to another location.

Status Display

While most JET-X troubleshooting sessions can be short, those involving multiple alarms or engine performance can be more complicated. During these longer sessions, or at the conclusion of any troubleshooting session, it is desirable to provide the user with feedback on the status of the session. To accomplish this, a "C" program was written to display the current contents of the engine file (as described above) on the JET-X monitor.

Numerical Manipulations

Several procedures were developed to assist the user in performing calculations using data available from the CEMS IV terminal; these usually require the determination of an average value and the comparison of that average to a threshold. While the skill of a typical JET-X user is adequate to enable him or her to perform these kinds of calculations, it was considered a nuisance to have to do so. Consequently, a support program was provided that enables all possible calculations to be incorporated into the JET-X environment. In general, the user is required to enter the appropriate data into the JET-X terminal; calculations and comparisons are performed by the support program, and the result is returned to the JET-X troubleshooting procedure in a form that is interpretable by GEN-X. In a fully integrated system, these functions (access to data, numerical manipulation, and returning information to the knowledge base) would be automatic.

Expanded Explanations

In order to include maintenance recommendations in the aforementioned engine file (see Record Keeping), it was necessary that they be as brief as possible. While a brief message can be sufficient to convey the follow-up actions to be performed, it does not offer the user an explanation of why the recommendation was made, or any specific details that may be relevant to the follow-up action. In order to circumvent this problem, after the engine file is displayed on the screen (Session Summary), the user is given the option to view and/or print an expanded explanation of each maintenance recommendation appearing in the engine file. In addition to the more detailed explanation of a recommendation, the option to view supplementary video Help displays is also provided.

Table 4 lists and summarizes all external "C" programs, excepting JRS.

Table 4. JET-X External "C" Programs.

Program Name	Description
SETFACTS.EXE	Allows the user to input information in the form of facts set either True or False. Input is made by menu selection from a table or list of items.
REPORTS.EXE	Maintains a record of all "Work" sessions and allows the user to view summary information, as well as records of individual troubleshooting sessions. Automatically maintains no more than five records in memory for an engine. Also enables the user to delete records from memory if no longer needed.
CLOSE.EXE	Creates and maintains a summary session record during the session and permits screen display and printing of that record if desired. The record created by "Close" is made permanent if a "Work" session is run, and is deleted if the session was "Training".
AVG.EXE	Enables the user to make form entry of numerical data into JET-X and, in general, will calculate an average (mean) for those numbers, compare the mean to a threshold, and return a True or False value to the knowledge base, depending on the result of the comparison.
PRINTREC.EXE	Enables the display and print-out of expanded explanations for maintenance recommendations made by JET-X. Many of these expanded explanations are accompanied by additional video "help" facilities. The user has the option to view, print, or both, or to skip the expanded description completely.
SESSION.EXE	Provides an interface to the user for form entry of information that goes into the session record; such as, user name and/or engine serial number. Performs transparent functions as well; such as, updating the master catalog of session records.
TEMPCREA.EXE	Manages the display and recording of information associated with the TEMPER simulation.
CURALARM.EXE	Provides the session summary record with information on which alarms were set to True for the current troubleshooting session. Its functions are transparent to the user.

Table 4. JET-X External "C" Programs (Concluded).

Program Name	Description
ALARMS.EXE	Determines if any alarms were set for the current session and passes a global fact to the knowledge base if alarms were set or not. Execution is transparent to the user.
PREVCHK.EXE	Enables the value of alarms, which have been treated in combination with other alarms, to be set "False" so they are not then analyzed individually. Program activity is invisible to the user.
ALARMTR.EXE	Determines the number of alarms remaining to be analyzed in a given session and returns either "none," "one," or "more," which is used as a path name for the calling JET-X module.
WFP3CALC.EXE	Requests the user to input numerical information; then calculates and displays a value for WFP3, a parameter used in determining the health of the main fuel control.
WFK68.EXE	Requests the user to input numerical information; after which it calculates and displays a value for corrected fuel flow (WFK.68), which is used to make a judgement about the level of minimum fuel flow.

11.0 JET-X FIELD EVALUATION SUMMARY

11.1 Test Description

JET-X field testing was conducted at Barksdale AFB, Louisiana from January through June 1988. During this time, the TEMS and CEMS IV alarms generated on A-10's operated by the Air Force Reserve's 917th Tactical Fighter Wing were analyzed using JET-X. While the diagnostic depth of JET-X is centered in the 3 CEMS performance alarms, all 21 CEMS IV performance alarms and the 33 TEMS events were analyzed with JET-X as they occurred. However, only 23 of the possible 54 alarms actually occurred during the test period.

Due to the demonstrator nature of the JET-X project, the field evaluation was carried out on a noninterference basis with the day-to-day maintenance operations at Barksdale. The GE field technical representative at Barksdale conducted the evaluation; Barksdale personnel were invited to utilize the system at their convenience and provide evaluation and commentary. JET-X analysis results were not used as the basis for maintenance actions; however, JET-X results were compared to those produced by the conventional manual (T.O. No. 1A-10A-2-71MS-5) diagnostic method.

11.2 JET-X Diagnostic Performance

Analysis was conducted on a total of 128 CEMS IV alarms and TEMS events during the test period. A majority (80) of these alarms consisted of two of the CEMS performance alarms:

- "Trended Performance Forecasted Low" (Fan Speed Forecasted Below Limit)
- "Step Loss in Trended Performance" (Step Loss: Trended Fan Speed).

In general, results of the JET-X evaluation indicate that its performance met or exceeded the diagnostic T.O. (1A-10A-2-71MS-5) results. Table 5 qualitatively summarizes the results of the field evaluation in terms of its performance relative to the referenced T.O. This summary is based on the findings and observations made by the GE field representative. Because most of the diagnostic procedures included in the JET-X knowledge base were entered with only minor modification from their format in the T.O., it is not surprising that significant differences in performance were not identified. Using the capabilities of GEN-X to fully develop troubleshooting procedures for all alarms would be expected to enhance alarm analysis beyond that capable in a publication.

For the two CEMS alarms that occurred frequently during the field trial, gun-gas contamination of the compressor was the cause of deteriorating performance in most cases. While this is normal for the TF34, analysis performed by JET-X and T.O. resulted in the same conclusion which, unfortunately, did not illustrate the diagnostic depth of JET-X. This would have become apparent only for more unusual causes of performance loss.

An enhanced validation procedure for the "Step Loss in Trended Performance" alarm in JET-X demonstrated more refined methodology for determining the authenticity of this alarm

when it occurs. On at least one occasion, JET-X identified a valid step loss that the T.O. procedures would have rejected.

Table 5. Field Evaluation of JET-X Diagnostics.

TEMS ALARM EVALUATION SUMMARY				
TEMS ALARM	NUMBER OCCURRING	JET-X COMPARISON TO THE T.O.		
		WORSE	EQUAL	BETTER
ITT > 945/1000	1		X	
ITT > 900/1000 (START	1		X	
OIL PRESSURE > SCHEDULE	1		X	
VG OFF SCHEDULE	1			X
FLAMEOUT	1			X
ROLLBACK	2			X
FLUCTUATING CORE SPEED	2		X	
FLUCTUATING FUEL FLOW	3		X	
VIBRATION (AT CORE FREQUENCY)	4			X
VIBRATION (AT FAN FREQUENCY)	6			X
LOW IDLE SPEED	2			X
SHIFT IN MAX T5	5			X

CEMS ALARM EVALUATION SUMMARY				
CEMS ALARM	NUMBER OCCURRING	JET-X COMPARISON TO THE T.O.		
		WORSE	EQUAL	BETTER
TAKEOFF FAN SPEED < LIMIT	3			X
FAN SPEED FORECASTED < LIMIT	43		X	
STEP LOSS: TRENDED FAN SPEED	37			X
STEP CHANGE: TRENDED OIL PRESS.	1		X	
STEP CHANGE: TRENDED VG SCHED.	3			X
RISING CORE VIBRATION TREND	5			X
IRON IN OIL FORECASTED > LIMIT	3		X	
SILVER IN OIL FORECASTED > LIMIT	1		X	
CHROMIUM IN OIL FORECASTED > LIM	1		X	

While only modest enhancement of the existing diagnostic procedures for most alarms was built into JET-X, one of the significant features implemented was the ability to check for related alarm combinations and make recommendations based solely on this criteria. Shortly before the close of the field test, an engine event occurred which underscored the value of this technique. Failure of a CIT (compressor inlet temperature) sensor (control only) produced three TEMS alarms: flameout, VG off schedule, and low idle speed. T.O. analysis might have identified the real cause of these alarms, but because of the multiple alarms, interpretation was difficult. Only after much discussion, did Barksdale experts decide on the cause. When these alarms were input into JET-X, an immediate recommendation was produced, identifying the faulty CIT sensor as the most probable cause.

While the 6-month field test was too short to fully evaluate JET-X diagnostic performance, many indications of its capabilities were noted.

11.3 System Evaluation

Aside from JET-X's diagnostic performance, many observations were made by Barksdale personnel and the GE representative regarding system architecture, methods, user interface, and other human-factors issues. These observations are summarized and explained below.

1. Perhaps the most distracting feature of JET-X is dealing with three separate machines. Hardware consolidation was considered the primary requirement before such a system could be deployed. The engine data/history base should reside on the same computer with the expert-system inference engine, knowledge base, help facilities, and associated software drivers. The end-user screen should provide for display of engine data, help images, and text, and should interface with the expert system.
2. Instead of going through an entire troubleshooting procedure for the CEMS performance alarms before concluding a water wash is required to remove compressor contamination, JET-X should immediately present the user with all relevant data required to determine if contamination is likely. This would save considerable time in using JET-X, since this is the most frequent cause of these alarms. If a question arises as to the cause of performance loss, the full JET-X troubleshooting procedures could then be invoked.
3. Video help examples have a place in engine diagnostics. Providing sample data displays that illustrate symptoms can be a valuable training tool, as well as an aid to the novice diagnostician. Diagrams or photos to aid in the location of engine hardware would be of little value for the TF34; most components are easy to find and identify. Engine cross sections showing primary sources of wear-metals could also serve some users, but would not add significantly to the diagnostic depth of JET-X. Borescope photographs typifying engine damage were not considered to contribute to engine troubleshooting.
4. The history window portion of the end-user screen was not used. Instead of documenting most steps in the troubleshooting process, the inclusion of brief text messages, summarizing significant findings was thought to be a better approach to documenting the troubleshooting session. Printing these "flags" as the session proceeded instead of putting them on the screen was also suggested. *In summary: the current history window format contributes nothing to JET-X.*
5. The order of menu selections presented should be consistent throughout JET-X. Instead of ordering menu selections on the basis of likelihood, using the same order was felt to be more important. Thus, YES and NO menu selections should always appear in the same sequence.
6. The ability to "back up" one step and readdress a question was needed. Frequently, the user wanted to go back to a previous decision point and choose a different path, but this feature is not permitted in GEN-X. Consequently, the

only option was to restart the entire session once the opportunity to do so was presented; this proved to be very annoying, as well as time consuming.

7. While every effort was made by developers to ensure a smooth and distraction-free transfer to external support programs, the Barksdale users found the exchange to be noticeable and mildly bothersome. A more complete integration of GEN-X and its external programs is required.
8. Experienced diagnosticians at Barksdale commented that they would prefer a system that performed all evaluations of engine data and history without involving the user. At the conclusion of such a session, the user would then be asked to accept or modify the recommendation of the expert system. Eliminating the drudgery of data examination was considered one of the features an expert diagnostic system should have. Providing this capability would require that considerable effort be spent in the development of analytical software, independent of expert-system technology.
9. Several JET-X processes were considered slow. Module swapping and calls to external programs (particularly displaying the session summary) were examples. While not a significant detriment, such delays could be reduced by accelerating disk-access time.
10. The module title appearing over the end-user procedure window was found to be too cryptic to be of assistance to the JET-X operator. Module titles should be more descriptive of the module function. Also titles which flash and then disappear were annoying.
11. JET-X question/answer sessions should be streamlined and consolidated. There are several areas where this could be accomplished. Eliminating unnecessary decision tree nodes, and where possible, combining nodes which present a statement or explanation with those that pose a required question would be beneficial. Also, procedure trees (consisting of single-child nodes), which require making a menu selection to index through each step should be combined into one or two large nodes to reduce the number of times the menu is exercised.
12. The format of the "Session Summary" screen should be made easier to read and understand. Avoiding wrapping to the next column, and using color to distinguish between headings and session data are specific suggestions for improvement. Separating recommendations by alarm is also preferred.
13. In some instances, system architecture results in the awkward reappearance of the Session Summary when no additional troubleshooting information has been added to the record; this should be avoided.
14. Allowing the user to view the Session Summary when desired, rather than requiring it be viewed at specific points in the analysis, is more appealing.
15. The "Restart" option, available at the conclusion of a session, allows the operator to begin a session over with the same or a new engine; however, this process is too slow.

16. The Symptom method of troubleshooting, available for the three CEMS performance alarms, did not match the troubleshooting techniques actually used by experienced Barksdale troubleshooters and, consequently, was not preferred. Symptom troubleshooting of performance alarms could be enhanced by providing the user with summary information about the engine at the beginning of the session. This would enable an educated entry into the Symptom session.

11.4 Overall System Assessment

One of the key messages returning from the field evaluation is the need to ensure that an expert-diagnostic system is fast, easy to use, and does not introduce annoying features into the troubleshooting process. Such a system has a place in the field; however, it must be capable of meeting the needs of the expert, as well as the novice.

The experience of the Air Force Reservists who accepted the invitation to operate JET-X ranged from that of the unit's prime diagnostician down to several troops who had been associated with aircraft and engine maintenance for several years, but were, themselves, not mechanics. Those with little experience found JET-X to be an educational tool. After familiarizing themselves with JET-X by means of the Introduction, most commented that with the imbedded explanations and the JRS Help displays, they felt comfortable performing data analysis. They not only reached the same conclusions as troubleshooters with more experience, but felt they understood what they were doing, and why and how they did it. In this respect, JET-X proved itself a far more useful training tool than the T.O.

Those with considerable experience in using TEMS and CEMS IV felt that JET-X was the "right direction" to go for ground-based diagnostics. JET-X enables more complete analysis and offers considerable potential. However, a direct link with CEMS IV is necessary for JET-X to begin to reach this potential. Also, seasoned diagnosticians favored automated data analysis over user-examination of data.

12.0 INTEGRATION OF TEMPER WITH JET-X

One of the original objectives for the design of JET-X involved the development of a gas-path analysis routine which could be integrated with the remainder of the expert system to provide a "smart Kalman filter." The approach envisioned an adaptation of GE's TEMPER (Turbine Engine Module Performance Estimation Routine) to the TF34 engine and a subsequent integration with the GEN-X routines to add expert guidance for the evaluation of the TEMPER results. For reasons that will be given in this section, this plan has not been carried through to fruition. Nevertheless, there is a role for a smart Kalman filter for military engines, provided the engines have sufficient gas path instrumentation and provided that the proper data base integration is available.

12.1 Description of Current TEMPER Algorithm

A fairly extensive discussion of gas-path analysis techniques, in general, and TEMPER, in particular, is provided in Appendix C of this report. The basic problem that is addressed by gas-path analysis techniques is a determination of engine performance (thrust, specific fuel consumption, etc.) relative to normal levels and an assignment of performance deficiencies to the faulted modules (fan, compressor, turbine, etc.). Original techniques that were used made no attempt to account for measurement error in their analysis. This led to sizable errors in their results, particularly with regard to the modular diagnosis; and generally, the results were largely unusable.

Once we recognized that measurement error was the cause of the poor reliability of the gas-path analysis techniques, we made considerable progress using two techniques which responded to the different parts of the measurement error. The first observation was that the measurements generally included large bias errors which, in turn, produced unrealistic values for component efficiencies. Although these bias errors tended to be somewhat different from one engine to another, they were relatively stable for a particular engine. It was noted that there was not a pressing requirement to be sure of the absolute performance level of a module, provided we could determine the amount of deterioration of the module from its initial installation on wing. To correct this problem, the initial data acquired on wing were used to establish a baseline for the engine. The subsequent level of deterioration of the engine and its modules was evaluated based on the deviation of the measurements from these initial levels.

The second technique addresses both the noise in the measurements and the possibility of drift in the measurement biases. This method utilizes weighted "least-squares" to simultaneously estimate component health parameters and measurement errors. This technique relies on the existence of redundant measures of both engine health and component health, in order to distinguish true changes in the engine performance from measurement errors. Generally, there is very little redundancy in a jet engine's measurement set; hence, this technique can easily be eliminated from possibility by the deletion of one or two key measurements. One problem introduced by the low level of redundancy in jet engine measurements is that the method is less sensitive to measurement errors and true performance changes. True changes

are only partially recognized, with the remainder of the changes being incorrectly attributed to other factors.

An important augmentation to the weighted, least-squares algorithm (in the case of repetitively acquired on-wing data), is the ability to properly recognize and account for a sudden shift in performance. This capability strengthens the technique, because it is reasonable to expect that a step-change in the measurement(s) is attributable to a single fault. If only one measurement jumps, it is reasonable to suppose that that specific measurement has had a bias shift. However, if several measurements shift at the same time, it is likely that a real engine problem was responsible.

TEMPER utilizes each of the three above-described techniques, as follows:

- The use of baselines to eliminate large bias errors that produce unrealistic levels for efficiencies and pumping capacities
- Use of weighted least-squares to reduce the errors introduced by measurement noise and sensor drift
- The use of a "fault logic" to recognize sudden changes to the performance of the engine and to identify the most likely cause for the observed shift (measurement error, or hardware problem).

These advancements have made TEMPER significantly more effective than its predecessors in analyzing on-wing-acquired data.

12.2 Need for Improvement in TEMPER Gas-Path Analysis

Despite the improvements, there are still a number of problems associated with the gas-path analysis of jet engines. There are frequently engine faults that are nearly indistinguishable, given the measurements available on jet engines. For example, an increase in the parasitic cooling flow produces the same changes to the measurements as a combination of decreased high pressure turbine efficiency and increased turbine flow function (in the proper proportion). All of these changes are typical of engine deterioration with time; consequently, no algorithm can separate these faults without additional sensors. There are many other equivalency sets of this kind which can be confused, based on the measurements.

There is often additional qualitative data that could be applied to the analysis. For example, when some maintenance is performed on wing (such as, when a sensor is replaced) it is reasonable to look for a shift in the behavior of the associated component. An obvious example for the TF34 is to expect improved fan and compressor performance as a result of a water wash. Also, there are rules of thumb for interpreting the TEMPER output. One example involves interpreting indicated measurement errors in fuel flow and control temperature sensors; when the indicated measurement errors are of opposite sign, the cause is likely to be a measurement error in one of the two measurements; however, when the errors have the same sign, there is likely to be a hardware cause for the deviation.

Fault analysis is often limited in scope, because of the rapid growth of processor cost associated with exhaustively searching for multiple-fault combinations. Some of these fault combinations can be lumped together, based on qualitative clues that could be interpreted by an expert system. Alternatively, the fault searches may select an inappropriate fault because a sensor is missing for a reading, and in the absence of the sensor, there is no logic to eliminate consideration of the fault.

For these and other reasons, the marriage of expert systems to the TEMPER gas-path-analysis algorithm seems to be a logical next step to improve the ability to fault-isolate performance problems. The INTERFACE IIL Advanced Diagnostic Software effort seemed to be a good opportunity to pursue this integration of technologies. The TF34 engine was selected for study, in part because of its gas-path-sensor complement which exceeds many of the newer military engines, such as the F101 or the F110.

12.3 Application of TEMPER to the TF34 Engine

Like most military engines, the TF34 does not have a repeatable, steady-state flight condition that can be used for TEMPER analysis. The only high power condition that is available on all flights is takeoff. Therefore, takeoff data were used for the TEMPER analysis. This can be expected to increase the scatter in the data (particularly when the aircraft is used for training purposes), because some of the takeoffs are made with the engine cool; while others are made with a hot engine.

Figure 64 demonstrates a sample data record for an A-10A flight. The data record contains data for both engines on the aircraft and includes environmental data that are obtained from the aircraft systems. Table 6 lists the parameters that are used for the TF34 TEMPER analysis. The measurement inputs are subdivided into "setting parameters" and "independent" measurements. The setting parameters are those measurements that are used to define the actual flight condition and engine power level; thus, they are not used to critique the performance of the engine or its modules. The independent measurements are used as the basis for determining overall engine health and for modular diagnosis.

Table 6 makes reference to the fact that the total air temperature is computed from both the fan speed and the compressor inlet temperature; of these, the latter measurement is clearly more significant in the determination of total air temperature. Normally, this calculation is performed utilizing the data from each engine, and the average of the two calculated values is used. If the data is not available for one of the engines, then the value from the other engine will be used. The impact of this is that the compressor inlet temperature is independent of the setting parameters only to the extent that total air temperature is also being computed from the other engine.

Thus, there are somewhere between five and six independent measurements available for the TF34 engine gas-path diagnosis. Most of these measurements provide data that can be used to evaluate the performance of the core engine. The only measurement available to assess the fan is the compressor inlet temperature, which is suspect for the previously stated reasons.

A/C AF 780382	FLIGHT 2.00	MONDAY 06 JUNE 1988
EPU SERIAL NUMBER 111	DDU SERIAL NUMBER 513	SOFTWARE VERSION 56.1
MRPDATA	(25-Nov-88)	FL0810561-001
ENGINE SERIAL NUMBER	LEFT 205363	RIGHT 205822
ITT TRIM	816	812
ENGINE TIME - START, TIME OF DAY	11:39:54	11:42:02
- RUN THIS FLIGHT	01:52:38	01:50:12
- EOT	2483.42	2839.85
- ENG FLIGHT HOURS	2067.25	1654.24
AC TIME - FLIGHT DURATION	01:19:52	
- TOTAL FLIGHT HOURS	2:89.27	
EPU TIME - TOTAL ON TIME	47.60	

EVENT:	TAKE OFF	CRUISE	R NG flameout 2	R-VG open 2
DURATION-QUANT			00:00:32-1	00:02:46-5
ALTITUDE FEET	-309	9994	380	-249
AIR SPEED KCAS	130	256	146	
MACH #	0.196	0.453	0.223	0.000
OAT DEG C	25.5	4.9	20.9	29.8
NZ C'S	1.0	2.0	1.1	1.0
AOA UNITS	14.6	16.1	20.9	10.7
TIME OF DAY HMS	11:56:39	12:23:27	13:15:53	13:17:27

TRIM DATA:	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT
B1 NF TRIM			-1.1	-2.3				
B2 NF RED LINE			2.3	1.8				
B3 NG TRIM IDLE								-5.3
B4 VG SCHEDULE			0.5	0.5				-3.6
B5 ITT DEG C	814	809	816	812	805	844	539	501
B6 T2C DEG C	60.8	60.3	52.9	51.3	60.5	32.8	34.3	33.1
B7 PLA DEG	84.6	84.2	99.3	99.5	82.9	79.5	20.2	19.8
B8 PLA STABLE	00:00:14	00:00:14	00:00:52	00:00:52	00:00:00	00:00:00	00:00:04	00:00:24
C1 NG MECH I	93.9	93.3	94.9	93.9	94.6	44.8	65.6	59.1
C2 NF MECH I	80.9	80.5	85.1	84.1	81.5	15.3	24.7	22.4
C3 VG DEG	28.2	28.5	26.1	26.7	27.7	61.9	59.3	57.9
C4 WF ACT FPM	2768	2675	2382	2281	2649	224	369	326
C5 PS3 PSIA	231.5	230.5	198.9	196.9	229.2	24.9	44.8	41.5
C6 PT5 PSIA	52.9	52.0	45.5	44.2	52.5	15.3	17.7	16.9
C7 POIL PSIG	81.7	80.0	79.3	78.5	76.9	37.0	56.9	48.4

COMPUTED DATA:	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT
D1 NFG T5 CORR	-0.4	-0.4	0.1	-0.8				
D2 NGC NFG	0.4	0.1	1.2	0.8				
D3 PT5 PS3	-1.0	-2.3	-2.1	-4.2				
D4 PS3C T5C	-3.0	-2.5	-0.9	-1.6				
D5 NGC	87.2	86.8	89.2	88.5	87.9	43.5	63.5	57.3
D6 T5 AMP	114	113	113	114	115	113	114	113
TT2 DEG C	27.8	27.6	17.2	16.4				
PT0 PSIA	15.27	15.27	11.71	11.71	15.01	15.01	14.84	14.84
PAMB PSIA	14.87	14.87	10.11	10.11	14.50	14.50	14.83	14.83

VIBRATION:	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT
E3 AGB G MILS	0.2	0.1	0.4	0.1	0.1	0.3	1.2	0.4
E4 FR HOR G MILS	1.2	0.4	1.5	0.3	0.4	0.3	0.7	0.6
E5 EXE HOR G MILS	0.4	0.1	0.5	0.1	0.2	0.2	0.1	0.3
E6 AGB F MILS	1.2	1.0	1.3	0.8	1.2			
E7 FR HOR F MILS	1.7	1.7	2.2	1.6	1.1			
E8 EXE HOR F MILS	1.0	1.6	1.5	1.9	1.4			
AGB BB IPS	0.4	0.3	0.5	0.2	0.5	0.3	0.7	0.3
FR HOR BB IPS	1.1	0.6	1.3	0.6	1.2	0.1	0.2	0.2
EXE HOR BB IPS	0.4	0.5	0.4	0.6	0.4	0.1	0.2	0.0

STATUS:	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT
G3 WF OVERRIDE	NORMAL	NORMAL	NORMAL	NORMAL	NORMAL	NORMAL	NORMAL	NORMAL
G5 COOLDOWN								
G6 START MODE								
G7 WOW	GROUND MODE		AIRBORNE		AIRBORNE		GROUND MODE	
G9 SLATS	RETRACT		RETRACT		RETRACT		RETRACT	
GUN	NO GUN LAST 16 S		NO GUN LAST 16 S		NO GUN LAST 16 S		NO GUN LAST 16 S	

Figure 64. Sample of TF34 TEMS Data.

Table 6. TF34 Gas-Path-Analysis Measurements.

Setting Parameters	
Pressure Altitude, ZALT	Feet
Mach Number, ZXM	Dimensionless
Total Air Temperature, ZT1A*	Degrees, Celsius
Fan Speed, PCN2	Percent of Reference Speed
Variable Guide Vane Position, ZVSV	Degrees
Independent Measurements	
Inter-Turbine Temperature, T45	Degrees, Celsius
Compressor Inlet Temperature, T21T*	Degrees, Celsius
Core Speed, PCN22	Percent of Reference Speed
Fuel Flow, WF36	Pounds per Hour
Compressor Discharge Static Pressure, PS3	Pounds per Square Inch, Absolute
Inter-Turbine Pressure, P45	Pounds per Square Inch, Absolute
* Note that for the TF34 engine, the total air temperature is computed from the compressor inlet temperature and the fan speed	

Further, there is no temperature measurement available to separate compressor performance from high pressure turbine performance. Hence, the most that can be hoped for from a TEMPER analysis is some appraisal of the core engine performance level and an evaluation of the compressor and turbine-pumping capacities. Also, the overall low pressure system performance is weakly determined based upon the observed power level relative to the ideal gas horsepower supplied by the core engine.

Once the basic TEMPER code for the TF34 had been generated, measurement baselines and standard deviations were determined. Data from 6 A-10A aircraft, encompassing a total of 11 different engines, was used to generate the statistical parameters. The standard deviations are assumed to be representative of the expected measurement errors and, thus, are used as input to the TEMPER calculation.

The weighted, least-squares analysis also requires input of standard deviations for the hardware variability (for example, HP turbine flow function or HP compressor efficiency). These standard deviations are not selected to necessarily represent the true expectation of the hardware behavior but are chosen to produce a desired output characteristic, given known inputs. At the simplest level, the hardware standard deviations could be chosen to make all faults equally detectable. This basic approach is modified to some degree to reflect experience of where problems typically occur. For example, the LP turbine flow function is generally not changed as a result of normal deterioration. Thus, a smaller standard deviation is assigned to this hardware parameter to reflect this knowledge. Several of the hardware standard deviations were set to extremely small values because there was no way to estimate the corresponding parameters given the limited measurement set.

The resulting standard deviations for the measurement errors and hardware variables are listed in Table 7, which indicates that a correlation is assumed between the HP compressor flow and efficiency. This is a normal practice for TEMPER codes. The physical mechanisms (such as tip rubs, erosion, corrosion, leading edge blunting, etc.) which cause a reduction to the compressor efficiency also tend to reduce the compressor flow. The inclusion of the correlation tends to force the algorithm to place fault in the HP compressor only if the evidence supports a flow change and an efficiency change. This mechanism normally helps in the process of distinguishing between a compressor and a turbine problem. In the case of the TF34, where there is no HP compressor discharge temperature, the correlation offers the only hope for separation of the core components. Normally, similar correlations would be used for the fan tip and hub; however, these variables are not all observable for the TF34, and so no correlations have been used.

Figure 65 provides a "response" analysis for the TF34 TEMPER program. This analysis demonstrates the response of TEMPER to 1 percent changes to various component efficiencies and flow capacities, and to 1 percent measurement errors. In the first two pages of Figure 65, the rows represent the results that TEMPER will produce in response to a 1 percent change in the indicated input variable. For example, the second row indicates that the change to fan efficiency will be 0.490 percent, if data consistent with a 1 percent fan efficiency change

Table 7. Standard Deviations for TF34 TEMPER.

Measurement Errors	Percent
Inter-Turbine Temperature, T45	0.493
Compressor Inlet Temperature, T21T	0.158
Core Speed, PCN22	0.153
Fuel Flow, WF36	1.188
Compressor Discharge Static Pressure, PS3	0.559
Inter-Turbine Pressure, P45	0.539
Hardware Parameters	Percent
Fan Flow, W2R	0.0001
Fan Tip Efficiency, SE21DT	1.000
Fan Hub Flow, SW22R	0.0001
Fan Hub Efficiency, SE21D	0.0001
HP Compressor Flow, SPN22R*	0.200
HP Compressor Efficiency, SE25D*	0.600
HP Turbine Flow Function, SW41R	0.200
HP Turbine Efficiency, SE41D	0.550
LP Turbine Flow Function, SW45R	0.100
LP Turbine Efficiency, SE5D	0.500
* An 80 Percent Correlation Coefficient is Assumed Between the HP Compressor Flow and Efficiency	

TF34 ON-WING TEMPER RESPONSE ANALYSIS FOR INTERFACE

SX VECTOR

W2R SE21DT SW22R SE21D SPN22R SE25D SW41R SE41D SW45R SE5D
0.0001 1.0000 0.0001 0.0001 0.2000 0.6000 0.2000 0.5500 0.1000 0.5000

SZ VECTOR

T21T PCN22 WF36 PS3 P45 T45
0.1580 0.1530 1.1880 0.5590 0.5390 0.4930

COVARIANCE TERMS

HPC CORRELATION COEFFICIENT BETWEEN SPN22R AND SE25D = 0.800

TEMPER RESPONSE

	W2R	SE21DT	SW22R	SE21D	SPN22R	SE25D	SW41R	SE41D	SW45R	SE5D	T21T	PCN22	WF36	PS3	P45
W2R	0.000	-1.476	-0.000	-0.000	0.019	0.035	-0.037	0.013	0.002	-0.680	0.062	0.015	0.357	0.276	0.291
SE21DT	-0.000	0.490	0.000	0.000	-0.008	-0.013	0.008	0.010	-0.001	0.170	-0.047	-0.008	-0.107	-0.063	-0.060
SW22R	-0.000	0.095	0.000	0.000	-0.129	-0.302	-0.013	0.035	-0.005	0.003	0.005	-0.079	0.001	0.136	0.140
SE21D	-0.000	0.067	0.000	0.000	-0.022	-0.039	-0.002	0.040	-0.001	0.048	0.005	-0.019	-0.000	0.018	0.026
SPN22R	0.000	-0.342	-0.000	-0.000	0.462	1.078	0.050	-0.152	0.017	-0.229	-0.018	0.287	0.005	-0.500	-0.522
SE25D	-0.000	0.055	0.000	0.000	0.073	0.286	-0.012	0.340	0.002	0.035	0.002	-0.010	-0.123	0.050	0.112
SW41R	-0.000	0.210	-0.000	-0.000	0.021	0.011	0.082	-0.216	-0.013	0.141	0.015	0.033	0.317	-0.621	0.244
SE41D	0.000	0.034	0.000	0.000	0.088	0.357	-0.029	0.460	0.004	0.020	-0.001	-0.018	-0.200	0.163	0.099
SW45R	0.000	-0.065	-0.000	-0.000	0.081	0.215	-0.051	0.123	0.018	-0.057	-0.012	0.036	0.308	0.359	-0.525
SE5D	-0.000	0.681	0.000	0.000	-0.023	-0.037	0.023	0.024	-0.002	0.511	0.067	-0.021	-0.313	-0.168	-0.171
T21T	0.000	-1.895	0.000	0.000	-0.023	-0.048	0.024	-0.008	-0.005	0.667	0.823	-0.017	-0.386	-0.180	-0.178
PCN22	0.000	-0.341	-0.000	-0.000	0.450	1.025	0.056	-0.227	0.016	-0.228	-0.018	0.291	0.043	-0.541	-0.529
WF36	0.542														

Figure 65. TF34 TEMPER Response Analysis.

WF36	0.000	-0.076	0.000	-0.000	0.008	-0.031	0.009	-0.043	0.002	-0.055	-0.007	0.001	0.774	-0.064	-0.068
	-0.125														
PS3	0.000	-0.200	0.000	0.000	-0.049	-0.096	-0.080	0.158	0.011	-0.135	-0.014	-0.041	-0.288	0.612	-0.237
	-0.011														
P45	0.000	-0.205	0.000	0.000	-0.035	-0.033	0.034	0.103	-0.018	-0.147	-0.015	-0.043	-0.330	-0.254	0.743
	-0.051														
T45	0.000	-0.135	-0.000	0.002	-0.104	0.005	-0.324	-0.002	-0.110	-0.016	0.052	-0.724	-0.014	-0.061	
	0.447														
ZT1	0.000	-1.435	-0.000	-0.000	0.186	0.329	0.078	-0.437	0.002	0.777	0.757	0.166	-0.743	-0.620	-0.708
	0.385														
ZXM	0.000	-0.001	0.000	0.000	-0.006	-0.010	-0.002	0.015	-0.000	-0.010	-0.004	-0.005	0.005	0.016	0.020
	-0.014														
PCN2R	0.000	-1.586	-0.000	-0.000	0.009	0.007	-0.033	0.001	-0.001	-0.713	0.075	0.010	0.334	0.248	0.384
	0.109														
WMB02D	0.000	-0.086	-0.000	-0.000	0.072	0.069	0.035	-0.346	0.003	-0.088	-0.018	0.092	0.440	-0.266	-0.253
	0.251														
DWC250	-0.000	0.153	-0.000	-0.000	-0.036	-0.209	0.087	-0.468	-0.013	0.104	0.013	0.040	0.395	-0.324	0.141
	0.186														
DP4P31	-0.000	0.230	-0.000	0.000	0.071	0.213	0.067	0.044	-0.011	0.154	0.015	0.023	0.206	-0.534	0.303
	-0.120														
SAB	-0.000	0.071	0.000	0.000	0.001	0.004	0.001	0.008	0.000	0.053	0.007	-0.001	-0.022	-0.005	-0.038
	-0.013														
SA28	-0.000	0.453	-0.000	0.000	0.001	0.002	0.008	-0.012	0.000	0.112	-0.066	0.001	-0.049	-0.058	-0.068
	-0.004														
HPC	0.000	-0.038	-0.000	-0.000	0.194	0.564	0.002	0.287	0.006	-0.027	-0.003	0.067	-0.118	-0.085	-0.031
	-0.063														
RESPONSE IN EIGENVECTOR DIRECTION FOR HPC EIGENVALUE # 1 = 0.595															
NOTE: EIGENVECTOR COMPONENTS ARE: SPN22R 0.267 SE25D 0.984 EQUIVALENT SIGMA = 0.622															
HPC	0.000	-0.344	-0.000	-0.000	0.426	0.962	0.051	-0.238	0.016	-0.230	-0.018	0.279	0.038	-0.495	-0.533
	0.538														
RESPONSE IN EIGENVECTOR DIRECTION FOR HPC EIGENVALUE # 2 = 0.154															
NOTE: EIGENVECTOR COMPONENTS ARE: SPN22R 0.964 SE25D -0.267 EQUIVALENT SIGMA = 0.116															

Figure 65. TF34 TEMPER Response Analysis (Continued).

TEMPER RESPONSE AND FAULT THRESHOLD SUMMARY

PARAMETER	RESPONSE	RESIDUAL	PROBABILITY	THRESHOLD
W2R	0.0000	4.905	55.608	1.597
SE21DT	0.4900	0.490	99.796	5.053
SW22R	0.0000	0.917	98.856	3.694
SE21D	0.0000	0.066	99.900	13.734
SPN22R	0.4624	12.152	5.865	1.015
SE25D	0.2863	0.861	99.034	3.813
SW41R	0.0823	2.057	91.440	2.467
SE41D	0.4603	1.522	95.804	2.868
SW45R	0.0183	1.832	93.452	2.614
SE5D	0.5106	2.043	91.575	2.475
T21T	0.8230	32.968	0.001	0.616
PCN22	0.2905	12.410	5.343	1.004
WF36	0.7741	0.548	99.720	4.777
PS3	0.6120	1.958	92.348	2.528
P45	0.7435	2.559	86.179	2.211
T45	0.4467	1.838	93.400	2.609
ZT1	-	34.281	0.001	0.604
ZXM	-	0.007	99.900	41.081
PCN2R	-	5.644	46.424	1.489
WMBQ2D	-	1.898	92.881	2.567
OWC25Q	-	2.787	83.507	2.119
DP4P31	-	1.767	93.984	2.661
SA8	-	0.025	99.900	22.428
SA28	-	0.458	99.831	5.227
HPC #1	0.5949	1.539	95.688	2.852
HPC #2	0.1539	11.475	7.477	1.044

Figure 6S. TF34 TEMPER Response Analysis (Concluded).

are input to the program. This row also shows that there will be an indication of a change to the LP turbine efficiency of 0.170 percent. The response analysis is computed using the input derivative matrix and the statistical inputs.

All of the hardware parameters and measurement errors are evaluated in this analysis, and in addition, some parameters that are not part of the weighted, least-squares analysis are also considered. The parameters that are denoted as "HPC" on the second page of the figure represent changes to the compressor behavior that correspond to a combined change to flow capacity and efficiency. Both eigenvector directions are computed in the analysis, but only the first (where flow and efficiency changes are compatible) is considered.

The third page of Figure 65 summarizes the response characteristics and indicates the necessary magnitude of change in the given parameter to trigger the fault-search logic. The column designated "threshold" indicates the required size change (in percent) for the given parameter to activate the fault logic. The residual and probability columns are for a 1 percent change. Note that the fault logic is triggered when the basic solution probability is below 5 percent.

The response analysis is the device that is used to select the hardware standard deviations. The response of any particular variable (either hardware parameter or measurement error) can be improved by increasing its standard deviation relative to the other variables. To improve the response of all hardware parameters without jeopardizing the ability to detect measurement errors, it is necessary to add redundant measurements to the system. This, of course, is not practical for the TF34. Thus, the response to those hardware factors which are observable is in the neighborhood of 50 percent at best.

12.4 Comments on Idealized Integration of TEMPER with JET-X

When designing the JET-X system, it became clear that TEMPER was not logically a part of the JET-X expert system at all. To clarify this remark, it should be pointed out that TEMPER is a pro-active tool which identifies faults that have occurred and sounds the alarm for further analysis. TEMPER also, by virtue of its trending of the data, provides assistance in locating the source of the problem. In this latter function, it can be made more effective through the infusion of qualitative information. However, TEMPER should be running continuously on all engines, whether they have a problem or not.

The ideal role for TEMPER would be as a part of the CEMS data management system. In this design, TEMPER would be run routinely on all data input to CEMS. This processing might include expert-system features that would not involve interaction with the user, but only to simplify the representation of logic that is difficult to code in a conventional language, such as FORTRAN. As part of its routine processing, TEMPER would evaluate its outputs against "alert levels." When an alert level was exceeded, TEMPER would produce an alarm comparable to those now generated by TEMS or CEMS.

The design of TEMPER should accommodate some interaction with the other data that is available to CEMS. In particular, maintenance input could be used to identify that a shift in

some module or measurement error is likely. For example, the noting of a water wash would force TEMPER to reevaluate the performance of the fan and compressor. Other events, such as a sudden shift in the vibration level or an indication of change in one of the control schedules, might induce a TEMPER fault search to try to find supporting evidence in the gas-path data.

It seems quite desirable to manage the TEMPER fault search by an expert system. The principal reason for this is that it permits more complex strategies than can be comfortably coded in FORTRAN. The specific faults to be searched can be readily tied to clues in the data, eliminating the need for at least some of the exhaustive searches. Knowledge of missing measurements can also be used to eliminate certain faults from consideration because they cannot be fairly evaluated from the available sensors.

Commercial experience with TEMPER suggests that wild points can cause problems by producing erroneous faults in the solution that for one reason or another do not get corrected on the succeeding reading. Although strategies to cope with this problem have been developed, they are awkward to code in conventional languages, and could probably be more general if expressed using an expert system.

There is also a role for user interaction with TEMPER. When an alarm has been produced (whether by TEMPER or by another source), TEMPER can be utilized to help understand the problem. This could involve reprocessing some of the prior data through TEMPER, using different assumptions to see if an alternate fault selection provides a better fit to other data that have been observed. It might involve interactive exploration of a group of faults that cannot be distinguished based on the gas-path data alone; for example, other data might be used to attempt to distinguish a turbine problem from a cooling flow increase. TEMPER can be used to examine specific explanations for observed behaviors; for example, can a loss of performance margin be attributed to an ITT measurement error or to an overboard leak? This latter approach has been addressed in the current work.

To summarize, there are two principal reasons why TEMPER does not play a larger role in the JET-X system. These are:

1. There are too few available measurements on the TF34 engine to accomplish an adequate TEMPER modular diagnosis
2. TEMPER should be a part of the CEMS system to adequately fulfill its proper role in a diagnostic system (the scope of the current contract precluded this type of arrangement).

One of the goals for future development is to revisit this work within a more favorable setting. Current plans involve attempting a TEMPER/JET-X type of integration for the ATFE (Advanced Tactical Fighter Engine).

12.5 JET-X TEMPER Interface

Since TEMPER could not be included in the CEMS IV data base, an attempt was made within JET-X to model a possible end-user interface with TEMPER. In this implementation, TEMPER was used to analyze a single point of data in an effort to determine what might be the cause of an observed shift in engine performance. The approach consisted of running TEMPER in a fault-search mode to identify the solution probability for each of a set of faulted solutions. For each solution, the magnitude of the indicated change and the solution probability are displayed on a summary screen for perusal by the user (Figure 66). Because this type of display is most likely to be usable by an experienced user, a filtered display is available upon request. The filtered display shows only those faults that are adjudged to be most likely by JET-X. The three criteria used to eliminate prospective faults from the filtered display are:

- Any components showing performance changes in "impossible" directions (such as improving efficiencies) are eliminated
- Any solutions that have very low probabilities are deleted
- Solutions that include excessively large component deviations are eliminated, since they should be obvious without the TEMPER analysis and, therefore, are deemed to be incorrect.

An example of the filtered TEMPER output is provided in Figure 67. The user may select to view only the filtered TEMPER results, or if the full TEMPER output is selected, the filtered output is also displayed.

Once the filtered TEMPER output is displayed, the user is asked to select one or more faults which are considered to be most likely, based on other information at his or her disposal. Once these selections have been made (on the filtered TEMPER display), follow-up maintenance recommendations are made based solely on this input. Figure 68 presents an example of two maintenance recommendations based upon two selected TEMPER symptoms.

As was indicated above, there has been no effort to integrate TEMPER with the other knowledge-base functions. For the present application, a stand-alone mode was appropriate, since there was no link between CEMS IV and TEMPER. This type of integration should be addressed in another program, such as the previously mentioned ATFE.




Figure 66. JET-X Screen for "Raw" TEMPER Output.





Figure 67. JET-X Screen for "Filtered" TEMPER Output.



1. The following information was received from the FBI on 11/11/81.

13.0 REFERENCES

1. Winograd, T., and Flores, F., "Understanding Computers and Cognition," Addison-Wesley, July 1987.
2. "CEMS IV Preliminary Maintenance Procedures, Revision 2," incorporated into U.S. Air Force T.O. (Technical Order) No. 1A-10A-2-71MS-5, 15 May 1984.

APPENDIX A

APPENDIX A

AI (ARTIFICIAL INTELLIGENCE) TOOLS FOR DIAGNOSTICS AND GEN-X OVERVIEW

Diagnosis of machines and systems is aimed at interpreting the nature of malfunction or loss of function and finding the underlying root causes. The process can be highly complex and time consuming as the machines or systems to be diagnosed get more complex. While jet engine health monitoring systems and engine data bases provide critical data for diagnostics, it is often the expertise of a human diagnostician and his or her ability to interpret and reason with the data that leads to diagnostic success.

Computer-assisted diagnostics tools range from automated health monitoring systems to the use of AI (artificial intelligence) techniques, such as expert systems. In a general sense, the diagnostic procedures built into an expert system can also be represented using conventional programming techniques (utilizing higher level languages, such as FORTRAN, or even using machine language). However, the expert-system route to diagnostics offers some unique features. For example, probability and cost data can be entered to establish the most effective, least expensive testing sequence for diagnosing a problem. Once a problem is diagnosed, the technician can then be guided through a fixed, repair procedure.

Rule-based expert systems consist of a body of knowledge (the knowledge base) and a mechanism for interpreting this knowledge (the inference engine). The knowledge base is further subdivided into a rule base, which contains information about the general problem domain, and a fact base, which contains information about the specific problem being solved. The inference engine is an interpreter of these facts and rules. The inference engine monitors the fact base and controls the order in which the rules are tested.

In many rule-based expert systems, the inference engine operates in either the forward chaining mode (event- or data-driven) or the backward chaining mode (goal-driven). The goal-driven mode is useful in deciding what information to seek next during a problem-solving session, while the event-driven mode is useful in responding to user/sensor inputs such as requests for help. Expert systems also differ in the choice of logic types and the procedure and format for representing the logic in these different forms. Two basic logic types are required to represent procedural logic (such as decision trees) and nonprocedural logic (such as, If/Then rules or And/Or trees).

GEN-X, the GENeric eXpert system, was used in the development of JET-X. GEN-X is a software package (a rule-based expert system shell), developed at the GE-CRDC (Corporate Research and Development Center). GEN-X is designed to facilitate the process of building expert systems (that is, applications). The usefulness and versatility of GEN-X software for building expert systems to achieve diagnostic goals can be attributed primarily to two factors: the extremely user-friendly graphics interface for building the knowledge base, and the packaging of this capability into a low cost hardware environment.

The knowledge base of the JET-X, rule-based expert system consists of various modules. The inference engine (of GEN-X) allows the JET-X system application developer to select either

the forward- or backward-chaining mode of operation appropriate for each logic module. GEN-X also offers three methods of representing the logic: Decision tree (Figure A-1), And/Or tree (Figure A-2), and If/Then table (Figure 11, presented in Section 7.3). The application developer has a free choice in the selection of any of these three logic types in creating a logic module. The modules of the knowledge base can be directly linked in GEN-X to realize desired inferencing order. The graphical format of GEN-X for creating the three logic types and inserting the rules and procedures (the expert knowledge) offers an efficient mechanism for creating, reviewing, updating, and extending the knowledge base.

The specific method of representing the logic in GEN-X is determined by its function. If/Then tables and And/Or trees are employed for formulating the fault hypotheses and for directing the search. Decision trees are used to validate the hypotheses. These logic implementation details are shielded from the end user; instead, only the relevant information is presented to the end user in a multiple-window format.

GEN-X can be used in one of two major modes: the development environment or the end-user environment. In the development environment, the actual expert-system application is constructed using any combination of If/Then tables, And/Or trees, or Decision trees, as described earlier. The end-user environment is for the person actually seeking the assistance of the expert system; in this report, the terms "user" and "operator" are used to describe this person.

Availability of various types of Help facilities plays a crucial role in any practical expert system. Some of the Help facility items are: graphs and tables capturing historical or uncommon events, photographs and animation sequences for explaining or interpreting more difficult maintenance procedures, electrical-wiring diagrams, and the location and illustration of uncommon components. These should be made available whenever the user might need them, but usually only if requested. The expert-system developer then chooses an appropriate medium (such as, video disk players or digitized images) suitable to application. While not directly available in GEN-X, access to such external devices as video disk players, digitized images, etc., is available from the knowledge base through a back-plane command facility in GEN-X. Additional details of the software used for displaying the graphical images, called JRS (JET-X Retrieval System), are provided in Section 10.0.

In applying GEN-X to jet engine troubleshooting, the broad scope inherent in the analysis of 54 TEMS and CEMS IV alarms demands that the knowledge-base architecture address such issues as: the efficiency of search (while retaining interest), usefulness to experts and novices alike, and communicating to the user a sense that he or she has the overall control of the troubleshooting process. The modular structure of GEN-X is inherently suitable to meeting this challenge, as discussed in Section 7.1 on the JET-X architecture.

This application operates in a microcomputer (IBM PC or compatibles) environment, supplemented by any peripheral device, such as JRS, desired for graphical/pictorial help facility. The integrated, diagnostics expert systems of the future may be linked to external computers (for access to a database), sensors (for on-line data), or other devices. The use of GEN-X as a shell provides JET-X with these capabilities.




Figure A-1. GEN-X Decision Tree Development Screen.




Figure A-2. GEN-X and/or Tree Development Screen.

APPENDIX B

APPENDIX B

KNOWLEDGE SOURCE(S) AND EXTRACTION

The development of the JET-X knowledge base was not handled in the conventional manner of using a knowledge engineer to capture an expert's experience through a series of detailed interviews. Two TF34 engine diagnostic experts were part of the JET-X development team: one was a field technical representative, experienced in troubleshooting day-to-day engine problems; the other was a factory engineer, with a background in the analysis of engine performance data.

The task of building the knowledge base was left largely in the hands of these experts. The GEN-X rule-based expert system shell, which utilizes a graphical approach for logic representation, is easily learned and applied by one who is neither an expert in computer science nor artificial intelligence. The modular nature of the knowledge base and the ease of linking the modules in any desired manner made the task of updating and extending the logic rapid and straightforward. Initial training in the use of the GEN-X shell was provided to both experts. Material presented in a 3-day training session provided adequate GEN-X fluency to begin the construction of the knowledge base. All troubleshooting "expertise" was entered into GEN-X by the factory engineer, using an IBM PC-XT configured with 640K RAM, a hard disk, and color monitor.

The existing engine monitoring systems, TEMS and CEMS IV, identify a combined total of 54 engine alarms, each of which indicates a malfunction or discrepancy in engine operation or health. A separate, unique, troubleshooting procedure was created in JET-X for each alarm. Initially, these procedures were based on past work performed by the two experts, which was documented in logic tree format. However, the necessity of limiting each tree to a single document page (with readable text) imposed a restraint on the procedures developed for the troubleshooting manuals. The removal of this physical constraint in the GEN-X system, plus additional diagnostic experience acquired since the completion of the original work, enabled JET-X troubleshooting methodologies to become more sophisticated and flexible.

Troubleshooting procedures were extensively developed for three of these alarms which identify abnormal changes in engine trended performance. Because considerable experience and data were available, the analysis of these three alarms was complex; expert-system technology provided the only practical method of utilizing this information. Of the remaining 51 alarms, significant enhancements were made to the procedures for 5 alarms, and modest improvements were made to those of another 10 alarms.

Typically, the development of a specific troubleshooting procedure began with the factory engineer incorporating his experience and knowledge into a GEN-X "algorithm." This initial procedure was then modified to reflect any additional experience provided by the field rep. While the procedure could be considered complete at this point, experience proved otherwise. Frequently, in the process of building the knowledge base, a better method of executing already

completed procedures was identified and incorporated. These unanticipated enhancements altered the knowledge base in various ways: some improved troubleshooting techniques; while others simplified communication with the end user. Several changes streamlined and more fully integrated the entire knowledge base.

Having the experts actually construct the knowledge base gave them opportunity to be creatively involved in development, rather than just being periodic dispensers of information. This involvement triggered a high level of commitment from the experts that resulted in a constantly improving knowledge base.

The role normally filled by the knowledge engineer was affected by this process. Instead of extracting and interpreting expert knowledge, the knowledge engineer utilized his experience to coordinate the design of the overall architecture of the knowledge base, and to identify and design support capabilities needed to supplement GEN-X features. Principal responsibilities of the knowledge engineer included the task of designing an architecture to avoid brittle behavior within search-space limits, incorporating a strategy for display of interim results while the troubleshooting process was still in progress, and providing record-keeping capabilities to permit generation of reports in various formats.

While the benefits of having the technical expert construct the knowledge base are many, there are certain limitations. Because of training and hardware requirements, development of the knowledge base by the expert may be limited to the use of expert-system shells that are easy to learn, and which do not need costly and specialized hardware or software. Because of lack of experience, knowledge-base architecture may be awkward, cumbersome, and brittle if left only in the hands of the expert. The disposition of the expert to becoming involved as the knowledge-base developer (from both an attitude and availability point of view) is perhaps the most significant factor when considering this option.

APPENDIX C

APPENDIX C

BACKGROUND ON GAS-PATH ANALYSIS TOOLS AND TEMPER

C.1 Gas-Path Analysis Tools

One of the more difficult diagnostic problems associated with jet engines has been the determination of the cause of an engine performance loss, which is characterized by reduced thrust margin and increased fuel consumption. The problem may be simply stated as:

- Determining that the performance loss is real (measurement error is frequently large and cannot be ignored)
- Identifying which engine component(s); for instance, the HP compressor and/or LP turbine, is (are) responsible for the drop in performance.

There are three principal factors that contribute to the difficulty of this problem. The first of these is the inherent complexity of a jet engine. A jet engine includes large asymmetries in the pressure and temperature profile characteristics; many secondary flow paths that influence these temperature/pressure profiles, as well as the overall performance of the engine; and alternate rotating and static components which continually modify these temperature and pressure patterns. The second complicating factor is the mechanisms that cause performance degradation. These include changes in rotating machinery and seal clearances, erosion of (and deposits on) compressor and turbine airfoils, deformation of nozzle vanes, foreign or internal object damage, stall-induced damage to compressor blades, and many more. Most of these performance-loss mechanisms are primarily dependent on the operating environment in ways that are not well-understood; hence, it is not easy to predict losses in performance. Finally, the ability to instrument the gas path in order to directly sense the performance changes on a component-by-component basis is severely limited by the weight and expense of performance instrumentation, not to mention the attendant performance loss associated with the presence of the actual instrumentation. Consequently, the performance instrumentation on a jet engine is usually very limited, relative to the need for modular fault isolation.

The initial attempts to monitor on-wing performance levels were based on the application of development test codes to on-wing measurements. These development test codes are successfully used to evaluate the performance of factory test engines, based on extensive pressure and temperature instrumentation. It is not unusual for a development test engine to have four to six temperature rakes at an axial location, with each of the rakes having five or more individual probes. Thus, the development test measurements are very accurate, as well as quite reliable. When these codes are applied in a revenue-service environment, the accuracy of the on-wing measurements is so poor that reliable modular diagnosis is impossible. It is not unusual to see 5 to 10 percent errors in the computed efficiencies of the core engine components (the low spool efficiency errors may be as large as 50 percent or more). These large errors preclude the use of the development test codes for on-wing modular diagnosis.

The next advance in on-wing gas-path analysis was pioneered by Lou Urban in his appropriately named GPA (gas-path analysis) program. The trick in GPA is to make use of additional

information available in the measurements and in the engine model. For example, the gas-path information contained in the fuel flow measurement is essentially the same as that contained in the turbine discharge temperature. Thus, these measurements are redundant, and only one can be used in a conventional analysis. The development test codes address this situation by using the fuel flow to calculate the engine performance and, in addition, compute a temperature measurement error that defines the inconsistency between the fuel flow measurement and the turbine discharge temperature measurement. The problem of resolving this inconsistency is left to the analyst, although the conventional codes frequently do not provide a ready means for implementing the conclusions of the analyst.

Lou Urban noted that there were many at least partially redundant measurements in the limited on-wing instrumentation set. For example, if compressor efficiency is below normal, this fact should be reflected not only in the temperature and pressure measurements surrounding the compressor, but also in the fuel flow (which should be increased) and in the downstream temperatures (which should be elevated), relative to normal operation. If a mathematical technique could be found to take advantage of this observation, a more powerful modular diagnostic technique should result. The required mathematical technique was weighted least-squares (or the Kalman filter); this algorithm does make use of all of the instrumentation according to its presumed accuracy. The details of the weighted least-squares technique are described in Sections C.3 and C.4 herein.

One other observation that proves to be key to the success of this algorithm is that there are frequently substantial biases in the on-wing measurements. These biases often are the major cause of the huge efficiency errors previously reported. This problem can be partially corrected through the use of "a priori estimates" (or baselines) that are based on a sample of data from the initial service of the engine. This baseline can be referred to the engine cycle model and then used to give a better prediction of the expected measurements. The deviation from this expectation then is either due to a hardware performance change or a shift in the measurement error. The use of these baselines is a major reason for the improvement of these weighted least-squares algorithms relative to the development test codes.

The GE Aircraft Engines TEMPER (Turbine Engine Module Performance Estimation Routine) program is similar in its basic design to the GPA program. The significant difference between TEMPER and the original GPA program is the addition (to TEMPER) of "fault" logic. Frequently, the performance of an engine will change suddenly from one reading to the next. When this happens, the basic weighted least-squares algorithm tends to assign the shift to multiple causes, including both measurement errors and hardware changes. Engineering judgement suggests it is more likely that a single fault was responsible for the sudden change in performance. This assumption forms the basis of the TEMPER fault logic which seeks to assign the shift to a single cause (either a component performance change or a measurement error). The mathematical details are described in Section C.4 hereof.

One consequence of this design approach is that the algorithm must be able to recognize that a sudden change in the performance level has occurred. This is accomplished by maintaining

current estimates of the status of each of the hardware deviations and each of the measurement errors. These can be used to predict an expected value for each of the measurements for the current reading. The size of the sum of squares of residuals between the new measurements and those predicted from the current status allows the recognition of a sudden shift in the measured performance of the engine.

Substantial improvements in the ability to diagnose engine performance problems have resulted from the transition from development test codes to the TEMPER-type of algorithms. These improvements are largely attributable to the use of more information to solve the problem. This additional input to the solution includes the following elements:

1. The utilization of all measurements in the analysis to take advantage of redundant information
2. Use of historic measurement data to eliminate the effect of bias errors on the solution
3. The use of the engine model to define normal engine operating characteristics, including more obscure effects such as clearance changes, etc.
4. Use of expected measurement error and hardware shift limits to evaluate the relative likelihood of competing explanations for observed performance shifts
5. The utilization of trended information to identify sudden changes in the observed performance data, so that an alternate-solution algorithm can be invoked.

A sixth source of data that is used in some weighted least-squares algorithms (but not in TEMPER) is a deterioration model based on prior experience with the engine model.

Although the gas-path analysis algorithms based on the weighted least-squares method are significantly better than their predecessors, there is still room for improvement in the tool. The algorithms are still unable to correctly diagnose all performance problems; particularly those that are not uniquely identified by the available measurements. For example, it is not possible to distinguish normal HPT (high pressure turbine) deterioration, which is characterized by a simultaneous decrease in the efficiency and increase in the flow function, from an increase in many of the internal cooling flows (for example, HPT rotor cooling). There are many such equivalent fault sets that may be confused with one another. No single technique has been found to be fully reliable in diagnosing engine performance problems. (There have been instances where an on-site inspection by experienced design engineers failed to achieve the correct diagnosis; there are even cases where TEMPER was apparently superior to these designers' inspections.)

Field experience indicates that a better diagnosis is possible, when additional data is considered. Some examples of additional data items that can be helpful in achieving an effective performance diagnosis include:

- Maintenance history for the deficient engine
- Comparison of the problem engine with the other engine(s) on the same aircraft

- Vibration trend data
- Oil metal content analysis results
- Pilot reports (for example, that a stall occurred)
- Engine/component life usage data, expressed in hours or cycles
- Special usage data (such as, the number of gun fires since last water wash for the A-10A)
- Engine control data
- Inspection results (for example, borescope).

Most of these data items are qualitative rather than quantitative; hence, they are not readily integrated with the gas-path analysis algorithm. (Some integration has been achieved for TEMPER with the maintenance data; however, this has been awkward at best and represents only a crude beginning.) At the start of this program, we believed that this integration could be effected through the use of knowledge-based systems technology. Consequently, this has been one of the guiding objectives for this effort.

C.2 Turbine Engine Module Performance Estimation Routine (TEMPER)

As indicated in Section C.1, GE Aircraft Engine uses a gas-path analysis algorithm known as TEMPER. TEMPER was developed to support commercial engines which, historically, have been better instrumented (at least for some operators) than their military counterparts. TEMPER was originally developed to process overhaul acceptance test data in the late 1970's. It was extended to on-wing application in the early 1980's. Both test cell and on-wing versions of TEMPER are currently in use as part of the GE-provided GEM (Ground-based, Engine Monitoring) package. This portion of the report gives a general description of TEMPER.

C.3 General Description

On-wing TEMPER uses the weighted least-squares technique as its central element in determining the cause of engine performance loss. This basic algorithm has been modified in a number of ways to increase the effectiveness of on-wing TEMPER. Many of these modifications simulate the features of a Kalman filter; others are deliberate changes to the design to make it more effective in evaluating jet engine performance. This section will provide a general description of the on-wing TEMPER calculation; whereas, Section C.4 will describe the mathematical details. Figure C-1 shows a general flow chart for the on-wing TEMPER calculation.

The objective of on-wing TEMPER, or any similar gas-path analysis tool, is to compute a number of "state variables," that define the well-being of the engine. These state variables include component efficiencies and flow capacities and, potentially, other characteristics of the engine (such as, cooling flow levels or nozzle areas). The state variable list should include all engine characteristics that are apt to change as a result of deterioration. The state variable list must be a minimum set in the sense that they are mathematically independent of one another. For

example, we could not use both a nozzle area and the associated flow coefficient, since there is no way to distinguish these state properties based upon their effects on measurable quantities. The set of state variables should also be complete in the sense that if we know the status of each of the state variables, the status of the engine can be directly inferred.




Figure C-1. On-Wing TEMPER Flow Chart.

On-wing TEMPER operates with routinely acquired engine data; typically, in commercial applications, these data are obtained at a steady-state cruise condition where transient effects can be ignored. For most military aircraft (the exception being military transport aircraft), there is no flight condition that can be characterized as being steady-state on a repeatable basis. The A-10A fits this latter pattern; hence, on-wing TEMPER uses takeoff data as the basis for its analysis. The takeoff data are acquired at a fairly repeatable condition and, thus, can be used for an on-wing TEMPER analysis. No attempt is made to interpret the variable transient effects; consequently, they are lumped with the other noise sources.

The available gas-path measurements are subdivided into two sets. The first set is denoted as "setting parameters;" this constitutes a minimal set needed to define the actual flight condition and engine power level. Included in this set are measurements such as altitude, mach number, fan speed, and total air temperature. These measurements are used as inputs to the engine cycle model. There must be a sufficient set to completely specify the inputs to the cycle calculation.

The remainder of the measurements form the basis for the evaluation of the overall performance of the engine, and the diagnosis of performance problems to the deficient component(s). Typical of these measurements are core speed, fuel flow, HP compressor discharge static pressure, HP turbine discharge temperature, etc. These measurements can be compared to the predictions from the engine cycle model. The results of this comparison define the degree to

which the actual engine deviates from its desired performance. This second group of measurements shall be referred to as the "independent measurements" in this discussion.

There are at least two minimal constraints on the group of independent measurements. The first of these constraints is that there be a sufficiently large set of measurements to facilitate a complete determination of all of the hardware state variables (previously defined). This requirement means that it should be possible to develop an algorithm that directly computes the values of all of the state variables, given a complete set of independent measurements (the setting parameters may be used as well). Further, it is desirable that this algorithm be numerically well behaved in the sense that the solution is not overly impacted by measurement error. In control theory, this requirement would be rendered as saying that the state variables are observable. The possibility of measurement error as an explanation for the deviation of the measurements from expectation cannot be ruled out. Therefore as a second requirement, it is desirable that these measurements include some redundancy to permit simultaneous evaluation of the hardware condition and the measurement errors. One specific example of this redundancy is the presence of both a fuel flow measurement and at least one turbine discharge temperature. These measurements are essentially redundant; only rare hardware changes can cause them to deviate from one another.

In practice, these constraints are seldom (if ever) satisfied for jet engine diagnosis. For example, even in the best of circumstances, there is not sufficient instrumentation to reliably distinguish between an increase in the turbine cooling flow and normal deterioration of the HP turbine. Unless the instrumentation is very complete, there usually are not sufficient measurements to allow full evaluation of the LP system. Thus the selection of state variables involves some compromise between the previously stated hard and fast rules and the limits of practicality. When two state variables are indistinguishable in the measurement set, one or the other can be selected for analysis; and, it is hoped, the determination of which was truly at fault can be accomplished subsequent to the on-wing TEMPER analysis by other means.

One of the most important elements of the on-wing TEMPER algorithm is the use of "initial installation baselines." There are many reasons that the measurements obtained on-wing would be biased relative to the engine cycle model.

1. The cycle deck normally defines plane average values for pressure and temperature measurements. These averages are normally based on factory testing using a large number of probes that are equally spaced at the appropriate engine station. In contrast, the in-service measurements are generally from single probes that are normally mounted near the wall to minimize the possibility of internal engine object damage from a fatigue failure of the sensor. Since there are generally sizable pressure and temperature gradients within the engine, these single-element measurements are usually biased, with respect to the plane average.
2. The cycle deck is generally matched to some engine quality level other than the initial installation level. Usually, the model represents an average of a group of production acceptance runs or possibly the result of a flight test.

3. The cycle deck may include modeling errors. Generally considerable emphasis is placed on accurately portraying the overall engine performance parameters, such as thrust and fuel flow. There is a lack of corresponding emphasis on some of the internal pressures and temperatures. This is not surprising, since the data may not be available to do as good a job on these internal measurements. The data acquired during a development test program often are based on engines that contain "slave" hardware that does not conform to the production parts list, or may be acquired on hardware that is subsequently improved during the design process. If the cycle model is based upon production acceptance runs, the internal pressure and temperature measurements are no better than those available on-wing; hence, they cannot be reliably used to determine component efficiencies, etc. In this context, selection of the appropriate performance level for the internal components becomes problematic. Further, virtually all testing occurs at sea level, so that altitude effects are not always correct, and are frequently based on analysis rather than solid test experience.
4. It is not always clear what pressure or temperature in the cycle model corresponds to a given probe. What location in the fan duct provides the predictor for a given static pressure? What is the appropriate location for a temperature probe to be modeled? Usually in the cycle model, flow additions or subtractions, or pressure losses are arbitrarily defined to occur between specific numeric engine stations. This is a consequence of the one-dimensional character of the cycle deck calculation. This makes it difficult to select the "correct" location that corresponds to a particular measurement probe.
5. In the case of the TF34, yet another problem is that the on-wing data are acquired during the takeoff transient. Accordingly, there are significant transient effects on the measurements that are not modeled in the cycle deck.

For these and other reasons, it is important to utilize initial service baselines (computed using initial installation data, when available) in the on-wing TEMPER calculation to represent the differences between the cycle model and on-wing measurements. The procedure implemented is to acquire data from a number of newly installed engines and to develop correlation factors relative to the cycle model; so that, with the baselines applied to its outputs, the cycle deck correctly predicts the initial service measurement levels. The baseline determination process involves close scrutiny of the data to determine that the engine is not deteriorating and that the measurement probes are performing properly. The comparison of data acquired from several engines helps in the latter judgement. Usually, the initial service baselines are expected to vary with power setting; however, variations in altitude, Mach number, or total air temperature are not normally expected.

In addition to the initial service baselines, on-wing TEMPER also uses a second type of baseline which shall be referred to as a priori estimates. These estimates are a further adjustment to the initial service baselines whose purpose is to provide a good prediction for the new data, based on the data which has already been acquired. The derivation of these a priori estimates will be described later in this discussion; for now, it is important only to know that they are used in the initial calculation in the same way as the initial service baselines.

With this discussion as background, the initial steps in the on-wing TEMPER calculation is described (Figure C-1). Incoming data are divided into setting parameters and independent measurements. The setting parameters are used as input to the cycle deck calculation to produce predicted values for the independent measurements. These predicted measurements are adjusted using the initial service baselines and the a priori estimates. The resulting predictions for the independent measurements are then compared to the actual measurements, and the differences are computed for input to the weighted least-squares analysis.

In addition to the measurement deviations, the on-wing TEMPER weighted least-squares analysis needs three other inputs. These are: cycle derivatives, measurement error statistics, and state variation statistics.

The weighted least-squares analysis is attempting to determine the most likely explanation for the observed measurement variations. To accomplish this, it must be able to hypothesize a solution in the form of a specification of the changes to the state variables and, from this, compute the probability of the hypothesized solution given the observed measurements. In order to compute this probability, the algorithm must know the measurement errors that are implied by the hypothesized hardware state. This calculation could be accomplished by entering the hypothesized state variable changes into the cycle deck and recomputing the predicted measurements. A comparison of these predicted measurements to the actual measurements would define the implied measurement errors. The use of the cycle deck as an intermediary in this process is not practical, because the cycle deck is a nonlinear, iterative calculation that is relatively expensive computationally; and because there are an infinite number of possible hypotheses to be evaluated.

To get around this problem, assume that the state variations are sufficiently small so that a linear approximation of the cycle deck can be used to perform this analysis. This is the motivation for the input of the cycle derivatives. The calculation of the cycle derivatives is not a trivial process. The "canned" approaches do not yield adequate results for this application. Instead, it is necessary to run a range of points at a given flight condition and use a linear-curve-fit analysis to compute the slopes. Frequently, it is necessary to plot and review the data to identify and resolve anomalies. The derivatives are generally computed at several power settings and flight conditions so that variations with these setting parameters can be included. It is also important to recognize the impact of variable geometry on the derivatives. In regions where the geometry varies with power, the slope of the variable geometry is superimposed on the basic engine slopes; in other power regimes where the geometry is fixed, these superimposed geometry slopes are absent.

In addition to cycle derivatives, the weighted least-squares analysis uses statistical models for the measurement errors and the hardware deviations in its analysis. The statistical models provide the additional equations to the weighted least-squares process that permit simultaneous calculation of the measurement errors and the hardware states; (without the statistical input, there are only enough equations to compute the measurement errors). Once again, we could consider the use of arbitrary statistical models for both the measurement errors and the

hardware states. For example, we might be tempted to select an asymmetric statistical model, for the component efficiencies, that favors a loss in efficiency over an increase in efficiency. Unfortunately, the use of any model, other than the Gaussian (or Normal) distribution, results in a nonlinear problem. Thus, assume that both the measurement errors and the hardware state variations are Gaussian. This is not an especially troubling approximation, because the Gaussian distribution has been shown to be the correct limiting distribution for statistically independent variables.

The procedure used to determine these statistical inputs to the weighted least-squares analysis is worthy of some discussion. The process begins with the determination of a first estimate for the measurement errors. This estimate is obtained from the same type of data used to identify the initial service baselines (as described in the preceding paragraphs). For each individual engine, the scatter about the mean is computed. In principal, this calculation could include a determination of the covariances as well as the variances; however, experience shows that these are negligible. The calculation of variances is done on a per engine basis (with a distinct average for each engine), because the engine-to-engine variation of the measurements is to be accommodated in another way in the on-wing TEMPER algorithm. Variances for the individual engines are averaged to obtain an ensemble mean for the variance of each of the individual measurements. The same type of careful review that was applied in determining the initial service baselines is also appropriate here in order to eliminate obviously bad measurements from consideration.

The technique used for determining the hardware state variations for input to the weighted least-squares algorithm is based on an interesting mathematical property of this algorithm. Like the Kalman filter, the weighted least-squares algorithm is a linear estimator. This means that the assessment of the hardware state and the measurement errors is a linear function of the input measurement deviations. If each of the incoming measurement deviations were doubled, then each of the derived hardware state deviations and measurement errors would also be doubled.

A consequence of this property is that it is possible to analyze the response of the weighted least-squares algorithm to known input problems; such as, a specific measurement error or a specific hardware problem. For example, a specific measurement error's response could be explored by running the weighted least-squares algorithm with a vector that consists of a deviation only in the measurement whose error is to be investigated (all other input measurement deviations would be zero). Similarly, a specific hardware deviation could be evaluated by simulating the measurement deviations that are associated with a change to that specific state variable, while all other state variables are at their nominal levels, and all measurement errors are zero. This type of analysis to explore the response of the weighted least-squares algorithm to known input problems will be referred to as a "response analysis" in this discussion.

The response of the weighted least-squares algorithm to an input vector is a function of the supplied values of the cycle derivatives and the measurement error and hardware state variation statistical models. The first two may be considered to be fixed by the cycle deck and the

data input, respectively. Experience with jet engine gas-path analysis (given the amount of instrumentation available) shows that it is reasonable to expect to achieve a correct identification of about 60 percent of an incoming problem. Thus, if a pure turbine efficiency problem is input to the algorithm, the output will correctly ascribe approximately 60 percent of the problem to turbine efficiency. The remaining 40 percent will be incorrectly attributed to other state variables and to measurement errors. Response to a turbine problem could be improved by increasing its input variance, but only at the cost of lowering the correct responses of the other hardware states and the measurement errors. The value of 60 percent as a reasonable value for the ultimate response depends on the amount of redundant instrumentation available; the value represents the best GE experience for current applications.

In theory, the statistical model for the hardware state variations could be derived from data. In fact, the measurement covariances derived from the observed scatter in the test data may include the effects of hardware state change. For this reason, it is important to restrict the calculation of these measurement error covariances to regions where there is no apparent deterioration; however, it is not necessarily desirable to use the actual hardware state covariances as input to the weighted least-squares analysis. Instead, it may be desirable to make the algorithm equally capable of detecting each of the potential problems that it could encounter. Thus, experience might suggest that compressor problems are extremely rare in comparison to turbine problems. If this experience were reflected in the input to the weighted least-squares analysis, the response to a true compressor problem would be much poorer than to a true turbine problem. It may be desirable instead to adjust the hardware state variances to permit the equal detection of either problem.

Both factors should be considered in selecting the hardware state variation statistics. There are certain types of problems that the algorithm must successfully detect. These include: HP compressor efficiency and flow problems, HPT efficiency and flow changes, and at least one of the LP shaft (either fan or LPT) efficiency variables. It is also very desirable to be able to determine the status of fan flow. Conversely, service experience suggests that there is little, if any, gradual deterioration in the LPT flow characteristic. Furthermore, the knowledge of the LP turbine flow function helps to offset measurement errors in the interturbine pressure measurement (when one is available) and, therefore, contributes to the ability to compute HPT efficiency. Thus, it is desirable to select a very small variance for the LP turbine flow function and accept the associated poor response to LP turbine flow problems.

Another consideration in selecting the hardware state variation statistics is the possibility of correlation between variables. For example, most of the mechanisms that can cause deterioration to the fan, affect both the flow and the efficiency; thus clearance change, surface erosion, leading edge distortion, etc., all effect both fan efficiency and fan flow capacity adversely. Further, the ratio of fan efficiency change to fan flow change for each of these mechanisms is roughly similar. Statistically, this means that we should expect a high correlation between the fan flow change and the fan efficiency change. In the limit of 100 percent correlation, only one variable would be needed; the other would be uniquely determined from the knowledge of the first. In practice, however, it appears that the true correlation is less than this limit;

perhaps only about 85 percent. A similar observation holds for the HP compressor, although the expected correlation is somewhat less (perhaps 80 percent). There does not seem to be any significant correlation between flow function and efficiency in the turbines. This knowledge of correlation can enhance the fault isolation because the instrumentation may do a better job of identifying the flow change than the efficiency (or vice versa); the correlation helps to achieve a correct estimate for the less observable state variable.

It should be noted, that the introduction of nonzero hardware correlations into the process of response analysis can produce some rather surprising results. The response to a pure flow or efficiency change in a correlated component can be greater than 100 percent, or less than zero (meaning that a loss of efficiency is interpreted as an increase, or vice versa). This behavior can be somewhat disconcerting; however, a more correct test in this case is to test the response to a properly correlated input problem with both the efficiency and the flow perturbed at the same time in the proper ratio.

A special tool has been developed to test the response characteristics of the weighted least-squares analysis given a set of input statistics. This tool is used iteratively to select the hardware state variances and includes the correlation effects, as well as all of the given data (cycle derivatives and measurement error variances). Once the hardware variances have been derived, a data sample is processed through the weighted least-squares algorithm. The resultant deviations are compared to the predicted histogram, based on the Chi-Squared distribution. The results of this comparison are used to scale all measurement error and hardware state variances to assure reasonable levels for subsequent probability calculations. This scaling does not alter the response characteristics of the algorithm.

With this additional background, the description of the on-wing TEMPER algorithm (Figure C-1) will be continued. A more detailed view of the weighted least-squares portion of the on-wing TEMPER algorithm is provided in Figure C-2. The first step in the process is to eliminate obviously bad measurements from the weighted least-squares analysis. Each incoming measurement deviation is compared to a limit to decide whether the measurement has experienced a hard failure. The limit is expressed as a multiple of the standard deviation for the measurement error (typically, a number in the range from 15 to 30 is used). When a measurement is judged to be in error, or if it is unavailable, its deviation is set to zero in the weighted least-squares input, and the corresponding measurement error sigma is increased by a factor of 100 to eliminate the measurement as a constraint on the weighted least-squares solution. In order to avoid ridiculous solutions, a minimum number of measurements is required before the weighted least-squares analysis will be performed.

To provide motivation for the next step in Figure C-2, some digression to discuss Kalman filters is required. The Kalman filter behaves much the same as a weighted least-squares analysis, with the exception that the Kalman filter is used to analyze a continuous stream of incoming data (as will be discussed, several characteristics of a Kalman filter have been adapted for the On-Wing TEMPER). When a Kalman filter begins an application, it has no initial knowledge of the state of the system that is being analyzed. As time passes, the incoming data provide a

more certain knowledge of the state of the system. In fact, if the system is one which does not change with time, the Kalman filter will become increasingly certain in its knowledge of the system; so that in the limit, new data will not alter the assessment of the system's state. Most problems of interest involve some gradual, unpredictable change of state, so that the Kalman filter approaches a steady-state condition where the new data information is offset by the process noise. When this occurs, the Kalman filter is in a steady-state condition, and we can consider the initial uncertainty in the system to be eliminated.




Figure C-2. Weighted, Least-Squares Flow Chart.

The mathematical formulation of the Kalman filter automatically reflects this qualitative behavior. The analogue of the hardware state variance is computed in the course of the analysis of the Kalman filter. Initially, the implied hardware variances are very large, suggesting little knowledge of the hardware state. As more data become available for the estimation, the initial uncertainty disappears, and the Kalman filter approaches a steady-state condition where the equivalent hardware variances represent only the process noise, with no contribution from the initial uncertainty.

In on-wing TEMPER, this increasing certainty in the knowledge of the system's state (including measurement biases) has been simulated by a device that scales the measurement error

and hardware state variances depending on the amount of data that have been utilized to establish their a priori estimates. Scaling is performed using the student distribution, which uses as an input the number of readings that are included in the specific a priori estimate baseline. In practice, the change in the standard deviation is negligible after 20 samples are included; consequently, the adjustment disappears after 20 readings are included in the a priori baseline.

From time to time, a measurement will fail and be out of service for some length of time. When the measurement is corrected, the on-wing TEMPER algorithm assumes that the a priori baseline for that measurement must be reestablished; therefore, the counter for that measurement is reset to zero. This has the effect of diminishing the impact of the measurement on the weighted least-squares analysis until confidence in its baseline has been reestablished.

Having completed these preliminary steps, weighted least-squares analysis is performed using the potentially modified measurement deviations and standard deviations. The mathematical description of this calculation is given in Section C.4. Qualitatively, the algorithm determines the hardware state vector that is most probable, given the observed measurement deviations. This hardware state solution also implies a value for each of the measurement errors; the calculation of the probability uses both the hardware state and the measurement errors. The input standard deviations tend to define the possible excursions for the measurement errors and the hardware state variations; a multi-sigma change to any specific parameter is unlikely.

Once the solution has been determined, an associated solution probability is also computed. This solution probability can be thought of as the likelihood of occurrence of the observed deviations in hardware state and measurement errors given the statistical model for each of these variables. The probability calculation uses the Chi-Squared distribution. An input to this calculation is a number of degrees of freedom, which represents the number of valid measurements used to obtain the solution (any measurement that is judged to be failed is not included in this count). Typically, this probability will be relatively high for a good solution. A small value for this solution probability can be interpreted as an indication that some element of the statistical model was invalid for the particular reading.

When the probability of the solution is beneath a certain threshold, the on-wing TEMPER algorithm attempts to determine a plausible explanation for the observed deviation. Because on-wing TEMPER maintains a set of a priori baselines which represent the accumulated knowledge about the state of the system and the measurement errors, the low probability is an indication that a significant shift in one or more of the measurements has occurred since the last reading was obtained. Given that the large shift occurred within a relatively short time interval, it is reasonable to suppose that the change is due to a single event. This assumption is, in fact, the basis of the on-wing TEMPER fault-search logic.

There are many candidates that could be considered for this single event. There may have been some moderately hard failure in one of the sensors; although not sufficiently hard to be detected by the limit check logic, hard enough to result in a low solution probability. Sensor

failure need not be one of the independent measurements; it might be a setting parameter failure. There may have been some hardware event that resulted in the sudden shift in performance. A compressor stall could have caused tip rubs in the compressor, as well as aluminum deposits on the compressor blades, resulting in a substantial change in the efficiency and flow capacity of the compressor. The sudden change need not affect one of the state variables that is normally included in the on-wing TEMPER analysis. For example, a combustor liner collapse would behave somewhat like a turbine performance problem.

On-wing TEMPER includes a list of fault candidates that are considered any time there is a sudden shift in performance. Each of these potential faults is evaluated by the on-wing TEMPER fault-search algorithm. When the fault to be evaluated represents an error in one of the dependent measurements, the corresponding standard deviation is increased by a factor of 100, and the data are resubmitted to the weighted, least-squares algorithm. The large increase in the standard deviation frees the specific measurement error to assume a very large value without penalizing solution probability. Similarly, when the prospective fault is one of the usual hardware state parameters, its standard deviation is increased by a factor of 100, and the weighted, least-squares analysis is repeated. The hardware state faults are done both individually and in component pairs (both flow and efficiency for a given component); the latter hypothesis indicates that a significant occurrence in a component is likely to affect both the efficiency and the flow capacity.

When the prospective fault represents an error in one of the setting parameters or a hardware state variable that is not normally included in on-wing TEMPER analysis, the additional fault is included in the state vector array for the fault test. The standard deviation assigned to this extra fault is large (comparable to the values given to the normal parameters in their fault tests). The augmented state set is then processed through the weighted, least-squares analysis to evaluate the prospective solution.

The solution probability is used to determine which of the prospective faults provides the best explanation for the observed data. Unless this best solution exceeds a required probability minimum (currently, equal to 25 percent), none of the fault-search solutions will be accepted as valid. If more than one of these solutions exceeds this minimum threshold, the solution with the highest probability will be chosen. One of the problems with this analysis in the current design is that some of the faults may be nearly indistinguishable in the data; therefore, the one selected is not necessarily the correct choice.

One brief note on the calculation of the solution probability in the fault-search analysis should be mentioned. The number of degrees of freedom is reduced by one for each "free" variable in the fault analysis. Accordingly, if a measurement error is being evaluated, the number of degrees of freedom is lowered by one; if the flow and efficiency of a component are being simultaneously evaluated, the number of degrees of freedom is reduced by two. The effect of this reduction in the number of degrees of freedom is a corresponding reduction in the computed probability.

C.4 On-Wing TEMPER Development

The most challenging element of the GEM (ground-based engine monitoring) design was the development of a gas-path-analysis algorithm, which has been named on-wing TEMPER. The on-wing gas-path analysis is a key element of the condition monitoring system. The importance of which is emphasized by the high cost of fuel since the decision not to overhaul a deteriorated engine can prove to be quite expensive. However, even before the rapid increase in fuel prices, the on-wing gas-path analysis provided a basis for identifying incipient failures and, thus, avoiding costly secondary damage. Another important goal of the on-wing gas-path analysis is to pinpoint which engine components are in need of overhaul in order to restore performance; the ability to perform this function reliably can considerably reduce overhaul costs and also can reduce or eliminate second overhauls.

In the past, engine/component analyses based on on-wing measurements have not been particularly successful. The principal cause of this lack of success has been the limited amount of gas-path instrumentation available on-wing, and the inaccuracy of that available data. The deterministic algorithms which were used interpreted the measurement errors as component faults since they had no way of estimating measurement error. Thus, measurement error might be manifested as an absurdly low HPT efficiency rate (indicative of a fault in the temperature or pressure measurement at the compressor discharge). At the beginning of the GEM project, newer, statistically based algorithms were successfully being used to analyze acceptance test data for jet engines. The extension of these techniques to the on-wing problem appeared promising; particularly since these techniques cope with measurement error, and because they took advantage of the limited redundancy available.

The statistical approach most commonly applied to this type of problem is the Kalman filter. (There is an enormous body of literature available on Kalman filters; Reference 1 provides a good fundamental description.) There are a number of ingredients needed for a successful application of the Kalman-filter technique. These include:

1. A good model of the physical process to be analyzed. This model should include a set of state variables (which are to be estimated by the Kalman filter, or else are assumed to be predictable) and the ability to predict the available measurements from the state variables. In a time-varying application, it is desirable that the model provide a reasonable prediction of the change of the state variables with time.
2. A set of measurements is needed, which in the time-varying case, are to be repeated at some interval. It should be possible, in the absence of measurement error, to use these measurements to calculate all of the state variables which are to be estimated by the Kalman filter.
3. The measurements are assumed to have a Gaussian error component. A model of this error component is needed, consisting of standard deviations (means are assumed to be an element of the basic physical model) and correlation coefficients.

4. In order to reflect the fact that the physical process can not be perfectly predicted, a "process noise" model is needed. This process noise accounts for the variability of the physical phenomena; it is, once again, assumed to be Gaussian in nature.

In the absence of process noise, the only unknown to the Kalman filter is the initial state of the system; in this case the filter uses the repeated measurements to continuously improve the initial state estimate, and the model defines the variation with time. The residual measurement deviation from this "fitted" model is attributed to measurement error. The Kalman filter maintains a running estimate of the certainty of its knowledge about the state of the system. As more measurements continue to be taken, the calculated uncertainty of the initial state estimate is, of course, reduced; accordingly, each new measurement has a smaller and smaller influence on the state estimate. In the limit, any deviation from the state estimate and model will be attributed to measurement error.

The addition of process noise to a Kalman filter estimation adds a new dimension to the problem; the state at each time step is independent (to some degree) from the past and, therefore, is unknown. This process noise counters the increased certainty associated with repeated measurements, so that there is always some uncertainty about the state of the system. As time progresses, this element of uncertainty associated with the initial system state gradually disappears. Thus, in the limit, a steady-state condition is reached, where each new estimate represents a balance between the process noise and the measurement uncertainty.

The Kalman filter exhibits a number of interesting mathematical properties. One of the most significant of these properties is that the Kalman filter is an unbiased, minimum variance, estimator. This means that if we consider a large class of estimators and apply each to a large number of problems, the Kalman-filter estimates will exhibit less error than any of the other estimators. This property is an important reason for the popularity of the Kalman filter.

Another significant mathematical property of the Kalman filter is that it is a linear estimator, and as such, the state estimate is supplied as a linear function of the observed measurement deviation from the a priori model. The matrix which expresses this proportionality between the measurement deviations and state changes is known as the Kalman-gain matrix and is computed from model slopes, measurement error and process noise covariance matrices, and a matrix defining the predicted time variation of the state. The Kalman gain does not depend upon present or past measurements, or upon the current estimate of the state. Thus if desired, the Kalman gain for each time step can be computed prior to any data acquisition.

The Kalman-gain matrix approaches a limiting value for large times. This steady-state limit is a consequence of the increasing certainty of the initial state, as described in the preceding paragraphs. In the absence of process noise, the limiting Kalman gain is the zero matrix; thus in this case, the estimate of state approaches independence of the measurements. It is necessary to introduce a process noise matrix in order to have a non-zero, limiting Kalman gain.

An important factor that should always be remembered when dealing with Kalman filters, or other similar algorithms, is that they are attempting to evaluate many more unknowns than there are equations. The unknowns are state variables and the measurement errors for each of the measurements. However, there is legitimately only one equation per measurement. The additional equations needed to produce a solution are developed from the statistical data (measurement covariances and process noise) and model constraints. One consequence is that measurement redundancy is highly desirable, since this reduces the ratio of unknowns to equations.

Because the Kalman-filter state and measurement error estimates are a linear function of the measurement deviations from the a priori model, it is possible to study the Kalman-filter response to known input problems. For example, a state variable in the jet engine, gas-path analysis example is HPT efficiency. The Kalman filter can be tested by simulating the input associated with a known change to turbine efficiency (for example, one percent) and then examining the Kalman filter's interpretation of this input. Since the Kalman filter is linear, the ratio of the solution vector (state and measurement error) to the input problem is independent of the amplitude of the input problem.

To continue the turbine efficiency example, experience shows that the Kalman filter will usually indicate a change in the turbine efficiency of less than the one percent input value. The exact response depends upon the statistical inputs; such as, the process noise and measurement error covariances. Typically, the output turbine efficiency change will be approximately 60 percent of the input change. *The remainder of the input measurement deviations associated with turbine efficiency will be incorrectly attributed to other components and to measurement error.* The addition of redundant measurements, or the improvement of the precision of the measurements, could improve the response to a pure hardware problem; however, these options are not practical for on-wing gas-path analysis.

The response of the Kalman filter to other state variable changes and to pure measurement errors is analogous to the turbine efficiency example. Generally, the correct element of the response will be a fraction between zero and one; although, it is possible (when correlation is present between the state variables) to have a response greater than unity or less than zero, with the remainder of the input being incorrectly attributed to other state variables and measurement errors. An obvious consequence of this observation is that the solution error is linear with the magnitude of the input deviation. As the problem becomes larger, the Kalman filter becomes less effective. This behavior is not inconsistent with the minimum-variance property, since the Gaussian distribution guarantees that the majority of the data will be near the mean and that the probability decreases rapidly as the deviation magnitude increases.

While the algorithm used in the on-wing TEMPER module is not a rigorous Kalman filter, it shares many of the properties attributed to the above-described Kalman filter. The principal reason for not using a Kalman filter is convenience; the design permits experimentation with various modifications to the basic concept, using relatively simple program changes. Many of the modifications will be described in the following paragraphs.

The on-wing TEMPER module has as its central element a weighted, least-squares algorithm. The weighted, least-squares algorithm may be thought of as a snapshot version of the Kalman filter; that is, a Kalman filter applied to a single set of data. A brief description of the weighted, least-squares algorithm is provided here as a basis for what follows (Reference 2 provides more detail).

The analysis is centered on a model of the measurement process in the form:

$$Z = h(x) + v \quad (1)$$

The elements of this equation are as follows:

1. "Z" is the measurement vector, consisting of independent measurements upon which the analysis is to be based. The number of measurements (and, therefore, the number of components in Z) is taken to be p. In the condition monitoring analysis, the elements of Z would be measurements; such as core speed, fuel flow, compressor discharge temperature, LP turbine inlet pressure, etc.
2. "x" is the state vector which includes those characteristics of the system which are needed to completely specify its condition. It is assumed that x is a minimum set in the sense that there is no duplication of information within x. The vector, x, is assumed to have n components. In the condition monitoring example, the elements of x would include those characteristics of the engine which are expected to vary with time; for example, compressor efficiency, HP turbine flow function, cooling flow, etc. Other attributes, such as nozzle area, which are not expected to change with time are excluded from x.
3. "h(x)" is a model of the process which accurately describes the effects of the state variables on the measurements. In the absence of measurement error, $Z = h(x)$ would be precisely true. In the condition monitoring example, h(x) would be represented by a cycle model which predicts the effect of the state variables on the measurements.
4. "v" is the measurement error vector consisting of the random element of the measurement error (any bias in the measurement is assumed to be incorporated in the model). In a practical sense, v will include, in addition to the random measurement errors, the effects of any secondary state variables not included in the state vector, x, and any random imperfections in the model.

In order to make the problem mathematically tractable, several simplifying assumptions are made; these are:

1. The measurement error, v, is assumed to follow the Gaussian (normal) probability distribution. Further, it is assumed that the mean, or expected, value of v is zero. The variance and covariance of v are given by:

$$E(v * v^T) = R \quad (2)$$

where the operator "E" denotes the expected value function, and the superscript, T, denotes the matrix or vector transpose. The diagonal elements of R are the variances of the measurement errors; any nonzero off-diagonal elements represent the existence of a correlation between two measurement errors (as would occur if there were a common element to the measurement process; such as, the use of the same barometer).

2. The measurement error, v, is assumed to be statistically independent of the state, x; that is,

$$E(x * v^T) = 0$$

3. The state vector, x, is also assumed to obey the Gaussian distribution with mean (or expected value), $\bar{x}$, and covariance:

$$E[(x - \bar{x}) * (x - \bar{x})^T] = M \quad (3)$$

The off-diagonal elements of M represent correlations between state variables. For instance, it is quite reasonable to expect that there would be a correlation between fan efficiency and fan flow since a common mechanism is likely to cause changes in these state variables.

4. The process model is linearized with respect to the state variable, x. This can readily be accomplished by performing a Taylor series expansion about $(x - \bar{x})$; this yields:

$$h(x) = h(\bar{x}) + dh/dx|x = \bar{x} (x - \bar{x}) + \dots$$

The terms neglected are of the order $(x - \bar{x})^2$ and higher; define:

$$dh/dx|x = \bar{x} = H \quad (4)$$

It should be noted that H is a matrix of dimension $p * n$. The introduction of Equation 4 into Equation 1 yields the following:

$$Z = h(\bar{x}) + H * (x - \bar{x}) + v$$

This equation may be simplified by redefining the origin of the measurement Z. Accordingly, let:

$$Z' = Z - h(\bar{x}) + H * \bar{x}$$

Now drop the prime (') for convenience, and Equation 1 becomes:

$$Z = H * x + v \quad (5)$$

Given the preceding definitions and assumptions, the objective is to develop an algorithm that estimates the current state of the system, X , given a set of measurements, Z . The selected approach considered x and Z to be stochastic variables, as defined in the foregoing, and then chose X to be the state vector estimate which maximizes the probability density function for x given Z , $P(x|Z)$. Reference 1 demonstrates that $P(x|Z)$ is a decreasing, monotonic function of the quadratic form:

$$J = 1/2 \{ (x - \tilde{x})^T M^{-1} (x - \tilde{x}) + (Z - Hx)^T R^{-1} (Z - Hx) \} \quad (6)$$

Employing Equation 6, if "J" is minimized with respect to the state vector, x , the result will be the state vector estimate with the highest conditional probability. To find this solution, evaluate the differential of Equation 6:

$$dJ = dx^T \{ M^{-1} (x - \tilde{x}) - H^T R^{-1} (Z - Hx) \}$$

To find the minimum, dJ must be zero for arbitrary dx^T ; therefore, the solution is:

$$X = \tilde{x} + (M^{-1} + H^T R^{-1} H)^{-1} H^T R^{-1} (Z - H\tilde{x}) \quad (7)$$

Equation 7 is used to determine the maximum likelihood estimate for the state vector. The corresponding estimate for the measurement error, $V = Z - Hx$, is:

$$V = \{ I - H(M^{-1} + H^T R^{-1} H)^{-1} H^T R^{-1} \} (Z - H\tilde{x}) \quad (8)$$

The weighted, least-squares algorithm can be understood as a single-point Kalman filter. In this context, most of the parameters share the same meaning for both algorithms. The state covariance matrix, M , may not be incorporated in the Kalman filter formulation, although it serves a role similar to that played by the process noise. "M" represents the a priori uncertainty in the estimate of state. The linearity properties of the Kalman filter apply equally well to the weighted, least-squares algorithm. In particular, this means that the algorithm response to known input problems can be studied and that the algorithm errors are proportional to the amplitude of the input problem.

One advantage for the weighted least-squares, in practice, is that there is a fixed Kalman-gain matrix for the algorithm, and thus, only one response characteristic need be studied; (the Kalman gain for a Kalman filter varies continuously until the initial-condition effects are damped out.) Also, it is somewhat easier to adjust the response characteristics of the weighted, least-squares filter by means of the M and R inputs than to perform the same adjustment for a Kalman filter.

A number of performance changes that are of interest in jet engine condition monitoring occur suddenly; for example, a rapid engine transient or a series of transients can open up compressor and/or turbine clearances, and can produce a significant change in performance. Other events which can alter engine performance by a significant amount in a short time are a stall, foreign

object damage, or a part failure. There are also a number of mechanisms which can cause a sudden, significant shift to a measurement bias. When a Kalman filter or weighted, least-squares algorithm is used to analyze one of these step change problems, it tends to attribute a fraction of the change to the correct explanation; but the remainder of the shift is incorrectly ascribed to measurement errors and hardware state variables that have not changed.

In the case that there is a sudden, large shift to a measurement between two consecutive readings, it is possible to use a technique to improve the algorithm results. Basically the approach consists of three steps:

- Perform the standard, weighted, least-squares analysis of the new data
- Evaluate the residual error from this analysis; if this residual error exceeds some threshold, it is assumed that a sudden shift has occurred
- When a sudden shift is detected, a "fault search" is performed to seek a single explanation for this shift.

The solution residual is determined by substituting X into Equation 6 and computing the associated J . Reference 2 shows that the expected value of this J is " $p/2$," where p is the number of independent measurements. For present purposes, the parameter $2J$ is used to compute a probability (which is determined by treating $2J$ as being a Chi-Square variable with p degrees of freedom). The probability computed may be thought of as the chance of observing a residual as large as (or larger than) $2J$, given the assumption that the model represents the mean of the data and given that the statistical model implied by the M and R matrices is also valid. The smaller the value of J , the higher this computed probability will be. This probability is tested against a threshold (currently 1 percent) to decide whether a fault search is justified. If the computed probability is less than the threshold, the fault search is performed.

The assumption of the fault search is that the shift is caused by a single event which might be a measurement bias shift or a component performance change. A menu of possible causes was developed. This list includes bias shifts in the measurements (including such setting parameters as fan speed or altitude), large changes in any individual component (which might effect the efficiency and/or flow capacity), and special hardware changes which are not routinely sought by the weighted, least-squares algorithm (such as, a bleed flow leak or a nozzle area change). This menu of possible problems can be augmented by other events which are observed in service. For example, a compressor stall would be expected to alter the efficiency and pumping capacity of the compressor and, in addition, might alter other engine components as well. If a repeatable model for a compressor stall can be observed in service, then this can be added to the menu of faults to be tested.

The fault-search algorithm tests each possible hypothesis to evaluate its ability to explain the observed shift. To accomplish this, the appropriate elements of " M " or " R " are modified to greatly enlarge the expected range of the particular fault being studied. For example, if the usual turbine efficiency variation is expected to be within a 1 percent range, the fault search will be performed with a possible turbine efficiency variation of 100 percent. The weighted, least-squares analysis is repeated with the modified M or R and the associated J , and the

corresponding probability is computed (with the number of degrees of freedom reduced by the number of free parameters associated with the particular fault being studied). In those instances where the fault represents a phenomenon which is not normally included in the standard analysis, the M and H matrices are augmented to include the new phenomenon.

Each possible fault in the menu of possible faults is tried, and the particular fault which yields the highest solution probability (based on J) is taken to be the prospective explanation to the problem. The philosophy was adopted that this best fault solution must provide a substantial improvement to the probability in order to be accepted as a valid explanation for the observed events (the current threshold for believing the fault explanation is a probability of 25 percent). If the best fault solution fails to achieve this threshold, the no fault solution is taken to be the best explanation of the data.

The fault search is limited to single faults. This is consistent with the assumption of a single failure within a short time span. It is also required to avoid excessive computer charges since the fault search uses significant computer resources. There are, however, situations where more than one measurement fails completely between readings. To cope with this situation, measurements are reviewed prior to the weighted, least-squares analysis. This review consists of a comparison of the measured value to the model-predicted value. If the absolute difference is larger than a threshold (based upon the standard deviation of the measurement error), the measurement is assumed to have failed. When a measurement is determined to be failed, the weighted, least-squares analysis is altered to ignore the measurement by:

- Setting the input measurement deviation for the failed measurement to zero
- Increasing the standard deviation of the failed measurement by a factor of 200.0
- Reducing the number of degrees of freedom for the probability calculation by one.

This has the effect of allowing the measurement error to "float" to a value which essentially makes it consistent with the other measurements. To protect the fault-search logic from reaching ridiculous conclusions, it is disabled from testing for certain types of faults when there are a large number of failed measurements.

In order to provide proper meaning to the probability calculation, as described above, it is necessary to maintain a continual estimate of the state of the engine and of the current levels of the measurement errors. This running status represents an a priori estimate of the engine state and the measurement error for each new reading, based upon prior results. In the weighted least-squares analysis, the new measurements are compared to this a priori estimate, and the difference forms the basis for the weighted, least-squares calculation (and the fault search, if appropriate). To understand the need for this moving frame of reference, consider an engine that gradually deteriorates over a long period of time. As deterioration continues, the residual error, relative to the initial performance of the engine, will gradually increase until eventually the probability is below the fault-search threshold. At this point, a fault search will be performed; the results would reflect the change from the initial engine state, rather

than from the most recent reading. Hence, the assumption upon which the fault search is based would be violated; that is, that the fault is likely to be attributable to a single cause. Further, the fault search would presumably be repeated on every subsequent reading, since the drift away from the initial state would continue.

To bookkeep the a priori estimate, the on-wing TEMPER (Turbine Engine Module Performance Estimation Routine) maintains a status estimate for each state variable and each measurement error. These status estimates may or may not be updated following each reading, based upon the following rules.

1. If the solution probability for a new reading is less than a preselected threshold (5 percent), the status estimate is not changed. The assumption here is that the data are not sufficiently well understood to justify the upgrade to the baseline.
2. When either a measurement error or a hardware state variable is detected to be a fault, the associated status estimate is adjusted to the value of the current reading (that is, a full adjustment is made). Note, that when the fault search selects a fault parameter that is not normally included in the weighted, least-squares analysis (such as, fan speed measurement error, or compressor discharge bleed flow) there is no status estimate associated with the parameter to update.
3. Status estimates other than those associated with a fault are upgraded using an exponential smoothing algorithm (Reference 3). The first 10 readings are used to establish a base value for the exponential smoothing algorithm; these readings are then reprocessed starting from this base value. The exponential smoothing coefficient used is 0.2. The purpose of the exponential smoothing is to reduce the impact of point-to-point measurement noise.

Another feature employed in on-wing TEMPER is a technique to properly handle the reintroduction of a measurement which has been failed for a period of time. Associated with each of the status estimates is a counter which represents the number of readings used to obtain the current value of the status estimate. When a measurement is determined to have failed, the counter is reset to zero. Statistical inputs to the weighted least-squares algorithm are scaled using this counter, together with a test scalar. Thus, when a failed measurement is reintroduced, its standard deviation is scaled by a factor of 200 (this choice is arbitrary, since the correct value for zero prior readings is infinity). On subsequent readings, the scalar decreases so that the measurement gradually achieves an influence comparable to the other measurements.

An important philosophy in the design of on-wing TEMPER is that as much information as possible be applied to determine the correct solution. Often, the desire for elegance is used as a justification for ignoring certain types of information. However, in the on-wing analysis arena, all available information should be applied in order to obtain the best possible solution. In keeping with this approach, on-wing TEMPER has been designed to utilize the maintenance information which is routinely recorded by the airlines. For instance, if the ITT harness is replaced, it is reasonable to expect that there may be a subsequent shift in the ITT measurement bias. The various maintenance actions are encoded in a form that is

recognizable to the on-wing TEMPER algorithm, and a fault-logic type of calculation is performed to assess the improvement to the solution which will be realized by assuming a shift to the particular parameter indicated. For the above-described ITT harness replacement, the ITT bias error is tested by the fault logic to see if an improved solution results. Unless there is a significant improvement to the solution, the maintenance information is ignored. (The ITT harness is often replaced as the first attempt to correct a high ITT problem; however, the harness may or may not be the cause. If it is not the cause, harness replacement is not likely to affect the ITT bias error. In this case, the extra latitude afforded to the fault logic is not appropriate.)

Figure C-1 is a block diagram of the on-wing TEMPER algorithm. Those on-board measurements which are needed to define engine environment and power level (such as, altitude, Mach number, total air temperature, fan speed, etc.) are used to run the engine model which predicts values for the remainder of the on-board measurements (fuel flow, ITT, core speed, etc.). Note that the model is modified through the application of in-service baselines to reflect the actual behavior observed in service. These and the actual measurements are used to compute the measurement deviations from the model (in percent). The baselines represent the sum total of three discrete phenomena.

1. There is a difference between a typical engine in the fleet when initially introduced to service and the engine model. This difference includes bias errors in the measurements and inconsistencies between the model quality and the initial on-wing engine performance. This portion of the baseline is obtained through analysis of the initial revenue service data for a number of engines (by means of comparisons with the engine model).
2. There is also typically a difference between any particular engine when it is newly installed on wing and the average of newly installed engines. The calculation of this difference is accomplished in on-wing TEMPER by means of the averaging of the initial 10 readings to establish the start-up levels for the a priori estimates (for the exponential smoothing).
3. A third difference is the result of the particular engine's deterioration, subsequent to its installation on wing. The calculation of this deterioration (and the corresponding drift in the measurement biases) has been described above; the deterioration is assumed to be independent of all setting parameters (such as, altitude, fan speed, etc.).

The difference between the measurements and the a priori estimate (from the model as modified by the baselines) forms the input to the weighted, least-squares calculation.

The filter solution is computed using a weighted, least-squares algorithm as described above. Statistical inputs for this computation are determined using a three-step process.

1. On-wing data are reviewed to determine standard deviations of the data during regions where there is apparently no deterioration. This short-term scatter is assumed to be indicative of the measurement error and is used to obtain a first approximation for the measurement error matrix, R .

2. The first estimate for the state variation matrix, M , is deduced by using the "R" matrix from Step 1, and iterating the M to achieve desired Kalman-gain response characteristics. As a first approximation, it is desirable to make all state changes and measurement errors equally detectable. However, as a refinement to this goal, those state changes which are judged to be unlikely in service (such as, LP turbine flow function change) are assigned a reduced response relative to other state changes (such as, HP turbine efficiency) that are considered more likely to change. No relative juggling of measurement error responses is performed, since the R matrix is assumed to be correctly indicated (on a relative basis) from the on-board data as described above.
3. A data sample is processed through the on-wing TEMPER algorithm. Resulting residuals are plotted in histogram form and compared to the Chi-Square distribution for the appropriate number of degrees of freedom. A scalar adjustment to both M and R is defined to provide an approximate fit between the observed frequency distribution of residuals and that predicted by the Chi-Square distribution. This scaling does not alter the response characteristics developed in Step 2, but only serves to limit the use of the fault search to appropriate situations.

If there is encoded maintenance input, it is utilized to provide one or more extra solutions to determine whether the indicated maintenance has altered an associated state variable or measurement error. The solution probabilities are computed for all solutions, and if they are sufficiently high, the state of the system is assumed to be known. If the solution probability is lower than the desired threshold, a fault search is performed in an effort to identify a single fault explanation for the deviant behavior.

Once the final solution has been obtained, initial baseline data are used to translate it to an absolute basis. (The internal solution is computed relative to a baseline which is a moving frame of reference; the output solution is quoted relative to the typical newly installed engine, thus providing a fixed frame of reference.) Then if the solution probability is sufficiently high, the a priori estimates are upgraded to reflect the experience gained from the latest data.

C.5 References

1. Gelb, A., et al., "Applied Optimal Estimation," The MIT (Massachusetts Institute of Technology) Press, 1974.
2. Bryson, A.E. Jr., and Ho, Y.C., "Applied Optimal Control," Hemisphere Publishing Corporation, 1975.
3. Bowerman, B.L., and O'Connell, R.T., "Time Series and Forecasting," Duxbury Press, 1979.

Nomenclature

CAS	Close Air Support
CEMS IV	Comprehensive Engine Management System, Increment IV
CND's	Can Not Duplicate's
EPU	Electronic Processing Unit
JRS	Jet-X Retrieval System
OCM	On-Condition Maintenance
RCM	Reliability-Centered Maintenance
RTOK's	Retest OK's
TEMS	Turbine Engine Monitoring System
TEMPER	Turbine Engine Module Performance Estimation Routine
T.O.	Technical Order
UDU	Umbilical Display Unit